Interfaces

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Physician Scheduling for Continuity: An Application in Pediatric Intensive Care

Hannah K. Smalley, Pınar Keskinocak, Atul Vats

Interfaces
Vol. 45, No. 2, March–April 2015, pp. 133–148
ISSN 0092-2102 (print) | ISSN 1526-551X (online)
http://dx.doi.org/10.1287/inte.2014.0774

Physician Scheduling for Continuity: An Application in Pediatric Intensive Care

Hannah K. Smalley, Pınar Keskinocak
Georgia Institute of Technology, Atlanta, Georgia 30332
{hksmalley@gmail.com, pinar@isye.gatech.edu}

Atul Vats
Children's Healthcare of Atlanta, Emory University School of Medicine, Atlanta, Georgia 30329, atul.vats@choa.org

A physician schedule that maximizes continuity (i.e., reduces instances of patients being treated by multiple physicians) could improve the efficiency of handoffs—the transfer of patients from the care of one physician to another. We present a modeling and solution approach for assigning physicians to service and call shifts in the pediatric intensive care unit (PICU) at Children's Healthcare of Atlanta at Egleston (Children's). We developed the handoff continuity score (HCS) for measuring the continuity of a schedule. We combined the HCS with a mixed-integer programming model (MIP) with the objective of maximizing the HCS, while minimizing violations of physician preferences. For a 51-week horizon and a physician pool of 16 physicians, no feasible solution to this MIP is found within 48 hours using CPLEX 12.4. However, an iterative heuristic incorporating modified versions of the MIP produces a schedule (3.42 percent optimality gap) for the scheduling instance faced by Children's for this period. Our solution approach facilitates resource optimization, and automated scheduling requires less time than manually constructing such a schedule. We generated six-month schedules that were implemented in the PICU at Children's in 2011, 2012, and 2013. Such automated schedule construction allows for creation of schedules that maximize continuity.

Keywords: mixed-integer programming; continuity of care; healthcare; physician scheduling.
History: This paper was refereed.

In 2003 and again in 2011, the Accreditation Council for Graduate Medical Education (ACGME) instituted work-hour restrictions that greatly limit the hours that residents can work (ACGME 2011b); its goal was to improve performance and reduce medical errors and sleep deprivation, while ensuring time for essential training and experience (Ulmer et al. 2008). One such restriction was a reduction in the length of allowed duty periods. Landrigan et al. (2004, p. 6) link extended duty periods of 24 hours or more to increased medical errors, and Wu et al. (1991) present results of a survey of 114 physicians, 41 percent of whom reported fatigue to be a contributing factor in past medical errors. Fatigue among medical residents also negatively impacts the personal lives of those medical residents (Papp et al. 2004).

In an inpatient setting in which the average patient length of stay exceeds one day, patients are likely to be treated by multiple physicians, including attending physicians and residents. Although they are beneficial in many ways, increased duty-hour restrictions contribute to fragmentation of care—the treatment of individual patients by multiple physicians. This occurs because medical residents may be on duty for shorter periods than were allowed prior to the implementation of the new duty-hour restrictions, and therefore may not be able to treat a patient during that patient's entire stay in the hospital. Increased duty-hour restrictions also force a trend toward night shift work in inpatient hospital settings and create concerns regarding negative impacts on education of residents, schedule flexibility, and continuity of care (Fortuna et al. 2009).

Continuity of care can impact the quality of patient care. Epstein et al. (2010) observe that more fragmented care may lead to longer patient lengths of stay. Saultz and Albedaiwi (2004) show that continuity of care increases patient satisfaction. Rodriguez et al. (2010) investigate the impact of the length of time surgical residents spend on individual rotations on continuity of care, and find that rotations of only one month in duration, which are often the norm, are insufficient. In a follow-up study, Turner et al. (2012) find that simply increasing the length of surgical rotations beyond one month is not enough to improve continuity of care. The authors offer other suggestions for improving continuity, including an apprenticeship model that assigns residents to one or two supervising physicians for the duration of their assignments to a rotation.

Fragmentation of care caused by increased duty-hour restrictions means increased frequency of handoffs of patients from one physician's care to another, increasing the risk that communication errors will occur during handoff (Van Eaton et al. 2011, Willis et al. 2009). Such communication errors can place patients at risk. Flawed communication at handoff has become "widely recognized as a leading safety hazard in healthcare" (Landrigan 2007, p. 6). For example, Pickering et al. (2009) developed a tool for measuring the level of corrupted information shared at handoff. The quality of communication between physicians at handoff largely impacts patient care, and efforts can be made to improve communication and identify best practices (Arora et al. 2009, Riesenburg et al. 2009, Solet et al. 2005).

A team of researchers and physicians from Georgia Institute of Technology and Children's Healthcare of Atlanta at Egleston (Children's) collaboratively developed a mixed-integer programming model (MIP) that constructs physician schedules that maximize continuity and familiarity by utilizing an objective scoring method for measuring continuity at each handoff. In the *Problem Description and Model* section, we outline the physician scheduling problem faced by many institutions at which multiple physicians treat hospitalized patients, and we discuss the scoring method we developed for measuring continuity. We present the MIP we developed, and discuss our solution approach. In the *Results and Discussion* section, we provide results from an application of our methods at the pediatric intensive care unit (PICU) at Children's. We discuss the implementation of our methods in the *Implementation* section, and summarize our findings in *Conclusions*.

Problem Description and Model

Joint scheduling of specialty residents or subspecialty fellows (we use the term residents to refer to physicians in both groups) and other physicians (e.g., attendings) is a complex problem because of ACGME duty-hour restrictions and individual physician preferences. ACGME restrictions include the following. When averaged over four weeks, residents (1) may be scheduled a maximum of 80 duty hours per week, (2) must have one day per week free of duty, and (3) may be scheduled for in-house call (i.e., the resident must remain on-site) at most once every third night (ACGME 2011a). Residents may be scheduled for a maximum of 24 hours of continuous duty, with at most four additional hours to ensure an effective transfer of care. An eight-hour layoff is required between scheduled duty periods (10 hours is recommended), and residents must have at least 14 hours free of duty after 24 hours of in-house duty. In addition to these requirements, a feasible schedule must satisfy the expected patient demand for medical doctors (MDs); note that we will use the terms physician and MD interchangeably.

We investigate the problem of scheduling staff under such restrictions for the PICU at Children's. This PICU, which is part of the largest pediatric healthcare system in the country, is a 30-bed multidisciplinary medical-surgical quaternary care unit that cares for acutely ill patients. A typical daily schedule for this PICU includes two attendings and two fellows on service during the day, one attending and one fellow on call at night, up to four residents on service during the day, and two residents on call at night. According to the virtual PICU performance system (VPS), a national PICU database (VPS 2012), of 26 PICU's around the country that submitted data for at least one quarter in 2011, only seven (including Children's) have 24 or more beds, and only Children's schedules fellows to work night call shifts in-house. Thus, the scheduling problem faced at Children's is challenging because of its size and complexity, as compared to other PICUs.

When scheduling staff, each situation often involves unique institutional or individual preferences. For example, in the PICU at Children's, a service-block schedule structure has been implemented with the goal of creating continuity in the unit. Specifically, attending physicians are scheduled to be on service for one week at a time (i.e., seven-day service block) and prefer not to exceed that limit, and fellows are on service for overlapping 14-day periods (i.e., one fellow starts a 14-day service block each Monday). In addition, individual preferences with regard to specific days on or off service or night call are considered when constructing a MD schedule for the unit, a process that was previously performed manually. For previous schedules, manual construction of a six-month attending-only schedule, which the ACGME does not regulate highly, required several hours plus additional time to accommodate schedule changes on an ongoing basis.

In light of these scheduling constraints and preferences, we developed an efficient and effective decision support tool for assigning physicians to day and night shifts, such that physician and patient continuity is maximized, within the boundaries of hard feasibility constraints and soft physician preferences. To measure the continuity of a schedule, we developed the handoff continuity score (HCS).

Handoff Continuity Score

Intuitively, a physician who treats a patient for multiple consecutive days becomes familiar with that patient's progress and current state; returning to duty after a multiple-day break may require a readjustment period for the physician to again become familiar with the patient's condition. Therefore, given a physician duty schedule, we measure the continuity at each handoff based on two assumptions:

- A physician's familiarity with a patient increases with multiple (possibly successive) on-duty days on which the physician cares for that patient.
- A physician's familiarity decreases as the number of recent off-duty days increases.

