Adaptive Optimization of Hospital Resource Calendars

I.B. Vermeulen[1], S.M. Bohte[1], S.G. Elkhuizen[2], J.S. Lameris[2], P.J.M. Bakker[2], and J.A. La Poutré[1]

[1] Centre for Mathematics and Computer Science (CWI), Amsterdam, the Netherlands, I.B.Vermeulen@cwi.nl,
[2] Academic Medical Center, University of Amsterdam, the Netherlands

Abstract. As demand for health care increases, a high efficiency on limited resources is necessary for affordable high patient service levels. Here, we present an adaptive approach to efficient resource usage by automatic optimization of resource calendars. We describe a precise model based on a case study at the radiology department of the Academic Medical Center Amsterdam (AMC). We model the properties of the different groups of patients, with additional differentiating urgency levels. Based on this model, we develop a detailed simulation that is able to replicate the known scheduling problems. In particular, the simulation shows that due to fluctuations in demand, the allocations in the resource calendar must be flexible in order to make efficient use of the resources. We develop adaptive algorithms to automate iterative adjustments to the resource calendar. To test the effectiveness of our approach, we evaluate the algorithms using the simulation. Our adaptive optimization approach is able to maintain overall target performance levels while the resource is used at high efficiency.

1 Introduction

Hospitals continuously aim to improve their patient-oriented care. They want to provide their patients with high service levels. However, the demand for health care is increasing, and more patients must be treated with the same capacity. High efficiency on resources is necessary for high service levels.

Traditional approaches to logistical improvement are usually not suited to the medical domain. The distributed authority in hospitals [1] makes improvements involving many departments difficult to implement. Furthermore, scheduling decisions must be made depending on the individual patient's specific attributes. Efficient scheduling of patient appointments on expensive resources a complex and dynamic task.

Hospital resources are many: ranging from CT and MRI scanners, to hospital beds, to attending staff. A resource is typically used by several patient groups with different properties [2]. There are groups of inpatients (admitted to the hospital) and outpatients (not admitted), with different levels of urgency [3]. The total hospital resource capacity is allocated to these groups, explicitly or implicitly. Either way, due to fluctuations in demand, this allocation must be flexible to make efficient use of the resources.

To allocate hospital resources, electronic calendar-systems are widely applied. However, they are mostly just storing the patient appointments. The actual scheduling approach largely influences whether resources are efficiently used. Efficient resource usage and short waiting times are of great importance to the hospital departments. Department managers can control options such as changing the capacity, or adjusting the performance goals. It is typically hard to determine the effect of such changes a priori.

In this paper, we study the case of scheduling Computer Tomography scans (CT-scans) at the radiology department within the Academic Medical Centre Amsterdam (AMC). The AMC is a large hospital with approximately one thousand beds, and treating around 350,000 outpatients annually. Currently the radiology department makes more than 15,000 CT-scans per year. Diagnostic resources such as the CT-scanners are literally central in the clinical pathways of many patients. Long access times to such resources are immediately felt as bottlenecks for health care processes in the hospital.

In recent years the logistical process around the CT-scan has improved already substantially [8]. The actual scheduling of appointments is (still) done manually. A calendar supervisor determines a long time in advance how the scanner capacity is allocated to different groups of patients. This allocation is determined based on experience, future expectation, and in cooperation with medical experts.

The calendar supervisor also monitors and adjusts the allocations in the calendar on a regular basis (at least daily) to maintain its efficiency. Often, the actual realization of patient arrival does not match the allocation. This results in inefficient use of the capacity and/or increased waiting time for patients. The calendar supervisor can adjust the calendar to counter such problems. With constant active – and time consuming – supervision of the calendar, the manual scheduling practice performs satisfactory.

However, making good adjustments is critically dependent on the supervisor's expertise: even a short vacation or illness of the calendar supervisor leads to immediate significant deterioration of the resource efficiency. From a planning and sustainability perspective, this is highly unsatisfactory.

