# Survey of Dynamic Medical Resource Allocation Methods

## Essay for CPSC 502 – Artificial Intelligence

**Matthew Dirks**
mcdirks@cs.ubc.ca

## Abstract

To better utilize health care resources, more time-efficient and cost-effective schedules are needed. This survey provides an overview and comparison of a variety of papers implementing dynamic resource allocation for patients in medical clinics or hospitals. Methods include Markov Decision Processes, Approximate Dynamic Programs, and Value Iteration.

## 1 Introduction

Health care resources need to be better utilized to save time and money. Medical equipment and staff is expensive so must be used efficiently and effective. Furthermore, patient health requirements rely on prompt care and high-quality service. However, most resources are not used to their fullest because of inefficient scheduling on varying types and needs of patients. To solve this problem, the goal is to dynamically allocate patients of varying priorities or classes to available resources without unnecessary cost; meanwhile keeping wait-times below desired targets and costs as low as possible [1] [2]. Until recently, the allocation of medical resources for multiple patient classes has had limited attention [2]. This survey will show some of the work in this area.

**Table 1** Timeline of papers surveyed, ordered by when first drafts were received.

| 2003 | 2004 | 2005 | 2006 | 2007 | 2008 | 2009 | 2010 | 2011 |
|---|---|---|---|---|---|---|---|---|
| Green et al. [3] | | | Patrick et al. [2] | Kolisch and Sickinger [1] | Nunes et al. [4] | Vermeulen et al. [5] | | Gocgun et al. [6] |

## 2 Why Allocation Methods?

Health care systems are being challenged more and more to provide medical services to an increasing population while doing it more efficiently and cheaper than they've ever done it before [5]. For instance, the US (2006) is giving increasingly more attention to reducing the costs of health care and finding ways to use health resources more efficiently [3]. Furthermore, Vancouver Coastal Health Authority management was concerned that OP wait times for CT scans were excessive [2]. Managers, clerks, and supervisors are under great pressure to manage these facilities in a more efficient and effective manner

There are many hospital resources which can be made more efficient, such as: CT and MRI scanners, hospital beds, attending staff, and operating rooms [5]. Medical services, and imaging equipment in particular, is often underutilized while patients have lengthy and increasing wait times [2] [3]. Analysis shows that many scheduled appointments for outpatients exceed target wait-times [2].

Of particular interest is diagnostic-imaging equipment. MRI installations are a critical component of a health care system [3] and CT scanners are central in the clinical pathways of many patients [5]. These scanners are expensive, running at about $2 million for a new MRI in 2006 in the US [3].

Current practice relies on the expertise of a booking clerk or calendar supervisor to efficiently plan when appointments should be based on their predictions of demand [2]. One reason for this is that acceptance for computer-based decision rules in medical environments is low [1]. However, as the area is researched and made better known, we will likely see more acceptance of computer-aided decision makers. In the meantime, the easiest workaround are simple decision rules which can be applied manually.

## 3 Domain Information

There are various patient classes. In general they are: inpatients (IP), patients who are staying at the hospital; outpatients (OP), patients outside the hospital scheduled for an appointment; and emergency patients (EP), patients in critical condition who need an immediate scan. Some facilities have a dedicated scanner for EPs (such as in [5] and [6]), and others are only for scheduled appointments with no IPs ([4] and [5]).

See Table 1 for a list of abbreviations used in the papers surveyed.

**Table 2** Abbreviations used in the papers surveyed. Grouped by category.

| | |
|---|---|
| **Patient Types** | |
| OP | Outpatient (scheduled) |
| IP | Inpatient |
| EP | Emergency Patient |
| CEP | Critical EP |
| NCEP | Non-Critical EP |
| Add-on OP | Outpatient with increased urgency (must be scanned within the day) |
| IVC | Intravenous Contrast needed |
| **Methods** | |
| MDP | Markov Decision Process |
| LP | Linear Program |
| ALP | Approximate Linear Program |
| ADP | Approximate Dynamic Program |
| VIA | Value Iteration Algorithm |
| LA [3] | Linear Approximation |
| **Rules** | |
| FCFS | First-Come First-Serve |
| FCRS | First-Come Randomly-Served |
| RS | Random Selection |
| FAS [3] | Fill All Slots (all appointments filled with OP) |
| LCA [3] | Linear Capacity Allocation |

## 4 Solutions

All the papers echo the feeling that modeling as an MDP results in so big a state space as to be unsolvable (intractable). However, they each find a way to overcome this problem either by simplification, approximation, or avoiding MDPs entirely.

