OPTIMIZATION METHODS FOR PHYSICIAN SCHEDULING

A Thesis
Presented to
The Academic Faculty

by

Hannah Kolberg Smalley

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Industrial and Systems Engineering

Georgia Institute of Technology
December 2012

OPTIMIZATION METHODS FOR PHYSICIAN SCHEDULING

Approved by:

Professor Pinar Keskinocak, Advisor
School of Industrial and Systems
Engineering
Georgia Institute of Technology

Professor Turgay Ayer
School of Industrial and Systems
Engineering
Georgia Institute of Technology

Professor David Goldsman
School of Industrial and Systems
Engineering
Georgia Institute of Technology

Professor Julie L. Swann
School of Industrial and Systems
Engineering
Georgia Institute of Technology

Dr. Atul Vats
Children's Healthcare of Atlanta
Emory University School of Medicine

Date Approved: 13 August 2012

To my parents, Ole and Betty Kolberg.

ACKNOWLEDGEMENTS

"I can do all things through Christ who strengthens me." *Philippians 4:13*

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

SUMMARY

This thesis considers three physician scheduling problems in health care systems. Specifically, we focus on improvements to current physician scheduling practices through the use of mathematical modeling. In the first part of the thesis, we present a physician shift scheduling problem focusing on maximizing continuity of care (i.e., ensuring that patients are familiar with their treating physicians, and vice versa). We develop an objective scoring method for measuring the continuity of a physician schedule and combine it with a mixed integer programming model. We apply our methods to the problem faced in the pediatric intensive care unit at Children's Healthcare of Atlanta at Egleston, and show that our schedule generation approach outperforms manual methods for schedule construction, both with regards to solution time and continuity.

The next topic presented in this thesis focuses on two scheduling problems: (i) the assignment of residents to rotations over a one-year period, and given that assignment, (ii) the scheduling of residents' night and weekend shifts. We present an integer programming model for the assignment of residents to rotations such that residents of the same type receive similar educational experiences. We allow for flexible input of parameters and varying groups of residents and rotations without needing to alter the model constraints. We present a simple model for scheduling 1st-year residents to night and weekend shifts. We apply these approaches to problems faced in the Department of Surgery Residency Program at Emory University School of Medicine. Rotation assignment is made more efficient through automated schedule generation, and the shift scheduling model allows us to highlight infeasibilities that occur when shift lengths exceed a certain value, and we discuss the impact of duty hour restrictions

under limitations of current scheduling practices.

The final topic of this thesis focuses on the assignment of physicians to various tasks while promoting equity of assignments and maximizing space utilization. We present an integer programming model to solve this problem, and we apply this model to the physician scheduling problem faced in the Department of Gynecology and Obstetrics at Emory University Hospital and generate high quality solutions very quickly.

CHAPTER I

INTRODUCTION

Operations research and mathematical modeling are increasingly applied to health care operations, particulary to problems involving resource allocation. The set of supply, demand, and feasibility requirements inherent in such problems often lend themselves well to a mathematical optimization model. While problems focused at the hospital level often seek to minimize costs, the ultimate goal is to deliver the best level of patient care.

This thesis focuses on three topics related to the allocation of resources in health care systems. The first two topics focus on the development of flexible models for assigning physicians to shifts or rotations to satisfy patient demand, physician preferences and educational requirements, while maximizing patient care through improved continuity or adequate physician coverage. The third topic is focused on the assignment of physicians to varying tasks so as to maximize space utilization and promote equity of assignments. For each of these topics, optimization models provide a means of developing scheduling strategies which consider physician preferences, work hour restrictions, and the expected demand for physicians' time.

1.1 Shift Scheduling for Physician and Patient Continuity

The Accreditation Council for Graduate Medical Education (ACGME) has instituted work hour restrictions which limit allowable duty hours for medical residents [1]. Increased duty hour restrictions placed on residents have resulted in more fragmented care, or the treating of patients by multiple physicians. Being treated by the same physician improves a patient's sense of continuity. It has been shown that continuity of care improves patient satisfaction [40], while fragmented care is reported to increase

patient length of stay [16].

With increased fragmentation of care comes more handoffs, or the transfer of a patient's care from one physician to another. Each handoff requires careful communication between physicians to relay all pertinant information regarding a patient's care. However, one impact of increased duty hour restrictions is an increase in handoff errors, simply due to the increased frequency of handoffs required by more fragmented care [50, 52].

1.1.1 Objectives

Patient care is largely impacted by the quality of communication between physicians at handoff. Miscommunications at handoff have become "widely recognized as a leading safety hazard in health care" [28]. However, communication can be improved and best practices can be identified [3, 37, 43]. As one approach to improving handoffs, our goal is to improve physician scheduling methods to maximize continuity for physicians and patients by considering the expected familiarity a physician may have with patients transferred to their care upon starting each shift. A physician schedule which maximizes continuity, while adhering to duty hour restrictions and physician and institutional preferences, could improve the efficiency of handoffs.

1.1.2 Methodology

We present a modeling and solution approach for assigning attending physicians, medical residents, and fellows to service and call shifts in the pediatric intensive care unit (PICU) at Children's Healthcare of Atlanta at Egleston (Children's) over a one-year period. We developed the Handoff Continuity Score, or HCS, for measuring the continuity of a schedule. The HCS measures the familiarity of oncoming physicians at each shift change based on previous days worked, using familiarity factors which rate the familiarity that a physician feels at handoff following a break of 1 up to 5 days, and considering the shifts worked in the previous five days. We combined

the HCS with a mixed integer programming model, the Children's PICU Physician Scheduling MIP (CPPS-MIP), which includes three sets of constraints to enforce feasibility requirements, institutional and physician preferences, and calculate the HCS. The objective of CPPS-MIP is to maximize the HCS while minimizing violations of physician preferences.

For a 51-week time horizon and a physician pool which includes 9 attending physicians and 7 fellows, no feasible solution to this MIP is found within 48 hours using CPLEX 12.4. However, a feasible schedule can be constructed using an iterative heuristic which incorporates modified versions of CPPS-MIP. We show that this heuristic produces a physician schedule which achieves an optimality gap of 3.42% for the scheduling instance faced by Children's for this time period.

We tested our solution approach using problems of different size and make, and investigated the benefits of variations to the heuristic incorporated into our solution approach. We also developed an alternative model formulation with the goal of improving on the efficiency of the heuristic.

1.1.3 Contribution

Attempting to satisfy physician preferences when faced with increased duty hour restrictions creates a complex scheduling problem. Our solution approach facilitates resource optimization, and automated schedule construction requires significantly less time than manually constructing such a schedule. We are the first to consider familiarity among oncoming physicians at shift change during schedule construction. Schedules which maximize continuity, in combination with other methods for improving communication, have the potential to improve handoff efficiency.

We generated 6-month schedules for attending physicians which were implemented in the PICU at Children's Healthcare of Atlanta in 2011 and 2012.

1.2   Medical Resident Rotation and Shift Scheduling

Most teaching hospitals experience two physician scheduling problems: (1) the assignment of residents to rotations (i.e., "rotation assignment") to both meet patient care demand and satisfy educational requirements and (2) the scheduling of day-to-day shifts (i.e., "shift scheduling"). In order to maximize their education, medical residents are often assigned to rotations in various hospitals and/or medical services to receive instruction during their training. These rotations require varying numbers of residents of different levels of experience to provide coverage to meet demand. Manually constructing a schedule which assigns residents to rotations can be a cumbersome task, particularly if there are a large number of rotations to cover with complicated demands with regards to the number and types of residents preferred for assignment. We discuss this problem as faced in the Department of Surgery Residency Program at Emory University School of Medicine (EUSOM).

Day-to-day shift scheduling of these residents is made additionally difficult due to increased duty hour restrictions [1]. One of the most recent restrictions implemented is the requirement that 1st-year residents be on duty for no more than 16 consecutive hours. This prevents the scheduling of 1st-year residents to back-to-back day and night shifts, a scenario which was previously standard practice. Training hospitals must find ways to overcome scheduling shortages when duty hour restrictions make past scheduling practices infeasible. EUSOM is one such program that has identified shortages while trying to construct day-to-day shift schedules.

1.2.1   Objectives

Resident physician rotation assignment and shift scheduling must adhere to rules governing resident education, work hours, and rest periods, but must also balance the objectives of maximizing resident experience and patient care. In hopes of alleviating the burden of manually constructing a rotation assignment, which is often a time

consuming task, our goal is to create a user-friendly automated tool which can be used in current and future years with varying numbers of residents and rotation demands.

The combination of ACGME duty hour restrictions and physician preferences complicate day-to-day shift scheduling of residents. One common preference among attending physicians at EUSOM is that 1st-year residents be given responsibility for all night and weekend shifts. Our analysis of the impact of duty hour restrictions and these preferences on the feasibility of scheduling inform decision-making with regards to resident shift scheduling.

1.2.2 Methodology

We developed two integer programming (IP) models with the goals of creating feasible assignments of residents to rotations over a one-year period (Resident Rotation Assignment Model (RRA-IP)), and constructing night and weekend call-shift schedules for the individual rotations (Surgical Resident Shift Scheduling Model (SRSS-IP)). These models create the ability to capture various duty-hour rules and constraints, test multiple what-if scenarios, and largely automate the process of schedule generation. We tested these models using scheduling constraints faced in the Department of Surgery at EUSOM.

We performed a detailed analysis of the performance of RRA-IP with varying objective functions. The complete objective function includes four objectives, namely (i) minimize weighted demand violations, (ii) minimize deviation from equality of assignments for residents in the same group, (iii) minimize denied resident requests for service assignments during specific periods, and (iv) maximize assignments to desirable services, of which the first two have higher priority. We limited the objective function to only the highest priority objectives to determine the impact on solution time and quality of the rotation schedule. We also investigated variations to education

requirements for residents, namely the requirements that residents of a given level and type require similar (or identical) experiences in the course of a year. Finally, we modified the constraints which enforce these education requirements in hopes of improving solution times.

We developed an Excel-based decision support tool for entering parameter values such as the numbers of residents, services, demands, and education requirements, without needing modifications to the constraints in RRA-IP. We constructed a 1-year rotation schedule using RRA-IP and find that in comparison to a manual schedule constructed for the same time period, it performs better with respect to solution time and equality of assignments for residents in the same group.

We used SRSS-IP to investigate how shift lengths and specific duty hour restrictions impact schedule feasibility given physician preferences (i.e., nights and weekends staffed by 1st-year residents). For education purposes, it is important that residents work as many daytime shifts as possible. We tested variations to the duty hour restriction limiting residents to 80 duty hours per week, when averaged over four weeks, and we considered varying weekday, night, and weekend shift lengths to determine the impact on resident education (specificially, the number of daytime shifts that can be assigned).

1.2.3 Contribution

Allocating residents to rotations more efficiently could have positive impacts on patient care [27], resident education [20], and compliance with duty-hour restrictions [44]. The Resident Rotation Assignment Model we have developed is general and can be used to generate schedules at EUSOM as well as in other residency programs, even if the number of residents, number of rotations to cover, and demands of those rotations change over time. The model constructs schedules faster than can be done

6

manually, and there is a sense of fairness among residents when a schedule is constructed by an objective model. An easy-to-use automated tool which incorporates these models could be rapidly adopted by programs throughout the country.

The shift scheduling model we present provides a means of understanding the impact of shift lengths and duty hour restrictions on resident education. It is possible to construct day-to-day shift schedules given current scheduling practices (i.e., 1st-year residents staffing night and weekend shifts), but with potential negative impacts on resident education. Relaxing ACGME duty hour restrictions limiting the number of allowable duty hours worked per week by a small margin could provide a means of improving resident education by increasing the number of daytime shifts assigned to each resident. Alternatively, reducing daily workloads or changing the preferred scheduling practices could provide an additional solution.

1.3 Task Assignment for Equity and Maximal Space Utilization

Manually constructing a staff schedule can be a difficult process, particularly if faced with conflicting objectives, non-homogeneous staff members, multiple tasks, and space availability limitations. Mathematical modeling could provide a tool for constructing such a schedule, balancing the complex factors at play.

1.3.1 Objectives

In hospital settings, it is not uncommon for physicians to be assigned to different tasks in a given week and from week to week, where each task may have different demands for physicians, and physician preferences and availability, both by day and task, complicate schedule construction. In the Department of Gynecology and Obstetrics at Emory University Hospital (Emory OB/GYN), a group of heterogeneous physicians are required to staff a number of assignments, including Labor and Delivery, the Emergency Room, two outpatient clinics, and Surgery. Some of these assignments

have fixed demands for physicians, while others should receive coverage if possible given physician and space availability. Physicians in the scheduling pool have varying availabilities, requirements, and preferences. Any assignment of physicians to tasks should maximize fairness with regards to a number of different metrics.

Our goal is to provide an automated tool for constructing physician schedules which satisfy demand and meet all objectives within feasibility requirements.

1.3.2 Methodology

We present an integer programming model, the Physician Scheduling model (PS-IP), for assigning physicians with non-homogeneous personal preferences and requirements to a set of tasks with varying demand and space availability. The model considers a large number of objectives from maximizing fairness of assignments to minimizing undesirable scheduling scenarios (e.g., assignments to surgery and night call in the same day). We test PS-IP by attempting to construct a 6-month schedule using constraints and vacation requests from Emory OB/GYN. After showing that no feasible solution exists due to a shortage of physicians to meet all demand (even when removing vacation requests), we relax a constraint preventing assignments to consecutive night call shifts. We seek to minimize assignments to consecutive night call shifts, in addition to the other objectives, and show that a feasible solution is possible in this case. We investigate alternatives to the objective function to determine the best option for meeting scheduling goals.

1.3.3 Contribution

We developed a model to solve a physician scheduling problem which includes a non-homogeneous physician pool, various tasks, and multiple objectives, while also considering availability of space. Using PS-IP, we solve this problem faced by Emory OB/GYN within minutes.

CHAPTER II

PHYSICIAN SCHEDULING FOR CONTINUITY: AN APPLICATION IN PEDIATRIC INTENSIVE CARE

When the care of a patient is transferred from one physician (MD) or team of physicians to another, a handoff takes place, whereby all rights and responsibilities regarding the patient change hands. In many cases, the first introduction an MD has with a patient's case is from another MD during a handoff. The communication that occurs between MDs at this time is a "vital link in the continuity of patient care" [43].

While many issues require careful discussion during a handoff, familiarity with a patient prior to taking responsibility of their care may help reduce the impact of a less than effective handoff and lead to a sense of continuity for the physician and patient alike. An MD schedule which meets all feasibility requirements (such as compliance with duty hour restrictions), satisfies physician preferences (when possible), and maximizes the likelihood of familiarity between physicians and their patients, is one approach to improving continuity and potentially handoff efficiency.

Optimized physician scheduling can provide a means of maximizing continuity within the confines of duty hour restrictions. In this chapter, we present a modeling and solution approach for assigning physicians to service (daytime) and call (night) shifts over a one-year period. Our approach is unique in that it seeks to make assignments which maximize continuity by considering expected familiarity of physicians with patients transferred to their care upon starting each shift.

2.1 Introduction

In an inpatient setting where the average patient length of stay (LOS) exceeds one day, patients are likely to be treated by multiple physicians including attending physicians and residents. In 2003 and again in 2011, the Accreditation Council for Graduate Medical Education (ACGME) instituted work hour restrictions greatly limiting the hours that residents can work [2]. These increased restrictions further contribute to fragmentation of care, or the treating of individual patients by multiple physicians. This is due to the fact that medical residents may be on duty for shorter periods of time than was allowed prior to the implementation of the new duty hour restrictions, and therefore may not be able to treat a patient during his entire stay in the hospital. Increased duty hour restrictions also force a trend towards night shift work in inpatient hospital settings and create concerns regarding negative impacts on education of residents, schedule flexibility, and continuity of care [18].

Continuity of care can have impacts on the quality of patient care. It has been shown that patient satisfaction is improved by continuity of care [40]. In one study, Epstein et al. (2010) observe a negative relationship between fragmentation of care and patient length of stay [16]. Rodriguez et al. (2010) investigate the impacts that the length of time surgical residents spend on individual rotations has on continuity of care, and find that rotations lasting only one month in duration are insufficient [38], though this is often the norm. 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 which assigns residents to one or two supervising physicians for the duration of their assignment to a rotation [49].

Fragmentation of care caused by increased duty hour restrictions means more handoffs of patients from one physician's care to another, increasing the risk for communication errors occurring during handoff due to increased frequency [50, 52].

Such communication errors can place patients at risk. Flawed communication at handoff has become "widely recognized as a leading safety hazard in health care" [28]. For example, Pickering et al. (2009) developed a tool for measuring the level of corrupted information shared at handoff [36]. The quality of communication between MDs at handoff largely impacts patient care, and efforts can be made to improve communication and identify best practices [3, 37, 43].

In this chapter, we present a mixed integer programming (MIP) model which constructs physician schedules that maximize continuity and familiarity by utilizing an objective scoring method for measuring continuity at each handoff. In Section 2.2, we outline the physician scheduling problem faced by many institutions where hospitalized patients are treated by multiple physicians due to their length of stay in the hospital, and we discuss the scoring method we developed for measuring continuity. We present the MIP we developed, and discuss our solution approach. In Section 2.3, we provide results from an application of our methods to the Pediatric Intensive Care Unit (PICU) at Children's Healthcare of Atlanta at Egleston (Children's). We also discuss applications to units and institutions with different preferred schedule structures. In Section 2.4, we present results from modifications to our solution approach as well as an analysis of our methods applied to problems of different size. In Section 2.5, we present an alternative model for solving this scheduling problem. We discuss implementation of our methods in Section 2.6 and summarize our findings in Section 2.7.

2.2   *Problem Description and Model*

Joint scheduling of specialty residents and/or subspecialty fellows (we use the term "residents" to refer to physicians in both of these groups) along with other physicians (e.g., attendings) is a complex problem due to ACGME duty hour restrictions as well as individual physician preferences. To abide by ACGME restrictions, residents may

be scheduled a maximum of 80 duty hours per week when averaged over four weeks, and may be scheduled for a maximum of 24 hours of continuous duty, with at most 4 additional hours to ensure an effective transfer of care [1]. An 8-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. Residents must have one day per week free of duty when averaged over four weeks, and may be scheduled for in-house call (i.e., the resident must remain on site) at most once every third night when averaged over four weeks. In addition to these requirements, a feasible schedule must satisfy the expected patient demand for MDs.

