Evolutionary Multiobjective Optimization for Dynamic Hospital Resource Management

Anke K. Hutzschenreuter[1,2], Peter A.N. Bosman[2], and Han La Poutré[1,2]

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

[2] Center for Mathematics and Computer Science, Amsterdam, The Netherlands
Peter.Bosman@cwi.nl, Han.La.Poutre@cwi.nl

Abstract. Allocating resources to hospital units is a major managerial issue as the relationship between resources, utilization and patient flow of different patient groups is complex. Furthermore, the problem is dynamic as patient arrival and treatment processes are stochastic. In this paper we present a strategy optimization approach where the parameters of different strategies are optimized using a multiobjective EDA. The strategies were designed such that they enable dynamic resource allocation with an offline EDA. Also, the solutions are understandable to health care professionals. We show that these techniques can be applied to this real-world problem. The results are compared to allocation strategies used in hospital practice.

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. In order 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.

Hospital resource management is concerned with the efficient and effective deployment of resources, i.e. operating rooms and beds, when and where they are needed. In many hospitals, this is a major managerial issue, especially due to the complex relationship between resources, utilization and patient throughput for different patient groups[1]. Moreover, the problem is stochastic as resource usage at a hospital unit behaves like a stochastic process. Emergency patients arrive in urgent need for care, complications require patient transfers and the patients' length of stays are stochastic. Different patient treatment processes need to be considered that typically involve several hospital units. Often, resources (e.g. at the Intensive Care unit) are shared by multiple treatment processes. Thus, hospital resource management is a complex and highly dynamic problem.

M. Ehrgott et al. (Eds.): EMO 2009, LNCS 5467, pp. 320–334, 2009.

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. Previous work [2] showed that a trade-off is needed between these conflicting objectives.

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, the simulation tool described in [2] is important to be used for the evaluation of a resource allocation. Moreover, 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 have been shown to be very powerful for multi-objective (MO) optimization [3,4,5].

For optimizing hospital resource management, we apply strategy (or policy) optimization, as advocated in [6]. Policies are parameterized functions that return an allocation decision for any given situation. The strategies' parameters are optimized using the EA. The advantage of using policies to solve stochastic dynamic optimization problems is that only one strategy has to be optimized that can be applied to a set of scenarios in the simulation. In cooperation with domain experts from the Catharina Hospital Eindhoven (CHE), the Netherlands, we designed strategies that enable dynamic resource allocations. The strategies can be easily understood by health care professionals which is important for the implementation and understanding in practice.

Thus, hospital resource management is a complex and dynamic problem that requires state-of-the-art techniques from dynamic MO research. Specifically, we combine strategy optimization with the SDR-AVS-$\mathbb{MIDEA}$ algorithm [7], an Estimation-of-Distribution (EDA) algorithm. The algorithm uses mixture distributions to stimulate the search for a broad Pareto-front and additionally contains techniques to prevent premature convergence (SDR-AVS). We demonstrate the applicability of these techniques to a real-world problem and their effectiveness.

Only few papers have addressed dynamic MO optimization, especially in stochastic environments. The approach presented in [5] is developed for seldom random changes of the environment and requires optimization from scratch if a change in the environment is detected. Our approach uses strategy optimization and therefore does not need to be re-optimized for each situation. Moreover, it can handle also frequent changes of the environment because the strategies describe what to do in any situation. In [4] the performance of the Non-dominated Sorting Genetic Algorithm version 2 (NSGA2) is evaluated for artificial objective functions. In our work, we use objective functions for a real-world application.

Work on hospital resource management can be found in the Operations Research and Operations Management literature. The models mainly focus on aggregated resource allocation policies, e.g. [1,8], or allocation policies for single units, e.g. [9,10]. Our approach allows for an in-depth analysis of allocation strategies also on the level of different hospital units. Furthermore, their work solely addresses static allocations whereas we consider also the optimization of dynamic resource allocation. The work in [11] provides theoretical results for hospital bed utilization. Our approach is more flexible and can easily be adopted to other settings. Moreover, earlier work considered hospital resource management as a single-objective optimization problem. In [12], the MO optimization problem is addressed. The model, however, is restricted to deterministic patient treatment processes. In our approach, we incorporate stochastic treatment processes that can be flexibly adjusted to other settings.

The remainder is organized as follows. First, we provide a model of the hospital domain and a description of the resource allocation problem in Section 2. Next, the allocation strategies and the algorithm used in our approach are presented in sections 3 and 4. The experiments are reported in Section 5. Finally, in Section 6 we provide our conclusions and an outlook on future work.

2 Simulation Model and Optimization Problem

2.1 Simulation Model

The simulation tool for hospital resource allocation is based on a case study at the CHE. The following features are included: patient characteristics influencing the patients' priority and pathway in the hospital and uncertainty related to the pathways. The model is described below. For a more detailed description, the reader is referred to [2].