The MDP model fits well to this scheduling problem. Firstly, it is a sequential decision process with uncertainty elements. Secondly, the full system state can be observed at each decision point in time. Lastly, the Markovian assumption that the current state depends only on one previous state is appropriate here [4]. These characteristics lend well to the MDP model.

MDPs have been used to model the scheduling problem in all but one of the papers surveyed in this essay ([1] [2] [3] [4] [6]), leaving Vermeulen et al. [5] as the only one to take an alternate approach – they developed a custom dynamic approach instead.

All MDPs can be modeled as a set *S* of states, where we can be in state *s* at time *t*. Then let *R(s, a, s')* be the reward function describing any rewards or costs associated with some state and action. Let *A* be the set of possible actions, and *P(s'|s, a)* is the probability of ending up in state *s'* given the previous state *s* and the action *a* taken (also known as the Transition or Movement Model) [7].

Next we can find the optimal action by solving Bellman's equation:

$$V(s) = \max_a (C(s,a) + \gamma \sum_{s'} p(s'|s,a)\, V(s'))$$

For instance, Green et al. model the scheduling problem as an MDP, and solve it using a Linear heuristic instead of a more difficult optimal real-time service policy [3]. As seen in lecture, MDPs are based on the assumption that the current states only depends on the previous state. As such, in Green's paper we have: State consisting of number of IPs and OPs; actions at each state depending on whether or not an EP has arrived in the current state; and a reward consisting of revenue from each patient type, wait cost per period, and penalty for patients not served by the end of the day. Then they formulate the profit maximization problem as a finite-horizon dynamic program, and end up with V* representing the optimal expected daily profits [3].

Gocgun et al. [6] base their method on Patrick et al.'s method. Patrick et al. solve the MDP differently because of the state space issues with MDPs. In an MDP we have to loop over all states. This is known as the curse of dimensionality. For this reason Approximate Dynamic Programming (ADP) is used. ADP is a technique for solving (via approximation) stochastic optimization problems such as MDPs.

The basic idea of the ADP algorithm is to iterate forward changing the value function approximation and changing our policy for making decisions. ADP also goes under different names, such as the "forward pass" algorithm. When we add expectations to ADP, then the reinforcement learning community would call this Temporal Difference learning (TD). TD is essentially an approximation of Value Iteration. ADP successfully overcomes the curse of dimensionality present in MDPs. However, it has its own problems to overcome as well, such as the problem of developing a good statistical approximation of the value of being in each state. More details on ADP can be found in the work by Powell [7].

# 5 Evaluation Methods

In general, each paper compares their optimal policy to various decision rules (or heuristics) and then compares them under some scenarios or simulation.

Kolisch and Sickinger [1] compared their optimal policy with three decision rules, and assess the impact of each on various scenarios. The decision rules used are: Linear Capacity Allocation (LCA) as introduced by Green at al. [3], First-Come First-Served (FCFS), and Random Selection (RS).

Patrick et al. [2] ran three scenarios to simulate a small outpatient clinic, a large outpatient clinic, and the Vancouver hospital. Their Approximate Optimal Policy performs well in each scenario regardless of the size of the hospital.

Green at al. [3] compared three heuristic appointment policies to the optimal policy in order to find out which policy, or combination of policies, performed best as an estimate of the optimal policy. They ran their simulation for 50,000 days for each policy. One notable conclusion is that two of the policies are significantly affected by the uncertainty in the durations of diagnostic exams.

Nunes et al. [4] compare the optimal policy with two others: greedy policy and fixed policy. The result was an optimal policy that performed best for lowering costs, and performed about the same under all other metrics.

Vermeulen et al. [5] also performed a patient scheduling simulation. Performances averaged over 70 runs with patients arrive during 20 weeks. The number of patients per week was modeled by via a random walk. They simulate reality by using a First Come Randomly Served (FCRS) rule – a variation of the FCFS rule as seen in other papers. Here patients are scheduled as they arrive, but then he is assigned a slot randomly from all the free slots within his planning window. This simulates the effect of manually assigning slots according to patient preferences and other considerations.

Gocgun et al. [6] used a computer simulation on 32 different scenarios. Each scenario had a different variety of parameters. Using tornado charts, one way sensitivity analyses were performed to evaluate the impact of specific parameters on model outcomes. They compared their optimal policy to commonly used decision rules, including FCFS, two rules similar to Random Selection, and two priority-based rules. They generated 100,000 independent day-long sample paths of random events, and then implemented each decision method.