To capture physician familiarity at handoff, we developed the familiarity factors reported in Table 1; Smalley et al. (2011) provide additional details. These familiarity factors are based on a five-day period (i.e., the average patient length of stay in the PICU at Children's) on which familiarity decreases (almost exponentially) as the number of days recently worked by the physician decreases. We developed these factors based on informal discussions and surveys at multiple meetings of the attending physician group at Children's. We considered linear familiarity factors (testing revealed similar results with linear factors); however, this attending physician group suggested that a linear relationship between the familiarity felt after having worked one up to five days prior to a handoff did not exist. Several formats for the familiarity factors were proposed and discussed, and we chose those we considered to be most appropriate for the physician group in this PICU. We also investigated the use of factors that consider periods shorter and longer than the average patient length of stay; a five-day period was identified as most appropriate because it is a good global marker for this patient population. Thus, although the factors in Table 1 are specific to the group at Children's, they are also flexible and can be adjusted for other units and institutions based on their perceptions of continuity.

Previous days worked	Familiarity factor
< 1 day ago	0.5
$[1, 2)$ days ago	0.25
$[2, 3)$ days ago	0.15
$[3, 4)$ days ago	0.075
$[4, 5)$ days ago	0.025

Table 1: Familiarity factors capture the familiarity physicians feel at handoff.

To determine the continuity of a schedule, we assign a continuity score between 0 and 1 to each oncoming physician (i.e., physician who is beginning a shift) at shift change. This score is equivalent to the summation of familiarity factors corresponding to each previous day worked by the physician. For example, if a physician is starting a shift after having worked two recent shifts, one two days prior and the other four days prior, that physician receives a score of 0.175 ($0.15 + 0.025$). A score of 1 implies the greatest familiarity. The score for each handoff is calculated as the average continuity score over all oncoming physicians; for a complete physician schedule, we average these handoff scores to determine the overall HCS.

A Mixed-Integer Programming Model

We developed a MIP for automated physician schedule generation. This model seeks to maximize the HCS (i.e., continuity) of the schedule, while conforming to feasibility constraints and satisfying MD preferences when possible.

Applying optimization methods to scheduling staff in a hospital setting is not a new concept. The problem of scheduling nurses to shifts in a hospital has been studied extensively; see Burke et al. (2004) and Cheang et al. (2003) for comprehensive reviews. Beaulieu et al. (2000) present a mathematical model for scheduling emergency room physicians to shifts over a six-month period. Sherali et al. (2002) use a mixed-integer program for scheduling residents to night shifts over a 4–5 week period. Rousseau et al. (2002) develop a flexible solution approach applying constraint programming, local search, and genetic algorithms to the physician scheduling problem, which various units and institutions face, with minimal customization. Topaloglu (2006) assigns emergency medicine residents to day and night shifts using goal programming, and in a later paper, applies sequential and weighted methods to a multiobjective optimization model for assigning residents to night call shifts, while considering levels of seniority (Topaloglu 2009). Ovchinnikov and Milner (2008) develop a user-friendly spreadsheet model to assign first- through fourth-year residents to night call and emergency rotation shifts in a radiology department, and Cohn et al. (2009) solve multiple nested integer programming (IP) models to assign 10–20 residents to various types of night call shifts (e.g., primary, backup) in three hospitals over a one-year period. Gunawan and Lau (2013) present mathematical models for solving a local hospital's physician scheduling problem, which includes assignment of physician tasks to time slots over a period of five days, with the goal of optimizing resource allocation. Brunner et al. (2010) present a branch-and-price algorithm for constructing daily physician schedules with flexible shift start times and lengths, for a scheduling horizon of up to six weeks, and Stolletz and Brunner (2012) use a reduced set-covering formulation for flexible shift scheduling of physicians, while considering fairness. Topaloglu and Ozkarahan (2011) combine mixed-integer programming and column-generation techniques to schedule resident duty hours for a four-week period within a matter of seconds. Rinder et al. (2012) provide a literature review of work that applies various industrial engineering methods to address physician scheduling problems, with a particular focus on patient impact on scheduling.

Other relevant scheduling problems addressed in the literature include airline-crew scheduling (Gopalakrishnan and Johnson 2005, Kohl and Karisch 2004). Similar to the physician scheduling problem that we address, crew scheduling involves extensive duty-hour restrictions, although continuity is less crucial. Ernst et al. (2004) compile a detailed list of previous work in personnel scheduling, including problem types and solution approaches.

The Children's PICU physician scheduling MIP (called CPPS-MIP in this paper) that we developed shares some characteristics with previous work. Turner (2011) develops optimization models for assigning surgical residents to individual patients to maximize continuity of care and education of residents by considering expected surgical cases for each resident, and Turner et al. (2013) present a heuristic-based system for assigning patients to physicians, with the goal of partially improving continuity of care. Kazemian et al. (2013) present an IP approach to shift scheduling, which seeks to minimize the number of patient handoffs over a four-week period in an intensive care unit based on the average patient census at shift change. The authors consider 12 possible times for shift changes each day and varying shift lengths. To the best of our knowledge, however, we are the first to consider the benefits of optimized MD shift scheduling on continuity beyond a single shift change. That is, our approach acknowledges that a prior familiarity may exist among oncoming physicians at shift change, thereby preserving some sense of continuity for the patients.

CPPS-MIP is general and can be easily applied to various units and institutions with different preferences, where each day is divided into two periods that start at the same time each day and do not overlap. This means, for example, that on any given day, a MD can be scheduled to be on service sometime between 8 AM and 4 PM or on call sometime between 4 PM and 8 AM the following day; however, different shift types are possible in each period (e.g., the 8 AM–4 PM or 8 AM–12 PM service shift).

Constraints incorporated in CPPS-MIP are listed in Appendix A. Some constraints in this model are applicable to a majority of institutions; see Equations (A1)–(A11). Others are more specific to the problem Children's faces because of physician preferences; see

Equations (A12)–(A20). Equations (A21)–(A23) calculate the HCS, and Equations (A23) and (A24) form the objective function.

The first set of constraints enforces the feasibility of the MD schedule. Each day, there is a demand for physicians during each period. The schedule must satisfy ACGME requirements, as we discuss in the *Problem Description and Model* section. Given that multiple shift types are available during each period, an additional feasibility constraint requires that no MD can be scheduled for overlapping shifts. Additional constraints are required if first-year residents are available for scheduling. Physicians in this group may work a maximum of 16 consecutive hours; therefore, they may not be assigned to duty shifts in two consecutive periods. Equations (A1)–(A11) in Appendix A represent these constraints.

As we mention previously, MDs in the PICU at Children's prefer a block structure to their service schedule; therefore, we added preference constraints in our MIP to ensure that a MD assigned to a service block is scheduled to work the appropriate shifts corresponding to that block. A predetermined number of MDs of each type must be assigned to each block, and MDs may not work overlapping blocks. Although assigning a MD to consecutive service blocks is possible, it generates a penalty. For additional details, see Equations (A12)–(A15).

Some constraints in our MIP are specific to the PICU at Children's, but may be modified for other units; these are represented by Equations (A16)–(A19) and include the following. Attending physicians prefer not to work two consecutive call shifts if they are on service in between those shifts, and over the scheduling horizon, the number of days on service for each MD should be close to average for each physician type. Of the three types of shifts, two are service shifts; shift 1 is 8 AM–4 PM; shift 2 is 8 AM–12 PM. Shift 3 (4 PM–8 AM) is a call shift. MDs in the PICU prefer that a fellow works shift 3 on weekend days only if that fellow was scheduled for shift 1, and fellows may only work shift 2 on any day if they were assigned to shifts 1 and 3 the previous day. Shift 2 is designed to allow time for effective transfer of care following a 24-hour shift.

When not scheduled in the PICU, fellows could be doing research or be assigned to a rotation in anesthesia (four weeks during the three-year fellowship), the cardiac intensive care unit (12 weeks), or an elective (four weeks). Fellows do not take night call during these rotations. Thus, assuming that these assignments to various rotations are made in advance of daily scheduling in the PICU, Equation (A20) specifies that fellows can only be assigned to service and call shifts in weeks they are available (i.e., are not assigned elsewhere).

The HCS is calculated based on previous days worked by oncoming MDs (Smalley et al. 2011). Specifically, each MD receives a continuity score between 0 and 1 (1 = most familiar) computed as the summation of a subset of familiarity factors (see Table 1) corresponding to each day worked in the previous five days. Each handoff score is based only on the previous days worked by oncoming MDs—not all MDs; therefore, we incorporate additional constraints to capture only the scores for oncoming MDs. The score for each handoff is the average continuity score over all oncoming MDs. The HCS for the entire schedule is calculated as the average of all handoff scores; see Equations (A21)–(A23).

To optimize continuity, our goal is to maximize the HCS. Because preferences of physicians in the PICU at Children's must be considered, a penalty is incurred if MDs are scheduled to consecutive service blocks, or if physician requests are violated; see Equation (A24).

Then our objective is: Maximize (HCS—Penalty). Note that because the penalty increases by increments of 1 and the HCS is between 0 and 1, the MIP will not seek to increase the HCS by increasing the penalty. That is, the penalty is determined exactly by physician requests; it exists to ensure that requests are granted if feasibly possible and are denied otherwise.