We investigate the problem of scheduling staff under such restrictions in the PICU at Children's. This PICU is a 30-bed multidisciplinary medical-surgical quaternary care unit which cares for acutely ill patients, and is part of the largest pediatric healthcare system in the country. A typical daily schedule for the PICU includes 2 attendings and 2 fellows on service during the day, and 1 attending and 1 fellow on call at night (in addition to up to 4 residents on service during the day, and 2 on call at night). According to the Virtual PICU Performance System (VPS), a national PICU database [51], of 26 PICU's around the country that submitted data for at least 1 quarter in 2011, only 7 (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 due to size and complexity, as compared to other PICU's.

When scheduling staff, there are often institutional or individual preferences that are unique to each situation. In the PICU at Children's for example, a block service 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., 7-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 every Monday). Additional individual preferences with regards to specific days on or off service or night-call also are taken into account when constructing an MD schedule for the unit, a process which was previously performed manually. For past schedules, manual construction of a 6-month attending-only schedule (which is not highly regulated by the ACGME) required several hours, with additional time invested on an ongoing basis to accommodate schedule changes.

In light of these constraints and preferences in scheduling, our goal is to develop an efficient and effective approach 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, or HCS.

2.2.1 Handoff Continuity Score

Intuitively, a physician has more familiarity with a patient's progress and current state if she treated the patient for multiple consecutive days. Conversely, returning to duty after a multi-day break may require a readjustment period for the physician to re-familiarize herself with the patient's condition. So given a physician duty schedule, we measure the continuity at each handoff based on two basic assumptions:

- A physician's familiarity with her patients increases with multiple (possibly successive) on-duty days where the physician cares for the patients.

- A physician's familiarity decreases as the number of recent days off-duty increases.

To capture the familiarity felt by physicians at handoff for each previous day worked within a specific time period, we developed the familiarity factors reported in Table 1. These familiarity factors are based on a five-day period (i.e., the average patient length of stay in the PICU at Children's) and have a seemingly inverse exponential relationship. These factors were developed based on informal discussions and

Table 1: Familiarity Factors

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

surveys of the attending physician group at Children's at multiple meetings. Linear familiarity factors were considered (and testing revealed similar results with linear factors), but this particular attending physician group suggested that there was not a linear relationship between the familiarity felt after having worked one up to five days prior to a handoff. Several formats for the familiarity factors were proposed and discussed, but those we have chosen were considered to be "most appropriate" for the physician group in this PICU. We also investigated the use of factors which consider time periods shorter and longer than the average patient length of stay, but a five-day period was identified to be most appropriate as it is a good "global marker" for this patient population. Thus, while 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 own perceptions of continuity.

To determine the continuity of a schedule, we assign a continuity score between 0 and 1 to each oncoming physician 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 2 days and 4 days previously, they receive a score of $0.15 + 0.025 = 0.175$. A score of 1 implies the greatest familiarity. The score for each handoff is calculated as the average continuity score over all oncoming physicians, and for a complete physician schedule, we average these handoff scores to determine the overall HCS. For more

details on the development of the HCS, see [42].

2.2.2 A Mixed Integer Programming Model

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

Optimization methods applied 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 [9, 10] for comprehensive reviews. Beaulieu et al. (2000) present a mathematical model for scheduling emergency room physicians to shifts over a 6-month period [4]. Sherali et al. (2002) use a mixed-integer program for scheduling residents to night-shifts over a 4-5 week period [41]. Rousseau et al. (2002) develop a flexible solution approach applying constraint programming, local search, and genetic algortihms to the physician scheduling problem faced by various units/institutions with minimal customization [39]. Topaloglu (2006) assigns emergency medicine residents to day and night shifts using goal programming [45], and in a later paper, applies sequential and weighted methods to a multi-objective optimization model for assigning residents to night-call shifts while considering levels of seniority [46]. 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 [35], and Cohn et al. (2009) solve multiple nested IP models to assign 10-20 residents to various types of night-call shifts (e.g., primary, backup) in three different hospitals over a 1-year period [11]. Brunner et al. (2010) present a branch and price algorithm for constructing daily physician schedules with flexible shift start times and shift lengths, for a scheduling horizon of up to 6 weeks [8].

Other relevant scheduling problems addressed in the literature include airline crew scheduling [21, 26]. Similar to the physician scheduling problem that we address, crew

scheduling involves extensive duty hour restrictions. Ernst et al. (2004) compile a detailed list of previous work in personnel scheduling, including problem type and solution approach [17].

The Children's PICU Physician Scheduling MIP (CPPS-MIP) 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 [48]. However, to the best of our knowledge, we are the first to consider the benefits of optimized MD shift scheduling on continuity. CPPS-MIP is general and can be easily applied to various units and institutions with different preferences where each day is divided into two non-overlapping time periods, each period starting at the same time every day. This means that on any given day, an MD can be scheduled to be on service sometime between the hours of 8am and 4pm, for example, and/or on call sometime between 4pm and 8am the following day. However, different shift types are possible in each time period (e.g., 8am - 4pm vs. 8am - 12pm service shift).

The notation we use in CPPS-MIP is given in Tables 2, 3, and 4.

Some constraints in CPPS-MIP are applicable to a majority of institutions (constraints (1a) through (1k)), while others are more specific to the problem faced at Children's due to physician preferences (constraints (2a)-(2i)). Constraints (3a)-(3c) calculate the HCS, and (3c) and (3d) form the objective function.

The following constraints enforce the feasibility of the MD schedule. Each day, there is a demand for physicians during each time period, denoted by constraints (1a), which can be written for different physician groups such as residents of all levels, attendings only, fellows only, etc. ACGME requirements must be satisfied by the schedule for residents, corresponding to a maximum of 80 duty hours per week, averaged over four weeks (constraints (1b)), and a maximum of 28 hours worked consecutively, where we assume that 4 of those hours are used to ensure an effective

Table 2: Children's PICU Physician Scheduling MIP (CPPS-MIP) - Sets and descriptions

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,..., R_n=fellows), $R = \bigcup_{i \in \{1,2,...,n\}} R_i$
J	Set of days in the planning horizon, $J = \{1, 2, ..., N\}$
L	Set of scheduling time 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, $W = \{1, 2, ..., N/7\}$
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., $\wp = \{A, R, R_1 \cup R_2, ...\}$)
B_P	Set of possible service block lengths (in days) for physicians in group P, $P \in \wp$ (e.g., $B_A = \{7\}$)

Table 3: Children's PICU Physician Scheduling MIP (CPPS-MIP) - Decision Variables and descriptions

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. $i \in I$, $j \in J$, $n \in \{1, 2, ..., N - j + 1\}$
D_{ijl}	1 if physician i works shifts in two consecutive time periods beginning with time 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 time 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 time 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. $i \in I$, $j \in J$

transfer of care and are not considered duty hours (constraints (1c)). A 10-hour rest is recommended between scheduled duty periods (constraints (1d)), and residents must have at least 14 hours free of duty after 24 hours of in-house duty (constraints (1e)). Residents must have one day per week free of duty, when averaged over four weeks (constraints (1f)-(1h)), and may be scheduled for in-house call at most once every third night (constraints (1i)). Given that multiple shift types are available during each time period, there is an additional feasibility constraint requiring that no MD be scheduled for overlapping shifts (constraints (1j)). Constraints (1k) are required if 1st year residents are available for scheduling. Physicians in this group may work a maximum of 16 consecutive hours, and therefore may not be assigned to duty shifts in two consecutive time periods.

Table 4: Children's PICU Physician Scheduling MIP (CPPS-MIP) - Parameters and descriptions

Parameter	Description
N	Length of scheduling horizon, in days
$\hat{N}$	Number of past 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 time 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 that must be scheduled to a block of length n which starts on day j. $P \in \wp$, $j \in J$, $n \in \{1, 2, ..., N-j+1\}$
F_m	familiarity factor corresponding to a shift worked $[m-1, m)$ days ago, $m \in \left\{1, 2, ..., \hat{N}\right\}$
Z_P	the maximum allowable deviation from the number of service shifts a physician in group P is scheduled from the amount they 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 time 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 time period l on day j, 0 otherwise. $i \in I$, $j \in J$, $l \in \{1,2\}$

(1a)
$$\sum_{k \in K_l} \sum_{i \in C} X_{ijk} = v_{ljP} \quad \forall j \in J, l \in L, P \in \wp$$

(1b)
$$\sum_{l \in L} \sum_{k \in K_l} \sum_{j \in \{m, m+1, ... m+27\}} \bar{L}_k X_{ijk}/4 \leq 80 \quad \forall i \in R$$
$$m \in \{1, 2, ..., N-27\}$$

(1c)
$$X_{ijk^1} + X_{ijk^2} + X_{i(j+1)k^3} \leq 2 \quad \forall i \in R, (k^1, k^2, k^3) \in \hat{C}_1$$
$$j \in \{1, 2, ..., N-1\}$$
$$X_{ijk^1} + X_{i(j+1)k^2} + X_{i(j+1)k^3} \leq 2 \quad \forall i \in R, (k^1, k^2, k^3) \in \hat{C}_2$$
$$j \in \{1, 2, ..., N-1\}$$

$$X_{ijk^1} + X_{ijk^2} \leq 1 \quad \forall i \in R, \, j \in J, \, (k^1, k^2) \in T$$

(1d)
$$X_{ijk^1} + X_{i(j+1)k^2} - \sum_{k \in K_2} X_{ijk} \leq 1 \quad \forall i \in R, \, (k^1, k^2) \in T_1$$
$$j \in \{1, 2, ..., N-1\}$$

$$X_{ijk^1} + X_{i(j+1)k^2} - \sum_{k \in K_1} X_{i(j+1)k} \leq 1 \quad \forall i \in R, \, (k^1, k^2) \in T_2$$
$$j \in \{1, 2, ..., N-1\}$$

$$X_{ijk^1} + X_{ijk^2} + X_{i(j+1)k^3} \leq 2 \quad \forall i \in R$$

(1e)
$$j \in \{1, 2, ..., N-1\}$$
$$(k^1, k^2) \in C_1, \, (k^2, k^3) \in \hat{T}_1$$

$$X_{ijk^1} + X_{i(j+1)k^2} + X_{i(j+1)k^3} \leq 2 \quad \forall i \in R$$
$$j \in \{1, 2, ..., N-1\}$$
$$(k^1, k^2) \in C_2, \, (k^2, k^3) \in \hat{T}_2$$

$$X_{ijk} - D_{ij1} \leq 0 \quad \forall i \in I, \, j \in J, \, k \in K_l, \, l \in L$$

(1f)
$$X_{ijk} - D_{ij2} \leq 0 \quad \forall i \in I, \, j \in J, \, k \in K_2$$
$$X_{i(j+1)k} - D_{ij2} \leq 0 \quad \forall i \in I, \, j \in \{1, 2, ..., N-1\},$$
$$k \in K_1$$

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

(1h)
$$\sum_{j \in \{m, m+1, ... m+27\}} D_{ij1}/4 \leq 6 \quad \forall i \in R$$
$$m \in \{1, 2, ..., N-27\}$$

(1i)
$$\sum_{k \in K_2} X_{ijk} + X_{i(j+1)k} + X_{i(j+2)k} \leq 1 \quad \forall i \in R, \, j \in \{1, 2, ..., N-2\}$$

(1j) $$\sum_{k\in K_l} X_{ijk} \leq 1 \quad \forall i \in I, j \in J, l \in L$$

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

As previously mentioned, MDs in the PICU at Children's prefer a block structure to their service schedule, so there are additional "preference constraints" in our MIP. Constraints (2a)-(2c) are necessary for scheduling MDs according to the preferred block structure. Constraints (2a) ensure that an 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 (constraints (2b)), and MDs may not work overlapping blocks (constraints (2c)). It is possible to be assigned to consecutive service blocks, but with a penalty (constraints (2d)).

(2a) $$\sum_{k\in K_1} X_{isk} \geq Y_{ijn} \quad \forall i \in P, s \in \{j, j+1, ..., j+n-1\}$$
$$j \in \{1, 2, ..., N-n+1\}, n \in B_P, P \in \wp$$

(2b) $$\sum_{i\in P} Y_{ijn} \geq e_{Pjn} \quad \forall j \in \{1, 2, ..., N-n+1\}, n \in B_P, P \in \wp$$

(2c) $$Y_{ijn} + Y_{ikr} \leq 1 \quad \forall i \in P, n \in B_P, k \in \{j, j+1, ..., j+n-1\}$$
$$r \in \{1, 2, ..., N-k+1\}$$
$$j \in \{1, 2, ..., N-n+1\}, (j,n) \neq (k,r)$$
$$P \in \wp$$

(2d) $$Y_{ijn} + Y_{i(j+n)r} \leq 1 + \Omega_{ij} \quad \forall i \in P, n \in B_P, j \in \{1, 2, ..., N-n+1\},$$
$$r \in \{1, 2, ..., N-j-n+1\}, P \in \wp$$

Some constraints in our MIP that are specific to the PICU at Children's, but may be modified for other units, include the following. Attending physicians prefer not to

work two consecutive call shifts if on service in between those shifts (constraints (2e)), and over the scheduling horizon, the number of days on service for each MD should be close to average for each physician type (constraints (2f)). There are two types of service shifts (1: 8am-4pm; 2: 8am-12pm) and 1 type of call shift (3: 4pm-8am). MDs in the PICU prefer a fellow work shift 3 on weekend days only if he or she was scheduled for shift 1 (constraints (2g)), and fellows may only work shift 2 on any day if they were assigned to shifts 1 and 3 the previous day (constraints (2h)). Shift 2 is designed to allow time for effective transfer of care following a 24-hour shift.

$$(2e) \qquad \sum_{k \in K_2} \left(X_{ijk} + X_{i(j+1)k} \right) + X_{i(j+1)l} \quad \leq \quad 2 \qquad \forall i \in A, \, k \in K_1$$
$$j \in \{1, 2, ..., N-1\}$$

$$(2f) \quad \sum_{k \in K_1} \sum_{j \in J} X_{ijk} - \left(\sum_{j \in J} v_{1jP} / |P| \right) \quad \geq \quad Z_P \quad \forall i \in P, \, P \in \wp$$

$$(2g) \qquad\qquad\qquad\qquad\qquad X_{ij3} \quad \leq \quad X_{ij1} \quad \forall i \in R_n, \, j \in \{r-1, r\}$$
$$r \in \{7, 14, ..., N\}$$

$$(2h) \qquad\qquad 0.5 \cdot \left(X_{i(j-1)1} + X_{i(j-1)3} \right) \quad \geq \quad X_{ij2} \quad \forall i \in R_n, \, j \in \{2, 3, ..., N\}$$

When not scheduled in the PICU, fellows could be doing research or be assigned to a rotation in Anesthesia (4 weeks during the 3 year fellowship), the Cardiac Intensive Care Unit (12 weeks), or an elective (4 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, constraints (2i) specify that each fellow can only be assigned to service and call shifts in weeks they are available (i.e., are not assigned elsewhere).

(2i) $\quad X_{ijk} \leq Q_{ir} \quad \forall i \in R_n,\, r \in W,\, j \in \{7r-6, 7r-5, ..., 7r\},\, k \in K_l,\, l \in L$

The HCS is calculated based on previous days worked by oncoming MDs [42]. Specifically, each MD receives a continuity score between 0 and 1 (1 = most familiar) computed as the summation of a subset of familiarity factors (Table 1) corresponding to each day worked in the last 5 days (constraints (3a)). Each handoff score is only based on the previous days worked by oncoming MDs, not all MDs, so we incorporate constraints (3b) to capture the scores for oncoming MDs only. 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 (constraints (3c)).

$$
\text{(3a)} \quad R_{ijl} = \sum_{m \in \{1,2,...,\hat{N}\}} F_m D_{i(j-m)l} \quad \forall i \in I,\, l \in L
$$
$$
j \in \left\{\hat{N}+1, \hat{N}+2, ..., N\right\}
$$

$$
\hat{R}_{ijl} \leq \sum_{k \in K_l} X_{ijk} \qquad\qquad \forall i \in I,\, l \in L
$$
$$
j \in \left\{\hat{N}+1, \hat{N}+2, ..., N\right\}
$$
$$
\text{(3b)} \quad \hat{R}_{ijl} \geq R_{ijl} + \sum_{k \in K_l} X_{ijk} - 1 \qquad \forall i \in I,\, l \in L
$$
$$
j \in \left\{\hat{N}+1, \hat{N}+2, ..., N\right\}
$$
$$
\hat{R}_{ijl} \leq R_{ijl} \qquad\qquad\qquad \forall i \in I,\, l \in L
$$
$$
j \in \left\{\hat{N}+1, \hat{N}+2, ..., N\right\}
$$

$$
\text{(3c)} \quad \mathbf{HCS} = \left(\sum_{j \in \{\hat{N}+1,..,N\}} \sum_{l \in L} \left[\sum_{i \in I} \hat{R}_{ijl} \Big/ \sum_{P \in \{A,R\}} v_{ljP} \right] \right) \Big/ \left(2N - 2\hat{N} \right)
$$

To optimize continuity, our goal is to maximize the HCS. Because preferences of physicians in the PICU at Children's must be taken into account, a penalty is incurred if MDs are scheduled to consecutive service blocks, or if physician requests are violated (constraints (3d)).

$$
\text{(3d)} \quad \mathbf{Penalty} = \sum_{i \in I} \sum_{j \in J} \Omega_{ij} + \sum_{i \in I} \sum_{j \in J} \sum_{l \in L} \left[(1 - O_{ijl}) \cdot X_{ijl} + U_{ijl} \cdot X_{ijl} \right]
$$

Then our objective is: **Maximize (HCS - Penalty)**. Note that since 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. In other words, the penalty is exactly determined by physician requests, and exists to ensure that requests are granted if feasibly possible and denied otherwise.

2.2.3 Solution Approach

Physicians at Children's provided us with a manually constructed schedule for the 51-week period from Monday, July 5, 2010 - Sunday, June 26, 2011. The schedule included only attendings and fellows, so we limit our solution approach discussion to these two groups. To illustrate the advantage of optimized MD scheduling with regards to physician and patient continuity, we attempted to use CPPS-MIP to generate MD schedules for these groups for the same time period. However, a significant amount of running time was required (no feasible solution was found within 48 hours using CPLEX 12.4). 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 CPPS-MIP with a reduced set of decision variables:

Heuristic 1 - Iterative Schedule Construction with Modified CPPS-MIP:

- **Step 1**: Iteratively assign attending physicians to shifts on a week-by-week basis, ignoring requests for time off and/or on duty.

- **Step 2a**: Assign fellows to service blocks using the following simple integer program, with sets, parameters, and decision variables as defined in Table 5.

