Operations Research for Health Care 3 (2014) 200–214

Contents lists available at ScienceDirect

Operations Research for Health Care

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

A network-based approach for monthly scheduling of residents in primary care clinics

Jonathan F. Bard [a,*], Zhichao Shu [a], Luci Leykum [b]

[a] *Graduate Program in Operations Research and Industrial Engineering, The University of Texas, Austin, TX 78712, United States*

[b] *School of Medicine, The University of Texas Health Science Center, San Antonio, TX 78229, United States*

ARTICLE INFO

Article history:
Received 5 November 2013
Accepted 21 August 2014
Available online 1 September 2014

Keywords:
Resident scheduling
Soft constraints
Network model
Goal programming
Internal medicine

ABSTRACT

This paper presents a network model with the objective of maximizing the number of interns and residents (collectively called *housestaff*) who are assigned clinic duty each month during their training in internal medicine. A complexity analysis is provided that demonstrates that the basic problem can be modeled as a pure network. When team considerations, on-service, and other hard clinic constraints are taken into account, it is shown that the problem can be modeled as a network with gains. This result is unaffected when a series of soft constraints related to the number of clinic assignments per week per housestaff, the number of required faculty, the ratio of housestaff to faculty, and clinic assignment requirements during the month are taken into account. To evaluate the effectiveness of the model, a comparative study for the 2012–13 academic year was undertaken in conjunction with the Internal Medicine Department at the University of Texas Health Science Center in San Antonio. The results indicated that an average increase of 5.3% in the number of clinic sessions assigned over the year is possible when compared to those obtained manually by the Chief Resident. In addition, the quality of the solutions with respect to the soft constraints was notably better.

1. Introduction

Upon completing medical school, physicians typically enter a three-year residency program for specialized training. In the first post-graduate year (PGY 1), they are referred to as interns, and in PGY 2 and 3, as residents. Collectively, they are referred to as *housestaff*. As part of the training, each member of the housestaff is expected to spend a minimum number of hours in clinic each year. This requirement varies by service or department but generally equates to one or two half-day sessions per week. Planning is done monthly with the overall goal of providing as many clinic hours as possible subject to the available resources and procedural rules of the program. The difficulty faced by those responsible for constructing the monthly schedules is that there are a variety of hard and soft constraints that must be taken into account which are impossible to deal with simultaneously without the aid of a decision support system. Compounding the problem is the absence of documentation on most of those constraints and the fact that their relative importance can shift over time. To date, there has been little effort to document the rules and requirements governing clinic assignments that would allow for the development of a mathematical model. Although information systems exist for managing housestaff, there has been no effort that we are aware of to automate the scheduling process.

In this paper, the Internal Medicine Department at The University of Texas Health Science Center in San Antonio (UTHSCSA) serves as the model for our work. A primary goal of their residency program and virtually all such programs across the United States is to maximize the number of hours that housestaff spend in clinic each month. Individual assignments are subject to team, call, rotational, day of the week, day of the month, post-graduate year, and supervisory constraints, to name the most prominent. Some of these can be viewed as soft constraints or goals that are desirable to meet but can be violated if no other option is available. The goals are ordered by priority and given weights in the form of user-supplied parameters that reflect their relative importance. This construct permits schedule violations of higher priority goals only if an equivalent number of lower priority goal violations are present.

The basic structure of each post-graduate year centers on monthly blocks or rotations, which are designed to provide a

* Corresponding author. Tel.: +1 512 471 3076.
E-mail addresses: jbard@mail.utexas.edu (J.F. Bard), shuzhichao@gmail.com (Z. Shu), leykum@uthscsa.edu (L. Leykum).

http://dx.doi.org/10.1016/j.orhc.2014.08.002
2211-6923/

variety of experiences in a controlled setting. Typical rotations include general hospital wards, specialty wards such as hematology and geriatrics, critical care units, and night admissions. In the US, all training falls under the purview of the Accreditation Council for Graduate Medical Education (ACGME), a national accrediting body that sets the guidelines and expectations for each residency program. In some cases, the guidelines are equivalent to hard constraints while in others, they can be considered goals.

The construction of monthly clinic schedules is the responsibility of the chief resident who serves in that position during his or her "fourth" year. It is a complex process that must take into account a conflicting set of ACGME requirements, the individual preferences of each housestaff, supervising faculty availability, and a varying list of training and procedural goals. For the vast majority of residency programs, clinic schedules are constructed manually in a haze of confusion and frustration that may consume up to a week of the chief resident's time. To compound matters further, the derived schedules are often suboptimal and in violation of one or more hard constraints.