Solution Approach

Physicians at Children's provided us with a manually constructed schedule for the period from July 1, 2010 to June 30, 2011. The schedule included only attendings and fellows; therefore, we limit our solution approach discussion to these two groups. To illustrate the advantage of optimized MD scheduling with regard to physician and patient continuity, we attempted to use CPPS-MIP to generate MD schedules for these groups for the 51-week period beginning Monday, July 5, 2010 and ending Sunday, June 26, 2011. Note that we developed our model to accommodate time horizons that begin on a Monday and end on a Sunday, primarily because Children's prefers this service-block structure; therefore, we attempted to generate schedules for this slightly shortened period. These schedules required a significant amount of running time; using CPLEX 12.4, we found no feasible solution within 48 hours. Note that this and the remaining computational experiments reported in this chapter were performed on one of two systems: (1) 2.27 GHz Xeon quad-core processor and 48 GB RAM, or (2) 2.33 GHz Xeon quad-core processor and 12 GB RAM. The poor performance of CPLEX on CPPS-MIP motivated us to develop the following heuristic, which finds a feasible solution (assuming enough physicians are available for scheduling) by fixing some assignments and then solving this model with a reduced set of decision variables:

Heuristic: *Iterative schedule construction with a modified CPPS-MIP*:

- Step 1: Iteratively assign attending physicians to shifts on a week-by-week basis, ignoring requests for time off or on-duty time.
- Step 2a: Assign fellows to service blocks using a simple integer program, which we call the fellows' service block assignment integer program (FSBA-IP); see Appendix B for details.

This integer program fairly assigns fellows to service blocks (which is fair with respect to the total number of block assignments given to each fellow over the schedule horizon), disallowing assignments to overlapping or consecutive service blocks.

- Step 2b: Given the solution to FSBA-IP, we next assign fellows to call shifts: In each week, assign the Monday, Thursday, and Sunday night call shifts to the fellow who is on service duty during the first week of a two-week service block, and assign the Wednesday and Saturday night call shifts to the other fellow who is on service that week. Assign the Tuesday and Friday night call shifts to one of the fellows who is not on service (arbitrary). Note that we will consider fellows' requests for time off at night in a later step.
- Step 3: Run CPPS-MIP with fellows' shifts fixed, optimizing for attending physicians using the schedule found in Step 1 as a starting point.
- Step 4: Run CPPS-MIP again with fellows' service shifts and attending physicians' service and call shifts fixed to solution values from Step 2, optimizing for fellows' call shifts.

We used this heuristic to generate MD schedules for the 51-week period. An optimal solution to FSBA-IP in Step 2a can be found in less than one second. Step 3 requires running the model with some variables fixed. After running for two hours, the penalty is minimized and only marginal improvement in the HCS is found in up to 48 hours. Therefore, for Step 4, we use the best solution found after running Step 3 for two hours. An optimal solution to CPPS-MIP with variables fixed according to Step 4 requires only a few seconds for this instance.

Using the schedule found by the heuristic to warm start CPLEX, CPPS-MIP could not find an HCS-improved schedule after running for 48 hours, and CPLEX reported a 25 percent optimality gap based on the LP-relaxation upper bound. To close the gap, we need to improve either the heuristic solution, or the upper bound, or both. Proposition 1 (see Appendix C) helps us identify a better upper bound, and shows that the optimality gap of the heuristic solution is much smaller than the one reported by CPLEX.

Fellows in the PICU at Children's work two-week service blocks, which is equivalent to two consecutive one-week service blocks, and these blocks coincide with the start and end days of attending service blocks. Therefore, we can apply this proposition to both fellows and attendings. Thus, a 51-week schedule for attendings and fellows at Children's cannot achieve a higher HCS than 1-, 3-, and 17-week schedules.

Because of the format of the HCS calculation and the service block structure at Children's, the one service block (i.e., one-week period) leading up to the start of a schedule is considered when determining the HCS of that schedule. Therefore, given Proposition 1, we used CPPS-MIP to generate a four-week schedule, optimizing the HCS over the last three weeks of the schedule to determine the best possible three-week HCS. In the next section, we refer to this three-week schedule as our three-week test problem, which CPLEX solves in less than one minute.

	Manual schedule	Heuristic-generated schedule	Percentage increase*
HCS	0.6627	0.7919	19.50
HCS: service shifts only	0.7578	0.8070	6.49
HCS: night call shifts only	0.5676	0.7769	36.87
HCS: attending physicians only	0.5961	0.7980	33.86
HCS: fellows only	0.7293	0.7859	7.76
Call shifts without an on-service physician ($N = 357$)	15% (52)	1.4% (5)	

Table 2: We calculate the HCS for manual and heuristic-generated schedules constructed for July 5, 2010–June 26, 2011.

*All statistically significant improvements in HCS, $p < 0.001$ using a z-test.

Results and Discussion

Table 2 shows the HCS for the manual and heuristic-generated schedules for the 51-week period from July 5, 2010 through June 26, 2011. We also report the HCS for daytime service shifts, night call shifts, attendings, and fellows, respectively. For each score presented in Table 2, the heuristic-generated schedule achieves a statistically significant improvement over the score of the manually generated schedule, determined using a z-test (computed using the HCS at each handoff over the schedule horizon). The manual and heuristic-generated schedules each incorporate the preferred service-block structure; therefore, the HCS for service shifts shows relatively little improvement (6.49 percent). The largest increase in the HCS (36.87 percent) is attributed to the night call shifts. The HCS for attendings increased significantly more than the HCS for fellows. The 14-day service blocks for fellows provide concentrated clinical time, which is important for education, but also provides more continuity for daytime service shifts than the seven-day service blocks worked by attendings. Therefore, the manually generated schedule had a high HCS for fellows compared to attendings as a result of the service block structure. In addition, ACGME duty-hour restrictions provide rigorous constraints that limit much variability of night call shift assignments for fellows compared to attendings. Therefore, the heuristic could only find marginal (but still statistically significant) improvements in the HCS for fellows. Note that all physician requests were satisfied by manual and heuristic-generated schedules.

Problem description	Optimal solution	Heuristic solution	Optimality gap (%)
Three-week schedule	0.8190	—	<0.01
51-week schedule	—	0.7919*	3.42+

Table 3: We compare the heuristic-generated schedule to a three-week test problem.

***Solution does not improve after running the MIP for 48 hours.**

+Optimality gap based on Proposition 1.

In addition to the HCS by physician type and time of day, Table 2 also reports the percentage (and number) of night call shifts over the schedule horizon without an on-service physician assigned. Although not part of the MIP's objective, it is intuitive that a schedule that maximizes continuity would likely have very few night call shifts on which an on-service physician was not scheduled. As reported in Table 2, the heuristic-generated schedule includes only five night call shifts on which an on-service physician is not assigned, compared to 52 in the manually generated schedule.

The heuristic-generated schedule requires no special considerations when scheduling for the December holidays, but maintains the service-schedule structure adopted for the rest of the year. Entering the manually generated holiday schedule (December 24–January 2) into the heuristic-generated schedule, the HCS for this new schedule is 0.7826, still an 18.09 percent improvement over the manual schedule.

Given that a 51-week block schedule cannot achieve a higher HCS than a three-week block schedule based on Proposition 1, we conclude that the HCS for the heuristic-generated schedule is within 3.42 percent of optimality. Table 3 shows a comparison of the three-week test problem and the 51-week heuristic solution.

Implementation

2011: Challenges and Results

We generated an attending-only schedule for the PICU at Children's, to be implemented from July 1, 2011 to December 23, 2011. Without needing to schedule fellows and with the shorter horizon, we used CPPS-MIP to create this schedule. Modifying the MIP and deciding on an appropriate schedule required many iterations; in each iteration, we presented a

Model version	Description	HCS
1	Call structure: On-service attendings on call either Monday or Thursday, then either Saturday or Sunday	0.70
2	Model version 1, plus the following constraints and (or) changes: (1) no consecutive call shifts, and (2) requests take priority over seven-day service blocks	0.64
3	Call structure: On-service attendings on call either Monday, Friday, and Sunday, or Thursday and Saturday, plus the following constraints and (or) changes: (1) no consecutive call shifts, and (2) requests take priority over seven-day service blocks	0.68
4	Model version 3, plus the following change: requests take priority over call structure	0.69

Table 4: To find an acceptable schedule, our model required revisions. In this table, we provide descriptions and the HCS for the modified MIP versions used to generate schedules at each iteration.

MIP-generated schedule to one or more physicians in the group, who identified issues that required the addition of new constraints in the MIP; see additional constraint definitions in Appendix D. CPPS-MIP could generate attending-only schedules for the six-month period in approximately two hours (no additional improvement in the HCS was found within 48 hours).