Hospital care units. In general, a hospital can be divided into several, medically specialized, care units [13]. The units like nursing wards provide treatment and monitoring and are typically dedicated to a medical specialty such as cardiothoracic surgery (CTS). The operating room (OR) is typically shared by different specialties which are assigned time slots for performing surgical procedures (indicated by a prefix). The intensive care unit (ICU) is often divided into several subunits providing patient care and monitoring with different intensity. We distinguish between intensive care (IC), high care (HC) and medium care (MC). The post anesthesia care unit (PACU) is dedicated to patients recovering from anesthesia. The set of care units relevant for the simulation model is denoted by U with U={CTS-OR, IC, IC-HC, MC, CTS-HC, CTS-PACU, CTS ward, o}. o denotes the possible destinations of a patient's discharge from the hospital which are home or other care facilities, but also mortality.

Patient pathways and scheduling. We distinguish between scheduled patients (i.e. elective surgical patients from the waiting lists) and non-scheduled patients (i.e. emergency patients in urgent need for surgical and/or intensive

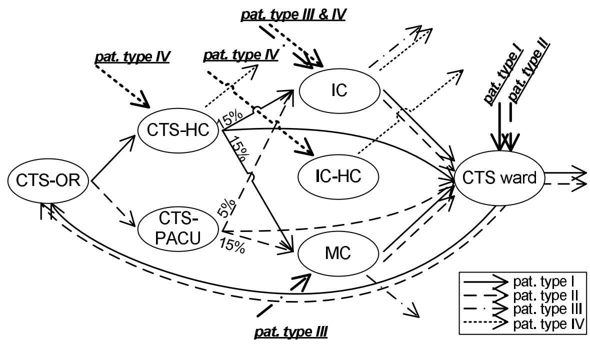

Fig. 1. Interference of CTS, other surgical and emergency patient pathways[1]

care). The set of patient groups is denoted by Θ. The Poisson arrival rate of patients is given as $\lambda_g, g \in \Theta$. We define a patient pathway of group $g \in \Theta$ as the sequence of actually required treatment operations and the respective length of stay (LoS). The patient process represents all possible pathways of patient type $g \in \Theta$ and is modeled by a probabilistic graph [14], $G^g = (V^g, A^g, P^g)$, where the set of nodes, $V^g \subset U$, represents the involved hospital units and the set of arcs, A^g, represents the possible adjacent treatment operations. The length of stay of a patient of group $g \in \Theta$ at hospital unit $u \in V^g$ is modeled as a random variable, LoS_u^g, that follows a probability distribution $P^{LoS_u^g}$. P^g is the set of conditional probability distributions defined on A^g with

$$P^g = \{Pr(v|u, g, t)|u \in V^g, (u, v) \in A^g, t \geq 0\} \text{ for } g \in \Theta. \tag{1}$$

$Pr(v|u, g, t)$ represents the probability that care provided at unit v is required given that a patient of group g has been admitted to unit u for t time units.

Resources are required to perform treatment operations at a hospital unit. Here, relevant resources are ORs and hospital beds. Often, hospital units operate autonomously which means that schedules and resources are managed locally by the units each applying their own (medical) priorities and preferences. The simulation model reflects the distributed decision making by representing each relevant unit by an agent. The policies for scheduling patient admissions and transfers implemented in the agent-based simulation tool were derived from the CHE case study. A detailed description is given in [2].

For the simulation we use four types of patient pathways (type I to IV) that were identified in the CHE case study for the CTS. Type I and II patients are CTS patients, for whom the immediate postoperative care is a priori indicated as CTS-HC and CTS-PACU, respectively. The type III pathway corresponds to the treatment process of emergency patients who arrive unexpectedly. The type IV patient path represents the inflow of other surgical patients in the system.

Figure 1 shows the four types of patient pathways[1]. The routing probabilities are indicated on the arrows where differing from 100%. Type III and IV pathways are restricted to their possible interference with type I and II patients at IC, IC-HC, CTS-HC and MC. The preceding and successive treatment steps are not considered because other dedicated resources are used.

2.2 Optimization Problem

In the following we consider a time horizon T with discrete time units t and n equidistant decision moments denoted by $t_i \in T' \subset T$ with $t_{i-1} < t_i$ for $i = 1, \dots, n-1$. Typically, t would be in steps of hours and t_i would be in steps of days.

Decision variables & parameters. In the simulation model described above, we consider the number of allocated resources as free decision variables (i.e. control variables that impact the performance of the system). Formally, an allocation policy, $\pi(t_i) = (\pi_u(t_i), u \in U)$, determines the number of resources, $r_u(t)$, allocated to hospital unit u at time $t \in T$. Thus, we have that

$$r_u(t) = \pi_u(t_i) \; \forall u \in U, \forall t \in T : t \in [t_i, t_{i+1}), \; t_i \in T'. \tag{2}$$

The model parameters (i.e. the variables whose values characterize the problem instance) are listed below:

P^g**:** the conditional routing probability distribution of patient group $g \in \Theta$
$P^{LoS_u^g}$**:** the length of stay probability distribution for type $g \in \Theta$ at unit $u \in U$
λ_g**:** the (daily) arrival rate of patients of type $g \in \Theta$
r_u^{min}, r_u^{max}**:** the lower and upper bound for the resource capacity allocated to unit $u \in U$; the values are imposed by the layout of a hospital unit, the available equipment, staff and funds
c_u**:** the cost[2] for a resource at hospital unit $u \in U$; specifically, costs for the OR are only accounted for if allocated OR capacity remains unused due to cancelations of surgeries resulting from unavailable postoperative care beds[3].