Here, we develop an approach to automatically determine effective optimizations of a resource calendar we derive from our case study. Our approach enables the calendar supervisor to quickly implement calendar adjustments, and anticipate – and remedy – the impact of current demand trends of future resource efficiency.

As a model of hospital resource scheduling, we present a precise model for our application case of CT-scan scheduling. Our model and its parameter values are determined from extensive case analysis. Sources include historical data and extensive discussion with experts at various levels in the organization.

We have implemented an extensive computer simulation of the application case. This allows us to study various problem scenarios and scheduling approaches. It can serve as a prototype as the final step towards application.

Furthermore we present our adaptive approach to automatic optimization of resource calendars. In our approach the allocation of capacity to different patient groups is flexible and adaptive to the current and future situation. To maintain high performance levels, our system exchanges capacity between different urgent and non-urgent patients groups. Additionally, resource openings hours can be reduced to increase capacity usage while maintaining high performance levels, or extended to counter increasing waiting time. We extensively evaluate our adaptive approach in our simulated environment,

Other approaches consider how to coordinate patient scheduling, such as [4], and [5]. In this work, human schedulers still make local decisions based on their experiences and knowledge of the individual patient. It supports the current process, making it suitable for fast application. In [6] the author considers updating the allocation of hospital resources to the departments; here, we provide a more operational approach at the level of scheduling patients. In [7] the authors present a first step towards a general model for solving resource conflicts represented as a constraint satisfaction problem, our approach directly focuses on maintaining high performance levels. Our approach can straightforwardly be applied to similar situations where resources are shared between patient groups for which an appointment calendar is used.

2 CT-scan Scheduling Model

In this Section, we define our CT-scan scheduling model. From the AMC electronic calendar system, we have collected the historical data of the appointments made (from October 2005 until March 2006). We have complemented this with data from actual production of CT-scans (November 2005 until January 2006). During this period, some scans were taken without an appointment. We derive patient distributions and scheduling practice from this data, site-visits, and extensive discussions with the human schedulers, the calendar supervisor, and resource manager.

Our model consists of three main parts, discussed in the following sections. One part is the set of arriving **patients** that need to be scheduled for an appointment. Secondly, we have the available resource, and the way the associated appointment **calendar** is structured. The third part is the **scheduling** process that determines how appointments are made by assigning patients to timeslots on the calendar.

2.1 Patients

An important issue is that there is a great variety in patients and scan attributes. We make the abstraction that a patient always needs to be scheduled for exactly one CT-scan. We therefore model the patient and his/her scan as a unity, which we from now on refer to as ’patient’. Patient attributes are listed in Table 1.

In practice, patients and their attributes are structured in different patient groups. Table 2 lists these groups and their specific properties. The group size is given relative to the total number of patients.

The largest group – **out+ivc** – is comprised of non-urgent outpatients who need intravenous contrast (ivc) injected before taking the CT-scan. Non-urgent outpatients who do not need intravenous-contrast are in the group **out–ivc**. All urgent outpatients form the group **urgent**. The fourth group – **clinic** – consists of all inpatients.

Besides these four groups, there are a number of smaller, highly specific groups: $special_n$. These include patients taking part in special programs, and patients who need a specific treatment for making a CT-scan. E.g., one special group is the group of patients, usually children, who need to be sedated while making the CT-scan.

Urgency of patients is defined in terms of planning windows (PLANWIN) with different sizes. Outpatients with normal urgency (out+ivc, and out–ivc) have a planning window of $(2, 14)$, which means that the appointment must be scheduled between 2 days and 14 days after the request for the scan is made. Urgent outpatients and clinic patients have high urgency and have planning windows of a few days. Patients from special groups are always scheduled to the first available timeslot of matching type.