The heuristic-generated schedule reported in Table 2 assigned attendings to alternating night call shifts in the weeks they are on service. Although this may be best for continuity, such a schedule would create fatigue; therefore, it was an unattractive solution. One standard practice when assigning attendings to call shifts in the PICU was to assign on-service attendings to work call either Monday or Thursday, as well as call on either Saturday or Sunday. These MDs were then not assigned to any other call shifts during the week. With constraints enforcing this preference, the optimal schedule assigned one off-service attending to the call shifts on Tuesday, Wednesday, and Friday. This option was undesirable because of the requirement for consecutive night call shifts. Therefore, we decided that no attending should be assigned to take call two nights in a row.

As an alternative to scheduling on-service attendings to call either Monday or Thursday and then once over the weekend, the physicians thought that scheduling one on-service attending to call shifts on Thursday and Saturday and the other on-service

	2011		2012		
Attendings	9		11		Percentage increase (per person)
	Requests	Average	Requests	Average	
On service	104	11.56	134	12.18	5
On call	22	2.44	34	3.09	26
Off service	217	24.11	723	65.73	173
Off call	212	23.56	687	62.45	165

Table 5: The number of attending physician requests varied from 2011 to 2012.

attending to call shifts on Monday, Friday, and Sunday would be an acceptable assignment and good for continuity. In some instances in both call models, physicians' requests were denied to maintain these preferred night call structures.

One physician in this group routinely requested to work call on Mondays and Saturdays and requested to be on service during the same week. This violates the alternative call structure. We decided that in weeks in which this is the case, the other on-service attending physician should be assigned to call on Thursday and Sunday.

Each physician in the group has a specific number of day and night shifts Monday–Friday and weekend day and night shifts that should be worked to satisfy various fellowship and other requirements. These requirements do not always coalesce well with the seven-day service blocks; however, these requirements and physician requests should take priority over the seven-day service blocks and preferred-call structure. Thus, we altered the constraints enforcing the seven-day service blocks and the preferred call structure to be soft constraints, by which we gave requests precedence in the objective function. Table 4 provides a description of the MIP versions used to generate a schedule at each iteration and the HCS for each schedule.

Children's implemented the final MIP-generated schedule, which is based on Model version 4 (see Table 4), with some minor manual modifications. For comparison, a physician at Children's, who has experience in scheduling and also has the goal of maximizing continuity, manually constructed an additional schedule for the same period; its result was an HCS of 0.63.

We used the heuristic with these new constraints and (or) modifications to CPPS-MIP to generate an additional schedule corresponding to the 51-week period between July 5, 2010 and June 26, 2011. The HCS for this new schedule is 0.7578, a 14.35 percent improvement over the HCS for the manually generated schedule constructed for the same period.

2012: Challenges and Results

Following the positive reception for the 2011 schedule, we generated an attending-only schedule for the PICU for July 1–December 23, 2012; however, we did not generate a parallel manual schedule. The pool of attendings available for scheduling increased from nine to 11 for this period. Using CPPS-MIP, we found a schedule (optimality gap < 1 percent) in approximately 35 minutes. This improvement in running time was largely the result of the extensive personal requests made by attendings for this period that limited possible variations to the schedule for purposes of HCS improvement. Table 5 reports the number of requests made for 2011 and 2012 for the same periods (i.e., July 1–December 23). Requests to be off service or off call increased per person on average by 173 percent and 165 percent, respectively.

This large increase in requests can be attributed partly to a format change for entering these requests. For the 2011 schedule, attendings made requests by month, with no predetermined format; Figure 1 shows an example of such requests. For the 2012 schedule, as a means of structuring how requests were entered, we created an Excel spreadsheet for entering requests for each individual day in the schedule horizon. We gave this spreadsheet to the attendings, and they used a slight variation of this form to communicate their requests (see Figure 2).

Further study is needed to better understand the reason for this large increase in requests for this period—whether it is the result of the change in request formats or for reasons more unique to each physician; however, automated schedule generation is clearly more efficient if a greater number of requests are made. Unfortunately, physician and patient continuity may suffer. When we remove all physician requests for this period, the HCS increases by 2.4 percent, which can be significant in terms of continuity. The number of night call shifts without an on-service physician scheduled decreases by 7.41 percent. If we randomly ignore 50 percent of the requests, the HCS increases by up to 1.24 percent, and the number of night call shifts without an on-service physician scheduled decreases by up to 5.56 percent (calculated using 10 instances, seven of which did not improve on the HCS with all requests considered within two hours of running time). For each of these cases, however, a large number of requests are denied, as expected.

Month	Physician							
	1	2	3	4	5	6	7	8	9
July	Off until July 2, Can do 7/4 week			No call 1-3, 14-31; Want service 4-10 with call 4, 9, 11	1st, 2nd, 3rd weeks off for vacation	No July 4 weekend; no service 18-22	25th is good week for service		No call 6; No service week of 6th
August	Can do Aug. 1 week; prefer week of 15th		18/19 no call; 20-21 off; 31 no service	No call 2-7, 15-28; Want service 8-14 & 29-9/2 with call 2,8,13,29		No call 2, 31; No service 5, 10, 19, 22-23, 31	No service or call 1-14		

Figure 1: This physician request entry form used in 2011 illustrates how physicians communicated their scheduling preferences.

Date	Service 1	Call 1	Service 2	Call 2	Service 3	Call 3	Service 4	Call 4	Service 5	Call 5	Service 6	Call 6	Service 7	Call 7	Service 8	Call 8
7/1/12									no	no call			no	no call		
7/2/12									want		want	want call	no		no	no call
7/3/12					no	no call			want		want		no		no	no call
7/4/12	no	no call			no	no call			want		want		no		no	no call
7/5/12					no	no call			want		want		no		no	no call
7/6/12					no	no call			want	no call	want		no		no	no call
7/7/12					no	no call	no	no call	want			want call	no	no call	no	no call
7/8/12					no	no call	no	no call	want				no	no call	no	no call
7/9/12	no	no call			no	no call	no	no call			no	no call	want	want call		
7/10/12	no						no	no call			no	no call	want			
7/11/12	no						no	no call	no	no call	no	no call	want		no	no call
7/12/12	no						no	no call			no	no call	want		no	no call
7/13/12	no						no	no call		no call	no	no call	want	want call	no	
7/14/12							no	no call	no	no call	no	no call	want			
7/15/12									no	no call	no	no call	want	want call		
7/16/12			no	no call			want		no	no call	no	no call				
7/17/12			no	no call			want		no	no call	no	no call		want call		
7/18/12			no	no call			want		no	no call	no	no call	no		no	no call
7/19/12		no call	no	no call			want		no	no call	no	no call			no	no call
7/20/12			no	no call	no	no call	want		no	no call	no	no call	no		no	no call
7/21/12					no	no call	want		no	no call	no	no call			no	no call
7/22/12					no	no call	want		no	no call	no	no call			no	no call
7/23/12	no	no call									no	no call	no			
7/24/12				no call				no call		no call	no	no call	no	want call		
7/25/12			no	no call			no		no		no	no call	no	no call		no call
7/26/12											no	no call	no			no call
7/27/12			no	no call						no call	no	no call	no			
7/28/12							no	no call	no	no call	no	no call				
7/29/12							no	no call	no	no call	no	no call				
7/30/12	no	no call	no	no call			no	no call			no	no call	no		no	no call
7/31/12	no	no call					no	no call			no	no call	no	want call	no	no call

Figure 2: We developed a structured Excel spreadsheet that physicians used to communicate their scheduling requests in 2012.

Conclusions

Physician preferences and increased duty-hour restrictions create a complex scheduling problem when attempting to satisfy all requirements in a manually generated schedule. The solution approach we present in this paper facilitates resource optimization; it constructs a feasible schedule in significantly less time than is needed to create a schedule manually. Further, by considering familiarity among oncoming physicians, the schedules produced maximize continuity. In conjunction with other methods for improving communication, schedules with greater continuity have the potential to enhance handoff efficiency and effectiveness.

The HCS provides a means of understanding the continuity of a physician schedule. With CPPS-MIP and the heuristic, it allows the construction of schedules that can improve handoff efficiency and effectiveness; however, the aforementioned analysis does not consider many factors (e.g., bed occupancy, new admissions, disease acuity, and fatigue) that can impact the handoff process. Physician fatigue cannot be eliminated completely, regardless of the stringency of duty-hour restrictions; therefore, communication failures created by exhaustion may still occur. Future work includes expanding the HCS to develop a score that considers both familiarity among oncoming physicians and fatigue of physicians signing off their shifts.

Although counterintuitive, one can argue that too much familiarity among oncoming physicians could negatively impact the handoff process. Increased familiarity may cause an oncoming MD to pay less attention during handoff, potentially missing important information about a patient's care. Therefore, optimized scheduling to maximize familiarity among oncoming physicians can only improve handoff efficiency if it is in combination with other steps for improving communication at handoff (e.g., a sign-out checklist).