Performance evaluation of resource allocations. In order to optimize resource allocation in hospitals a trade-off is needed between conflicting objectives, i.e. high patient throughput at low resource costs and back-up capacity usage. The outcome resulting from running the simulation applying allocation policy π is denoted by $F = F(\pi) = (F_0(\pi), F_1(\pi), F_2(\pi))$ with

[1] The actual patient routing may deviate from the medical indication depending on the available beds at the respective hospital care units. Patients may only be transferred to a higher care level than indicated. This gives more routing possibilities and thus makes the patient flows in Figure 1 more complex. The procedure is described in detail in [2].

[2] Costs for hospital resources relate to the daily costs for staff and materials and are expressed relative to the costs of a nursing ward bed.

[3] We assume that all fixed and variable costs for an operating room are covered by the surgical procedure that is to be performed. Therefore, only unused OR capacity is accounted for in the resource costs for the OR.

$F_0(\pi)$: the mean total throughput of patients under allocation π, defined as the number of patients discharged from the hospital after treatment.

$F_1(\pi)$: the mean total resource costs given by

$$F_1(\pi) = \sum_{u \in U \setminus \{CTS-OR\}} \sum_{t_i \in T'} c_u \cdot \pi_u(t_i) + c_{CTS-OR} \cdot uc_{CTS-OR}(\pi),$$

where $uc_{CTS-OR}(\pi)$ denotes the unused CTS-OR capacity due to canceled surgeries resulting from unavailable postoperative care beds given π.

$F_2(\pi)$: the mean total weighted back-up capacity usage under allocation π. The weighting factors correspond to the cost weights $c_u, u \in U$.

For optimizing resource management, $F_0(\pi)$ has to be maximized, while $F_1(\pi)$ and $F_2(\pi)$ have to be minimized. In the following we use as objective function

$$F'(\pi) = (-F_0(\pi), F_1(\pi), F_2(\pi)).$$

Optimization problem. The MO problem can thus be formulated as

$$\min F'(\pi) \tag{3}$$

where

$$\forall u \in U\, \forall t \in T : r_u(t) \in \mathbb{N} \cap [r_u^{min}, r_u^{max}]. \tag{4}$$

3 Strategies for Hospital Resource Management

As described in Section 2, hospital resource management is a highly stochastic and dynamic problem. In our approach, we use strategy optimization as advocated in [6]. Strategies are parameterized functions that return an allocation decision given the current situation. We thus have to optimize only one strategy that can be applied to a set of scenarios in the simulation because it describes what to do in any given situation. The strategies described below were developed in cooperation with domain experts from CHE. Therefore, the strategies can be easily understood by health care professionals which is important for the implementation and understanding under practical conditions.

In the following, the allocation strategies used in this study are described. Moreover, a mechanism for exchanging resources among units is described that enables the implementation of dynamic resource allocation in practice.

3.1 Static Resource Allocation Policies

Static allocation policies allocate a fixed number of resources to the different hospital units. We consider day-constant allocation policies, denoted by $\pi_u(t_i), u \in U$, given by $\pi_u(t_i) \equiv \mathfrak{r}_u \in \mathbb{N}$ for $t_i \in T'$. Day-constant policies are typically employed by hospitals and are also current practice at the CHE.

3.2 Dynamic Resource Allocation Strategy

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 strategies 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. Below, the state representation, the policy and its usage for dynamic resource allocation are described.

State description. The state at unit u at decision moment t_i, $s_u(t_i)$, is determined by the resource utilization rate at u, i.e. the ratio between the utilized capacity[4] at the start of day t_i and the resource capacity, $r_u(t_i^-)$, just before the adjustment at t_i, denoted by t_i^-. Formally, we have $s_u : T' \rightarrow \mathbb{R}_0^+$, $u \in U$, with

$$s_u(t_i) = \frac{\text{utilized capacity at unit } u \text{ at start of day } t_i}{r_u(t_i^-)}. \tag{5}$$

At the postoperative care units (CTS-PACU and CTS-HC) resources are available only for a couple of hours during the day. For these units the state at the start of day t_i defined in (5) may not be representative for the resource occupancy during the remainder of t_i, i.e. due to empty beds at the start of the day and canceled surgeries during the day. For these units, the *expected* resource utilization rate is used to determine $s_u(t_i)$. The *expected* utilized capacity for day t_i is calculated as the utilized capacity at time t_i minus the expected patient outflow plus the expected inflow (determined by the surgery scheme in the OR) for day t_i.