Table 1. Patient attributes

attribute	**description**
request time	date and time when request for CT-scan is made
in-/outpatient	is the patient an in- or outpatient?
contrast needed? ($\pm ivc$)	does intravenous-contrast need to be injected?
planning window ($planwin$)	expresses urgency of patient.
duration	of the needed appointment

Table 2. Patient groups

group	**urgency**	**planwin** (fraction)	**duration**	**size(%)**
out+ivc	normal	$(2, 14)$	15 mins	$52\% \pm 6\%$
out–ivc	normal	$(2, 14)$	15 mins	$23\% \pm 4\%$,
urgent	high	$(0, 1)(33\%), (0, 2)(33\%)$, or $(0, 3)$	15 mins	$10\% \pm 3\%$
clinic	high	$(0, 1)(40\%)$, or $(0, 2)$	30 mins	$6\% \pm 2\%$
$special_n$	n.a.	n.a.	$duration_n$	$9\% \pm 2\%$

2.2 Resource Calendar

Patients must be scheduled to a timeslot on the calendar. The total resource capacity is given by the number of actual CT-scanners m (in our case $m = 2$) and the opening hours. The hospital's emergency room has an additional CT-scanner, which is used as a walk-in facility for emergencies, which we do not consider in our model.

A standard calendar is used, structured in days and weeks. The time on the calendar is partitioned into timeslots of different sizes. All timeslots have a size of a multitude of the unit size us. (In our case us is 15 minutes, and we use timeslots of sizes $1us$ up to $4us$.) The parameters in Table 3 define the resource calendar. The parameters m and us are fixed for long periods of time, the remaining can vary. Openings hours must be known at least one week in advance to plan staff. In general we assume that the m actual resources are interchangeable.

Timeslot Type Specification CT-scan capacity is reserved for different patients groups and these allocations serve medical restrictions (e.g. due to preparation constraints for narcosis), as well as a scheduling goal (e.g. reserve timeslots for urgent patients). E.g., on the actual calendar, three timeslots are reserved on all Thursday mornings for patients from a $special_n$ group, who need to be sedated while making the CT-scan. During lunch time, radiologists schedule meetings and other activities. Therefore, out+ivc patients, who need to be injected with intravenous contrast for which a radiologist must be present, cannot be scheduled during lunch. In the afternoon of every day a number of timeslots is reserved for urgent outpatients.

We model this allocation by using a timeslot-type specification (TTS). A timeslot-type specifies which patient can be scheduled to a certain timeslot (Table 4). The TTS thus determines how much of the resource capacity is allocated to the patient groups. The TTS is not necessarily fixed as the capacity allocation can be dynamically altered.

The $TTspecial_n$ type of timeslots can only be used by very specific types of patients. For each of these types there is a rule which states that if there are still any free slots remaining r_n days in advance, these slots are changed to TTout type of timeslots. This rule is currently the only automatic TTS adjustments in operation at the hospital.

Table 3. Calendar parameters

parameter	description
m	number of resources
$o_{j,d}$	opening time of resource j on day d
$c_{j,d}$	closing time of resource j on day d
us	unit size timeslots
TTS	timeslot type specification

Table 4. Timeslot Types

Timeslot-Type	allowed patients	size
TTout	out+ivc, out–ivc, urgent	$1us$
TT-ivc (during lunch)	out–ivc, urgent(with no ivc)	$1us$
TTurgent	urgent	$1us$
TTclinic	clinic	$2us$
$TTspecial_n$	$special_n$	1–4us

2.3 Scheduling

Scheduling is the process of assigning patients to timeslots, i.e. making appointments. In the case we describe, scheduling performance of different approaches is influenced by two things: first, by how well the TTS matches the actual situation, and second, by the actual scheduling method (the selection of a timeslot per patient given the TTS).

As in many hospitals, for the AMC CT-scanners the actual scheduling of appointments is done manually. Human schedulers schedule patients in turn, by looking on the calendar for an available slot, or using the search function of the electronic calendar system. The search returns a list of available timeslots. Human schedulers have expertise in taking the individual patient's attributes into account. They can also use a patient's preference (e.g. for a specific day, or time). However, the human schedulers have little overview on how their local decisions will influence overall performance goals

The calendar supervisor can adjust the TTS in case there is a mismatch between the TTS and the actual realization of patient arrivals. The human schedulers are used to working with different types of timeslots on the calendar. They use their expertise and long-time experience to select timeslots within the scheduling rules.

