Computers & Operations Research 66 (2016) 1–11

Contents lists available at ScienceDirect

Computers & Operations Research

journal homepage: www.elsevier.com/locate/caor

Modelling and solving generalised operational surgery scheduling problems

Atle Riise [a,b,*], Carlo Mannino [a], Edmund K. Burke [c]

[a] SINTEF ICT, Department of Applied Mathematics, P.O. Box 124 Blindern, 0314 Oslo, Norway
[b] University of Oslo, Department of Informatics, P.O. Box 1080 Blindern, 0316 Oslo, Norway
[c] University of Stirling, Cottrell Building, Stirling FK9 4LA, UK

ARTICLE INFO

Available online 29 July 2015

Keywords:
OR in health services
Project scheduling
Surgery scheduling
Operating room management
Heuristics

ABSTRACT

The term 'surgery scheduling' is used to describe a variety of strategic, tactical and operational scheduling problems, many of which are critical to the quality of treatment and to the efficient use of hospital resources. We consider operational surgery scheduling problems. The exact problem formulation varies substantially between hospitals or, even, hospital departments. In addition, the level of detail varies between different planning situations, ranging from long term patient admission planning to the very detailed scheduling of a particular day's surgeries. This diversity makes it difficult to design general scheduling methods and software solutions that can be applied without extensive customisation for each application. We approach this challenge by proposing a new generalised model for surgery scheduling problems. We show how this model extends the multi-project, multi-mode resource constrained project scheduling problem with generalised time constraints, including some extensions that to our knowledge have not been previously studied. Furthermore, we present a search method for solving the proposed model. The algorithm uses on-line learning to balance computational loads between a construction and an improvement method, both working on high level solution representations. An adapted schedule generation scheme is used to map these to concrete schedules. We perform computational experiments using realistic problem instances from three surgery scheduling planning situations at a medium sized Norwegian hospital; day scheduling, week scheduling and admission planning. The results show that the algorithm performs well across these quite different problems without any off-line customisation or parameter tuning.

1. Introduction

Efficient surgery planning and scheduling is crucial for minimising patients' waiting time, reducing the number of cancellations, levelling staff work load and improving the overall performance of a hospital [1]. The focus of this paper is on the wide range of operational surgery scheduling problems that arise in the daily scheduling of actual individual patients. Constrained by higher level decisions, such as a defined Master Surgery Schedule [2], these problems may be informally described as the task of assigning start times to all surgery related activities for each patient, while reserving capacity for these on a chosen set of constrained resources. Such activities may include, for example, the preparation of a patient for surgery, preparation of equipment, removing unnecessary equipment, surgery, waking the patient, cleaning of the operating room and equipment, transporting the patient to the recovery room, and recovery. The involved resources can include operating rooms, operation teams, surgeons, anesthesiologists, equipment, post-operative beds, and others. Objectives are typically represented by resource overtime, hospitalisation costs, intervention costs, operating room utilisation, patient's waiting time, and patient or personnel preferences, among others. These scheduling problems are often NP-hard [3]. The level of detail varies between different planning situations: patient admission planning may consider only one or two kinds of resources, is mainly concerned with allocating a date of admission for each patient, and typically has a long time horizon. The number of activities, resources, and choices to make increase closer to the day of surgery. This might be when scheduling surgeries for the next one or two weeks. Finally, a very detailed schedule is made for the next day, considering all relevant activities and resources in full detail. Due to unforeseen events and activity durations, this schedule is usually subject to dynamic re-scheduling during the actual day of surgery.

General surgery planning software and associated scheduling algorithms must be able to handle this variety without excessive

* Corresponding author at: SINTEF ICT, Department of Applied Mathematics, P.O. Box 124 Blindern, 0314 Oslo, Norway. Tel.: +47 922 48 389.
E-mail addresses: atle.riise@sintef.no (A. Riise), carlo.mannino@sintef.no (C. Mannino), e.k.burke@stir.ac.uk (E.K. Burke).

http://dx.doi.org/10.1016/j.cor.2015.07.003
0305-0548/

customisation to each individual hospital and planning situation. So far, however, the surgery scheduling literature directly reflects the diversity of the problem domain (see recent surveys in [1,4–6]). Different authors use different problem definitions. They also consider different aspects of the problem, such as the intensive care unit as a bottleneck resource [7], the prediction of hospital bed availability [8], the use of mobile equipment [7], uncertainty in surgery durations [9,10,3] and 'no shows' [11], human resources availability and preferences [12], and combined resource allocation and sequencing in admission planning [13]. Ref. [14] formulates a problem that includes several real world constraints, including team preferences, for a day-scheduling problem. Bulgarini et al. [15] consider a combination of short term scheduling with a long term assignment of patients to weeks, while [16,17] consider the allocation of patients to days or operating sessions in combination with master surgery schedule creation. We are not aware of any common ontology for operational surgery scheduling problems.

In order to amend this situation, we propose a model for a 'generalised operational surgery scheduling problem' (GOSSP) that can express a wide range of real world surgery scheduling problems. This also answers the need expressed in [5] for the inclusion of more real world aspects in surgery scheduling models. We also propose a general purpose meta-heuristic that generates good quality schedules for GOSSP problem instances with different characteristics.