State-dependent allocation policy. A state-dependent allocation policy, denoted by $(\pi_u(t_i, s_u), t_i \in T', u \in U)$, is determined by five parameters: a base resource allocation, r_u^{base}, two adjustments, r_u^{decr} and r_u^{incr}, and two utilization thresholds, $\mathcal{UT}_u^{decr}$,$\mathcal{UT}_u^{incr}$ with $\mathcal{UT}_u^{decr} \leq \mathcal{UT}_u^{incr}$. We use an iterative step-function $\pi : T' \times \mathbb{R}_0^+ \rightarrow \mathbb{N}^{|U|}$ given as

$$\pi_u(t_i, s_u) = \begin{cases} \max\{r_u^{min}, r_u(t_i^-) - r_u^{decr}\} & \text{, if } s_u(t_i) < \mathcal{UT}_u^{decr} \\ r_u(t_i^-) & \text{, if } s_u(t_i) \in [\mathcal{UT}_u^{decr}, \mathcal{UT}_u^{incr}] \\ \min\{r_u^{max}, r_u(t_i^-) + r_u^{incr}\} & \text{, otherwise} \end{cases} \tag{6}$$

for $t_1, \ldots, t_{n-1}$ and

$$\pi_u(t_0, s_u) = r_u^{base}, \tag{7}$$

with $\pi_u(t_i, s_u) \in [r_u^{min}, r_u^{max}]$ $\forall t_i \in T', u \in U$. In (6) the current resources allocation, $r_u(t_i^-)$, is decreased by r_u^{decr} if the resource utilization rate is below the threshold $\mathcal{UT}_u^{decr}$. If the utilization rate is above $\mathcal{UT}_u^{incr}$, $r_u(t_i^-)$ is increased

[4] Note that due to the usage of back-up capacity the utilized capacity may exceed the allocated resources, thus s_u may be greater than 1.

by r_u^{incr}. Otherwise, the current allocation remains unchanged. Note that the policy specifies the allocation at the different units independently.

In the simulation the policy is applied at the start of every day after a warming-up period. Warming-up is necessary to avoid early convergence to minimal allocations due to the empty hospital in the start of a simulation run.

For the dynamic resource allocation problem (4) is changed to

$$\forall u \in U : r_u^{base} \in \mathbb{N} \cap [r_u^{min}, r_u^{max}], \tag{8}$$

$$\forall u \in U \, \forall t_i \in T' : s_u(t_i) \in \mathbb{R}_0^+, \tag{9}$$

$$\forall u \in U : r_u^{decr}, r_u^{incr} \in [0, 5] \tag{10}$$

$$\forall u \in U : \mathcal{UT}_u^{decr} \in [0, 1], \mathcal{UT}_u^{incr} \in [\mathcal{UT}_u^{decr}, \mathcal{UT}_u^{decr} + 1]. \tag{11}$$

As large adjustments are not desirable for hospital management, a maximal adjustment of 5 beds was chosen. Based on preliminary runs, a theoretical upper bound of 2 for $\mathcal{UT}_u^{incr}$ appeared to be more than sufficient.

3.3 Bed Exchange Mechanism for Dynamic Resource Allocation

In the state-dependent strategy described in Section 3.2, a large supply and stock of beds is assumed which enables 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 to 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 described in Section 3.2 is often not practically feasible. To enable dynamic resource allocation in hospitals, we propose an exchange mechanism that is based on fixed personnel resources. The resources are exchanged among the hospital units to meet the current local need.

Here, $\pi_u(t_i, s_u)$ denotes the number of resources *required* by unit u at time t_i, determined by (6). The fixed personnel resources are determined by $r_u^{base}, u \in U$. The actual resource allocation, $r_u(t)$, is set by the mechanism below and not by (2).

We classify hospital units into three care levels, level 1 to level 3, based on the intensity of care and monitoring and the skill level of the personnel. Here, level 1 is the intensive care (IC), level 2 comprises the IC-HC, MC, CTS-HC and the CTS-PACU unit. Level 3 is the CTS ward, or shortly referred to as ward. From the application domain three rules arise for feasible bed exchanges:

R1: Due to staff training and physical requirements (i.e. access availability to the isolated electric power system in the hospital), beds can only be exchanged within the same or between adjacent care levels.

R2: Due to the staff assigned to a bed, shifting one bed from level l to level $l+1$ yields two beds at $l+1$ for $l = 1, 2$.

R3: Due to the personnel required to operate a bed, only an even number of beds can be shifted from level l to level $l-1$ for $l = 2, 3$ (i.e. the reverse of **R2**).

For the sake of reproducibility the mechanism and the method `shiftBeds` are described in detail in Algorithm 1. The number of resources available for exchange, $\mathcal{E}_l$, in care level l at time t_i is determined by

$$\mathcal{E}_l = \sum_{u \in \, level \, l} \max\{0, r_u(t_i^-) - \pi_u(t_i, s_u)\}, \; l = 1, 2, 3. \tag{12}$$

First, beds are shifted from level 1 to level 2. Then, level 2 beds are shifted to level 1 if necessary. Subsequently, beds of level 2 are exchanged within level 2. Finally, beds are exchanged between level 2 and 3. All exchanges are performed only if additional resources are required. The order of the care levels is based on the resource costs associated with the different units (given in Section 5.2). In future work, also other orderings of care levels will be considered. The order of units within a care level is determined randomly.

Through the mechanism, the implementation of (6) is extended with the above adjustments at time $t_i \in T'$, depending on the interaction with other units. This complex interaction mechanism answers to reality, however, it further complicates the optimization of resource management. Therefore, a state-of-the-art technique is needed for this optimization, which is described in Section 4.