3 Simulation

Based on the model we have implemented a patient scheduling simulation (PSS). We use the PSS in the evaluation of different scheduling and resource management approaches. The PSS takes as input distributions of patients attributes, the standard resource openings hours and TTS, a scheduling method, a performance measure, and an adaptive model of how to adjust the TTS and openings hours. The PSS generates a patient stream, a filled-in calendar, the used openings hours and a performance value.

Patient arrival simulation With our model of patient properties and the relative request proportions, we can simulate the arrival of all patients during a week. In the simulation, we have structured the arrival process by the following steps for each week:

1. A standard random walk with a drift τ towards the average $\bar{n}$ fits the distribution over the number of patients arrivals per week. The number of patients for next week (n_{w+1}) is determined as a function of the current patient arrivals n_w as:

$$n_{w+1} = n_w + \mathcal{N}(0, \sigma) + \frac{\bar{n} - n_w}{\tau},$$

where $\mathcal{N}(0, \sigma)$ a normally distributed fluctuation of patient arrivals. We set: $\bar{n} = 250$, $n_0 = \bar{n}$, $\sigma = 30$, and $\tau = 3$.
2. Given n_w divide the patients over the groups, using the distribution from Table 2, where out+ivc will get the remainder of $n_w \approx 52\%$.
3. Per patient determine request date within week (see below).
4. Per patient determine request time on request day, using a uniform distribution over the opening hours.
5. Per patient determine the planning window using the distributions from Table 2.
6. Order patients by increasing request time within week.

Because of extra rounds for inpatients on Monday and Friday, patient arrival (step 3 above) is slightly structured during the week. On Monday and Friday twice as many requests for CT-scans of inpatients are ordered compared to the other three weekdays. Requests for outpatients arrive uniformly over the week. Note that as resource is closed on Saturday and Sunday, urgent and clinic patients requested on a Friday with a PLANWIN of (0,1) or (0,2) must also be scheduled to that Friday.

Resource Calendar and Scheduling Approach In our simulation we use a resource calendar, which is similar to the calendar used in practice. Opening time on the calendar is 8:30, while the resource closes at 16:45. To simulate current scheduling practice we use the following scheduling method:

First Come Randomly Served (FCRS) Patients are scheduled in order of arrival. A patient is assigned to a timeslot within his planning window, randomly selected from all the free timeslots of the allowed types. If there are no free timeslots within the planning window, the first free timeslot after the planning window is selected.

For non-urgent patients FCRS simulates the scheduling process where patient preferences are taken into account. We represent this by random allocation to free slots. Urgent and clinic patients have high urgency and thus patient preferences are of little importance. In current practice however, the human schedulers do not take the individual urgency of these patients into account and are therefore also scheduled randomly. We will present a dynamic approach to the scheduling of urgent and clinic patients in the next Section. Additionally our adaptive model for calendar adjustments is input for the scheduling approach used in the PSS.

Performance Measure Based on discussions with hospital experts, we want our performance measure to express that patients must be scheduled within their planning windows. It is important that each group (G) has a good service level. We define the minimum service level (MSL), over the four main groups of Table 2, as:

$$MSL = \min_G \left(\frac{|\text{patient } \in G = \text{ontime}|}{|G|} \right),$$

where ontime is defined as scheduled within the planning window.

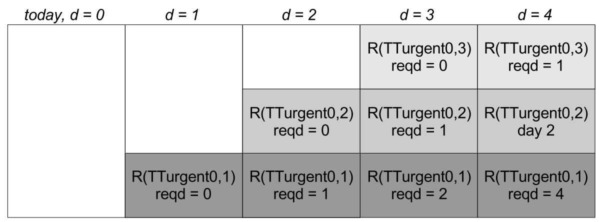

Fig. 1. Reservation within TTurgent.