In order to qualify the generality of the model (and thus the general applicability of the scheduling algorithm), it is useful to consider how it can model problem aspects that have been previously treated in the surgery scheduling literature. The literature review of Cardoen et al. [1] (and the recent update [6]) uses a classification of this literature along six dimensions: *patient characteristics*, *performance measures*, *decision delineation*, *research methodology*, *applicability of research*, and *uncertainty*. Regarding *patient characteristics*, the GOSSP model can express elective surgery scheduling problems, for both inpatient and outpatient clinics. It may also express on-line re-scheduling problems with a mix of elective and urgent, or emergency, patients.

Regarding *performance measures*, the presented model uses a weighted sum of a wide range of objective components. By the term *decision delineation* in [1], the authors refer to "indicating what type of decision has to be made (date, time, room or capacity) and whether this decision applies to a medical discipline, a surgeon or a patient (type)". Since the GOSSP concerns operational problems, it addresses decisions made at the *patient level*. At this level, the GOSSP framework enables the modelling of all the above decisions, and more. The choice of any renewable resource (such as operating rooms, preparation rooms, teams, recovery rooms, surgeon, and equipment) may be modelled, along with the choice of date and time for each surgery related activity. In the literature, surgery scheduling that cover more than one day is sometimes considered as a two-step process [7,18]. In the first step, operating rooms and dates of surgery are assigned to each patient [19–21]. In the second step, a schedule is created for each day. Some authors address both steps [7]. The advantage of this two-step approach is that it reduces the problem complexity. Both steps are, at least in the general version, NP-hard [3,22]. The drawback, however, is that the assignments from the first step may not permit a reasonably good, or even feasible, schedule to be found in the second step where more detailed constraints and preferences are considered [22,7]. In real life, the admission planner actually considers scheduling preferences when assigning days to patients. One example is the 'time-of-day' preference for certain patients, such as children or patients with diabetes. As in [13], we therefore choose to address the full problem, including both day assignment and scheduling. Only a few other papers take this approach, and then only in connection with short term planning (a day, or a week) [20,23–25,12]. In Section 4, we apply this integrated, one-step, approach also to long term admission planning problems.

From the reviews in [1,4–6], it can be seen that a wide range of *research methodologies* have been applied to various versions of surgery scheduling problems. In this work, we describe a meta-heuristic that through the generality of the GOSSP model can be applied efficiently to very different planning situations. Note that as such, it can also be used to implement daily (re-) scheduling in a simulation framework used for scenario analysis.

In terms of the *applicability of research*, we demonstrate the effectiveness of the proposed model and scheduling heuristic by testing it on real life problem instances from three different planning situations: admission planning, week-planning, and day-planning. The data are from a surgery department in Bærum Sykehus, a medium sized Norwegian hospital, and are available on-line [26].

The last of the classification dimensions in [1] is *uncertainty*. Operational surgery scheduling is inherently stochastic. The uncertainty has three main sources: stochastic activity durations, stochastic resource availability (staff illness), and stochastically occurring cancellations or arrivals of emergency care patients. The hospitals that we have interviewed take a pragmatic approach to these stochastic aspects. They treat activity durations as deterministic, but may add some slack to hedge for unforeseen delays. They may also consciously under-estimate activity durations to plan for over-booking. In long term admission planning, they may plan with reduced future capacity for certain resources, in order to reserve capacity for patients that may need treatment urgently, but that are not yet known at the time of scheduling. In detailed surgery scheduling, with a planning horizon of a few days, the scheduling is done without any regard to unforeseen arrivals or cancellations; these are simply handled by reactive re-scheduling as they occur. Although a more scientific approach to these stochastic scheduling problems may in principle provide better schedules, there is a pedagogical challenge to persuade a surgery planner to give up these easily understandable planning strategies. As noted in [1], the majority of hospitals that have been previously studied handle disruptions in a similar way.

The literature is sparse on stochastic surgery scheduling, and the problems that are treated are often unrealistically simple or apply simplified recourse actions. While the stochastic nature of the problem is obviously important, we have chosen in this paper to focus on creating a rich and realistic deterministic model for surgery scheduling problems, and on solving such a model to a reasonable level of quality within an acceptable time frame (minutes). Efficient algorithms working on such a general model will enable efficient reactive re-scheduling in a real world setting with stochastic disruptions and activity durations.

Some researchers have modelled surgery scheduling problems as extensions of the resource-constrained project scheduling problem (RCPSP) [12,23] or the job shop scheduling problem [24]. This is a useful approach, as it describes surgery scheduling problems in the context of well known operations research models. We follow the same approach, and aim for a high level of generality. We model the GOSSP as an extension of the multi-mode RCPSP with minimum and maximum time lags (MRCPSP/Max). This is sometimes also called the multi-mode RCPSP with generalised precedence relations (MRCPSP-GPR). The GOSSP extends this problem, both by being multi-project and by including several additional constraints. Some of these, the 'inter-activity mode compatibility constraints' and 'project disjunctions', have, to our knowledge, not been previously studied. The most relevant existing literature describe algorithms for the MRCPSP/Max, and even this literature is very limited. Previous studies have applied