4 EDA for Multi-objective Optimization

For the optimization of the dynamic and complex multi-objective resource allocation problem, we apply the SDR-AVS-$\mathbb{MIDEA}$ algorithm [7]. The algorithm was shown to be an efficient optimization technique for MO problems [7]. A brief outline of SDR-AVS-$\mathbb{MIDEA}$ is given in Section 4.1.

We use a strategy optimization approach with the policies defined in Section 3. The parameters of the strategies, specified in Section 5.1, are optimized using SDR-AVS-$\mathbb{MIDEA}$. The fitness is determined using the simulation tool described in Section 2.1.

The optimization of the strategies is performed in an offline fashion. As the strategies are used online in the simulation, the anticipation of time-dependency effects [6], i.e. the impact of decisions taken now on the future, is implicitly included in the optimization of the strategies' parameters. Thus, MO techniques can be applied in a straightforward fashion to solve this dynamic problem. Since designing online MO appears to be rather hard, this approach yields an additional advantage. The policies proposed can be easily understood by health care professionals, so this approach is also practically implementable.

4.1 Outline of SDR-AVS-$\mathbb{MIDEA}$

In this section, a brief outline of the evolutionary algorithm is given. For a detailed description the reader is referred to [7].

The algorithm divides the generated solutions into clusters of equal size that are kept separated in the objective space throughout a run. The use of clusters stimulates the search for a broad Pareto-front. New solutions are generated

Algorithm 1. Pseudo-code description of the bed exchange mechanism

Input: Set of hospital units, U, $\pi_u(t_i, s_u)\, \forall u \in U$, determined by (6), and $\mathcal{E}_l$, $l = 1, 2, 3$, determined by (12)
Result: $r_u(t_i)\, \forall u \in U$

for $l = 1$ **to** *3* **do** $EX_l \leftarrow \mathcal{E}_l$;
`/*exchange from level 1 to level 2 applying rules R1 and R2; if the exchange results in more beds at unit` u `than required, a bed is shifted to another unit of level 2 in need of additional resource capacity */`
while $EX_1 > 0$ *and* $\exists u \in level\,2$ *with* $\pi_u(t_i, s_u) > r_u(t_i^-)$ **do**
  if $(\pi_u(t_i, s_u) - r_u(t_i^-))$ *is an even number* **then**
    `shiftBeds(`*1,2,IC,u*,$\min\{EX_1, (\pi_u(t_i, s_u) - r_u(t_i^-))/2\}$`)`
  else
    `shiftBeds(`*1,2,IC,u*,$\min\{EX_1, \lceil(\pi_u(t_i, s_u) - r_u(t_i^-))/2\rceil\}$`);`
    Find a $v \in level\,2, v \neq u$ (if any) with $\pi_v(t_i, s_v) > r_v(t_i^-)$ and
    `shiftBeds(`*2,2,u,v,1*`);`
`/*exchange from level 2 to level 1 applying rules R1 and R3; if an exchange from one unit is not feasible (R3), the exchange is performed together with another unit of level 2 (if possible) */`
while $EX_2 \geq 2$ *and* $\pi_{IC}(t_i, s_{IC}) > r_{IC}(t_i^-)$ **do**
  Find $v_1 \in U$ with $\pi_{v_1}(t_i, s_{v_j}) < r_{v_1}(t_i^-)$;
  if $|\pi_u(t_i, s_u) - r_u(t_i^-)|$ *is an even number* **then**
    `shiftBeds(`*2,1,v_1,IC*,$\min\{|\pi_{v_1}(t_i, s_{v_1}) - r_{v_1}(t_i^-)|/2, \pi_{IC}(t_i, s_{IC}) - r_{IC}(t_i^-)\}$`);`
  else
    Find $v_2 \in U$, $v_2 \neq v_1$, with $\pi_{v_2}(t_i, s_{v_2}) < r_{v_2}(t_i^-)$ and
    `shiftBeds(`*2,2,v_2,v_1*,1`)` and `shiftBeds(`*2,1,v_1,IC*,$\min\{\lceil|\pi_{v_1}(t_i, s_{v_1}) - r_{v_1}(t_i^-)|/2\rceil, \pi_{IC}(t_i, s_{IC}) - r_{IC}(t_i^-)\}$`);`
    if *there is no such* v_2 **then**
      `shiftBeds(`*2,1,v_1,IC*,$\min\{\lfloor|\pi_{v_1}(t_i, s_{v_1}) - r_{v_1}(t_i^-)|/2\rfloor, \pi_{IC}(t_i, s_{IC}) - r_{IC}(t_i^-)\}$`);`
`/*exchange within level 2 applying rule R1 */`
while $EX_2 > 0$ *and* $\exists u \in level\,2$ *with* $\pi_u(t_i, s_u) > r_u(t_i^-)$ **do**
  Find a $v \in level\,2$, $v \neq u$, with $\pi_v(t_i, s_v) < r_v(t_i^-)$ and
  `shiftBeds(`*2,2,v,u*,$\min\{|\pi_v(t_i, s_v) - r_v(t_i^-)|, \pi_u(t_i, s_u) - r_u(t_i^-)\}$`);`
`/*exchange from level 2 to level 3 applying rules R1 and R2 */`
if $EX_2 > 0$ *and* $\pi_{ward}(t_i, s_{ward}) > r_{ward}(t_i^-)$ **then**
  Find a $v \in level\,2$, with $\pi_v(t_i, s_v) < r_v(t_i^-)$ and
  `shiftBeds(`*2,3,v,CTS ward*,$\min\{|\pi_v(t_i, s_v) - r_v(t_i^-)|, \lfloor(\pi_{ward}(t_i, s_{ward}) - r_{ward}(t_i^-))/2\rfloor\}$`);`
`/*exchange from level 3 to level 2 applying rules R1 and R3 */`
while $EX_3 \geq 2$ *and* $\exists u \in level\,2$ *with* $\pi_u(t_i, s_u) > r_u(t_i^-)$ **do**
  `shiftBeds(`*3,2,CTS ward,u*,$\min\{\lfloor|\pi_{ward}(t_i, s_{ward}) - r_{ward}(t_i^-)|/2\rfloor, \pi_u(t_i, s_u) - r_u(t_i^-)\}$`);`
`/*if no exchange is possible, the resource allocation remains unchanged */`
forall $u \in U$ *that were not yet considered* **do** $r_u(t_i) \leftarrow r_u(t_i^-)$;