4 Adaptive Model

The TTS and total capacity, and the actual method of scheduling, determine the performance of a scheduling approach. To cope with uncertainty in patient arrival, an additional surplus of capacity above the expected number of urgent and clinic patients must be available. This allocation of capacity in the TTS must be dynamically managed for maximum efficiency. We propose a three-part approach to scheduling and calendar adjustments, to best fit the calendar to current and future situations. Our approach is adaptive to, first, the current (partly filled-in) calendar, and second, the current expectation of the arrival of patients and their attributes.

Adaptive Urgent Scheduling Method To schedule urgent and clinic patients on time, the allocated capacity must be large enough. Patients with different PLANWINs use the same type of slots: urgent with (0,1); (0,2); (0,3) use TTurgent; clinic with (0,1); (0,2) use TTclinic. In practice, hospital schedulers do not distinguish between different PLANWINs; less urgent patients are regularly scheduled in place of high urgency ones.

To counter this problem we virtually divide urgent capacity while scheduling, by making reservations for different PLANWINs. Within the slots for TTurgent we make reservations (R): $R(TTurgent^{(0,1)}_{reqd})$, $R(TTurgent^{(0,2)}_{reqd})$, and $R(TTurgent^{(0,3)}_{reqd})$, for all $reqd$ ($reqd$ is the request day of the patients relative to today (0), 1 is tomorrow, etc.). To make sure patients are scheduled on time, these reservations are placed on the last day of the PLANWIN: $R(TTurgent^{(0,1)}_{0})$ on day 1, $R(TTurgent^{(0,2)}_{0})$ on day 2, etc., see Fig. 1. The same is done for clinic patients: within TTclinic reservations: $R(TTclinic^{(0,1)}_{reqd})$ and $R(TTclinic^{(0,2)}_{reqd})$ are added for all $reqd$. Note that since there can be no reservations on weekends, a large number of reservations are made on Friday; in practice, a corresponding large capacity allocated to urgent patients is found.

Given these reservations, patients are scheduled in first come first serve order, as long as enough timeslots are still available for patients with higher urgency (smaller PLANWIN). Algorithm 1 describes this method for urgent patients. We use a similar algorithm for clinic patients within TTclinic.

By dividing the total capacity as above, we also increase the variance in its usage. To deal with possible occurring problems, we allow for a reservation violation only if the patient is not scheduled on time otherwise. Specifically, in that case the patient is scheduled to the day within his PLANWIN with the most available timeslots regardless of reservations, see Algorithm 2 for urgent patients. We use a similar algorithm for clinic patients. This method makes the capacity division by reservations more flexible.

Algorithm 1 Reservations for urgent patient within TTurgent.

```
1: p is the current to be scheduled patient at day 0
2: R(TTurgent^{pw}_{reqd}) is the number of TTurgent slots reserved for patients with PLANWIN pw
   and have a request date of reqd.
3: FREE(TTurgent_d) is the number of free TTurgent timeslots on day d
4: TS is the first available timeslot within TTurgent
5: if planwin == (0,2) OR planwin == (0,3) then
6:    if (TS is on day 1) AND (FREE(TTurgent_1) ≤ R(TTurgent^{(0,1)}_{reqd=0}) then
7:       TS = the first available TTurgent timeslot after day 1
8: if planwin == (0,3) AND TS is on day 2 AND
   (FREE(TTurgent_2) ≤ R(TTurgent^{(0,1)}_{reqd=1} + R(TTurgent^{(0,2)}_{reqd=0}) then
9:    TS = the first available TTurgent timeslot after day 2
10: schedule p to TS
```

Algorithm 2 Additional steps to insert between line 9 and 10 in Algorithm 1.

```
1: if TS is outside planwin then
2:    D is day within planwin with most free TTurgent slots
3:    if FREE(TTurgent_D) > 0 then
4:       TS = the first available TTurgent timeslot on day D
```