Currently, the HCS is measured at each shift change, and no direct measures are in place to completely understand the impact that the HCS has on patient outcome measures, including mortality rates, complication rates, readmission rates, and lengths of stay. A simple analysis of rounding times recorded between 2008 and 2013 point to a trend toward reduced rounding times with increased HCS; unfortunately, attributing this reduction entirely to increased HCS is difficult. Although many other factors also impact rounding times, including patient acuity and day of the week, this downward trend is encouraging. Identifying measures to capture the impact of a single shift change or combination of shift changes on patient outcomes is the next step to a better understanding of the benefits of the HCS.

Moving forward, we will continue to utilize CPPS-MIP to assist with schedule construction in the PICU at Children's. We will adjust it as physician preferences change over time and (or) duty-hour restrictions change. Our next steps include developing a tool for physicians to use at the operational level to facilitate schedule construction by physicians in the PICU; this tool will allow a physician to create what-if scenarios to better understand the impacts of schedule changes on continuity.

Appendix A. Children's PICU Physician Scheduling MIP (CPPS-MIP)

Tables A.1, A.2, and A.3 show the notation we use in CPPS-MIP.

$$\sum_{k\in K_l}\sum_{i\in C} X_{ijk} = v_{ljP} \quad \forall j\in J,\ l\in L,\ P\in \wp \tag{A1}$$

$$\sum_{l\in L}\sum_{k\in K_l}\sum_{j\in\{m,m+1,\ldots,m+27\}} \bar{L}_k X_{ijk}/4 \le 80 \quad \forall i\in R,\ m\in\{1,2,\ldots,N-27\} \tag{A2}$$

$$\begin{aligned} &X_{ijk^1}+X_{ijk^2}+X_{i(j+1)k^3}\le 2 \\ &\quad \forall i\in R,\ (k^1,k^2,k^3)\in \hat{C}_1,\ j\in\{1,2,\ldots,N-1\} \\ &X_{ijk^1}+X_{i(j+1)k^2}+X_{i(j+1)k^3}\le 2 \\ &\quad \forall i\in R,\ (k^1,k^2,k^3)\in \hat{C}_2,\ j\in\{1,2,\ldots,N-1\} \end{aligned} \tag{A3}$$

$$\begin{aligned} &X_{ijk^1}+X_{ijk^2}\le 1 \quad \forall i\in R,\ j\in J,\ (k^1,k^2)\in T \\ &X_{ijk^1}+X_{i(j+1)k^2}-\sum_{k\in K_2} X_{ijk}\le 1 \\ &\quad \forall i\in R,\ (k^1,k^2)\in T_1,\ j\in\{1,2,\ldots,N-1\} \\ &X_{ijk^1}+X_{i(j+1)k^2}-\sum_{k\in K_1} X_{i(j+1)k}\le 1 \\ &\quad \forall i\in R,\ (k^1,k^2)\in T_2,\ j\in\{1,2,\ldots,N-1\} \end{aligned} \tag{A4}$$

Set	Description
A, R	Sets of attending physicians and residents, respectively.
I	Set of all physicians, $I = A\cup R$.
R_t	Sets of all residents of type t (e.g., $R_1 = $ 1st year residents, $\ldots$, $R_n = $ fellows), $R = \bigcup_{i\in\{1,2,\ldots,n\}} R_i$.
J	Set of days in the planning horizon, $J = \{1,2,\ldots,N\}$.
L	Set of scheduling periods in each day = {1, 2}, 1: day service, 2: night call.
K_l	Set of shift types for time period l in L.
W	Set of weeks in the planning horizon, $W = \{1,2,\ldots,N/7\}$.
C_1	Set of doubles (k^1,k^2), such that $k^1\in K_1$, $k^2\in K_2$, and working shifts k^1 and k^2, which start on the same day requires 24+ consecutive hours on duty.
C_2	Set of doubles (k^1,k^2), such that $k^1\in K_2$, $k^2\in K_1$, and working shift k^1 on some day j and shift k^2 on day $j+1$ requires 24+ consecutive hours on duty.
$\hat{C}_1$	Set of triples (k^1,k^2,k^3), such that k^1,$k^3\in K_1$, $k^2\in K_2$, and working shifts k^1 and k^2 on some day j, and then shift k^3 on day $j+1$ would require 29+ consecutive hours on duty.
$\hat{C}_2$	Set of triples (k^1,k^2,k^3), such that k^1,$k^3\in K_2$, $k^2\in K_1$, and working shift k^1 on some day j, and then shifts k^2 and k^3 on day $j+1$ would require 29+ consecutive hours on duty.
T	Set of doubles (k^1,k^2), such that $k^1\in K_1$, $k^2\in K_2$, and on any given day, there are less than 10 hours between the end of k^1 and the start of k^2.
T_i	Set of doubles (k^1,k^2), such that k^1,$k^2\in K_i$ and there are less than 10 hours between the end of k^1 on any given day and the start of k^2 on the following day, $i\in\{1,2\}$.
$\hat{T}_1$	Set of doubles (k^2,k^3) such that $k^2\in K_2$, $k^3\in K_1$, and there are less than 14 hours between the end of k^2 on any given day and the start of k^3 on the following day.
$\hat{T}_2$	Set of doubles (k^1,k^2) such that $k^1\in K_1$, $k^2\in K_2$, and there are less than 14 hours between the end of k^1 and the start of k^2 on the same day.
$\wp$	Set of physician groups (e.g., $\wp = \{A,R,R_1\cup R_2,\ldots\}$).
B_P	Set of possible service block lengths (in days) for physicians in group P, $P\in\wp$ (e.g., $B_A = \{7\}$).

Table A.1: The table presents the sets (and descriptions) used in CPPS-MIP.

Decision variable	Description
X_{ijk}	1 if physician i works shift k on day j; 0 otherwise; $i\in I$, $j\in J$, $k\in\bigcup_{l\in\{1,2\}} K_l$.
Y_{ijn}	1 if physician i is assigned to a service block of length n days starting on day j; 0 otherwise; $i\in I$, $j\in J$, $n\in\{1,2,\ldots,N-j+1\}$.
D_{ijl}	1 if physician i works shifts in two consecutive periods beginning with period l on day j; 0 otherwise; $i\in I$, $j\in J$, $l\in\{1,2\}$.
R_{ijl}	Continuity score for physician i at start of period l on day j; $i\in I$, $j\in J$, $l\in\{1,2\}$.
$\hat{R}_{ijl}$	R_{ijl} if physician i works during period l on day j; 0 otherwise; $i\in I$, $j\in J$, $l\in\{1,2\}$.
Ω_{ij}	1 if physician i is assigned to consecutive service blocks beginning with a block starting on day j; 0 otherwise; $i\in I$, $j\in J$.

Table A.2: The table presents the decision variables (and descriptions) used in CPPS-MIP.

Interfaces 45(2), pp. 133–148, © 2015 INFORMS

Parameter	Description
N	Length of scheduling horizon in days.
$\hat{N}$	Number of previous days to consider when determining HCS (e.g., $\hat{N} = 5 =$ average patient length of stay in the PICU at Children's).
v_{ljP}	Demand for physicians in group P during period l on day j; $P \in \wp$, $j \in J$, $l \in \{1, 2\}$.
$\bar{L}_k$	Length (in hours) of shift k; $k \in \bigcup_{l \in \{1,2\}} K_l$.
Q_{ir}	1 if resident i is available for scheduling in the PICU during week r; 0 otherwise; $i \in I$, $r \in W$.
e_{Pjn}	Number of physicians in group P who must be scheduled to a block of length n, which starts on day j; $P \in \wp$, $j \in J$, $n \in \{1, 2, \ldots, N - j + 1\}$.
F_m	Familiarity factor corresponding to a shift worked $[m-1, m)$ days ago, $m \in \{1, 2, \ldots, \hat{N}\}$.
Z_P	Maximum allowable deviation from the number of service shifts a physician in group P is scheduled from the amount that physician would work if all service shifts were assigned evenly among that group; $P \in \wp$;
O_{ijl}	0 if physician i requested not to work during period l on day j; 1 otherwise; $i \in I$, $j \in J$, $l \in \{1, 2\}$.
U_{ijl}	1 if physician i requested to work during period l on day j; 0 otherwise; $i \in I$, $j \in J$, $l \in \{1, 2\}$.