Table 5: Fellows' Service Block Assignment Integer Program (FSBA-IP) - Sets, Parameters, and Decision Variables

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 2-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 2-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

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

$$\text{Minimize d}$$

$$\text{s.t.}$$

$$\sum_{j \in \{1,2,...,W\}} Y_{ij} \leq d \quad \forall i \in F$$

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

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

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

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

FSBA-IP fairly assigns fellows to service blocks (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, next assign fellows to call shifts: In each week, assign the Monday, Thursday, and Sunday night call shifts to the fellow on week one of a two-week service block, and assign the Wednesday and Saturday night call shifts to the other fellow on-service that week. Assign the Tuesday and Friday night call shifts to one of the fellows not on service (arbitrary). Note that fellows' requests for time off at night will be taken into

account at 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 1 second. Step 3 requires running CPPS-MIP with some variables fixed. After running for 2 hours, the penalty is minimized and only marginal improvement in the HCS is found up to 48 hours. Therefore, for Step 4, we use the best solution found after running Step 3 for 2 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 Heuristic 1 to warm start CPLEX, CPPS-MIP could not find an HCS-improved schedule after running for 48 hours, and the optimality gap is reported by CPLEX at 25% 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 helps us identify a better upper bound, and shows that the optimality gap of the Heuristic 1 solution is much smaller than the one reported by CPLEX.

Proposition 1: Let H be the HCS of an optimal N-block schedule (equal-length, consecutive and non-overlapping blocks) and let H^1 = the best HCS of any k-block schedule, $k \leq n$ and N mod k = 0. Then $H^1 \geq H$.

Proof: Let H^1 = the best HCS of any k-block schedule. Let Δ be an optimal N-block schedule (N mod k = 0). Then Δ can be broken up into N/k k-block segments.

Let H = the HCS of Δ. Let $X_1, X_2, ..., X_{N/k}$ equal the summation of all handoff scores in block segments 1, 2, ..., N/k, respectively, for the optimal N-block schedule. Then

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

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

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

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

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

Thus, by (P3), $H^1 \geq H$.

$\square$

Fellows in the PICU at Children's work two-week service blocks, which is equivalent to 2 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.

By 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 taken into account when determining the HCS of that schedule. Therefore, given Proposition 1, we used CPPS-MIP to generate a 4-week schedule, optimizing the HCS over the last 3 weeks of the schedule in order to determine the best possible 3 week HCS. In the next section, we refer to this 3-week schedule as our "3 week test problem," which CPLEX solves in less than 1 minute.

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, $k \leq n$ and N mod k = 0. Then $H^1 \geq H$.

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

After first removing the constraints enforcing the service block structure implemented at Children's (constraints (2a)-(2d)), we used the modified MIP to generate 4 and 5-week schedules (again optimizing over the last 3 and 4 weeks of the schedules, respectively) to determine the best possible HCS of such schedules without service blocks.

2.3   Results and Discussion

The HCS for the manual and heuristic-generated schedules (for the 51-week period from July 5, 2010 through June 26, 2011) are reported in Table 6. We also report the HCS for daytime service shifts, night call shifts, attendings, and fellows, respectively. For each score presented in Table 6, the heuristic-generated schedule achieves a statistically significant improvement over the score incurred by the manually generated schedule, determined using a z-test (computed using HCS at each handoff over schedule horizon). The manual and heuristic-generated schedules each incorporate the preferred service block structure, and so there is understandably relatively little improvement (6.49%) in the HCS for service shifts. The largest increase in the HCS (36.87%) is attributed to the night call shifts. The HCS for attendings increased significantly more than the HCS for fellows. 14-day service blocks for fellows provide concentrated clinical time which is important for education, but

Table 6: Handoff Continuity Score (HCS) Results

	Manual Schedule	Heuristic-Generated Schedule	Percent 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)	

* All statistically significant increases, $p < .001$ using a z-test.

also provides more continuity for daytime service shifts than the 7-day service blocks worked by attendings. Therefore, the manually generated schedule had a high HCS for fellows compared to attendings due to the service block structure. In addition, ACGME duty hour restrictions provide rigorous constraints which limit much variability of night call shift assignments for fellows compared to attendings. Therefore the heuristic could only find marginal (but still statistically significant) improvement in the HCS for fellows. Note that all physician requests were satisfied by manual and heuristic-generated schedules.

In addition to the HCS by physician type and time of day, Table 6 also reports the percentage (and the number) of night call shifts over the schedule horizon without an on-service physician assigned. While not part of the MIP's objective, it is intuitive that a schedule which maximizes continuity would likely have very few night call shifts where an on-service physician was not scheduled. As reported in the table, the

Table 7: Test Problem Results and Comparison to Heuristic Solution

Problem Description	Optimal Solution	Heuristic 1 Solution	Optimality Gap
3 week schedule	0.8190	-	$< .01\%$
51 week schedule	-	0.7919^*	$3.42\%^+$

* *solution doesn't improve after running MIP for 48 hours*
$^+$ *optimality gap based on Proposition 1*

heuristic-generated schedule includes only five night call shifts where an on-service physician is not assigned, compared to 52 in the manually generated schedule.

The heuristic-generated schedule takes no special considerations when scheduling for the December holidays, but maintains the service schedule structure adopted for the rest of the year. Inputting the manually generated holiday schedule (Dec. 24 - Jan. 2) into the heuristic-generated schedule, the HCS for this new schedule is .7826, still an 18.09% improvement over the manual schedule.

Given the fact that a 51 week block schedule cannot achieve a higher HCS than a 3 week block schedule based on Proposition 1, we conclude that the HCS for the heuristic-generated schedule is within 3.42% of optimality. A comparison between the 3 week test problem and the 51 week heuristic solution is given in Table 7.

To test the applicability of CPPS-MIP to other units and institutions without a service block structure, as well as the efficiency of the model in such scenarios, we removed the constraints enforcing the service block structure (constraints (2a)-(2d)) and used Heuristic 1 with this modified MIP to generate a schedule for the same 51-week time period. Results from this analysis are reported in Table 8. We found that removing the requirements for the service block structure greatly increased the running time required to reach a good solution. After running the modified MIP for 24 hours for a time horizon of 51 weeks, no improvement beyond the warmstart solution provided by the heuristic is found, and the optimality gap reported by CPLEX is

Table 8: HCS Results for General Problem without Service Block Structure

Problem Description	Modified CPPS-MIP Solution	Optimality Gap (CPLEX)	Optimality Gap (Corollary 1)
3 weeks	0.9238	< 0.01%	-
4 weeks	0.9203	< 0.01%	-
12 weeks	0.8775	12.48%	4.88%
27 weeks	0.8599*	16.30%	7.43%
51 weeks	0.7719*	29.54%	19.68%

* *no improvement of warmstart solution*

29.54%. With Corollary 1, this gap is reduced to 19.68%, which is still quite a large gap. In light of this, we tested the model on problems with smaller time horizons, namely 27 weeks and 12 weeks. For the 27 week problem, an optimality gap of 7.43% can be achieved within 24 hours, using Corollary 1. For the smallest problem that we tested (i.e., 12 week time horizon), an optimality gap of 4.88% can be achieved in 24 hours.

From this analysis, we see that CPPS-MIP is more effective and Proposition 1 provides a tighter upper bound under service block constraints. The modified MIP used in the analysis presented in Table 8 is relatively efficient on small problems (e.g., 12 week time horizon), but doesn't approach optimality very rapidly for larger problems. To improve the efficiency of the modified MIP on larger problems, we can either (1) develop an alternative heuristic for constructing a better initial feasible schedule without the service block structure, or (2) search for methods to reduce the upper bound.

In hopes of reducing the upper bound on physician scheduling problems without the service block structure implemented at Children's, we considered the inclusion of additional constraints to limit assignments to multiple consecutive days. We observed

in solutions to the 3 week and 4 week test problems reported in Table 8 that physicians were assigned to many consecutive days (beyond the 7-day limit implemented at Children's). While this may be good for continuity, such assignments can have negative impacts on physician fatigue and morale. Our goal was to improve the solution time of these problems by limiting consecutive assignments. However, the large number of symmetric solutions prevent improvement in solution times. For example, if we limit the number of consecutive daytime assignments to 7, for example, there are many possible start days for each possible 7-day period (e.g., Monday, Tuesday, etc.), and the assignment of one physician is equivalent to the assignment of any other (unless specific requests for the relevant time periods are considered). Alternative methods for reducing the upper bound need to be found, which we leave to future work. However, we expect that for most practical instances, while no defined service block structure may be implemented, there will exist other rules with regards to consecutive days on service. Such rules would limit the range of feasible schedules, and thus we expect the performance of the modified MIP (with the addition of these new constraints) to achieve a performance level more in line with that observed for Children's.

2.4   *Heuristic and Problem Size Variations*

We tested alternatives to Heuristic 1 in order to determine if there was a different heuristic which achieved better results. Areas where we considered creating variations include Steps 2b, 3 and 4. Step 2b requires the assigning of fellows to night call shifts according to a prespecified structure. Step 3 of the heuristic, which seeks to find an optimal schedule for attendings while leaving fellow assignments fixed, is the step which requires the most running time and has the greatest optimality gap. Once a reasonable schedule is found (we stopped the MIP at 2 hours), these attending assignments are fixed, and Step 4 requires running CPPS-MIP again to optimize

Figure 1: Descriptions of Original and Alternate Heuristics

fellows' night call shift assignments.

We tested variations of Heuristic 1 by changing the order of assignments. Essentially, rather than first preassigning fellows to call shifts, we preassign attendings to call shifts. We then seek to optimize fellows' night call shift assignments, and then with these assignments fixed, solve for attendings' day and night shifts. Figure 1 describes these alternate heuristics in more detail. Heuristics 2 and 3 differ in the preassignment of night call shifts to attendings. Both of these call structures are models proposed by physicians in the PICU at Children's. We originally planned to test an alternate heuristic which kept the same order of assignments as Heuristic 1, but changed the fellows' night call assignments in Step 2b. However, due to ACGME rules requiring that no fellow be assigned to more than one night call shift in three days, we see that no improvement in this night call structure can be made in terms of continuity, and so alterations to it would not likely result in improvements over the original heuristic.

Results from testing these alternate heuristics with a running time of two hours

Table 9: Heuristic 1 Variations - Results

	Step 3			Step 4		
	HCS	**Running Time**	**Gap**	**HCS**	**Running Time**	**Gap**
Heuristic 1	0.7904	2 hours	8.28%	0.7919	< 10 sec.	< 0.01%
Heuristic 2	0.7224	< 10 sec.	< 0.01%	0.7899	2 hours	8.34%
Heuristic 3	0.7519	< 10 sec.	< 0.01%	0.7907	2 hours	8.24%

in Step 4 are reported in Table 9. Heuristic 1 outperforms the alternates, but only by a very slight margin; there is only a gap of 0.15% between the best and worst HCS values found following Step 4, and there are no noticable differences in the structures of the solutions. If we look at the solutions more closely in terms of the objectives achieved prior to 2 hours, we see that these heuristics perform almost identically (see Figure 2). Extending the running time to 24 hours reveals no significant differences (see Figure 3). The majority of improvement in objectives is achieved for each case prior to a running time of 1 hour (see Figure 4). However, there is some modest improvement between 1 hour and 2 hours. Thus, given the small difference between 1 hour and 2 hour running times in practice, we conclude that running CPPS-MIP in the relevant steps for two hours is logical.

We also investigated changing the problem size to determine if Step 3 of Heuristic 1 solves to optimality more quickly with more or less physicians in the scheduling pool. The PICU at Children's requires a minimum of 7 attendings, and can accommodate at most 13 attendings. Therefore, we tested the heuristic with these values, leaving the number of fellows the same. We also investigated the impact of doubling the number of attendings and fellows available for scheduling (i.e., 18 attendings, 14 fellows) to see if there was any impact on efficiency of this step of the heuristic. Figure 5 reports the LP and IP objectives for these variations to our problem size with a running time of 24 hours. The figure also reports these values for the original problem size which includes 9 attendings and 7 fellows. Each of these variations is initialized with
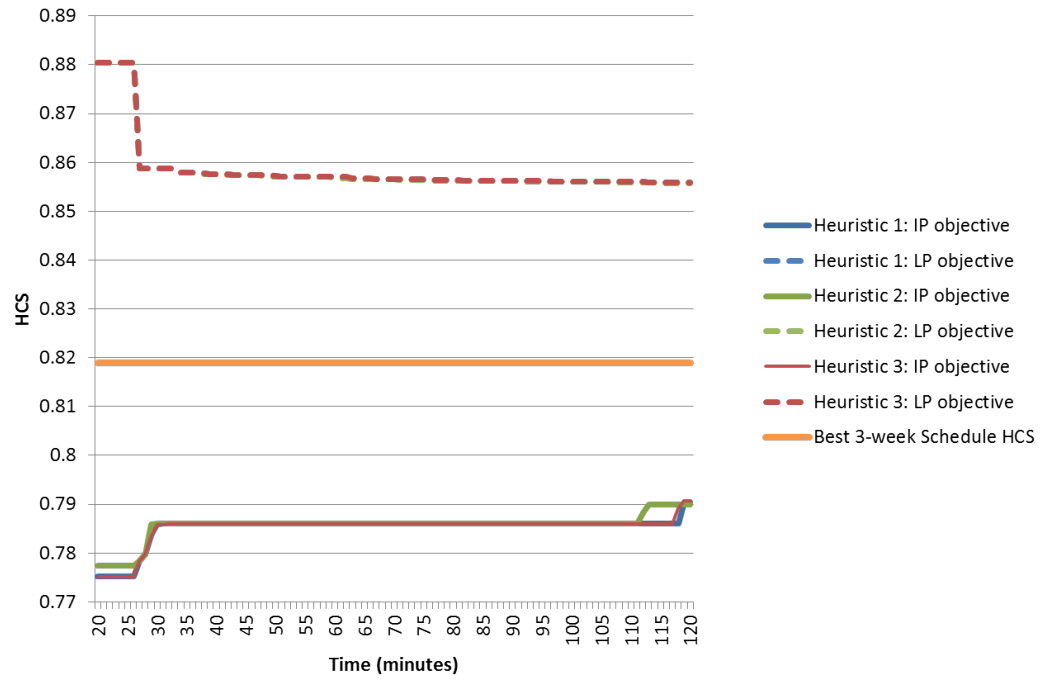
Figure 2: LP vs. IP objectives for variations to Heuristic 1 with a running time of 2 hours.
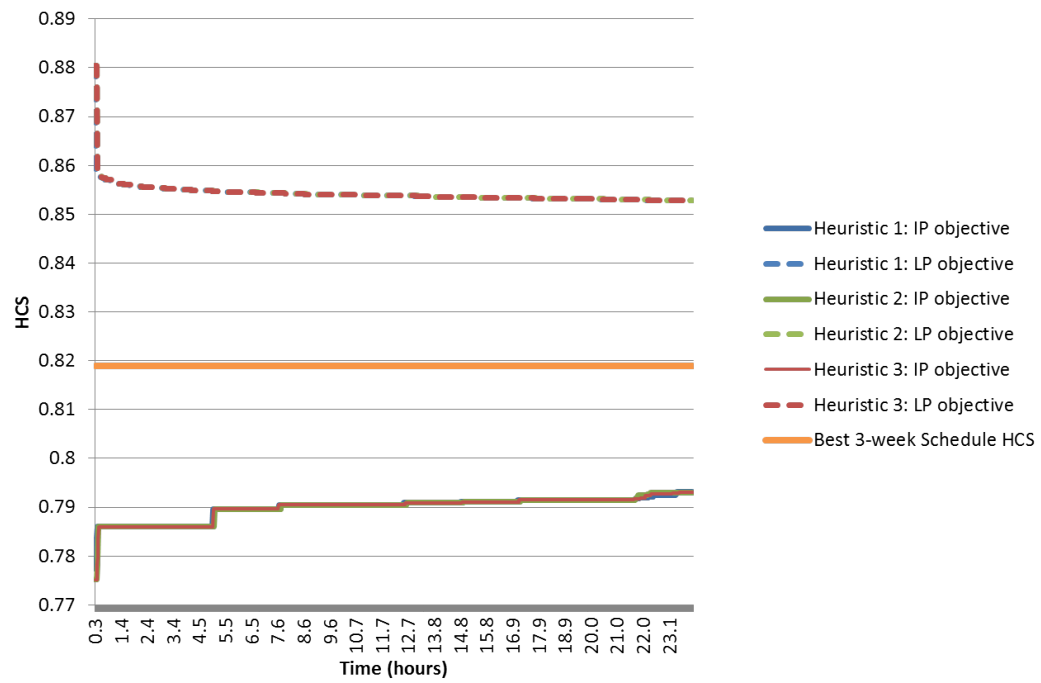
Figure 3: LP vs. IP objectives for variations to Heuristic 1 with a running time of 24 hours.
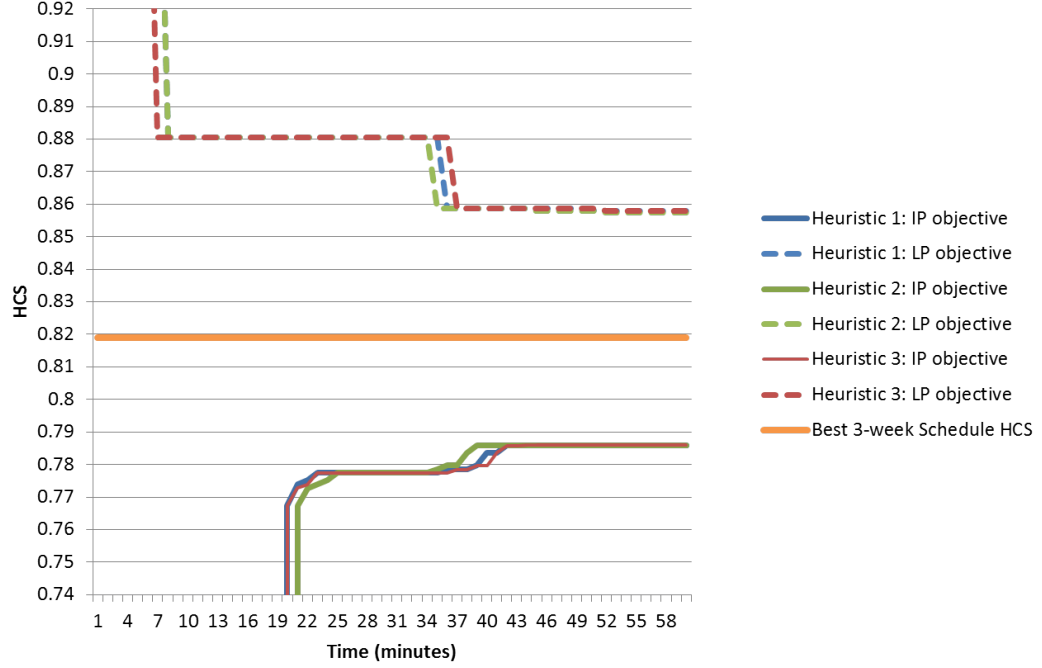
Figure 4: LP vs. IP objectives for variations to Heuristic 1 with a running time of 1 hour.

a similar feasible schedule. The variation which includes only 7 attendings achieves an HCS greater than 0.80 relatively quickly, and reaches an optimality gap of 6.67% after 24 hours. Variations which include 13 and 18 attendings, respectively, achieve similarly high HCS values as well, but require significantly longer running times to achieve those values (i.e., 4 and 18 hours, respectively), and optimality gaps at the end of 24 hours are 9.57% and 11.14%, respectively.