Function. `shiftBeds(`*int* l^{from}, *int* l^{to}, *unit* u^{from}, *unit* u^{to}, *int* n`)`

$k = 1$;
if $l^{from} < l^{to}$ **then** $k = 2$;
else if $l^{from} > l^{to}$ **then** $k = 0.5$;
$r_{u^{to}}(t_i) \leftarrow r_{u^{to}}(t_i) + k \cdot n$, $r_{u^{from}}(t_i) \leftarrow r_{u^{from}}(t_i) - n$, $EX_{l^{from}} \leftarrow EX_{l^{from}} - n$;

according to the EDA principle. In each separate cluster a single normal distribution is used. The algorithm uses adaptive variance scaling (AVS) in combination with standard-deviation ratio (SDR) triggers to prevent premature convergence. This means that if the best fitness in a cluster is improved in one generation and the average improvement is more than one standard deviation away from the estimated mean of the distribution, then the variance of the estimated distribution is scaled up to increase the area of exploration. If, however, the improvements are obtained near the mean of the predicted distribution, then the variance is scaled down to allow for a faster convergence.

5 Experiments

5.1 Settings of SDR-AVS-$\mathbb{MIDEA}$

The settings of the parameters in SDR-AVS-$\mathbb{MIDEA}$ are based on the guidelines reported in [7,15] with the percentile for truncation selection set to 0.3, $k = 4$ clusters. The guideline in [15] is used and results in a population size per cluster of 49 and 130 for the day-constant and the dynamic policies, respectively. The variance multiplier decreaser of AVS equals 0.9 and the SDR threshold is set to 1.0. As in [7], an elitist archive is maintained. To this end, the objective space is discretized in each objective with a discretization length of 10^{-3}. This provides sufficient granularity for the final Pareto-front approximations. We allowed 1600 generations for the different allocation policies.

In the EDA representation, the genes correspond to allocation policy parameters. For the day-constant policies described in Section 3.1, the genotype comprises the values for $\mathfrak{r}_u$, $u \in U$, with $\mathfrak{r}_u \in \mathbb{N} \cap [r_u^{min}, r_u^{max}]$. For the dynamic policies described in Section 3.2 and Section 3.3, a genotype is composed of values $r_u^{base} \in [r_u^{min}, r_u^{max}]$, $r_u^{decr}, r_u^{incr} \in [0, 5]$, $\mathcal{UT}_u^{decr} \in [0, 1]$, and $\mathcal{T}_u \in [0, 1]$ for $u \in U$. The parameter $\mathcal{T}_u$ is used to determine $\mathcal{UT}_u^{incr}$ by $\mathcal{UT}_u^{incr} = \mathcal{UT}_u^{decr} + \mathcal{T}_u$. The bounds, r_u^{min} and r_u^{max}, for the resource allocations were obtained from domain experts from CHE. These values are given in Table 1.

5.2 Settings of the Simulation Tool

Applying SDR-AVS-$\mathbb{MIDEA}$ to a real-world problem is associated with a large number of potential solutions using a complex simulation model. We run 10 simulation runs of 20 weeks including 8 weeks of warming-up to evaluate the allocation strategies. The warming-up period is not measured in the simulation outcomes. This setting results in a runtime of about 6 seconds per evaluation.