# 6 Survey

See Table 3 for a comparison of each paper.

## 6.1 Kolisch and Sickinger [1]

In the paper by Kolisch and Sickinger [1], they assume that only one IP can enter per slot. In reality this is a poor assumption as any number of IPs could become ready for a scan at the same time.

This paper differs from the others in that they model the problem and solve for the optimal solution using the Backward Induction Algorithm.

## 6.2 Patrick et al. [2]

This paper differs from the others in a few ways. Firstly, they consider an arbitrary number of patient priority classes, rather than just 1, 2, or 3 as the others have. Secondly, they specifically optimize for wait-time targets while being cost-effect. Most others maximize revenue. Thirdly, they were the first to adapt the Approximate Dynamic Programming method to this particular problem, thus solving a multipriority patient-scheduling problem which was previously intractable.

One particularly good strength of their approach is when patient demand fails to be taken care of in the current day it reappears in the next day’s demand. As opposed to starting fresh at each day as most approaches do. Furthermore, patients may also be diverted to other facilities, as a way of dealing with excessive demand, but at a penalty cost signifying lost business and poor service.

## 6.3 Green et al. [3]

This paper was the first of the papers surveyed here to work on the scheduling problem. Green at al. provided a solid groundwork for the research which came after and came to reasonable results themselves. Their work has been expanded on many times.

The biggest difference with their approach was using a Linear Heuristic to estimate the solution, whereas most other papers now have solved or approximated the MDP.

## 6.4 Nunes et al. [4]

One weakness of their approach, is the simplification they made of not including emergency patient cases, because the hospital they examined does not provide emergency care. This simplification helped make the problem easier to solve, however.

However, their approach stands out in a number of ways. Firstly, Nunes et al. implemented a hypothetical prototype and applied Value Iteration to it as their solution. No other papers have done this.

Secondly, their approach introduced the idea of treatment patterns. Handling the case where a patient

**Table 3** Comparison of papers surveyed

| Paper | Horizon | Model | Solved Via | # CT/MRI | Patient Types | Objective |
|---|---|---|---|---|---|---|
| **Kolisch and Sickinger [1]** | Finite | MDP | Backward Induction Algorithm | 2 CT | OP, IP, EP | Maximize expected net revenue over work day (revenue, waiting cost, penalty cost) |
| **Patrick et al. [2]** | Discounted infinite | MDP | ALP created by ADP | 1 CT | Arbitrary number of priority classes | Optimize for wait-time targets in cost-effect manner |
| **Green et al. [3]** | Finite | MDP | Linear Heuristic | MRI | OP, IP, EP | Maximize expected net revenue over work day |
| **Nunes et al. [4]** | Finite (7 or 15 days) | MDP | Value Iteration | *m* specialties | OP | Minimize the expected average cost over all possible policies |
| **Vermeulen et al. [5]** | N/A | Adaptive Approach to Automatic Optimization | N/A | 2 CT | OP+IVC, OP-IVC, IP | High service-levels |
| **Gocgun et al. [6]** | Finite | MDP | ALP created by ADP | 1 CT, 2 CT | OP (Add-on OP), IP, EP (Critical & Non-critical) | Maximize expected net revenue over work day |

needs a sequence of different treatments, not just CT or MRI scans.

Lastly, they plan admissions with fixed periods of time longer than just one day, looking instead at periods of one week or fifteen days.

### 6.5 Vermeulen et al. [5]

This paper develops an adaptive approach to automatic optimization of resource calendars. This approach is different from other papers in that they work on the operational level and they match the current scheduling procedure in the hospital. Their approach improves the current scheduling process in the hospital as opposed to replacing it – which they propose is most beneficial. Their adaptive approach with fixed capacity declines in performance in busy periods, as expected. However, their adaptive approach with adjustable opening hours can adjust capacity such that high performance is maintained over all twenty weeks of the simulation.

### 6.6 Gocgun et al. [6]

This paper has a few differences from the others. Firstly is the addition of non-stationary arrival probabilities. Secondly, they do extensive sensitivity analysis using tornado charts – no other papers do this. Thirdly, they also model 2 CT scanners and use the ADP technique. Other papers have done some of these things, but this paper was the first to combine everything.

One weakness of this paper is the restriction that at most one request from each patient-type can arrive during a slot.