We conclude that, regardless of the number of attendings available for scheduling, this problem remains difficult, though slightly less so with only the minimum number of attendings.
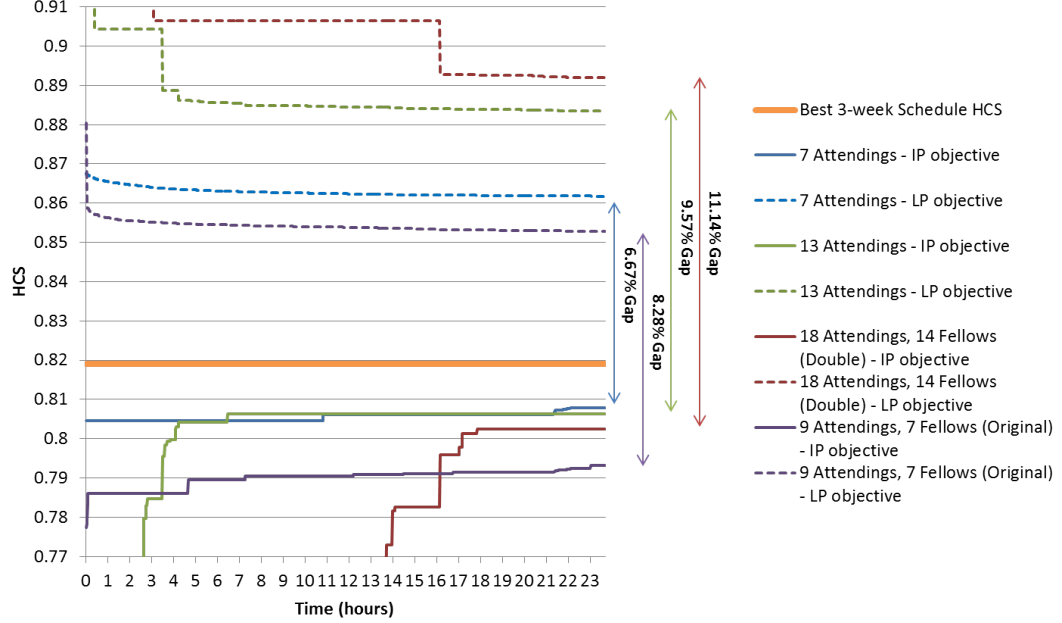
Figure 5: LP vs. IP objectives in Step 3 of Heuristic 1 for variations to the problem size with a running time of 24 hours.

2.5 Alternative Model

To determine if we could improve on the running time required by CPPS-MIP, we developed an alternative model using decision variables which consider the sequence of shift assignments for each physician, and compare this model to CPPS-MIP. By altering the formulation in this way, we hope to eliminate many constraints present in CPPS-MIP through preprocessing. New sets and parameters required by this model are listed in Table 10. The remaining sets and parameters are identical to those defined for CPPS-MIP and are listed in Tables 2 and 4.

Decision variables are listed in Tables 11 and 3. Primary decision variables $X_{i\theta}$ consider the sequence of assignments to shifts for each physician. Therefore, unlike CPPS-MIP, we assign a unique number to each possible shift in the schedule horizon. We number these shifts in order of the start time of each shift. For shifts that start at the same time, the longer shift is numbered first. For an example, see Figure 6.

We can eliminate many constraints present in CPPS-MIP thanks to the format

Table 10: Alternative Scheduling Model - Sets and Parameters

Set	Description
I, I_1, I_2	Set of all physicians, attending physicians, and residents, respectively
K_{lj}	Set of shifts during time period l on day j, $j \in J$, $l \in L$
K	Set of all shifts
Δ_t	Set of possible sequences for physicians in I_t, $t \in \{1,2\}$, such that each component of Δ_t corresponds to a 3-shift sequence
$\bar{\sigma}_{t\theta}$	Set of shifts in sequence θ for physicians in I_t, $\theta \in \{1,2,...,|\Delta_t|\}$, $t \in \{1,2\}$
Parameter	**Description**
$\sigma_{t\theta s}$	shift s in sequence θ for physicians in I_t, $s \in \{1,2,3\}$, $\theta \in \Delta_t$, $t \in \{1,2\}$

Table 11: Alternative Scheduling Model - Decision Variables and descriptions

Decision Variable	Description
$X_{i\theta}$	1 if physician i is assigned to shift sequence θ; 0 otherwise
S_{ik}	1 if physician i is assigned to shift k; 0 otherwise

of the primary decision variables in this alternative model. We do this by simply only allowing shift sequences in the sets Δ_t which do not violate these constraints. The eliminated constraints are (1c)-(1e), (1j)-(1k), (2e), and (2g)-(2i). The remaining constraints are included in this alternative model, with some slight modifications to accommodate the new variables. For easier comparison with CPPS-MIP, the constraints are numbered similarly, with each constraint label preceded by "A.".

Figure 6: Numbering of Shifts in Schedule Horizon for Alternative Model

$$\text{(A.1a)} \qquad \sum_{i \in P} \sum_{k \in K_{lj}} S_{ik} \;=\; v_{ljP} \qquad \forall l \in L,\, j \in J,\, P \in \wp$$

$$\text{(A.1b)} \qquad \sum_{k \in \bar{K}} \bar{L}_k S_{ik}/4 \;\leq\; 80 \qquad \bar{K} = \bigcup_{p \in \{m,...,m+27\},\, l \in L} K_{lp},\ \forall i \in I_2$$
$$m \in \{1, 2, ..., N - 27\}$$

$$\text{(A.1f)} \qquad D_{ij1} - S_{ik} \;\geq\; 0 \qquad \forall i \in I,\, k \in K_{lj},\, j \in J,\, l \in L$$
$$D_{ij2} - S_{ik} \;\geq\; 0 \qquad \forall i \in I,\, k \in K_{2j} \cup K_{1(j+1)}$$
$$j \in \{1, 2, ..., N - 1\}$$

$$\text{(A.1g)} \qquad D_{ij1} - \sum_{l \in L} \sum_{k \in K_{lj}} S_{ik} \;\leq\; 0 \qquad \forall i \in I,\, j \in J$$
$$D_{ij2} - \sum_{k \in \bar{K}} S_{ik} \;\leq\; 0 \qquad \bar{K} = K_{2j} \cup K_{1(j+1)},\ \forall i \in I,\, j \in J$$

$$\text{(A.1h)} \qquad \sum_{j \in \{m,...m+27\}} D_{ij1}/4 \;\leq\; 6 \qquad \forall i \in I_2,\, m \in \{1, 2, ..., N - 27\}$$

$$\text{(A.1i)} \qquad \sum_{j \in \{s,s+1,s+2\}} \sum_{k \in K_{2j}} S_{ik} \;\leq\; 1 \qquad \forall i \in I_2,\, s \in \{1, 2, ..., N - 2\}$$

$$\text{(A.2a)} \qquad \sum_{k \in K_{1s}} S_{ik} - Y_{ijn} \;\geq\; 0 \qquad \forall i \in P,\, s \in \{d, d+1, ..., d+n-1\}$$
$$d \in \{1, 2, ..., N - n + 1\},\, n \in B_P$$
$$P \in \wp$$

$$\text{(A.2b)} \qquad \sum_{i \in P} Y_{ijn} \;\geq\; e_{Pjn} \qquad \forall j \in \{1, 2, ..., N - n + 1\}$$
$$n \in B_P,\, P \in \wp$$

(A.2c) $Y_{ijn} + Y_{ikr} \leq 1$ $\forall i \in P, \, n \in B_P, \, k \in \{j, j+1, ..., j+n-1\}$

$$r \in \{1, 2, ..., N-k+1\}$$

$$j \in \{1, 2, ..., N-n+1\}$$

$$(j,n) \neq (k,r), \, P \in \wp$$

(A.2d) $Y_{ijn} + Y_{i(j+n)r} \leq 1 + \Omega_{ij}$ $\forall i \in P, \, n \in B_P, \, j \in \{1, 2, ..., N-n+1\}$

$$r \in \{1, 2, ..., N-j-n+1\}, \, P \in \wp$$

(A.2f) $\sum_{k \in K_{1_j}} S_{ik} - \left(\sum_{j \in J} v_{1jP} / |P| \right) - Z_P \geq 0$ $\forall i \in P, \, P \in \wp$

(A.3a) $R_{ijl} = \sum_{m \in \{1,2,...,\hat{N}\}} F_m D_{i(j-m)l}$ $\forall i \in I, \, l \in L$

$$j \in \left\{ \hat{N}+1, \hat{N}+2, ..., N \right\}$$

$\hat{R}_{ijl} \leq \sum_{k \in K_{l_j}} S_{ik}$ $\forall i \in I, \, l \in L$

$$j \in \left\{ \hat{N}+1, \hat{N}+2, ..., N \right\}$$

(A.3b) $\hat{R}_{ijl} \geq R_{ijl} + \sum_{k \in K_{l_j}} S_{ik} - 1$ $\forall i \in I, \, l \in L$

$$j \in \left\{ \hat{N}+1, \hat{N}+2, ..., N \right\}$$

$\hat{R}_{ijl} \leq R_{ijl}$ $\forall i \in I, \, l \in L$

$$j \in \left\{ \hat{N}+1, \hat{N}+2, ..., N \right\}$$

(A.3c)   **HCS** $= \left(\sum_{j \in \{\hat{N}+1,..,N\}} \sum_{l \in L} \left[\sum_{i \in I} \hat{R}_{ijl} / \sum_{P \in \{I_1, I_2\}} v_{ljP} \right] \right) / \left(2N - 2\hat{N} \right)$

(A.3d)   **Penalty** $= \sum_{i \in I} \sum_{j \in J} \Omega_{ij} +$

$$\sum_{l \in L} \sum_{j \in J} \sum_{i \in I} \sum_{k \in K_{l_j}} \left[(1 - O_{ijl}) \cdot S_{ik} + U_{ijl} \cdot S_{ik} \right]$$

The next set of constraints are new for this alternative model. Constraints (A.4a) ensure that variables S_{ik} correspond appropriately with variables $X_{i\theta}$. Constraints

(A.4b) guarantee that if a physician is assigned to a shift sequence that ends with shifts k and $\hat{k}$, then they must be assigned to another shift sequence which begins with k and $\hat{k}$. Finally, each physician must be assigned to a shift sequence which begins with dummy shift 0 (constraints (A.4c)) and each physician must be assigned to a shift sequence which ends with dummy shift $|\bar{K}|+1$ (constraints (A.4d)).

$$\text{(A.4a)} \quad \sum_{\theta \in \Delta_t : k \in \bar{\sigma}_{t\theta}} X_{i\theta} - S_{ik} \;\geq\; 0 \quad \forall i \in I_t,\, k \in \bar{K},\, t \in \{1,2\}$$

$$\text{(A.4b)} \quad \sum_{\theta \in \Delta_t : \sigma_{t\theta 2}=k,\, \sigma_{t\theta 3}=\hat{k}} X_{i\theta} - \;=\; 0 \quad \forall i \in I_t,\, k \in \{1,...,|\bar{K}|-1\},\, \hat{k} \in$$
$$\sum_{\theta \in \Delta_t : \sigma_{t\theta 1}=k,\, \sigma_{t\theta 2}=\hat{k}} X_{i\theta} \qquad \{k+1,...,|\bar{K}|\},\, t \in \{1,2\}$$

$$\text{(A.4c)} \quad \sum_{\theta \in \Delta_t : \sigma_{t\theta 1}=0} X_{i\theta} \qquad\;=\; 1 \quad \forall i \in I_t,\, t \in \{1,2\}$$

$$\text{(A.4d)} \quad \sum_{\theta \in \Delta_t : \sigma_{t\theta 3}=|\bar{K}|+1} X_{i\theta} \quad=\; 1 \quad \forall i \in I_t,\, t \in \{1,2\}$$

As in CPPS-MIP, the objective for this model is: **Maximize (HCS - Penalty)**.

We tested this alternative model for the same 51-week time horizon reported in Table 6. However, due to the large number of variables and constraints required by this problem, the model could not be constructed using CPLEX 12.4 due to memory issues. For a 10-week time horizon with the current pool of physicians, there are over 450,000 possible shift sequences per attending, and almost 800,000 possible shift sequences per fellow (after preprocessing).

To test the efficiency of this model on a small problem, we chose a time horizon of only 4 weeks, with 5 attending physicians and 3 fellows. However, no feasible schedule was found up to 24 hours of running the model, with or without the constraints enforcing the service block structure. When given a feasible schedule as a starting point (31.71% optimality gap), this alternative model can find an improvement to an optimality gap of 21.53% up to 24 hours with service blocks, and improves to an optimality gap of 18.70% up to 24 hours without service blocks. With this same problem,

41

CPPS-MIP can find an optimal solution (optimality gap < 0.01%) in approximately 2 minutes, with or without the service block structure. Thus, CPPS-MIP outperforms this alternative model, both on general problems as well as those with a service block structure.

One possible solution to this memory issue is using column generation techniques to reduce the number of shift sequences considered at any one time. The addition of constraints limiting assignments to overlapping sequences may strengthen the model as well. We leave these steps towards model improvement to future work.

2.6   Implementation

2.6.1   2011 - Challenges and Results

We generated an attending-only schedule for the PICU at Children's, to be implemented from July 1 - Dec. 23, 2011. Without needing to schedule fellows, and with the smaller time horizon, we simply used CPPS-MIP to create this schedule, rather than Heuristic 1. Modifying the MIP and deciding on an appropriate schedule required many iterations, where a MIP-generated schedule was presented to one or more physicians in the group, and issues were identified that required the addition of new constraints in the MIP (see additional constraint definitions in Appendix A). Attending-only schedules for the 6-month period could be generated by CPPS-MIP in approximately 2 hours (no additional improvement in the HCS was found up to 48 hours).

The heuristic-generated schedule reported in Table 6 assigned attendings to alternating night call shifts in the weeks they are on service. While this may be best for continuity, such a schedule would create fatigue and was therefore 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 due to the requirement for consecutive night call shifts. Therefore, we decided that no one 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 attending to call shifts on Monday, Friday, and Sunday, would be an acceptable assignment and good for continuity. With both of these call models, there were instances where a physician's requests were denied in order to keep with these preferred night call structures.

One physician in this group routinely requested to work call on Mondays and Saturdays during weeks they requested to be on-service. This violates the alternative call structure. We decided that in weeks where 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 on Monday through Friday, as well as weekend day and night shifts, that should be worked to satisfy various fellowship and other requirements. These requirements don't always coalesce well with the 7-day service blocks, but these requirements, as well as physician requests, should take priority over the 7-day service blocks and preferred call structure. Thus, the constraints enforcing the 7-day service blocks, as well as the preferred call structure, were altered to be soft constraints, where requests were given precedence in the objective function. Table 12 provides a description of the MIP versions used to generate a schedule at each iteration, as well as the HCS for each schedule.

The final MIP-generated schedule (based on Model Version 4, see Table 12) has

Table 12: Model Development to Find an Acceptable Schedule - descriptions and Handoff Continuity Score (HCS) for modified MIP versions used to generate schedules at each iteration

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/changes: (i) no consecutive call shifts, and (ii) requests take priority over 7-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/changes: (i) no consecutive call shifts, and (ii) requests take priority over 7-day service blocks	0.68
4	Model Version 3, plus the following change: requests take priority over call structure	0.69

been implemented, with some minor manual modifications. For comparison, an additional schedule was constructed manually for the same time period by a physician at Children's experienced in schedule construction, and with the same goal of maximizing continuity, resulting in an HCS of 0.63.

We used Heuristic 1 with these new constraints/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 .7578, still a 14.35% improvement over the HCS for the manually generated schedule constructed for the same time period.

2.6.2 2012 - Challenges and Results

Following the positive reception of the 2011 schedule, we generated an attending-only schedule for the PICU for July 1 - Dec. 23, 2012. No manual schedule was generated in parallel. The pool of attendings available for scheduling increased from 9 to 11 for this time period. Using CPPS-MIP, we found a schedule (optimality gap < 1%) in

Table 13: Attending Physician Requests: 2011 vs. 2012

	2011		2012		
Attendings	9		11		
	Requests	**Average**	**Requests**	**Average**	**Increase**[*]
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%

* Percent increase per person

approximately 35 minutes. This improvement in running time was largely due to the extensive personal requests made by attendings for this period, which limited possible variations to the schedule for purposes of HCS improvement. Table 13 reports the number of requests made for 2011 and 2012, for the same time periods (i.e., July 1 - December 23). Requests to be off service or off call increased per person on average by 173% and 165%, respectively.

This large increase in requests can be partly attributed to the change in formats that occurred for entering these requests. For the 2011 schedule, attendings made requests by month, with no predetermined format (see a subset of these requests in Figure 7). For the 2012 schedule, as a means of structuring the way 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 for communicating their requests. Figure 8 shows this form, and a sample of the requests that were made.

Further study is needed to better understand the reason for this large increase in requests for this period, whether due to the change in request formats or for reasons more unique to each physician. However, it is clear that with more requests, automated schedule generation is more efficient. Unfortunately, physician and patient continuity may suffer. When we remove all physician requests for this time period,

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 7: 2011 Physician Requests Entry Form - Subset

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	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 8: 2012 Physician Requests Entry Form - Subset

there is an HCS increase of 2.40% which can be significant in terms of continuity. The number of night call shifts without an on-service physician scheduled is reduced by 7.41%. If we randomly ignore 50% of requests, the HCS is increased by up to 1.24%, and the number of night call shifts without an on-service physician scheduled is reduced by up to 5.56% (calculated using 10 instances, 7 of which did not improve on the HCS with all requests considered within 2 hours). However, for each of these cases, a large number of requests are denied, as expected.