Table A.3: The table presents the parameters (and descriptions) used in CPPS-MIP.

$$X_{ijk^1} + X_{ijk^2} + X_{i(j+1)k^3} \le 2 \quad \forall i \in R,$$
$$j \in \{1, 2, \ldots, N-1\},\ (k^1, k^2) \in C_1,\ (k^2, k^3) \in \hat{T}_1$$
$$X_{ijk^1} + X_{i(j+1)k^2} + X_{i(j+1)k^3} \le 2 \quad \forall i \in R,$$
$$j \in \{1, 2, \ldots, N-1\},\ (k^1, k^2) \in C_2,\ (k^2, k^3) \in \hat{T}_2 \tag{A5}$$

$$X_{ijk} - D_{ij1} \le 0 \quad \forall i \in I,\ j \in J,\ k \in K_l,\ l \in L$$
$$X_{ijk} - D_{ij2} \le 0 \quad \forall i \in I,\ j \in J,\ k \in K_2$$
$$X_{i(j+1)k} - D_{ij2} \le 0$$
$$\forall i \in I,\ j \in \{1, 2, \ldots, N-1\},\ k \in K_1 \tag{A6}$$

$$D_{ij1} - \sum_{l \in L} \sum_{k \in K_l} X_{ijk} \le 0 \quad \forall i \in I,\ j \in J$$
$$D_{ij2} - \sum_{k \in K_2} X_{ijk} - \sum_{k \in K_1} X_{i(j+1)k} \le 0$$
$$\forall i \in I,\ j \in \{1, 2, \ldots, N-1\} \tag{A7}$$

$$\sum_{j \in \{m, m+1, \ldots, m+27\}} D_{ij1}/4 \le 6$$
$$\forall i \in R,\ m \in \{1, 2, \ldots, N-27\} \tag{A8}$$

$$\sum_{k \in K_2} X_{ijk} + X_{i(j+1)k} + X_{i(j+2)k} \le 1$$
$$\forall i \in R,\ j \in \{1, 2, \ldots, N-2\} \tag{A9}$$

$$\sum_{k \in K_l} X_{ijk} \le 1 \quad \forall i \in I,\ j \in J,\ l \in L \tag{A10}$$

$$\sum_{k \in K_1} X_{ijk} + \sum_{k \in K_2} X_{ijk} \le 1 \quad \forall i \in R_1,\ j \in J$$
$$\sum_{k \in K_2} X_{ijk} + \sum_{k \in K_1} X_{i(j+1)k} \le 1$$
$$\forall i \in R_1,\ j \in \{1, 2, \ldots, N-1\} \tag{A11}$$

$$\sum_{k \in K_1} X_{isk} \ge Y_{ijn} \quad \forall i \in P,\ s \in \{j, j+1, \ldots, j+n-1\},$$
$$j \in \{1, 2, \ldots, N-n+1\},\ n \in B_P,\ P \in \wp \tag{A12}$$

$$\sum_{i \in P} Y_{ijn} \ge e_{Pjn}$$
$$\forall j \in \{1, 2, \ldots, N-n+1\},\ n \in B_P,\ P \in \wp \tag{A13}$$

$$Y_{ijn} + Y_{ikr} \le 1 \quad \forall i \in P,\ n \in B_P,\ k \in \{j, j+1, \ldots, j+n-1\},$$
$$r \in \{1, 2, \ldots, N-k+1\},\ j \in \{1, 2, \ldots, N-n+1\},$$
$$(j, n) \ne (k, r),\ P \in \wp \tag{A14}$$

$$Y_{ijn} + Y_{i(j+n)r} \le 1 + \Omega_{ij} \quad \forall i \in P,\ n \in B_P,\ P \in \wp,$$
$$j \in \{1, 2, \ldots, N-n+1\},\ r \in \{1, 2, \ldots, N-j-n+1\} \tag{A15}$$

$$\sum_{k \in K_2} (X_{ijk} + X_{i(j+1)k}) + X_{i(j+1)l} \le 2 \quad \forall i \in A,\ k \in K_1,$$
$$j \in \{1, 2, \ldots, N-1\} \tag{A16}$$

$$\sum_{k \in K_1} \sum_{j \in J} X_{ijk} - \left(\sum_{j \in J} v_{1jP}/|P| \right) \ge Z_P \quad \forall i \in P,\ P \in \wp \tag{A17}$$

$$X_{ij3} \le X_{ij1} \quad \forall i \in R_n,\ j \in \{r-1, r\},\ r \in \{7, 14, \ldots, N\} \tag{A18}$$

$$\cdot(X_{i(j-1)1} + X_{i(j-1)3}) \ge X_{ij2} \quad \forall i \in R_n,\ j \in \{2, 3, \ldots, N\} \tag{A19}$$

$$X_{ijk} \le Q_{ir} \quad \forall i \in R_n,\ j \in \{7r-6, 7r-5, \ldots, 7r\},$$
$$r \in W,\ k \in K_l,\ l \in L \tag{A20}$$

$$\sum_{m \in \{1, 2, \ldots, \hat{N}\}} F_m D_{i(j-m)l} = R_{ijl}$$
$$\forall i \in I,\ l \in L,\ j \in \{\hat{N}+1, \hat{N}+2, \ldots, N\} \tag{A21}$$

$$\sum_{k \in K_l} X_{ijk} \ge \hat{R}_{ijl}$$
$$\forall i \in I,\ l \in L,\ j \in \{\hat{N}+1, \hat{N}+2, \ldots, N\}$$
$$R_{ijl} + \sum_{k \in K_l} X_{ijk} - 1 \le \hat{R}_{ijl}$$
$$\forall i \in I,\ l \in L,\ j \in \{\hat{N}+1, \hat{N}+2, \ldots, N\} \tag{A22}$$
$$R_{ijl} \ge \hat{R}_{ijl}$$
$$\forall i \in I,\ l \in L,\ j \in \{\hat{N}+1, \hat{N}+2, \ldots, N\}$$

$$\textbf{HCS} = \left(\sum_{j \in \{\hat{N}+1, \ldots, N\}} \sum_{l \in L} \left[\sum_{i \in I} \hat{R}_{ijl} \Big/ \sum_{P \in \{A, R\}} v_{ljP} \right] \right) \Big/ (2N - 2\hat{N}) \tag{A23}$$

$$\textbf{Penalty} = \sum_{i \in I} \sum_{j \in J} \Omega_{ij} + \sum_{i \in I} \sum_{j \in J} \sum_{l \in L} [(1 - O_{ijl}) \cdot X_{ijl} + U_{ijl} \cdot X_{ijl}] \tag{A24}$$

Set	Description
F	Set of fellows.
W	Set of weeks in schedule horizon.
Parameter	**Description**
G_{ij}	1 if fellow i is available to work a two-week service block beginning in week j; 0 otherwise; $i \in F$, $j \in W$.
Decision variable	**Description**
Y_{ij}	1 if fellow i is assigned to a two-week service block beginning in week j; 0 otherwise; $i \in F$, $j \in W$.
d	Maximum number of service blocks assigned to any fellow.

Table B.1: The table presents the sets, parameters, and decision variables used in FSBA-IP.

Appendix B. Fellows' Service Block Assignment Integer Program (FSBA-IP)

Table B.1 shows the notation used in FSBA-IP.

Fellows' Service Block Assignment Integer Program (FSBA-IP):

$$\text{Minimize} \quad d$$

$$\text{s.t.} \quad \sum_{j\in\{1,2,\ldots,W\}} Y_{ij} \le d \quad \forall i \in F$$

$$\sum_{i\in F} Y_{ij} = 1 \quad \forall j \in \{1,2,\ldots,W\}$$

$$Y_{ij} + Y_{i(j+1)} \le 1 \quad \forall i \in F, \forall j \in \{1,2,\ldots,W-1\}$$

$$Y_{ij} + Y_{i(j+2)} \le 1 \quad \forall i \in F, \forall j \in \{1,2,\ldots,W-2\}$$

$$Y_{ij} \le G_{ij} \quad \forall i \in F, \forall j \in \{1,2,\ldots,W\}$$

Appendix C. Proposition 1 and Proof

Proposition 1. *Let H be the HCS of an optimal N-block schedule (equal-length, consecutive, and nonoverlapping blocks), and let $H^1 =$ the best HCS of any k-block schedule, $k \le n$ and $N \bmod k = 0$. Then $H^1 \ge H$.*

Proof. Let $H^1 =$ the best HCS of any k-block schedule. Let Δ be an optimal N-block schedule ($N \bmod k = 0$). Then Δ can be broken up into N/k k-block segments. Let $H =$ the HCS of Δ. Let $X_1, X_2, \ldots, X_{N/k}$ equal the summation of all handoff scores in block segments $1, 2, \ldots, N/k$, respectively, for the optimal N-block schedule. Then

$$H = \left(\sum_{i\in\{1,2,\ldots,N/k\}} X_i \right) \Big/ (2NB) \quad \text{(P1)}$$

where $B =$ block length in days. Because $H^1 =$ the best HCS of any k-block schedule, each X_i, i in $\{1, 2, \ldots, N/k\}$, cannot exceed H^1 when divided by the number of days and handoff periods in that k-block period. That is,

$$H^1 \ge X_i/(2kB) \quad \forall i \in \{1, 2, \ldots, N/k\}. \quad \text{(P2)}$$

Calculating the summation of both sides of P2 over all i and then taking the average, we have the following inequality:

$$H^1 \ge \left(\sum_{i\in\{1,2,\ldots,N/k\}} X_i \right) \Big/ (N/k \cdot 2kB)$$

$$= \left(\sum_{i\in\{1,2,\ldots,N/k\}} X_i \right) \Big/ (2NB) = H. \quad \text{(P3)}$$

Thus, by (P3), $H^1 \ge H$. □

We can generalize Proposition 1 to instances without the service block structure.