Managing Urgent Capacity In the previous section we discussed how use reservations in scheduling urgent and clinic patients. If timeslots within the reservations are not used, these could be made available for other groups. In our approach we dynamically manage the surplus capacity allocated to deal with uncertain patient arrival. In general, to maintain high MSL, we can shift capacity between the groups urgent, clinic, and out+ivc. At the start of each day, the total surplus capacity is reallocated between groups, by the following four steps:

1. Change all remaining free TTout capacity of day 0 and 1 into TTurgent.
2. Change free TTurgent capacity on day 0 and 1 above the reservations into TTclinic.
3. Change free TTclinic capacity on day 0 and 1 above the reservations into TTurgent.
4. If free TTurgent capacity on day 0, 1, and 2 is above the capacity of the reservations, this amount of timeslots on day 2 is changed into TTout.

By using this specific order, capacity can be shifted between the three types of TTS. Empty slots of type TTout on day 0 and 1 – which can not be used by out+ivc patients (they have a PLANWIN of (2,14)) – can result in extra TTout timeslots on day 2.

Adjusting Opening Hours In busy periods, when the total demand reaches or exceeds resource capacity, waiting time can increase rapidly. With a little extra capacity this can usually be avoided. In addition to the adaptive scheduling method described above, we can use a directed search method on our Patient Scheduling Simulator (PSS) to find the required opening hours (OH) for a given desired MSL.

5 Experiments

We compare the performance of our fully adaptive approach to benchmark approaches with various levels of adaptivity. We conduct computer experiments to evaluate our

Table 5. Reservation sizes (expected n.o. patients).

	TTurgent	**TTclinc**	**TTclinc**
planwin	all days	mon, fri	tue, wed, thu
(0,1)	3 (1.6)	4 (2)	2 (1)
(0,2)	2 (1.6)	2 (2)	1 (1)
(0,3)	2 (1.6)		

Table 6. Average performance with 41h15min openings hours per week.

approach	**perf. (MSL)**	**cap. usage**
FCRS static calendar	0.77 ± 0.24	0.90 ± 0.04
Reservations	0.80 ± 0.27	0.90 ± 0.04
Flexible Reservations	0.83 ± 0.27	0.90 ± 0.04
Fully Adaptive	0.94 ± 0.15	0.91 ± 0.04
FCRS static, +2,5h	0.93 ± 0.15	0.86 ± 0.04

adaptive optimization of the scheduling process. In PPS simulations, realistic problem runs are generated. We average performances over 40 runs. Within each run patients arrive during 20 weeks. To avoid start-up effects, we start with a partially filled-in calendar, and measure average performance (MSL) over the last 10 weeks. We use a TTS optimized for an average arrival of patients. In this TTS, there are 18 TTclinic timeslots reserved for an average of 14 (± 4) clinic patients per week. There are 34 TTurgent timeslots reserved for an average of 25 (± 8) urgent patients per week. Note that patients also arrive randomly during the week.

In various experiments, we have determined the best sizes of the reservations, see Table 5. The shortest planning window needs the most surplus, and patients with lower urgencies can also make use of this surplus (as a result from Alg. 2). First we show results for scheduling methods and adaptive management of urgent capacity, for fixed openings hours. Secondly we show how opening hours can be adjusted to maintain high MSL or increases resource usage.

Fixed capacity In Table 6 we present the average performances for four approaches. The first benchmark is the baseline approach using FCRS for all patients, with a fixed resource calendar. This resembles the practical case where the calendar supervisor is absent. The second benchmark is the baseline plus the reservation blocks of Alg. 1. The third benchmark uses flexible reservations (Alg. 2). We compare this with our fully adaptive approach, that additionally manages the urgent capacity. For baseline performance to match our adaptive approach, 2.5 hours per week openings hours are required.