2.7 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 presented here facilitates resource optimization, constructing a feasible schedule in significantly less time than is needed to create a schedule by hand. Further, by considering familiarity among oncoming physicians, schedules are produced which maximize continuity. In conjunction with other methods for improving communication, schedules with greater continuity have the potential to enhance handoff efficiency.

The HCS provides a means of understanding the continuity of a physician schedule, and together with CPPS-MIP and Heuristic 1, allows for the construction of schedules which may improve handoff efficiency. However, there exist many factors not considered in the above analysis (e.g., bed occupancy, new admissions, disease acuity, and fatigue) which may impact the handoff process. Physician fatigue cannot be completely eliminated, regardless of the stringency of duty hour restrictions, and therefore communication failures created by exhaustion will still occur. Future work includes developing a score as an expansion of the HCS which accounts for familiarity among oncoming physicians as well as fatigue of physicians signing out.

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

CHAPTER III

AUTOMATED MEDICAL RESIDENT ROTATION AND SHIFT SCHEDULING TO ENSURE QUALITY RESIDENT EDUCATION AND PATIENT CARE

At academic teaching hospitals around the country, the majority of clinical care is provided by resident physicians. During their training, medical residents often rotate through various hospitals and/or medical services to maximize their education. Depending on the size of the training program, manually constructing such a rotation schedule can be cumbersome and time consuming. Further, rules governing allowable duty hours for residents have grown more restrictive in recent years [2], making day-to-day shift scheduling of residents more difficult. These rules limit lengths of duty periods, allowable duty hours in a week, and rest periods, to name a few.

In this chapter, we present two integer programming models (IPs) which (1) assign surgical residents to services over a 1-year time period, and given that assignment, (2) schedule the residents to night and weekend shifts over a one month period. These IPs solve these scheduling problems more effectively and efficiently compared to manual methods. The shift scheduling IP highlights the infeasibilities created by increased duty-hour restrictions placed on residents in conjunction with current scheduling paradigms.

3.1   Introduction

Resident physician rotation assignment is a challenging task which must balance the goals of maximizing resident experience and providing sufficient staff to provide excellent patient care, while adhering to the rules governing resident education. During

their five years of training after graduation from medical school, surgery residents at the Department of Surgery Residency Program at Emory University School of Medicine (EUSOM) spend time rotating through six hospitals so that they are exposed to the different areas of medicine that they may face as practicing surgeons. These hospitals include Emory University Hospital, Emory University Hospital Midtown, Grady Memorial Hospital, Children's Healthcare of Atlanta - Egleston Campus, the Veterans Affairs Medical Center in Atlanta, and Piedmont Hospital. Each hospital has one or more clinical services that must be staffed by residents in various levels of training (post-graduate years (PGY) 1-5), midlevel providers and attending surgeons.

Assigning residents to the different services in a way that meets patient care demands and satisfies educational requirements is a challenging and time-consuming task. Development of the surgery resident rotation schedule at EUSOM typically consumes 48 work-hours of a surgery residency program coordinator's time and 24 hours of the surgery residency program director's time. Additional time is often required throughout the year to adjust the rotation assignment schedule to accommodate unplanned events (e.g., absence of residents due to illness).

Day-to-day shift scheduling is made difficult by restrictions regarding work hours and rest periods. These rules and restrictions are constantly evolving, with the most recent changes in duty hours and rest periods instituted by the Accreditation Council for Graduate Medical Education [1] on July 1, 2011. While designed to reduce fatigue and improve patient safety, increased duty hour restrictions create challenges in scheduling to ensure that surgery residents meet educational requirements [12, 33, 52]. Current scheduling techniques coupled with ever more restrictive duty hour constraints limit surgical residents' exposure to the operating room [12] and to contact with their attending surgeon mentors in the operating room, the clinics, and on the wards. Further, staffing/work requirements for various surgical services are also made

more complex by these restrictions [33]. Restricted duty hours limit the amount of time senior and junior residents can spend mentoring those junior to them (such as medical students rotating on the clinical teams) in the training program, extends the burden of clinical care to teams with fewer members, and limits the pool of residents available for night and emergency coverage. The impact of this resource-constrained environment is potentially far-reaching.

As allowable duty hours for resident physicians continue to decrease, training hospitals must overcome scheduling shortages that are created. EUSOM has identified shortages in the clinical coverage that can be provided by resident physicians as a result of the new duty hours regulations, according to Keith Delman, Program Director for the General Surgery Residency Program in the Department of Surgery at Emory University School of Medicine. Resident physicians are trainees, and the number of trainees in a department is strictly regulated by the ACGME to ensure 1) that there is a sufficient amount of clinical exposure for training and 2) that resident physicians are not being employed solely for their function as providers of clinical service. The ACGME is unlikely to allow clinical departments additional residency positions; therefore, other strategies must be utilized to ensure coverage. Residency programs will need to redesign their schedules to accommodate duty hours regulations, and may need to employ midlevel providers (e.g., nurse practitioners and physician's assistants) to augment the staff needed to provide good patient care.

Idiosyncratic resident scheduling practices which cater to the individual preferences of attending surgeons create additional scheduling difficulties. For example, at EUSOM, 1st-year residents are responsible for their service's routine night and weekend shifts, according to Program Director Keith Delman. Trying to accommodate scheduling practices such as these, which are often not based on thoughtful study of the actual service needs, only adds to the complexity of scheduling.

To address the complex problem of assigning residents to services, and then creating day-to-day schedules which abide by ACGME duty hour restrictions, we have developed two integer programming (IP) models. The first IP, the Resident Rotation Assignment Integer Programming Model (RRA-IP) creates a feasible assignment of residents (PGY1-PGY5) to services over a one-year period. For day-to-day scheduling, we developed a simple integer programming model for the assignment of 1st-year residents to weekday, night, and weekend shifts, given the service assignments provided by RRA-IP. The goal in development of this IP, which we call the Surgical Resident Shift Scheduling Model, or SRSS-IP, was to better inform decisions regarding scheduling, as current scheduling paradigms may prove to be infeasible in light of constantly evolving duty-hour restrictions.

The remaining sections of this chapter are organized as follows. Section 3.2 provides a review of literature related to the problems we address. We provide details of RRA-IP in Section 3.3, followed by a discussion of efficiency results of this model in Section 3.4. We present a decision support tool we developed in Section 3.5, and discuss steps towards implementation in Section 3.6, as well as goals for future work. We present SRSS-IP in Section 3.7, followed by results and discussion in Section 3.7.1. We provide conclusions in Section 3.8.

3.2   *Literature Review*

To our knowledge, little work has been done to optimally assign residents to services, particularly the large number faced at EUSOM (28 services staffed by residents from EUSOM in 2012-2013). Franz and Miller (1993) present a solution approach for assigning medical residents to rotations over a 1-year period using a linear programming model and rounding procedure to find an integer solution [19]. Day et al. (2006) develop an integer programming model for assigning medical residents and fellows to services over a 2-week period [13]. Javeri (2011) develops an integer programming

model for assigning monthly rotations [24], but the discussion is limited to a presentation of the model definitions, and lacks validation and efficiency results. Unlike the prior work in this area, the main characteristics of the rotation assignment problem we face include:

- Assigning rotations over a 1-year period

- Education requirements:

    Required rotations

    Equivalent experience among residents of the same type

- Flexibility of model parameters

With the inclusion of extensive supply and demand constraints, as well as a user-friendly format for entering parameters specific to different institutions, our goal is to design a flexible, widely applicable scheduling system for rotation assignment which, given a set of inputs and preferences, can be modified to provide the output to a schedule of any complexity.

The impact of increased duty hour restrictions has been investigated extensively. Following the implementation of the 80-hour work week in 2003 (i.e., residents may work a maximum of 80 duty hours per week, when averaged over four weeks), Kort et al. (2004) conducted a survey of surgical residents and found that residents perceived negative impacts on continuity and safety of care, their operative experience, and their relationships with attending physicians [27]. Connors et al. (2009) review resident operative logs and find decreases in total cases as a result of increased duty hour restrictions, possibly implying inadequate operative experience [12]. A survey of surgery program directors reveals that duty-hour restrictions hinder clinical education opportunities and compromise patient safety [52]. Van Eaton et al. (2011) discuss challenges created due to the evolution of duty hour restrictions since 2003, as well as the new restrictions implemented in July, 2011 [50].

Increased duty hour restrictions complicate day-to-day shift scheduling of physicians, particularly medical residents, and mathematical modelling is a common approach to day-to-day scheduling problems [17], more increasingly those faced in a hospital setting. The most common application is nurse scheduling; see [9, 10] for comprehensive reviews. Applications in physician scheduling have also received extensive investigation. Beaulieu et al. (2000) develop a mathematical model for distributing shifts over a 6-month period to emergency room physicians [4]. Sherali et al. (2002) assign residents to night shifts using a mixed-integer program [41], and Rousseau et al. (2002) develop an easily customizable solution approach to the day-to-day physician scheduling problem [39]. Day et al. (2006) develop an integer programming model to generate a weekly shift schedule for a surgery residency program at a large academic program [13]. Topaloglu (2006) uses goal programming to assign emergency medicine residents to day and night shifts [45], and assigns residents to night shifts by applying sequential and weighted methods to a multi-objective optimization model which considers levels of seniority [46]. Ovchinnikov and Milner (2008) use Microsoft Excel and a solver program to create feasible schedules for radiology residents [35]. Cohn et al. (2009) assign residents to primary and backup night-call shifts in three hospitals using multiple nested IP models [11]. While considering flexible shift times and lengths, Brunner et al. (2010) schedule physicians over a 6 week period using a branch and price algorithm [8]. Topaloglu and Ozkarahan (2011) use column generation to minimize the sum of deviations from the desired service levels for a resident scheduling problem [47]. Javeri (2011) develops an integer programming model for scheduling day-to-day shift assignments and tests this model at 2 hospitals which included four different facilities [24].

The IP we developed for day-to-day shift scheduling solves the relatively simple problem of assigning 1st-year residents to weekday, night, and weekend shifts over a one-month period. This model allows us to test multiple scenarios very quickly,

to determine how current scheduling paradigms and duty hour restrictions impact schedule feasibility to determine if current scheduling practices are sufficient or if changes are needed. McCoy et al. (2011) find that it is possible to redesign a resident schedule to abide by the new restrictions without increasing staffing or decreasing patient admissions, but with possible negative impacts on continuity of care, quality of education, and time off between shifts [33]. Freiburg et al. (2011) conducted a survey to understand how residents perceive the benefits of 10 separate accommodation strategies used by training programs to adhere to resident work-hour restrictions [20]. The use of health IT is one of the most highly rated, as well as hiring of nurse practitioners and physician assistants, with respect to resident surgical education. The latter of these would be especially useful because a survey determined that surgical residents spend a large amount of time in noneducation activities [7] which could be redistributed if other staff were available, but this could require a large amount of human and fiscal capital, as determined by Mitchell et al. (2007) using a computer prediction model [34]. Our analysis of a one-month PGY1 shift schedule provides a means of understanding the impact of current scheduling paradigms combined with limiting duty hour restrictions, and reveals the possible need for midlevel providers to share daytime responsibilities with residents.

3.3 *Resident Rotation Assignment Model*

We developed an integer programming model, the Resident Rotation Assignment Model (RRA-IP), to assign residents to monthly rotations based on the demands of each service and the supply of residents during one year. Individual services request that a certain number of residents at certain levels in their training (i.e., PGY1 through PGY5) be assigned to that service for a period of time (typically one month at a time). Some of these requests represent a strong need for residents to cover the needs of the service, while others represent a desire for residents of a certain

Table 14: Emory Surgical Resident groups - 2012-2013

Resident Group	Number of Residents
PGY1-Categorical	9
PGY1-Prelim	6
PGY1-Urol.	3
PGY1-Neuro	3
PGY1-Ortho	5
PGY1-ENT	3
PGY1-CT	1
PGY1-Plastics	1
PGY2-Categorical	9
PGY2-Prelim	3
PGY2-GU	3
PGY2-CT	1
PGY3	9
PGY4	10
PGY5	7

training level to be assigned, but coverage is not necessary. The supply of residents available for assignment to each service is determined by a number of factors, the most important being educational requirements which specify the experiences residents of certain levels should receive to gain a well-rounded education. Further, it is important that residents of the same level and type receive equivalent (or similar) experiences throughout the year. Such resident groups are identified in Table 14.

Sets and parameters used by RRA-IP are defined in Tables 15 and 16. Decision variables in RRA-IP (listed in Table 17) represent the assignment of residents to specific services during a given time period. There are additional dummy variables to capture the case when there is no feasible assignment of residents to satisfy an individual service's demand.

The complete model formulation of RRA-IP is presented below. There are a series of both supply and demand constraints. On the supply side, there is a limit on the number of times that each resident may be assigned to an individual service

Table 15: Resident Rotation Assignment Model - Sets and Descriptions

Set	Description
L	Set of resident levels, e.g., $L=\{1,2,...,5\}$ to represent PGY1-PGY5
$R, R_1, ..., R_{|L|}$	Set of all medical residents, set of level 1 residents, ..., set of level $|L|$ residents, respectively. Note that $R = \bigcup_{l \in \{1,2,...,|L|\}} R_l$
K	Set of rotations
$S, S_1, ..., S_{|K|}$	Set of all services, set of services in rotation 1, ..., set of services in rotation $|K|$, respectively
P_l	Set of time periods for resident level l, $l \in L$. Note that residents of different levels may be assigned to services for time periods less than or greater than a one-month period, but possible time periods for each level are disjoint.
C	Set of service clusters such that each cluster contains services which represent equivalent experience
$\bar{C}$	Set of service clusters such that each cluster contains services for which a limit may be placed on the number of times that a resident may be assigned to that group of services
U_s	Set of services $s \in S$ such that a resident can not be assigned to any individual service in the set for two consecutive months
$\bar{U}_s$	Set of services $s \in S$ such that if a resident is assigned to any individual service in the set multiple times, then those assignments must be consecutive
I	Set of resident subcategories over all resident levels (e.g., $I = \{1,2,...\} = \{$PGY1-Categoricals, PGY1-Prelims,...$\}$)
G_i	Set of residents in resident subcategory i, $i \in I$
$\bar{I}_1$	Set of resident subcategories i such that all residents in group G_i require identical service assignments, $i \in I$. For each of these groups, if one resident in the group is assigned to a service, then all residents in that group must be assigned to that same service during some period in the year.
$\bar{I}_2$	Set of resident subcategories i such that all residents in group G_i require identical cluster assignments, $i \in I$. For each of these groups, if one resident in the group is assigned to a service in some service cluster, then all residents in that group must be assigned to a service in that cluster during some period in the year.
$\bar{I}_3$	Set of resident subcategories i such that all residents in group G_i require similar cluster assignments, $i \in I$. For each of these groups, if one resident in the group is assigned to a service in some service cluster, then all residents in that group should be assigned to a service in that cluster during some period in the year, if possible.

Table 16: Resident Rotation Assignment Model - Parameters and descriptions

Parameter	Description
Q_{rk}	Binary parameter indicating whether or not resident r must be assigned to a service in rotation k, $r \in R$, $k \in K$
$\bar{Q}_{rs}$	Binary parameter indicating whether or not resident r must be assigned to service s, $r \in R$, $s \in S$
$\bar{L}_{sl_1 l_2 p}$	Binary parameter indicating whether or not service s requests at least one resident from levels l_1 and l_2 in period p, $s \in S$, $p \in P_{l_1}$
$\bar{O}_c$	Integer paramater which sets a limit on the number of assignments of any resident to services in service group c, $c \in \bar{C}$
$W_{rsp}, \bar{W}_{rsp}$	Binary parameters indicating whether or not resident r requested to be assigned or not assigned to service s during period p, $r \in R_l$, $s \in S$, $p \in P_l$, $l \in L$
$\bar{D}_{si}$	Binary variable indicating whether or not a resident may be assigned to service s more than once in a time period of i months, $s \in S$, $i \in \{1, 2, 3, ..., 12\}$
D_{lsp}	Penalty incurred for not assigning the preferred number of residents of level l to service s during time period p, $s \in S$, $p \in P_l$, $l \in L$
B_{isp}	Bonus for assigning a resident from group i to service s during period p, $i \in I$, $s \in S$, $p \in P_l : R_l \cap G_i \neq \emptyset$, $l \in L$
Φ_{lsp}	Preferred number of residents of level l requested for service s during period p, $p \in P_l$, $l \in L$, $s \in S$
Γ_{lsp}	Maximum number of residents of level l that can be assigned to service s during period p, $p \in P_l$, $l \in L$, $s \in S$
V_{rp}	Binary parameter indicating the availability of resident r during period p, $r \in R_l$, $p \in P_l$, $l \in L$
M_{ls}	Maximum number of periods that a resident of level l can be assigned to service s, $l \in L$, $s \in S$
Ψ_i	Maximum number of periods that residents in group i can be assigned to services, $i \in I$
Θ_{lc}	Minimum number of periods that residents of level l must be assigned to some service in cluster c, $c \in C$, $l \in L$
$\tilde{P}_{p_1 l_1 p_2 l_2}$	Binary parameter indicating whether or not period p_1 for resident level l_1 intersects with period p_2 for resident level l_2, $l_1, l_2 \in L$, $p_1 \in P_{l_1}$, $p_2 \in P_{l_2}$

Table 17: Resident Rotation Assignment Model - Decision variables and descriptions

Decision Variable	Description
X_{rsp}	Binary variable, $= 1$ if resident r is assigned to service s during period p; 0 otherwise. $r \in R_l$, $s \in S$, $p \in P_l$, $l \in L$
Y_{lsp}	Integer variable ≥ 0, shortage below preferred number of residents of level l requested for service s during period p, $s \in S$, $p \in P_l$, $l \in L$
α_{rc}	Integer variable ≥ 0, for resident r and service group c, the maximum deviation between assignments of resident r to service group c from other residents in the equivalence group, $r \in R$, $c \in C$
β_{rsp}	Binary variable, $= 1$ if resident r's request to be assigned to service s during period p is denied; 0 otherwise. $r \in R_l$, $s \in S$, $p \in P_l$, $l \in L$
$\bar{\beta}_{rsp}$	Binary variable, $= 1$ if resident r's request to be assigned off service s during period p is denied; 0 otherwise. $r \in R_l$, $s \in S$, $p \in P_l$, $l \in L$

(constraints (1)). Residents may not be available for assignment during all periods (constraints (2)), and residents may be assigned to at most one service per time period (constraints (3)). For education purposes, it is required that some residents be assigned to a service in certain rotations at some point during the year (constraints (4)).