Table 1. Resource bounds, unit resource costs and benchmark day-constant policies obtained from CHE case study

	CTS-OR	CTS-HC	CTS-PACU	IC	IC-HC	MC	CTS ward
r_u^{min}	0	0	0	5	2	2	20
r_u^{max}	6	6	6	20	6	10	50
c_u	0.09	2	2	4	2	2	1
π_u^{CHE}	4	4	4	11	4	4	35

Table 2. Input parameters of patient pathways with LoS in hours (mean±stdev)

Type	Unit	LoS	P^g	Type	Unit	LoS	P^g
I	CTS-HC	15 ± 0	-	II	CTS-PACU	6 ± 0	-
	IC	48.48 ± 54	0.15		IC	42 ± 57.12	0.05
	MC	24.48 ± 38.52	0.15		MC	10.32 ± 22.08	0.15
	CTS ward	120 ± 22.08	0.7		CTS ward	120 ± 22.08	0.8
III	IC	89.48 ± 200.82	-	IV	IC-HC	34.94 ± 68.51	-

In a sensitivity analysis, the mean and variance of the relevant outcome values appear to be linear for increasing simulation run durations.

The settings for the patient pathways are based on the statistical data analysis conducted in the case study at the CHE. The relevant parameters of the different patient pathways introduced in Section 2.1 are given in Table 2. We use a Lognormal distribution for sampling patients' LoS. Arrivals of type III patients are Poisson with daily arrival rate $\lambda^{III} = 2$. Patients of type IV arrive daily in bulks between 2 and 4 patients. Costs for the different types of hospital resources related to the daily costs for staff and materials and are expressed in terms of relative costs of a nursing ward bed. The costs are given in Table 1. The OR costs account for the unused OR capacity due to cancelations of surgeries that result from unavailable postoperative care beds.

5.3 Results

One run of the EA takes approximately 10 hours for static strategy optimization and 30 hours for dynamic strategy optimization on a high-performance computer cluster. Specifically, we used 40 nodes running at speeds between 1.4Ghz and 2.2Ghz. We have run the EA for each strategy three times, yielding very comparable and stable results (within a strategy type).

In Figure 2 the results for the day-constant, the state-dependent strategies and the exchange mechanism are presented. In the application domain, the exact values for back-up capacity usage are of minor importance and a categorization of minimal (corresponding to $F_2 \in [0, 50)$), small ($F_2 \in [50, 100)$), medium, etc. is therefore sufficient for the representation of the optimization results. The results

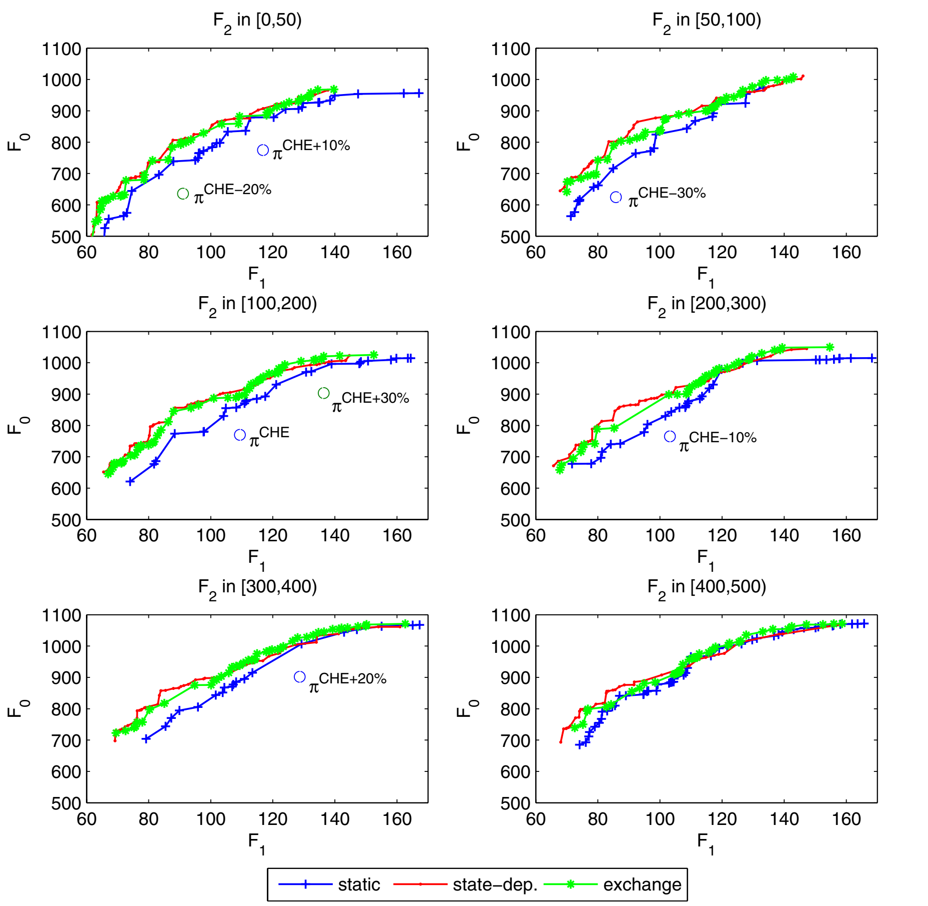

Fig. 2. Pareto-fronts for static and dynamic allocation policies including benchmarks from CHE

are confined to F_2-values below 500 as higher back-up capacity usage is not desirable for many hospitals. We depict the Pareto-fronts with respect to F_1 and F_0 for F_2-values in the predefined intervals. This in addition allows us better visibility of the results as opposed to 3D plots. To assess the performance of the policies the currently used day-constant resource allocation at the CHE and linearly scaled allocations are included as a benchmark for the relevant intervals. The currently employed policy of the CHE is denoted by π^{CHE} given in Table 1. Also, benchmarks determined by linearly scaled allocations are considered that are denoted by π^{CHE+i} with $\pi^{CHE+i} = (\lfloor \pi_u^{CHE} \cdot (1+i) + 0.5 \rfloor, u \in U, i = \pm 10\%, \pm 20\%, \pm 30\%)$.