Corollary 1. *Let H be the HCS of an optimal N-week schedule and let $H^1 =$ the best HCS of any k-week schedule, $k \le n$ and $N \bmod k = 0$. Then $H^1 \ge H$.*

The proof for this corollary is identical to that of Proposition 1, with the exception of the word "block" replaced with "week" wherever it appears, and block length "B" is replaced by "7." We can also prove a more general version of this corollary by replacing "week" with "day."

Appendix D. Children's PICU Physician Scheduling MIP (CPPS-MIP): Additional Constraint Definitions

We added additional constraints to CPPS-MIP during the creation of an attending-only schedule for the PICU for July–December 2011. These constraints, as well as additional set, parameter, and decision variable definitions (Table D.1), are presented next. Note that in the set definitions, we refer to a physician M. This physician routinely requested to be assigned to specific night call shifts that did not align with the preferred call structures of other physicians in the group at Children's. Thus, we define sets related to this physician's requests for clarity of constraint definitions.

We modified Equation (A12) in Appendix A to be soft constraints, to allow the possibility that attendings not be assigned to seven consecutive days if necessary to satisfy requests and (or) fellowship requirements:

$$Y_{ijn} - \sum_{k\in K_1} X_{isk} \le \Phi_{1is} \quad \forall i \in A,\ s \in \{j, j+1, \ldots, j+n-1\},$$

$$j \in \{1, 2, \ldots, N-n+1\},\ n \in B_p,\ P \in \wp. \quad (\overline{\text{A12}})$$

Attendings should not be assigned to night call two days in a row. This is a modification of Equation (A16) in Appendix A:

$$\sum_{k\in K_2} (X_{ijk} + X_{i(j+1)k}) \le 1 \quad \forall i \in A,\ j \in \{1, 2, \ldots, N-1\}. \quad (\overline{\text{A16}})$$

Each attending should be assigned to no more than three night call shifts per week.

$$\sum_{k\in K_2} \sum_{r\in\{0,1,\ldots,6\}} X_{i(j+r)k} \le 3 \quad \forall i \in A,\ j \in \bar{J} \quad \text{(D1)}$$

Set	Description
$\bar{J}$	Set of Mondays in schedule horizon.
$\tilde{J}, \hat{J}$	Set of Mondays in schedule horizon for which physician M requested or did not request to be assigned to night call shifts on Monday and Saturday, respectively.
$\bar{A}$	Set of attendings, excluding physician M.
Parameter	**Description**
Ψ_{ti}	Required weekday service ($t=1$), call ($t=2$), and weekend service ($t=3$) shifts for physician i over the schedule horizon, $i \in A$.
λ_p	Weight on penalty p in the objective function, $p \in \{1, 2, \ldots, 8\}$.
Decision variable	**Description**
Φ_{pij}	Binary variable = penalty p incurred for physician i in week beginning with day j; $i \in A$, $j \in \bar{J}$, $p \in \{1, 2, \ldots, 8\}$.

Table D.1: The table presents the additional sets, parameters, and decision variables used in CPPS-MIP.

Attendings should be assigned to at most two night calls of Tuesday, Thursday, and Saturday, and Tuesday, Thursday, and Sunday, respectively.

$$\sum_{k\in K_2}(X_{i(j+1)k}+X_{i(j+3)k}+X_{i(j+5)k})\le 2 \quad \forall i\in A,\ j\in\bar{J} \tag{D2}$$

$$\sum_{k\in K_2}(X_{i(j+1)k}+X_{i(j+3)k}+X_{i(j+6)k})\le 2 \quad \forall i\in A,\ j\in\bar{J} \tag{D3}$$

Monday and Thursday night call shifts should be assigned to on-service physicians if possible. We assume that a physician is an on-service physician for the week if that physician works the service shift on Monday.

$$\sum_{k\in K_2} X_{ijk} - \sum_{k\in K_1} X_{ijk} \le \Phi_{2ij} \quad \forall i\in A,\ j\in\bar{J} \tag{D4}$$

$$\sum_{k\in K_2} X_{i(j+3)k} - \sum_{k\in K_1} X_{ijk} \le \Phi_{3ij} \quad \forall i\in A,\ j\in\bar{J} \tag{D5}$$

A physician who is one of the on-service physicians for the week should be assigned to at least one weekend night call shift.

$$\sum_{k\in K_1} X_{ijk} - \sum_{k\in K_2}(X_{i(j+5)k}+X_{i(j+6)k}) \le \Phi_{4ij} \quad \forall i\in A,\ j\in\bar{J} \tag{D6}$$

The same attendings should work night call on Monday, Friday, and Sunday, and Thursday and Saturday, respectively, in weeks physician M did not request to be on call on Monday and Saturday.

$$2\cdot\sum_{k\in K_2} X_{ijk} - \sum_{k\in K_2}(X_{i(j+4)k}+X_{i(j+6)k}) \le \Phi_{5ij}+\Phi_{6ij} \quad \forall i\in A,\ j\in\hat{J} \tag{D7}$$

$$\sum_{k\in K_2}(X_{i(j+3)k}-X_{i(j+5)k}) \le \Phi_{7ij} \quad \forall i\in A,\ j\in\hat{J} \tag{D8}$$

In weeks physician M did request to be on call on Monday and Saturday, the same attending should work night call on Thursday and Sunday.

$$\sum_{k\in K_2}(X_{i(j+3)k}-X_{i(j+6)k}) \le \Phi_{8ij} \quad \forall i\in\bar{A},\ j\in\tilde{J} \tag{D9}$$

Fellowship requirements should be met within 90 percent for each physician. These include the number of weekday, night call, and weekend day shifts assigned over the schedule horizon.

$$\sum_{j\in\bar{J}}\sum_{r\in\{0,1,\ldots,4\}}\sum_{k\in K_1} X_{i(j+r)k} \ge 0.9\cdot\Psi_{1i} \quad \forall i\in A \tag{D10}$$

$$\sum_{j\in J}\sum_{k\in K_2} X_{ijk} \ge 0.9\cdot\Psi_{2i} \quad \forall i\in A \tag{D11}$$

$$\sum_{j\in\bar{J}}\sum_{r\in\{5,6\}}\sum_{k\in K_1} X_{i(j+r)k} \ge 0.9\cdot\Psi_{3i} \quad \forall i\in A \tag{D12}$$

The objective then becomes:

$$\textbf{Maximize } \text{HCS} - \text{Penalty} - \sum_{i\in A}\sum_{j\in\bar{J}}\sum_{p\in\{1,2,\ldots,8\}} \lambda_p\cdot\Phi_{pij}.$$

Note that requests are given precedence over both the service-block structure and the preferred call shift structure.

Acknowledgments

This research has been supported in part by the following Georgia Tech benefactors: William W. George, Andrea Laliberte, Joseph C. Mello, and Richard "Rick" E. and Charlene Zalesky.

References

ACGME (2011a) ACGME: Common program requirements. Accessed December 1, 2011, http://www.acgme.org/acgmeweb/Portals/0/PDFs/Common_Program_Requirements_07012011[2].pdf.

ACGME (2011b) Resident duty hours in the learning and working environment: Comparison of 2003 and 2011 standards. Accessed July 5, 2012, https://www.acgme.org/acgmeweb/Portals/0/PDFs/dh-ComparisonTable2003v2011.pdf.

Arora VM, Manjarrez E, Dressler DD, Basaviah P, Halasyamani L, Kripalani S (2009) Hospitalist handoffs: A systematic review and task force recommendations. *J. Hospital Medicine* 4(7):433–440.

Beaulieu H, Ferland J, Gendro B, Michelon P (2000) A mathematical programming approach for scheduling physicians in the emergency room. *Health Care Management Sci.* 3(3):193–200.

Brunner JO, Bard JF, Kolisch R (2010) Midterm scheduling of physicians with flexible shifts using branch and price. *IIE Trans.* 43(2):84–109.

Burke EK, De Causmaecker P, Berghe GV, Van Landeghem H (2004) The state of the art of nurse rostering. *J. Scheduling* 7(6):441–499.

Cheang B, Li H, Lim A, Rodrigues B (2003) Nurse rostering problems—A bibliographic survey. *Eur. J. Oper. Res.* 151(3):447–460.

Cohn A, Root S, Kymissis C, Esses J, Westmoreland N (2009) Scheduling medical residents at Boston University School of Medicine. *Interfaces* 39(3):186–195.

Epstein K, Juarez E, Epstein A, Loya K, Singer A (2010) The impact of fragmentation of hospitalist care on length of stay. *J. Hospital Medicine* 5(6):335–338.

Ernst AT, Jiang H, Krishnamoorthy M, Owens B, Sier D (2004) An annotated bibliography of personnel scheduling and rostering. *Ann. Oper. Res.* 127(1–4):21–144.