4th and 5th-year residents (PGY4's and PGY5's) must have identical assignments to other residents of their level throughout the year (constraints (5)). This means that if one PGY4 is assigned to service X at some point during the year, for example, then all PGY4's must be assigned to service X in some period during the year. 1st through 3rd year residents in the same resident subgroup (e.g., PGY1-Categorical, PGY2-GU, etc.) should have identical experience throughout the year (constraints (6)). This means that if one PGY1-Categorical resident is assigned to service X in service cluster Y, for example, then all PGY1-Categorical residents must be assigned to a service in service cluster Y at some point during the year. While experience

equality is strongly desired, there may be circumstances where this is not possible, and thus we have added the soft constraints (7). Parameters can be adjusted to allow for identical service cluster assignments, or just similar cluster assignments. Nine services are grouped into a general surgery cluster because they represent identical experience. These services are identified in Appendix B. The remaining services provide unique experiences, and thus there are no other equivalence group clusters.

There may be a limit on the number of times that a resident may be assigned to a specific groups of services (constraints (8)). For example, it may be undesirable that residents be assigned to too many night rotations. Some services are required for some residents (constraints (9)), and some service group clusters are required (constraints (10)). For example, PGY2-Urol residents should be assigned to some vascular surgery service and a specific general surgery service. In some cases, it may be preferred that if a resident is assigned to the same service more than once, then those assignments should be consecutive or not (constraints (11)-(12)). It may be required that a resident not be assigned to the same service in a 4-month period, for example, which is ensured by constraints (13). Residents may request that they be assigned to a specific service in a specific month (constraints (14)), or off a specific service during a certain month (constraints (15)).

Demand constraints include the following. Each service prefers a specific number of residents from each level during each period (constraints (16)). We have added the dummy variable Y_{lsp} to allow for the possibility that one of these preferences cannot be satisfied. If this is the case, a penalty is incurred in the objective function. There is a maximum number of residents from each level that may be assigned to each service in each period (constraints (17)). Constraints (18) ensure that if a service requests a resident from one of two different resident levels, then at least one of those resident levels is represented by residents assigned to that service.

The objectives are to:

i Minimize the violations of service demands for residents

ii Minimize the violations of resident preferences for services during/not during a specific month

iii Minimize the deviation from equality of experience for residents in the same equivalence group

iv Maximize the number of residents that are assigned to desirable services

It is not required that residents be assigned to some of the services for education purposes, but the experience they receive on some services is highly desirable, and thus there is a bonus when an assignment to one of those services is made. For example, there are three services that are highly desirable for PGY1-Categorical residents due to the education they gain on those services.

The complete model formulation is as follows:

$$\text{Minimize } \sum_{l \in L} \sum_{p \in P_l} \sum_{s \in S} D_{lsp} \cdot Y_{lsp} + \sum_{c \in C} \sum_{l \in L} \sum_{r \in R_l} \alpha_{rc} +$$
$$\sum_{l \in L} \sum_{p \in P_l} \sum_{r \in R_l} \sum_{s \in S} \left(\beta_{rsp} + \bar{\beta}_{rsp} \right) - \sum_{i \in I} \sum_{p \in \bar{P}_i} \sum_{r \in G_i} \sum_{s \in S} B_{isp} \cdot X_{rsp}$$

subject to:

$$(1) \qquad \sum_{p \in P_l} X_{rsp} \;\leq\; M_{ls} \quad \forall r \in R_l,\, l \in L,\, s \in S$$

$$(2) \qquad\qquad X_{rsp} \;\leq\; V_{rp} \quad \forall r \in R_l,\, p \in P_l,\, l \in L,\, s \in S$$

$$(3) \qquad \sum_{s \in S} X_{rsp} \;\leq\; 1 \qquad \forall r \in R_l,\, p \in P_l,\, l \in L$$

$$(4) \quad \sum_{p \in P_l} \sum_{s \in c} X_{rsp} \;\geq\; \Theta_{lc} \quad \forall r \in R_l,\, l \in L,\, c \in C$$

$$(5) \quad \sum_{p \in P_l : R_l \cap G_i \neq \emptyset, l \in L} (X_{rsp} - X_{\hat{r}sp}) \quad = \quad 0 \qquad \begin{array}{l} \text{for some } \hat{r} \in G_i,\ \forall r \in G_i \\ \text{such that } r \neq \hat{r},\ \forall i \in \bar{I}_1, \\ s \in S \end{array}$$

$$(6) \quad \sum_{p \in P_l : R_l \cap G_i \neq \emptyset, l \in L} \sum_{s \in c} (|G_i| \cdot X_{\hat{r}sp}) \quad \leq \quad 0 \qquad \forall \hat{r} \in G_i,\ i \in \bar{I}_2,\ c \in C$$
$$- \sum_{r \in G_i} \sum_{p \in P_l : R_l \cap G_i \neq \emptyset, l \in L} \sum_{s \in c} X_{rsp}$$

$$(7) \quad \sum_{p \in P_l : R_l \cap G_i \neq \emptyset, l \in L} \sum_{s \in c} (|G_i| \cdot X_{\hat{r}sp}) \quad \leq \quad \alpha_{\hat{r}c} \qquad \forall \hat{r} \in G_i,\ i \in \bar{I}_3,\ c \in C$$
$$- \sum_{r \in G_i} \sum_{p \in P_l : R_l \cap G_i \neq \emptyset, l \in L} \sum_{s \in c} X_{rsp}$$

$$(8) \quad \sum_{p \in P_l} \sum_{s \in c} X_{rsp} \quad \leq \quad \bar{O}_c \qquad \forall r \in R_l,\ l \in L,\ c \in \bar{C}$$

$$(9) \quad \sum_{p \in P_l} X_{rsp} \quad \geq \quad \bar{Q}_{rs} \qquad \forall r \in R_l,\ l \in L,\ s \in S$$

$$(10) \quad \sum_{p \in P_l} X_{rsp} \quad \geq \quad Q_{rk} \qquad \begin{array}{l} \forall r \in R_l,\ l \in L,\ s \in S_k, \\ k \in K \end{array}$$

$$(11) \quad X_{rsp} + X_{rs\hat{p}} - X_{rs(p+\tilde{p})} \quad \leq \quad 1 \qquad \begin{array}{l} \forall r \in R_l \\ p \in \{1, 2, ..., |P_l| - 2\} \\ \hat{p} \in \{p+2, p+3, ..., |P_l|\} \\ \tilde{p} \in \{1, 2, ..., \hat{p} - p - 1\} \\ l \in L,\ s \in \bar{U}_s \end{array}$$

$$(12) \quad X_{rsp} + X_{rs(p+1)} \quad \leq \quad 1 \qquad \begin{array}{l} \forall r \in R_l \\ p \in \{1, 2, ..., |P_l| - 1\} \\ l \in L,\ s \in U_s \end{array}$$

$$(13) \quad \sum_{i\in\{0,1,..,q\}} X_{rs(f+i)} + q \cdot \bar{D}_{sq} \quad \leq \quad q+1 \quad \forall r \in R_l, \ f \in \{1 + j \cdot (q+1)\},$$
$$q \in \{1,2,...\}, \ j \in \{0,1,...\}: f \leq$$
$$M_{ls}, \ \forall l \in L, \ s \in S$$

$$(14) \quad X_{rsp} + \beta_{rsp} \quad \geq \quad W_{rsp} \quad r \in R_l, \ p \in P_l, \ l \in L, \ s \in S$$

$$(15) \quad X_{rsp} - \bar{\beta}_{rsp} \quad \leq \quad \bar{W}_{rsp} \quad r \in R_l, \ p \in P_l, \ l \in L, \ s \in S$$

$$(16) \quad \sum_{r\in R_l} X_{rsp} + Y_{lsp} \quad \geq \quad \Phi_{lsp} \quad \forall p \in P_l, \ l \in L, \ s \in S$$

$$(17) \quad \sum_{r\in R_l} X_{rsp} \quad \leq \quad \Gamma_{lsp} \quad \forall p \in P_l, \ l \in L, \ s \in S$$

$$(18) \quad \sum_{r\in R_{l_1}} X_{rsp_1} + \sum_{r\in R_{l_2}} X_{rsp_2} \quad \geq \quad \bar{L}_{sl_1l_2p_1} \quad \forall l_1, l_2 \in L, \ s \in S, \ p_1 \in P_{l_1}, \ p_2 \in$$
$$P_{l_2} : \tilde{P}_{p_1l_1p_2l_2} = 1$$

3.4   Results

We attempted to use RRA-IP to generate a 1-year rotation schedule for July 2012 - June 2013, for all resident levels 1 through 5. Using CPLEX 12.4, a feasible schedule is found within 2 minutes, and after running the model for 24 hours, the optimality gap is 8.50% and does not improve rapidly (see Figure 9). Note that this and the remaining computational experiments reported in this chapter were carried out on either a system with a 2.27 GHz Xeon quad-core processor and 48 GB RAM or another with a 2.33 GHz Xeon quad-core processor and 12 GB RAM. While this solution may be considered acceptable for implementation, we find that the optimality gap is often higher when there is a slight change in equivalence requirements. The current equivalence requirements state that residents in groups PGY1-Categorical, PGY1-Prelim, up to PGY3's, all require similar experience to other residents within their
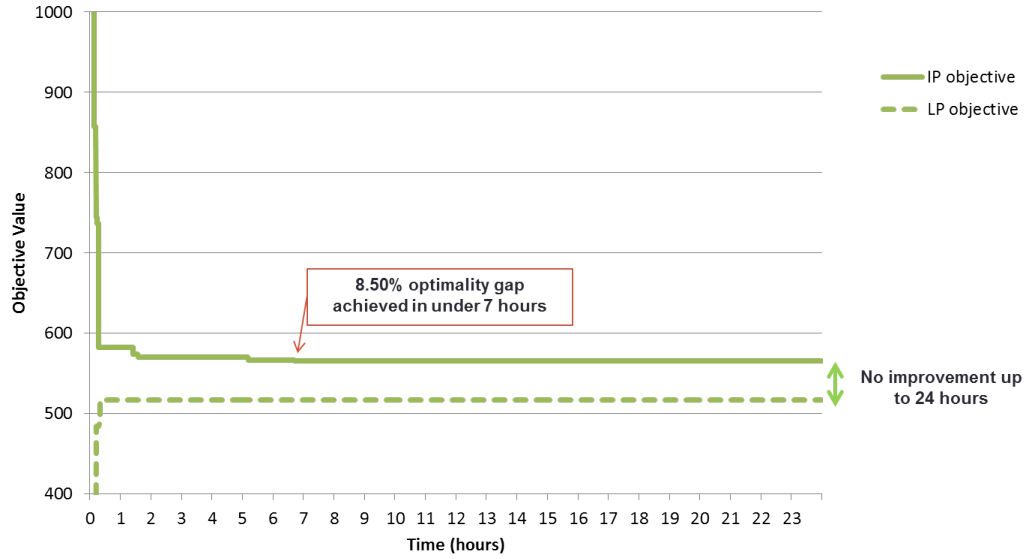
Figure 9: PGY1-PGY5 Resident Rotation Assignment Model with current equivalence requirements - LP vs. IP objectives up to 24 hours

respective groups. PGY4 and PGY5 residents require identical service assignments over the year to other residents within their respective groups. We modified this requirement to state that residents in all resident groups (i.e., PGY1-Categorical through PGY5) require similar experience to other residents within their group. Such a variation in the equivalence requirements may occur in other programs. With this slight modification, the model achieves an optimality gap of 15.65% in 2 hours, and does not improve up to 24 hours (see Figure 10).

We investigated how variations to the objective function might impact this solution time and the structure of the solutions. The rotation scheduling problem faced at EUSOM has four main objectives, with different priority levels:

1. Minimize weighted demand violations (highest priority)

2. Minimize deviation from equality of assignments for residents in the same group (high priority)

3. Minimize denied resident requests for service assignments during specific periods (low priority)

64
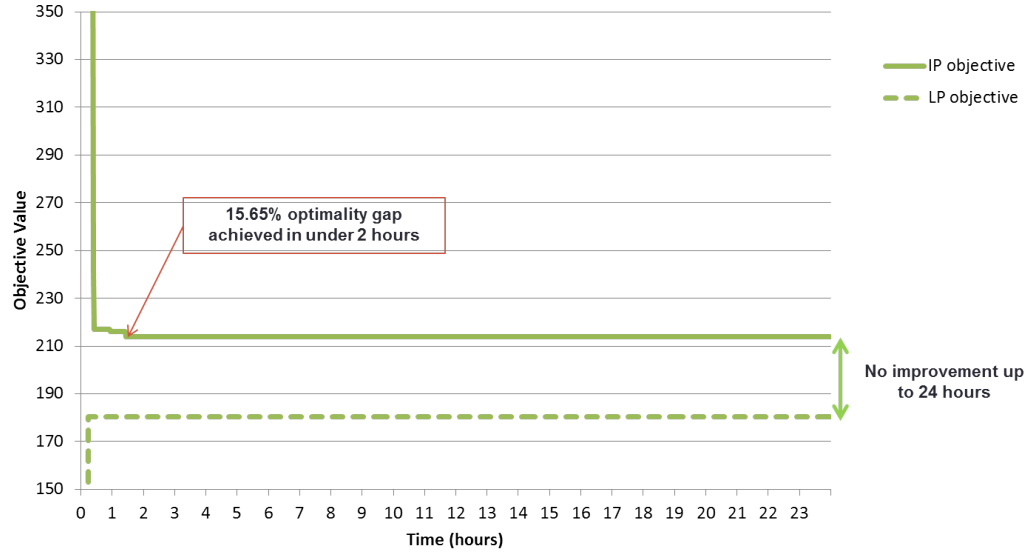
Figure 10: PGY1-PGY5 Resident Rotation Assignment Model with all resident groups requiring similar experience - LP vs. IP objectives up to 24 hours

4. Maximize assignments to desirable services (low priority)

Objectives 3 and 4 have approximately equal priority levels which are considerably lower than objectives 1 and 2. Different weights could be assigned to different components of the objective function, based on their relative importance. However, using only objectives 1 and 2, we generated 1-year schedules with RRA-IP assuming (1) current equivalence requirements, or (2) all resident groups requiring similar experience. For the case with the current equivalence requirements where some of the experiences have to be exactly the same, the optimality gap achieved by 24 hours is 4.51% (see Figure 11). However, if all resident groups require similar experience (i.e., exact equivalence is not required and we have some flexibility), the optimality gap after running RRA-IP for 24 hours is 33.38% (see Figure 12).

We investigated whether providing RRA-IP (with these two objectives) with an initial feasible solution would improve the running time for either case of equivalence requirements. To find such an initial solution, we used RRA-IP to generate a schedule with only objective one for each of the two sets of equivalence requirements. Optimal
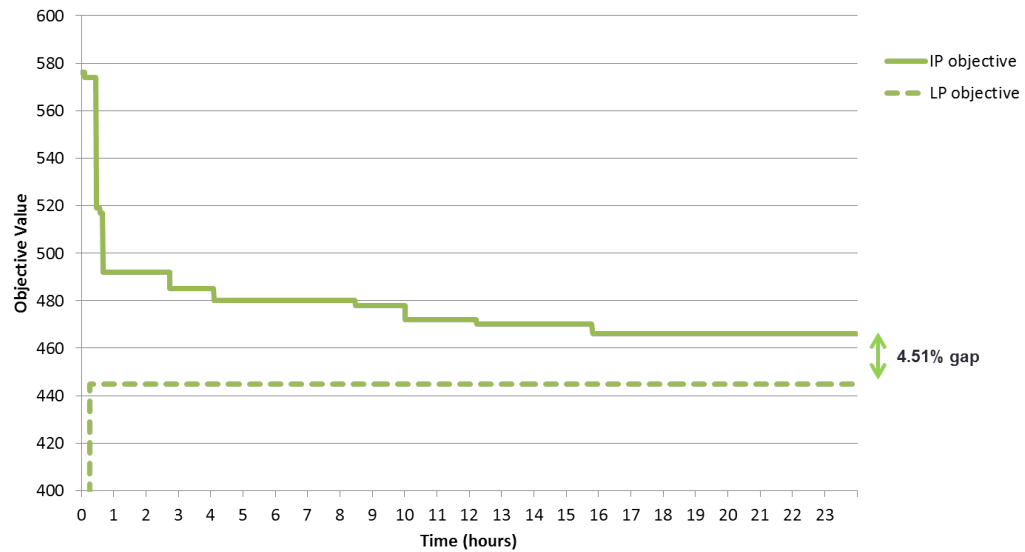
Figure 11: PGY1-PGY5 Resident Rotation Assignment Model with objectives 1 and 2 and with current equivalence requirements - LP vs. IP objectives up to 24 hours
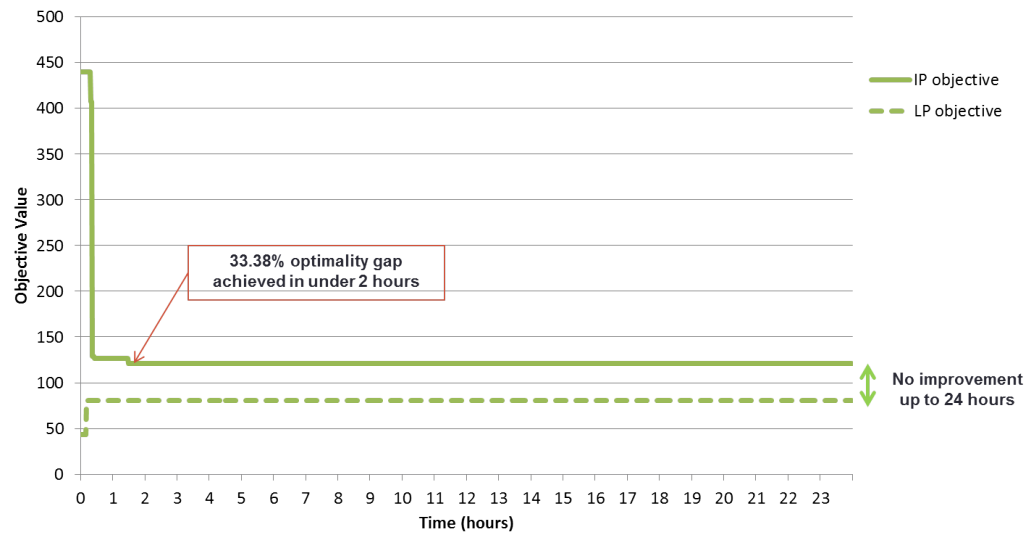
Figure 12: PGY1-PGY5 Resident Rotation Assignment Model with objectives 1 and 2 and with all resident groups requiring similar experience - LP vs. IP objectives up to 24 hours
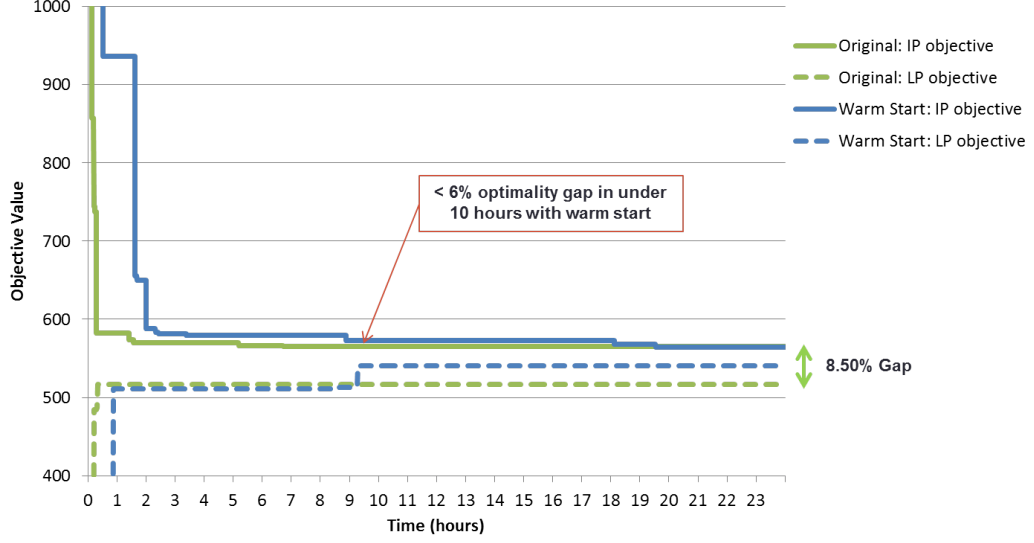
Figure 13: PGY1-PGY5 Resident Rotation Assignment Model with full objective function, with and without warm start, and with current equivalence requirements - LP vs. IP objectives up to 24 hours

solutions (optimality gap $< 0.01\%$) can be found with this modified objective function in less than one minute. These solutions have equivalent or fewer demand violations than the best solutions found within 24 hours using the full objective function, but include many more instances of uneven assignments to service clusters among residents in the same resident group (50 more instances with current equivalence requirements, 52 more if all groups require similar experience). Using these solutions as a starting point, RRA-IP, with objectives 1 and 2, cannot find improved solutions in 24 hours.