The results in Figure 2 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. For F_2-values of above 300 and F_1-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 F_1 and F_2 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 can be attributed to the interaction between the hospital units due to which required allocation adjustments cannot always be fully undertaken.

6 Conclusions

In this paper, multiobjective optimization for dynamic hospital resource management using evolutionary algorithms was addressed. We use a strategy optimization approach for which we designed policies that allow for the dynamic allocation of resources in hospital practice. Due to the complexity of the allocation strategies and the dynamic application domain, we used a state-of-the-art evolutionary MO technique, SDR-AVS-$\mathbb{MIDEA}$. The fitness of the solutions was determined using a simulation tool developed for this application domain. Our results show that the benchmark allocations obtained from a case study could be considerably improved using the optimized strategies. Furthermore, we showed that policies that incorporate more dynamic resource allocations result in further improvements. These improvements are made possible by the design of the policy. SDR-AVS-$\mathbb{MIDEA}$ then is powerful enough to detect and exploit the additional possibilities. We demonstrated that proper design in combination with state-of-the-art EAs can make an important contribution and achieve an improvement for complex real-world dynamic MO problems as in hospital logistics.

By using our strategy types, we can circumvent the online MO optimization and use offline MO techniques to optimize the parameters of the strategies. Furthermore, our approach is feasible for stochastic fitness functions obtained from a simulation model.

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 SDR-AVS-$\mathbb{MIDEA}$ in relation to the above extensions.

Acknowledgements

We thank Han Noot for his programming contributions, and Jan van Aarle, Ilona Blonk, Floor Haak and Dick Koning, for their expert feedback.

References

1. Harper, P.R., Shahani, A.K.: Modelling for the planning and management of bed capacities in hospitals. J. Oper. Res. Soc. 53, 11–18 (2002)
2. Hutzschenreuter, A.K., Bosman, P.A.N., Blonk-Altena, I., van Aarle, J., La Poutré, J.A.: Agent-based patient admission scheduling in hospitals. In: Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008) - Industry and Applications Track, pp. 45–52 (2008)
3. Bosman, P.A.N., Thierens, D.: Multi–objective optimization with diversity preserving mixture–based iterated density estimation evolutionary algorithms. Int. J. Approx. Reason. 31, 259–289 (2002)
4. Bui, L.T., Abbass, H.A., Branke, J.: Multiobjective optimization for dynamic environments. In: Congress on Evol. Comp., vol. 3, pp. 2349–2356 (2005)
5. Farina, M., Deb, K., Amato, P.: Dynamic multiobjective optimization problems: test cases, approximations, and applications. IEEE T. Evol. Comp. 8, 425–442 (2004)
6. Bosman, P.A.N., La Poutré, J.A.: Learning and anticipation in online dynamic optimization with evolutionary algorithms: the stochastic case. In: GECCO 2007: Proc. of the 9th annual conference on Genetic and evol. comp., pp. 1165–1172 (2007)
7. Bosman, P.A.N., Thierens, D.: Adaptive variance scaling in continuous multi-objective estimation-of-distribution algorithms. In: GECCO 2007: Proc. of the 9th annual conference on Genetic and evol. comp., pp. 500–507 (2007)
8. Vissers, J.M.H.: Patient flow-based allocation of inpatient resources: A case study. E. J. Oper. Res. 105(2), 356–370 (1998)
9. Ridge, J.C., Jones, S.K., Nielsen, M.S., Shahani, A.K.: Capacity planning for intensive care units. E. J. Oper. Res. 105(2), 346–355 (1998)
10. Kim, S., Horowitz, I., Young, K.K., Buckley, T.A.: Flexible bed allocation and performance in the intensive care unit. J. Oper. Manag. 18(4), 427–443 (2000)
11. Kusters, R.J., Groot, P.M.A.: Modelling resource availability in general hospitals design and implementation of a decision support model. E. J. Oper. Res. 88(3), 428–445 (1996)
12. Blake, J.T., Carter, M.W.: A goal programming approach to strategic resource allocation in acute care hospitals. E. J. Oper. Res. 140(3), 541–561 (2002)
13. Vissers, J., Beech, R. (eds.): Health Operations Management: Patient flow logistics in health care. Health Management Series. Routledge (2005)
14. Bollobàs, B.: Graph Theory: An Introductory Course. Springer, Heidelberg (1979)
15. Bosman, P.A.N., Grahl, J., Thierens, D.: Enhancing the performance of maximum-likelihood gaussian edas using anticipated mean shift. In: Rudolph, G., Jansen, T., Lucas, S., Poloni, C., Beume, N. (eds.) PPSN 2008. LNCS, vol. 5199, pp. 133–143. Springer, Heidelberg (2008)