Fortuna RJ, Palfrey JS, Shelov SP, Samuels RC (2009) Paediatric experiences with work-hour limitations. *J. Evaluation Clinical Practice* 15(1):116–120.

Gopalakrishnan B, Johnson EL (2005) Airline crew scheduling: State-of-the-art. *Ann. Oper. Res.* 140(1):305–337.

Gunawan A, Lau HC (2013) Master physician scheduling problem. *J. Oper. Res. Soc.* 64(3):410–425.

Kazemian P, Dong Y, Rohleder TR, Helm JE, Van Oyen MP (2013) An IP-based healthcare provider shift design approach to minimize patient handoffs. *J. Health Care Management Sci.* 17(1): 1–14.

Kohl N, Karisch SE (2004) Airline crew rostering: Problem types, modeling, and optimization. *Ann. Oper. Res.* 127(1–4):223–257.

Landrigan CP (2007) The handoff: A critical point of vulnerability. *CRICO/RMF Forum* 25(1):6–7.

Landrigan CP, Rothschild JM, Cronin JW, Kaushal R, Burdick E, Katz JT, Lilly CM, et al. (2004) Effect of reducing interns' work hours on serious medical errors in intensive care units. *New England J. Medicine* 351(18):1838–1848.

Ovchinnikov A, Milner J (2008) Spreadsheet model helps to assign medical residents at the University of Vermont's College of Medicine. *Interfaces* 38(4):311–323.

Papp KK, Stoller EP, Sage P, James E, Owens J, Avidan A, Phillips B, Rosen R, Strohl KP (2004) The effects of sleep loss and fatigue on resident-physicians: A multi-institutional, mixed-method study. *J. Assoc. Amer. Medical Colleges* 79(5):394–406.

Pickering BW, Hurley K, Marsh B (2009) Identification of patient information corruption in the intensive care unit: Using a scoring tool to direct quality improvements in handover. *Critical Care Medicine* 37(11):2905–2912.

Riesenburg LA, Leitzsch J, Massucci JL, Jaeger J, Rosenfeld JC, Patow C, Padmore JS, Karpovich KP (2009) Residents' and attending physicians' handoffs: A systematic review of the literature. *Acad. Medicine* 84(12):1775–1787.

Rinder MM, Weckman G, Schwerha D, Snow A, Dreher PA, Park N, Paschold H, Young W (2012) Healthcare scheduling by data mining: Literature review and future directions. *J. Healthcare Engrg.* 3(3):477–502.

Rodriguez H, Turner JP, Speicher P, Daskin MS, DaRosa D (2010) A model for evaluating resident education with a focus on continuity of care and educational quality. *J. Surgical Ed.* 67(6): 352–358.

Rousseau L, Pesant G, Gendreau M (2002) A general approach to the physician rostering problem. *Ann. Oper. Res.* 115(1):193–205.

Saultz JW, Albedaiwi W (2004) Interpersonal continuity of care and patient satisfaction: A critical review. *Ann. Family Medicine* 2(5):445–451.

Sherali HD, Ramahi MH, Saifee QJ (2002) Hospital resident scheduling problem. *Production Planning Control* 13(2):220–233.

Smalley HK, Keskinocak P, Vats A (2011) Development of a handoff continuity score to improve pediatric ICU physician schedule design for physician and patient continuity. *Critical Care* 15(5):R246.

Solet DJ, Norvell JM, Rutan GH, Frankel RM (2005) Lost in translation: Challenges and opportunities in physician-to-physician communication during patient handoffs. *Acad. Medicine* 80(12):1094–1099.

Stolletz R, Brunner JO (2012) Fair optimization of fortnightly physician schedules with flexible shifts. *Eur. J. Oper. Res.* 219(3): 622–629.

Topaloglu S (2006) A multi-objective programming model for scheduling emergency medicine residents. *Comput. Indust. Engrg.* 51(3):375–388.

Topaloglu S (2009) A shift scheduling model for employees with different seniority levels and an application in healthcare. *Eur. J. Oper. Res.* 198(3):943–957.

Topaloglu S, Ozkarahan I (2011) A constraint programming-based solution approach for medical resident scheduling problems. *Comp. Oper. Res.* 38(1):246–255.

Turner JP (2011) Overcoming obstacles to resident-patient continuity of care and equitable allocations of surgical experience. PhD dissertation, Northwestern University, Department of Industrial Engineering and Management Sciences.

Turner JP, Rodriguez HE, DaRosa DA, Daskin MS, Hayman A, Mehrotra S (2013) Northwestern University Feinberg School of Medicine uses operations research tools to improve surgeon training. *JAMA* 43(4):341–351.

Turner JP, Rodriguez HE, Daskin MS, Mehrotra S, Speicher P, DaRosa DA (2012) Overcoming obstacles to resident-patient continuity of care. *Ann. Surgery* 255(4):618–622.

Ulmer C, Wolman DM, Johns ME (2008) *Resident Duty Hours: Enhancing Sleep, Supervision, and Safety* (The National Academies Press, Washington, DC).

Van Eaton EG, Tarpley JL, Solorzano CC, Cho CS, Weber SM, Termuhlen PM (2011) Resident education in 2011: Three key challenges on the road ahead. *Surgery* 149(4):465–473.

VPS (2012) MyVPS Online. Accessed July 6, https://portal.myvps.org/Default.aspx.

Willis RE, Coverdill JE, Mellinger JD, Collins JC, Potts JR III, Dent DL (2009) Views of surgery program directors on the current ACGME and proposed IOM duty-hour standards. *J. Surgical Ed.* 66(4):216–221.

Wu AW, Folkman S, McPhee SJ, Lo B (1991) Do house officers learn from their mistakes? *JAMA* 265(16):2089–2094.

Verification Letter

James D. Fortenberry, MD, MCCM, FAAP, Pediatrician in Chief, Children's Healthcare of Atlanta, Professor of Pediatric Critical Care, Emory University School of Medicine, 1405 Clifton Road Northeast, Atlanta, GA 30322, writes:

"I am writing this letter to verify that the Handoff Continuity Score (HCS) and Children's PICU Physician Scheduling Model developed by Dr. Hannah K. Smalley, Dr. Pinar Keskinocak, and Dr. Atul Vats has been used at our hospital to produce attending shift schedules which maximize continuity of care. As staffing models evolve in intensive care units, one of the concerns of "shift models" relates to continuity of care. The HCS allows us an objective assessment tool for continuity as we consider alternate staffing models.

"The HCS developed by this team measures the level of familiarity among oncoming physicians at shift change, and correlated with physician perceptions of improved continuity, as determined by a survey. Our physician group has now been using the automated schedule generation for over two years, and it has been well received.

"We look forward to the new schedule currently being developed by this team for the 2nd half of 2013."

Hannah K. Smalley, PhD, is a recent graduate of Georgia Institute of Technology's Stewart School of Industrial and Systems Engineering. Her dissertation focused on optimization methods applied to real-world physician scheduling problems faced by physicians from Children's Healthcare of Atlanta and Emory University School of Medicine. Other projects she has contributed to include catch-up scheduling for childhood, adolescent, and adult immunizations, a collaborative effort with the CDC; analyzing patient overlap between hospitals in the Atlanta area to better understand the need for a health information exchange between hospitals; infectious disease modeling and immunization strategies during a vaccine shortage; and planning facility location and mass dispensing strategies for emergency response.

Pınar Keskinocak, PhD, is the William W. George Chair and Professor, and co-founder and co-director of the Center for Health and Humanitarian Logistics in the Stewart School of Industrial Engineering, and the associate director for research at the Health Systems Institute at Georgia Tech. Her research focuses on the applications of operations research and management science with societal impact, particularly health and humanitarian applications, supply chain management, and logistics/transportation. Her recent work has addressed infectious disease modeling; catch-up scheduling for vaccinations; hospital operations management; disaster preparedness and response; debris management; centralized and decentralized price and lead time decisions. She has worked on projects with companies, governmental and non-governmental organizations, and healthcare providers, including American Red Cross, CARE, CDC, Children's Healthcare of Atlanta, and Intel Corporation. She serves as a department editor for *Operations Research* (Policy Modeling and Public Sector area) and president of the INFORMS Health Applications Society.

Atul Vats, MD, FCCM, FAAP is an associate professor of pediatrics at Emory University School of Medicine, chair of peer review at Children's Healthcare of Atlanta, and a pediatric intensivist at Children's Healthcare of Atlanta. He is a fellow of the American College of Critical Care Medicine and a fellow of the American Academy of Pediatrics. His 10+ years as a medical director of one of the nation's highest acuity pediatric intensive care units, eight years of experience as a fellowship director, and decade-long experience in physician peer review has positioned him to lead efforts on physician workflow and efficiency. His research focuses on process improvement, enhancing quality of care, and improving patient continuity by specifically focusing on physician dependent workflow.