Using these same initial feasible schedules to warm start RRA-IP with the full objective function, we see that an improved solution can be found within 10 hours with the current equivalence requirements, with an optimality gap of 4.08% achieved by 24 hours (see Figure 13). If all resident groups require similar experience, providing the warmstart solution does not reduce the optimality gap by 24 hours.

We observe that objective 2 and related constraints greatly impact the solution time required by RRA-IP, but are crucial for producing solutions which have greater equality of assignments among residents in the same group. Therefore, we tested

a variation to the equivalence constraints which we will refer to as "balancing constraints" [24]. We define new variables Z_{icj} $\forall i \in I$, $c \in C$, and $j \in \{1,2\}$ which represent the maximum ($j = 2$) and minimum ($j = 1$) number of assignments of residents in group i to services in cluster c. We replace constraints (7) in the original formulation with constraints ($\bar{7}.1$) and ($\bar{7}.2$) to calculate these minimum and maximum values. Constraints (6) in the original formulation are removed, and replaced with constraints ($\bar{6}$) which force equality of upper and lowerbounds when identical clusters are required. Then we replace $\sum_{c \in C} \sum_{l \in L} \sum_{r \in R_l} \alpha_{rc}$ in the objective function with $\sum_{c \in C} \sum_{i \in \bar{I}_3} (Z_{ic2} - Z_{ic1})$, and we seek to minimize this value.

$$(\bar{6}) \qquad (Z_{ic2} - Z_{ic1}) \quad = \quad 0 \qquad \forall i \in \bar{I}_2,\ c \in C$$

$$(\bar{7}.1) \quad \sum_{p \in P} \sum_{s \in c} X_{rsp} \quad \leq \quad Z_{ic2} \quad \forall r \in G_i,\ i \in I,\ c \in C$$

$$(\bar{7}.2) \quad \sum_{p \in P} \sum_{s \in c} X_{rsp} \quad \geq \quad Z_{ic1} \quad \forall r \in G_i,\ i \in I,\ c \in C$$

We tested these new balancing constraints against the original equivalence constraints. Table 18 reports the solution time and optimality gaps for different combinations of the four objectives, for both equivalence requirement cases. Tests A and B can be solved very quickly (< 1 minute each) for each equivalence requirement, because deviation from equivalence is not a part of the objective function for these two tests. With the addition of objectives 2 or 2^+ (i.e., with balancing constraints and corresponding objective), we see an increase in solution time as expected. However, with current equivalence requirements, an optimal solution (gap $< 0.01\%$) can be found within 13 hours for Test F, which includes the full objective function. The remaining tests achieve optimality gaps below 5% by 24 hours.

When all resident groups require similar experience, however, the results are not as attractive for Tests C and E with regards to solution times. Fortunately, replacing the

original equivalence constraints with the new ones, an optimal solution (optimality gap $< 0.01\%$) can be found within 6 hours with the full objective (see Test F). With the original equivalence constraints, the optimality gap reaches 15.65% by 2 hours, but does not improve up to 24 hours.

Tests E and F, with current equivalence requirements, solve the problem faced at EUSOM. The efficiency of RRA-IP is much improved by inclusion of the balancing constraints, particularly for the case where residents in all resident groups require similar experience to other residents in their respective groups. This is due to the fact that, rather than attempting to minimize the deviation from equivalence for each individual resident in the scheduling pool, we are now looking at each resident group as a whole unit, with only two variables to represent the upper bound and lower bound on number of assignments of residents in that group to each service cluster.

Rules regarding equivalent experience may differ greatly for other programs, so we tested RRA-IP with various equivalence requirements. Results are reported in Table 19 with either equivalence constraints or balancing constraints.

We see that RRA-IP is relatively efficient in solving the problem in a majority of the scenarios. Test 3, which requires identical services for all resident groups, required the longest solution time, and does not rapidly approach optimality. Incorporating the balancing constraints rather than the original equivalence constraints produced improvements in solution times (or equivalent solution times) for all tests except Test 3. Fortunately, this is a scenario which is not likely to appear often in practice, due to the need to satisfy service demands which should take precedence over equivalence of assignments.

To better understand the impact of equivalence constraints vs. balancing constraints (and corresponding objectives) on the quality of the solutions produced, the following is a detailed discussion of the differences in the solutions reported in Table 19 for each set of equivalence requirements.

Table 18: Resident Rotation Assignment Model - Objective Function Analysis

Objective	A	B	C	D	E	F
				Tests		
1	X	X	X	X	X	X
2			X		X	
2^+				X		X
3		X			X	X
4		X			X	X

	Current Equivalence Requirements					
Solution Time	< 1 minute	< 1 minute	16 hours*	4 hours*	20 hours*	< 13 hours$^\alpha$
Optimality Gap	< 0.01%	< 0.01%	4.51%	4.49%	4.08%	< 0.01%

	All Resident Groups Require Similar Experience					
Solution Time	< 1 minute	< 1 minute	2 hours*	< 1 hour	2 hours*	< 6 hours
Optimality Gap	< 0.01%	< 0.01%	33.38%	< 0.01%	15.65%	< 0.01%

+ Balancing constraints and corresponding objective

* No improvement in optimality gap up to 24 hours

α < 1% optimality gap achieved in less than 1 hour

Table 19: Resident Rotation Assignment Model - Varying Experience Rules

Test	Description	Equivalence Constraints		Balancing Constraints	
		Solution Time	Optimality Gap	Solution Time	Optimality Gap
1	PGY1-3 require similar experience, PGY4-5 require identical services (current rules)	20 hours*	4.08%	< 13 hours[+]	< 0.01%
2	All resident groups require similar experience	2 hours*	15.65%	< 6 hours	< 0.01%
3	All resident groups require identical services	21 hours*	7.52%	22 hours*	12.58%
4	All resident groups require identical experience	7 hours*	5.98%	< 1.5 hours	< 0.01%
5	No similar/identical experience/service requirements	< 1 minute	< 0.01%	< 1 minute	< 0.01%
6	PGY1 require similar experience, PGY2-5 require identical experience	24 hours	0.63%	< 4 hours	< 0.01%

* No improvement in optimality gap up to 24 hours
+ < 1% optimality gap achieved in less than 1 hour

Test 1 PGY1-3 require similar experience, PGY4-5 require identical services (current rules)

The solution achieved with balancing constraints had 1 fewer demand violation than the solution from the model with equivalence constraints, and 4 fewer request violations. PGY1-Prelim residents had fewer total assignments, but the number of assignments per resident were equal as opposed to the solution with the equivalence constraints. PGY1-Urol residents and PGY2-Prelim residents had more total assignments, and assignments per resident were equal for each group with either balancing or equivalence constraints. PGY2-Categorical residents had fewer total assignments, and assignments per resident were not equal as opposed to the solution with the equivalence constraints. The remaining resident groups had equal total numbers of assignments (and equal numbers of assignments per resident) with balancing and equivalence constraints. There were some differences in assignments to service clusters for each resident group, other than PGY3, PGY4, and PGY5, and the solution produced with equivalence constraints included 5 instances of assignments to service clusters by residents in a resident group where equivalence was not achieved across all residents within that group, as compared to only 2 instances in the solution produced with balancing constraints.

Test 2 All resident groups require similar experience

The solution achieved with balancing constraints had an identical number of demand violations to the solution from the model with equivalence constraints, and 2 fewer request violations. PGY1-Categorical residents had fewer total assignments, but the number of assignments per resident were equal with balancing and equivalence constraints. PGY1-Prelim residents and PGY1-Urol

residents had more total assignments, and the number of assignments per resident (by group) were equal as opposed to the solution with the equivalence constraints. PGY2-Categorical residents had fewer total assignments, and assignments per resident were not equal as opposed to the solution with the equivalence constraints. The remaining resident groups had equal total number of assignments (and equal numbers of assignments per resident) with balancing and equivalence constraints. There were some differences in assignments to service clusters for each resident group, other than PGY3 and PGY5, and both solutions included 8 instances of assignments to service clusters by residents in a resident group where equivalence was not achieved across all residents within that group.

Test 3 All resident groups require identical services

The solution achieved with balancing constraints had 7 more demand violations than the solution from the model with equivalence constraints, and an identical number of request violations. PGY1-Prelim residents had fewer total assignments, PGY1-Urol residents had more total assignments, but assignments per resident were equal for each group with either balancing or equivalence constraints. The remaining resident groups had equal total numbers of assignments (and equal numbers of assignments per resident) with balancing and equivalence constraints. There were some differences in assignments to service clusters for most resident groups, and both solutions included 0 instances of assignments to service clusters by residents in a resident group where equivalence was not achieved across all residents within that group.

Test 4 All resident groups require identical experience

The solution achieved with balancing constraints had 5 fewer demand violations than the solution from the model with equivalence constraints, and an identical

number of request violations. PGY1-Categorical residents, PGY1-Urol residents, and PGY2-Prelim residents had more total assignments, PGY1-Prelim residents had fewer total assignments, but for each group, assignments per resident were equal with either balancing or equivalence constraints. The remaining resident groups had equal total numbers of assignments (and equal numbers of assignments per resident) with balancing and equivalence constraints. There were some differences in assignments to service clusters for some resident groups, and both solutions included 0 instances of assignments to service clusters by residents in a resident group where equivalence was not achieved across all residents within that group.

Test 5 No similar/identical experience/service requirements

The solution achieved with balancing constraints had an identical number of demand and request violations to the solution from the model with equivalence constraints, and demand violations were identical by service and time period. PGY2-Categorical residents had fewer total assignments, and assignments per resident were equal with either balancing or equivalence constraints. PGY2-Prelim residents had more total assignments, and assignments per resident were equal with balancing constraints unlike the solution found with equivalence constraints. The remaining resident groups had equal total numbers of assignments (and equal numbers of assignments per resident) with balancing or equivalence constraints. The number of assignments to each service cluster was almost identical for each resident group, but both solutions included 59 instances of assignments to service clusters by residents in a resident group where equivalence was not achieved across all residents within that group.

Test 6 PGY1 require similar experience, PGY2-5 require identical experience

The solution achieved with balancing constraints had an identical number of demand and request violations to the solution from the model with equivalence constraints. PGY1-Prelim residents had fewer total assignments, and assignments per resident were equal with either balancing or equivalence constraints. PGY1-Urol residents had more total assignments, and assignments per resident were equal with balancing constraints unlike the solution found with equivalence constraints. The remaining resident groups had equal total numbers of assignments (and equal numbers of assignments per resident) with balancing or equivalence constraints. There were some differences in assignments to service clusters for PGY1 resident groups, and the solution produced with equivalence constraints included 1 instance of assignments to service clusters by residents in a resident group where equivalence was not achieved across all residents within that group, as compared to 0 instances in the solution produced with balancing constraints.

For all cases except Test 3 (which did not achieve an improved solution time with the inclusion of the balancing constraints and corresponding objective function), we see either an equivalent number or a reduction in demand violations with the balancing constraints. Request violations are also identical or reduced with the balancing constraints. There is some variation for the solutions to each test with regards to the total number of assignments by resident group, but there were only two instances where assignments per resident within a group were not equal in the solution produced with balancing constraints, but were equal in the solution produced with equivalence constraints. For each test, the number of instances of assignments to service clusters by residents in a specific group where equivalence of assignments was not achieved was identical or improved in the solution produced with balancing constraints. We

conclude that the inclusion of balancing constraints and corresponding objective function produce solutions which are as good or better than solutions produced by the model with equivalence constraints for these instances with a maximum running time of 24 hours.

Requests made by residents impact the quality of any rotation assignment schedule. Thus, we conducted these 6 tests again, this time excluding residents' requests on or off specific services during specific time periods. We found that RRA-IP performed worse for Tests 1, 3, and 6, in terms of the time required to reach optimality, or the optimality gap achieved by 24 hours, whether we used equivalence constraints or balancing constraints (and the corresponding objective functions). The remaining tests performed better for one set of experience constraints (i.e. equivalence or balancing), but not for both. Demand violations were similar for each test, as compared between tests with equivalence constraints or balancing constraints, and also with the original tests which included resident requests. Instances of uneven assignments to service clusters by resident group were similar as well, but in some cases were increased when resident requests were not considered. These cases each had larger optimality gaps than when resident requests were considered. Despite these reductions in performance of RRA-IP without resident requests, it is important to note that for each test except Test 6, RRA-IP performed better with the inclusion of balancing constraints and corresponding objective function than with the equivalence constraints.

While inclusion of the balancing constraints and corresponding objective do improve efficiency, they do not produce identical solutions to the model with the equivalence constraints, assuming the model with equivalence constraints could be solved to optimality in a reasonable amount of time. Consider the following scenario of residents in some group G (containing 10 residents who require similar experience) assigned to services in some service cluster C:

- **Case 1:** 5 residents are assigned to service cluster C 2 times each, the other 5

residents are assigned to service cluster C 1 time each.

- **Case 2:** 9 residents are assigned to service cluster C 1 time each, the remaining 1 resident is assigned to service cluster C 2 times.

The impact of each of these cases on the objective function value is reported in Table 20, for both the model with equivalence constraints as well as balancing constraints. We see that the result is not identical, and the model with the equivalence constraints places higher value on having less deviation on an individual basis, rather than for a resident group as a whole.

Table 20: Resident Rotation Assignment Model - Balancing Constraints vs. Equivalence Constraints: Impact on Objective Function Value (Sample Scenario)

	Equivalence Constraints	Balancing Constraints
Case 1	25	1
Case 2	9	1

Whether this is a priority or not, we investigated the impact of using the IP with balancing constraints to produce a solution, and then using that solution to warm start the IP with equivalence constraints. Results from this analysis are presented in Table 21. We only considered Tests 1, 2, 4 and 6 since Test 5 performed similarly for both equivalence and balancing constraints, and Test 3 actually performed worse with balancing constraints than equivalence constraints. The right most column of the Table is the total solution time, which includes the time to generate the schedule with balancing constraints, as well as the time to improve on that schedule with equivalence constraints, up to a maximum running time of 24 hours for the latter case. We see that we approach optimality much faster with each test using the solution found with balancing constraints as an initial starting solution.

Table 21: Resident Rotation Assignment Model - Equivalence Constraints: Warm Start Analysis

Test	Description	RRA-IP		RRA-IP[+]		
		Solution Time	Optimality Gap	Solution Time	Optimality Gap	Total Solution Time
1	PGY1-3 require similar experience, PGY4-5 require identical services (current rules)	20 hours*	4.08%	1 hour*	0.18%	14 hours
2	All resident groups require similar experience	2 hours*	15.65%	< 4 hours*	11.60%	10 hours
4	All resident groups require identical experience	7 hours*	5.98%	< 10 minutes	< 0.01%	1.5 hours
6	PGY1 require similar experience, PGY2-5 require identical experience	24 hours	0.63%	< 10 minutes	< 0.01%	4 hours

+ with warm start solution from model with balancing constraints

* No improvement in optimality gap up to 24 hours

3.5   Decision Support Tool

We developed an Excel-based decision support tool for easy entry of all the parameters reported in Tables 15 and 16 to allow for flexibility in service requirements and demand from year to year. As an example, the set of residents in the scheduling pool for the 2012-2013 rotation schedule are shown as entered in the spreadsheet in Figure 14. The user simply enters the name of each resident available for scheduling, grouped by resident level and type.

Services needing staffing, as well as specific demands for each service, are entered into the tool, with the demand assumed to be identical for each month of the year (see Figure 15). If this is not the case, however, additional spreadsheets are available for entering month specific demand to allow for fluctuations in demand. If a service should be added or removed from this list, the user can simply enter or erase the name, respectively, and select the button labeled "Compile Data", and the table is automatically adjusted to allow for additional/fewer entries, with default values of "0" entered.