## 7 Progress

Kolisch and Sickinger [1] build on Green et al. [3] by adding an additional scanner (2 CT's as opposed to their 1 MRI) and appointments per period can exceed available time slots. Gocgun et al. [6] also builds on Green at al. [3], using 2 scanners and classifying patients into finer sub-categories to better reflect reality.
Patrick et al. [2] and Vermeulen et al. [5] both noted how modeling the multipriority patient-scheduling problem as an MDP results in a state space of unsolvable size. This previously intractable problem of solving the MDP has a reasonable solution by Patrick et al. [2] using a linear program equivalent to the MDP via approximate dynamic programming. Gocgun et al. [6] then uses this approximation technique in their solution.

The Linear Capacity Allocation decision rule has been introduced by Green et al. [3] for the single resource case. Kolisch and Sickinger [1] use this and extends it to two resources. They also expand on Green et al. [3] by considering appointment schedules where the number of appointments per period is arbitrary and not limited by the number of available resources – in an attempt to overcome outpatient no-shows by overbooking – and by modeling the probability of a patient showing up. Gocgun et al. [6] also models arrival probabilities.

Green et al. suggested that the time horizon for the analysis of capacity management decisions might more appropriately be extended to a week. This extension would also allow for a more detailed and accurate analysis of inpatient demand, which is sometimes pushed over from one day to the next [3]. Nunes et al. [4] independently – with no reference to Green et al. [3] – does implement a longer period of time of one week or fifteen days.

## 8 Future Work

An interesting idea for future work, suggested by Patrick et al., is to investigate having centralized booking system that services a number of different hospitals or clinics [2]. Of course this would also require taking into account patient preferences and travel costs. Vermeulen takes a simpler approach, proposing to scale the problem to multiple departments and research mechanisms for coordination between departments [5].

Patrick et al. also suggests applying this scheduling problem to other areas such as surgical scheduling or radiation treatments [2]. In Nunes et al. we see a general approach which can already be applied to any number of hospital services, other than just CT or MRI scans [4]. Vermeulen et al. also proposed generalizing the method [5]. Nunes et al. did work with the idea of a pattern (or sequence) of resources assigned to a patient [4]. A pattern could include, for example, medical-consultations, in-patient days, and MRI scans. A pattern could also be the sequence of appointments needed in radiation treatments, as Patrick et al. suggested [2]. Therefore, further studies in applying clustering methods to real data, in order to determine the treatment patterns and the related transition probability matrices, is a promising potential future work [4].

None of these papers have considered seasonal influences on the patient arrival probabilities. This is an area that has yet to be researched. However, Patrick et al. proposes that even if seasonal patterns are significant, the optimal policy is extremely robust to changes in the specific data and thus re-solving may not be necessary [2].

## 9 Conclusion

Medical resources are an expensive but very important part of medical care. More efficient and effective use of them can provide higher-quality care to patients without increasing the cost of medical bills.

The inefficient use of imaging equipment can be reduced by better scheduling admission for each patient class. Thus, equipment idle-time can be reduced, if not prevented entirely, while considering the varying patient types and treatments.

These papers have shown impressive progress in developing useful decision rules for efficiently and effectively dynamically allocating the medical resources to multiple patient groups. In time these solutions can be implemented into all clinics and hospitals. Plus, the promising future work should lead to even better and more robust solutions down the road.

## References


[1] Kolisch R, Sickinger S. Providing radiology health care services to stochastic demand of different customer classes. OR Spectrum 2008;30:375–95.

[2] Patrick J, Puterman ML, Queyranne M. Dynamic multipriority patient scheduling for a diagnostic resource. Operations Research 2008;56:1507–25.

[3] Green LV, Savin S, Wang B. Managing patient service in a diagnostic medical facility. Operations Research 2006;54:11–25.

[4] Nunes LG, de Carvalho SV, de R, Rodrigues CM. Markov decision process applied to the control of hospital elective admissions. Artificial Intelligence in Medicine 2009;47:159–71.

[5] Vermeulen IB, Bohte SM, Elkhuizen SG, Lameris H, Bakker PJ, La Poutre H. Adaptive resource allocation for efficient patient scheduling. Artificial Intelligence in Medicine 2009;46:67–80.

[6] Gocgun Y, Bresnahan BW, Ghate A, Gunn ML. A Markov decision process approach to multi-category patient scheduling in a diagnostic facility. Artificial Intelligence in Medicine 2011;53:73-81.

[7] Powell WB. Approximate dynamic programming: solving the curse of dimensionality. NJ, USA: Wiley, Hoboken; 2007.