In our first effort to provide a more rational and systematic approach to clinic scheduling, we developed a three-phase methodology centered on solving a mixed-integer nonlinear program (MINP) with the ultimate objective of maximizing the number of clinic sessions assigned to the eligible housestaff per month [1]. In the first phase, block data on each member of the housestaff are extracted from the Internal Medicine Department's information system, and those determined to have fixed assignments that month are removed from the model. Data structures are adjusted accordingly. In the second phase, a relaxed solution is found with a commercial optimizer; in the third phase, feasible schedules are obtained and a local optimum is found using a randomized exchange procedure. Somewhat surprisingly, the MINP solves quickly for instances with approximately 80 housestaff, 60 sessions, 40 rotations, 3 clinics, and 9 teams even though it contains many thousands of binary variables and constraints. The majority of the time was taken up in the third phase. These observations led to a complexity analysis of the problem and the realization of its general network structure. This motivated the development of a network flow-based model and additional analyses to determine if the overall results could be improved.

With this in mind, the purpose of this paper is threefold. First we provide a new model for scheduling residents at clinics during their monthly rotations; second, we examine the complicating constraints and determine the degree of difficulty they add to the problem; and third, we conduct a full analysis for the 2012–2013 academic year using data provided by UTHSCSA and compare the results obtained with the network model with those obtained manually by the chief resident. Although the previously developed MINP provides good results, the advantages of the network approach are fivefold: first, it proved to be more intuitive and hence easier for the chief resident to understand, second, it is simpler to code, third, it is more flexible in accommodating new constraints, fourth, it provides a globally optimal solution while all that can be said about the solutions provided by the MINP is that they are locally optimal, and fifth, it runs about 40% faster than the MINP which is critical when multiple scenarios are to be evaluated.

In the next section, we highlight the related literature. This is followed in Section 3 by the problem statement, including a discussion of the hard and soft constraints, the primary and secondary objectives, and how fairness is achieved in the schedules. The network concepts that underlie the scheduling problem are discussed in Section 4 accompanied by a complexity analysis of the hard constraints. This analysis provides the foundation for the network model used to find solutions and represents a major contribution of the paper. Section 5 contains a discussion of the preprocessing steps, the network formulation and a description of the input data. The code that embodies the model has been in use since 2013. The results from the comparative analysis are highlighted in Section 6. They indicate that increases of up to 11% in the number of sessions scheduled per month can be realized with the proposed approach. Conclusions are drawn in Section 7.

2. Literature review

There is a vast amount of literature on staff scheduling in the service industry, particularly related to the monthly tour scheduling problem investigated here. The interested reader is referred to [2,3]. Perhaps the area that comes closest to ours in terms of rules and requirements is airline crew scheduling. A good overview of models and solution techniques is given by Kohl and Karisch [4].

The majority of research on physician scheduling has concentrated on emergency room (ER) and operating room (OR) assignments (e.g., see [5–11]). There has also been some parallel work on scheduling anesthesiologists for surgeries [12,13]. With respect to healthcare workers in general, nurses have been the most widely studied (see [14,3,15]). Effective methods now exist that address both the midterm and short-term problems with an emphasis on satisfying individual preferences, and have been implemented in many facilities with some degree of success. Bard and Purnomo [16,17], for example, used a column generation scheme to solve the midterm scheduling problem for nurses. Individual preferences were addressed in their line-of-work construction subproblem. Jaumard et al. [18] used an exact branch-and-price algorithm to solve a similar problem, and were successful for hospital units with a few dozen nurses. Others, including Bester et al. [19], Burke et al. [20], Burke et al. [21], and Dowsland [22], have relied on intelligent heuristics to find good solutions. For work on daily nurse scheduling, see [17,23].

The problem of physician scheduling outside of the ER and OR has received much less attention, and in the vast majority of cases, is still done manually at great time and expense. The complex nature of the general problem, particularly in the case of housestaff scheduling where specialized constraints based on training requirements and departmental rules can be overwhelming, makes it difficult to define a set of widely accepted hard or even soft constraints, unlike the situation with nurses where the restrictions are similarly from one hospital to the next. By implication, it is difficult to formulate a generic problem that lends itself to common modeling approaches that rely on a set of basic shifts and start times (e.g., see [24,25]).

Nevertheless, there has been some progress in this area. Rousseau et al. [26] provide a list if rostering issues confronting the scheduler but offer little in the way of modeling or solutions. Gunawan and Lau [27] addressed the planning problem of assigning physicians daily tasks such as surgery, clinic, call and administration to defined time slots or shifts over a time horizon. Their model incorporated a large number of constraints and complex physician preferences, and had the objective of satisfying as many preferences and duty requirements as possible while ensuring optimal usage of available resources.

With respect to exact methods, Franz and Miller [28] developed a MIP for scheduling residents at a large teaching hospital, while Cohn et al. [29] employed optimization techniques to schedule the teaching phase of special training programs at Boston University School of Medicine. One-year schedules that consider on-call service and vacation requests were obtained. White and White [30] explored the similar problem of scheduling hospital rounds by specialty for teams and used tabu search to construct monthly rosters. Ovchinnikov and Milner [31] describe an implementation of a spreadsheet model for assigning radiology residents over a 1-year horizon in radiology at the University of Vermont's College of Medicine. They argue that spreadsheets are preferable to free