Adaptive opening hours When more patients arrive than expected, waiting time increases exponentially. Adding extra capacity temporarily can prevent this from happening. Our approach can then propose OH changes to resource managers to maintain high performance. In the following experiment we study a specific scenario of 16 weeks with a short busy period: $n_w = 200|w \leq 4, n_w = 300|6 \leq w \leq 11, n_w = 250|w = 5, w \geq 12$. In Fig. 2 we show the performance (averaged over 10 runs) of the baseline approach with fixed OH, our fully adaptive approach with fixed OH, and our fully adaptive approach with variable OH. We also plot the extra OH (in minutes) used by the fully adaptive approach with variable OH.

It is clear that a busy period would result in a great decline in performance for the baseline approach. Our fully adaptive approach with fixed OH does decline in performance but reaches good performance quickly after the busy period. The fully adaptive approach with variable OH can adjust the OH such that high performance is maintained over all weeks. Summed over all 16 weeks, it uses almost the same amount of OH as the approaches with fixed OH.

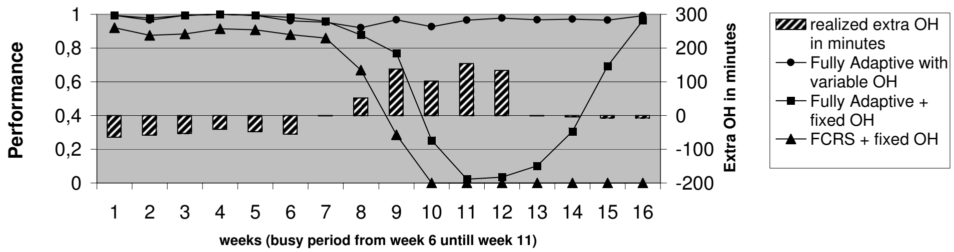

Fig. 2. Performance over weeks with variable and fixed OH.

6 Conclusions

We presented a detailed model for the CT-scan scheduling practice at the AMC. This case has similar scheduling problems as other places in the hospital. We describe how the resource calendar is structured and how various patient groups with different levels of urgency are scheduled. The resulting Patient Scheduling Simulator enables us to model various scenarios and to evaluate different allocation and scheduling approaches.

We developed an adaptive approach to the scheduling process and resource calendar management. We showed that this enables us to effectively schedule patients with different urgencies and make efficient use of capacity. By dynamically managing surplus capacity, overall, all patient groups benefit. In current practice this task requires constant attention and is critically dependent on the expertise of the calendar supervisor. Additionally we have shown that we can adjust the opening hours automatically to maintain high service levels. This is an important contribution, because currently it is very hard to determine when and by how much capacity should be extended or reduced to achieve certain patient service levels.

We are currently extending the presented work to cases where appointments must be coordinated between multiple departments, and we are looking into incorporating more patient preferences into appointment scheduling.

Acknowledgement We would like to thank Han Noot for his programming contributions, and Anieke Eikelenboom for all of her expert knowledge.

References

1. Vissers, J., Beech, R.: Health Operations Management. Routledge (2005)
2. Maruster, L., Weijters, T., de Vries, G., van den Bosch, A., Daelemans, W.: Logistic-based patient grouping for multi-disciplinary treatment. AI in Medicine 26(1-2) (2002)
3. Bowers, J., Mould, G.: Managing uncertainty in orthopaedic trauma theatres. Europ. J. of Operational Research 154 (2004)
4. Decker, K., Li, J.: Coordinating mutually exclusive resources. Proc. 3rd AAMAS (2000)
5. Vermeulen, I., Bohte, S., Somefun, D., La Poutré, J.A.: Improving patient activity schedules by multi-agent pareto appointment exchanging. In: Proc. IEEE CEC/EEE. (2006)
6. Vissers, J.M.: Patient flow-based allocation of inpatient resources: A case study. Europ. J. of Operational Research 105 (1998)
7. Oddi, A., Cesta, A.: Toward interactive scheduling systems for managing medical resources. Artificial Intelligence in Medicine 20 (2000)
8. Elkhuizen, S., van Sambeek, J., Hans, E., Krabbendam, J., Bakker, P.: Applying the variety reduction principle to management of ancillary services. Health Care Man Rev 32 (2007)