Additional necessary data is entered into the tool on a number of different spreadsheets, to include but not limited to:

- Number of time periods for scheduling of each resident level (default = 12)

- Maximum number of times any resident can be assigned to each service

- Service penalties for demand violations

- Resident bonuses for assignments to desirable services

These and the remaining spreadsheets include macros for automatically updating the forms when changes to the numbers of residents and services, or service demands, are made. Once a rotation assignment schedule is generated using RRA-IP given

Figure 14: Resident Rotation Assignment Decision Support Tool - Screenshot of Resident Data Entry Form

the entered parameters, the schedule is presented using an excel spreadsheet, with assignments listed by service and time period for each resident level.

3.6 Implementation

We used this decision support tool and RRA-IP to generate a rotation schedule for PGY1 residents for July 2012 - June 2013. Generating such a schedule required multiple iterations where we generated a schedule and experienced schedulers at EUSOM reviewed the schedule to identify any issues. Some of the issues that arose include:

- PGY1-Categorical residents should not have more than 2 night rotations in one academic year.

	A	B	C	D	E	F	G	H	I	J	K	L
1												
2	Compile Data											
3			PGY1		PGY2		PGY3		PGY4		PGY5	
4	Service	Rotation	preferred	max	preferred	max	preferred	max	preferred	max	preferred	max
5	GMHGenSA	1	1	1	0	0	1	1	0	0	1	1
6	GMHGenSB	1	1	1	0	0	1	1	0	0	1	1
7	EUHGenSA	1	1	1	0	0	1	1	0	0	1	1
8	EUHGenSB	1	0	1	1	1	0	0	0	0	1	1
9	EUHACS	1	0	0	1	1	0	0	1	1	0	0
10	EUHSurgOnc	1	1	1	0	0	1	1	1	1	1	1
11	EMHGenSA	1	1	1	1	1	0	0	1	1	0	0
12	PGHGenS	1	1	1	1	1	1	1	0	0	0	0
13	VAHGenS	1	0	0	1	1	0	0	1	1	0	0
14	EUHSelective	2	0	0	0	0	0	0	0	0	1	1
15	GMHTraumA	3	1	1	0	0	0	0	0	0	0	0
16	GMHTraumNight	4	0	0	0	0	1	1	1	1	0	0
17	EUHSICU	5	0	0	1	1	0	0	0	0	0	0
18	GMHSICU	6	1	1	2	2	0	0	0	0	0	0
19	EUHVascS	7	1	1	1	1	0	0	1	1	0	0
20	VAHVascS	8	0	0	0	1	1	1	0	0	0	0
21	HEHPedSA	9	3	3	0	0	0	0	1	1	0	0
22	GMHBurns	10	1	1	1	1	1	1	0	0	0	0
23	GMHPlastS	11	1	1	0	0	0	0	0	0	0	0
24	GMHUrol	12	1	1	0	0	0	0	0	0	0	0
25	EUHTranspS	13	1	1	0	0	0	0	1	1	0	0
26	EUHCTSurg	14	0	0	1	1	0	0	0	0	0	0
27	EUHNight	15	1	1	0	0	1	1	0	0	0	0
28	HEHPedSNight	16	1	1	0	0	0	0	0	0	0	0
29	VAHCTSurg	17	0	0	0	1	0	0	0	0	0	0
30	GMHQuadNight	18	1	1	0	0	0	0	0	0	0	0
31	EUHColorectal	19	1	1	1	1	0	0	0	0	1	1
32	GMHGenSC	20	0	2	1	1	0	0	1	1	0	0

Figure 15: Resident Rotation Assignment Decision Support Tool - Screenshot of Service Data Entry Form

- No one can have back-to-back night rotations

- If repeated by the same resident, some rotations should or should not be consecutive

- Misunderstandings of demands by services with regards to resident levels requested

- Some rotations absolutely must be covered

- Errors with fixed assignments of some residents

We corrected these issues which revolved around incorrectly entered data, and initial feedback regarding use of this model to create rotation schedules has been positive. Decision-makers at EUSOM hope to use this decision support tool for construction of the 2013-2014 rotation schedule.

3.6.1 Comparison to Manual Schedule

We received a manually constructed schedule for PGY2-PGY5 residents for July 2012 - June 2013. We compared manual and IP-generated PGY2 and PGY3 schedules (note that residents in groups PGY4 and PGY5 require identical service assignments, and thus there was little room for potential improvement). Both the manual and IP-generated schedules satisfied all demand. The manual schedule included much deviation from equivalence for residents in groups PGY2-Categorical, PGY2-Prelim, and PGY2-GU (there existed deviation from equivalence for 5, 4, and 6 service clusters, respectively). The IP-generated schedule included only two instances where assignments to service clusters were not equivalent across a resident group. Thus, RRA-IP can produce a schedule which satisfies all demand and meets all feasibility requirements, but in a faster time and with less deviation from equivalence for relevant groups than compared to manual methods.

One significant limitation here, however, is that requests for vacation time were not considered in RRA-IP. Including such requests could reduce this improvement in schedules, and we plan to include vacation requests in the next phase.

3.6.2 Future Work

After receiving individual resident requests for the 2012-2013 rotation schedule, we identified additional constraints necessary if we wish to accommodate requests with RRA-IP. Unfortunately, there is currently not a standardized format for residents of each level to make requests for vacation time, and not all requests have the same level of importance. Types of requests include resident A requests:

- to be or not to be assigned to service S in month M (included in RRA-IP).

- not to be assigned to specific rotations until later in the year.

- time periods for three vacation weeks, with different options for each vacation

listed in order of preference. Time periods can be for specific weeks, or just preferred months. Preferred vacation times possibly change with denial of vacation dates of highest preference.

- to be or not to be assigned to the same rotation with resident B.

We plan to provide a template for entering future request data which captures a majority of the types of requests listed above. We will allow the various types of requests to be prioritized by residents using a rating system to allow for greater flexibility in terms of preferences. Requests are gathered by multiple sources before being submitted to decision-makers, as residents represent multiple specialties. Therefore, the success of such a template relies heavily on the cooperation of the numerous groups providing data. If accepted by all groups, such a template could greatly improve the process of schedule generation due to the reduction in time required to decipher non-standardized request information.

The analysis reported herein relies on use of the commercial solver CPLEX. Next steps include the combining of the decision support tool we developed with a freely available solver such as GLPK [30] or OpenSolver [32] for regular use at the operational level. Further tool development will also be necessary to accommodate changes in availability and preferences that can occur throughout the year. This will require additional constraints in RRA-IP for recording a previous assignment schedule, as well as an addition to the objective function to minimize the level of change in any new schedule from the previous one. Such an update to the model will allow for constructing a new but similar schedule so as not to disrupt the assignments for all residents due to a perhaps small change in availability.

3.7   Surgical Resident Shift Scheduling Model

The solution given by the rotation assignment model RRA-IP provides input to the day-to-day scheduling model, which we will refer to as the Surgical Resident Shift

Scheduling Model, or SRSS-IP. Once residents are assigned to services on a monthly basis, their day-to-day schedule must be determined for the service he or she is assigned to. The day-to-day schedules for some subsets of the services are solved together, as some services share residents to staff night and weekend call shifts. SRSS-IP is identical for different groups of services, so we focus on one such group which we call the Emory group. Services in this group include EUH General Surgery A, General Surgery B, Surgery Oncology, ACS, Vascular Surgery and Night.

The preferred shift duration for residents assigned to services in the Emory group is 13 hours. Day shifts start at 6am and end at 7pm; night shifts start at 6pm and end at 7am. There is a one-hour overlap between night and day shifts dedicated for information exchange and care transition. In practice, it is often difficult for residents to leave on time due to workload, so we model this problem with flexible shift lengths in order to investigate the impacts that longer shifts would have on the ability to schedule residents so as to meet demand while complying with ACGME duty hour restrictions.

The objective of SRSS-IP is to maximize the minimum number of hours worked over all residents in order to maximize fairness in terms of hours scheduled, while staying within the duty hour regulations. Sets and parameters incorporated into the model are shown in Tables 22 and 23.

Table 22: Surgical Resident Shift Scheduling Model - Sets and descriptions

Set	Description
R	Set of all 1st-year medical residents assigned to the Emory group
S	Set of services in the group (i.e., $S = \{$1 (EUH General Surgery A), 2 (EUH General Surgery B), 3 (EUH Surgery Oncology), 4 (EUH Vascular Surgery), 5 (EUH ACS), 6 (EUH Night), 7 (EUH Weekend)$\}$)
K_s	Set of shifts required for service s, $s \in S$

Table 23: Surgical Resident Shift Scheduling Model - Parameters and descriptions

Parameter	Description
A_{rs}	1 if resident r is assigned to service s; 0 otherwise. $r \in R$, $s \in S$
N	Days in schedule horizon
H_t	Length of shift t, $t \in \{1, 2, ..., 2 \cdot N\}$

Decision variables in this model represent the determination that a resident is assigned to a specific shift. If a resident cannot be assigned to a certain shift, then a copy of the resident may be assigned. Defining decision variables in this way allows for clearer constraint definitions. An additional dependent variable represents the days in which residents are not scheduled for any shifts. We present formal descriptions of the variables in Table 24.

Table 24: Surgical Resident Shift Scheduling Model - Decision variables and descriptions

Decision Variable	Description
X_{rsk}	Binary variable, $= 1$ if resident r is assigned to service s for shift k; 0 otherwise. $r \in R$, $s \in S$, $k \in K_s$
Y_{rsk}	Binary variable, $= 1$ if a copy of resident r is assigned to service s for shift k; 0 otherwise. $r \in R$, $s \in \{1, 2, ..., 5\}$, $k \in K_s$
Z_{ra}	Binary variable, $= 1$ if day a of the month is free of duty for resident r; 0 otherwise. $r \in R$, $a \in \{1, 2, ..., N\}$
P	Integer variable, $=$ minimum number of shifts worked by any resident
$\bar{P}$	Integer variable, $=$ total number of shifts worked by all residents

The complete model formulation follows below. Constraints (1) ensure that variable P does not exceed the minimum number of shifts worked among all the residents. Constraints (2) set $\bar{P}$ equal to the total of all resident assignments, to reduce the

number of resident copies which are assigned. If a resident is assigned to a particular service, then that resident (or a copy of that resident) must be at their service for all weekday shifts (constraints (3)). For example, when the assignment model assigns residents to EUH General Surgery A, then they should be at that service for all weekday shifts from Monday through Friday (unless prevented due to a previous night shift assignment). One resident (and not a copy) should be assigned to each night and weekend shift (constraints (4)).

One PGY1 resident is assigned to EUH Night each period, which is a night float resident who covers the needs of all services in the group for a majority of the nights in the period. Standard practice in the Emory group is to assign other PGY1 residents assigned to services in the group to cover the remaining night call shifts, as well as weekend shifts, since duty hour restrictions prevent the resident assigned to EUH Night from working every night in the period. Weekend shifts are very similar to night shifts in that any resident assigned to a weekend shift must cover the needs of all services in the group.

Educational requirements and ACGME duty hour restrictions constitute the supply constraints. Residents must not be scheduled for more than six consecutive nights of night float (constraints (5)), hence the reason that the resident assigned to Emory Nights cannot work every night during the period. Duty periods of PGY1 residents must not exceed 16 hours in duration, a new restriction as of July, 2011 [1]. PGY1 residents must also have eight hours free of duty between scheduled duty periods, although 10 is desired. Since the preferred shift length is 13 hours, PGY1 residents cannot be assigned to two consecutive shifts (constraints (6)). Note that this constraint holds for shifts of longer lengths as well. All residents may work a maximum of 80 hours each week, when averaged over four weeks (constraints (7)). Residents must have at least one day free of duty each week when averaged over four weeks (constraints (8)). Constraints (9) and (10) force Z_{ra} equal to 1 if resident r has day

a free of duty, and equal to 0 otherwise.

$$\text{Maximize P}+\bar{P}$$

subject to:

(1) $\qquad \sum_{s\in S}\sum_{k\in K_s} X_{rsk} \geq \text{P} \qquad \forall r \in R$

(2) $\qquad \sum_{r\in R}\sum_{s\in S}\sum_{k\in K_s} X_{rsk} = \bar{P}$

(3) $\qquad X_{rsk} + Y_{rsk} = \text{A}_{rs} \qquad \forall r \in R,\ s \in \{1,2,3,4,5\},\ k \in K_s$

(4) $\qquad \sum_{r\in R} X_{rtk} = 1 \qquad \forall k \in K_t,\ t \in \{6,7\}$

(5) $\qquad \sum_{t\in\{k,k+2,...,k+12\}} X_{r6t} \leq 6 \qquad \forall k \in \{1,2,...,2N-12\},\ r \in R$

(6) $\qquad \sum_{s\in S}\sum_{t\in\{k,k+1\}} X_{rst} \leq 1 \qquad \forall k \in \{1,2,...,2N-1\},\ r \in R$

(7) $\quad \sum_{s\in S}\sum_{t\in\{k,k+1,...,k+55\}} H_t \cdot X_{rst} \leq 80\cdot4 \quad \forall k \in \{1,2,...,N-55\},\ r \in R$

(8) $\qquad \sum_{a\in\{t,t+1,...,t+27\}} Z_{ra} \geq 4 \qquad \forall r \in R,\ t \in \{1,2,...,N-27\}$

(9) $\sum_{s\in S}\left(X_{rs(2a-1)} + X_{rs(2a)}\right) + 2Z_{ra} \leq 2 \qquad \forall r \in R,\ a \in \{1,2,...,N\}$

(10) $\sum_{s\in S}\left(X_{rs(2a-1)} + X_{rs(2a)}\right) + Z_{ra} \geq 1 \qquad \forall r \in R,\ a \in \{1,2,...,N\}$

3.7.1 Results and Discussion

We used SRSS-IP to generate a one-month (31 day) schedule for the Emory group given varying values to the parameters H_t (i.e., the length of shift t). For each

scenario tested, the model solves to optimality within a matter of seconds. The Emory group is assigned 4 PGY1 residents each month (as requested) by RRA-IP. We found that it was possible to assign all night and weekend shifts to PGY1 residents assuming weekday service shifts, night call shifts, and weekend day shifts all have a duration of 13 hours. However, as mentioned previously, in practice, it is difficult for residents to complete a weekday service shift in 13 hours. The average duration of a weekday service shift for residents assigned to the Emory group is approximately 14.5 hours, but can be as high as 16 hours (the maximum allowed by ACGME duty hour restrictions). In light of this, we investigated the impact of variations to the preferred shift lengths of 13 hours. We also investigated the number of weekday service shift assignments that could be given to each resident with each parameter set tested (using Equation S.1).

$$(S.1) \quad \sum_{s \in \{1,2,3,4,5\}} \sum_{k \in K_s} X_{rsk} \quad \forall r \in R$$

While it may be possible to construct a feasible schedule which assigns PGY1 residents to all night call and weekend shifts, even with shift lengths up to 16 hours, if the result is a schedule which reduces resident education due to a small number of weekday service shift assignments given, then this is an undesirable result. Our goal in performing this analysis is to determine an answer to the following questions:

- At what shift lengths does it become infeasible to only schedule PGY1 residents to night and weekend shifts, assuming a preferred minimum number of weekday service shifts assigned to each resident?

- How do variations in ACGME duty hour restrictions impact these infeasibilities?

Figure 16 reports the largest number of weekday service shifts that can be assigned to each resident depending on the length of a weekday service shift (in hours). Note that night call and weekend shifts are assumed to be 13 hours in duration. Thus,
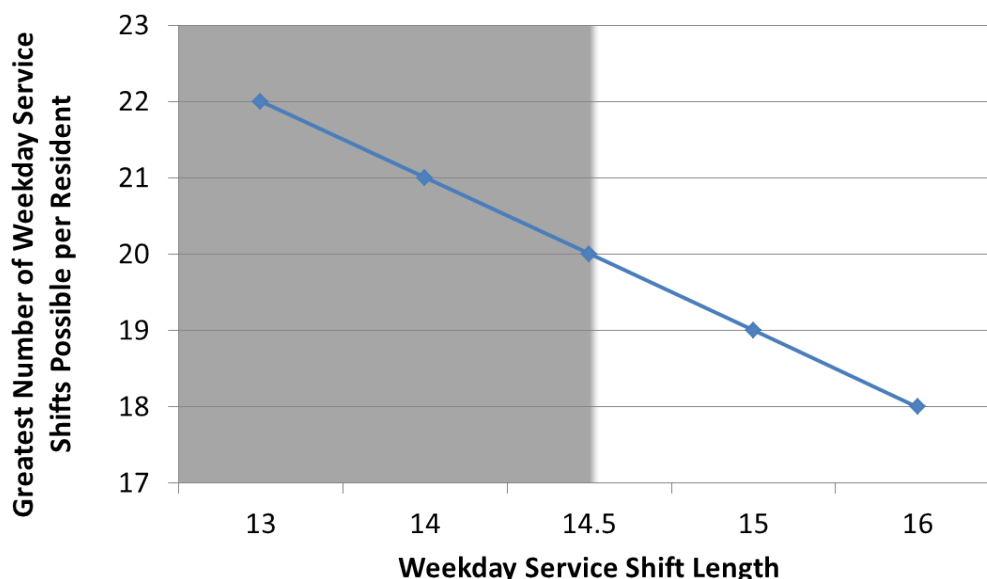
Figure 16: The maximum number of weekday service shift assignments that can be given to each resident depending on the length in hours of weekday service shifts. Night and weekend shifts are assumed to be 13 hours in duration. A typical weekday service shift has a duration of approximately 14.5 hours for residents assigned to the Emory group.

with 13 hour night and weekend shifts and 14.5 hour weekday service shifts (i.e., the average for the Emory group), it is possible to fill these night and weekend shifts with PGY1 residents and each resident can be assigned to 20 weekday service shifts. Shorter weekday service shifts are needed to assign a larger number of weekday service shifts to each resident.

ACGME duty hour restrictions mandate a maximum of 80 duty hours per week for all residents, when averaged over four weeks. In 2011, duty hour restrictions expanded to include a limit of 16 hours of consecutive duty for PGY1 residents. This limits PGY1 residents from working consecutive day and night shifts. Without this new restriction, night and weekend shift scheduling may be easier to accomplish with only PGY1 residents, but with an additional restriction limiting consecutive duty hours to 28 hours, assigning a PGY1 resident to a day and night shift in the same day would prevent them from being on duty the following day. If one of our goals with regards

to day-to-day shift scheduling is to keep residents on weekday service shifts as much as possible for education purposes, then removing this 16-hour restriction would not be an ideal solution. Thus, rather than further investigating the impact of this new 16-hour restriction, we analyzed the impact of the 80-hour duty week.

Given a specific number of weekday service shifts that should be assigned to each resident, Figure 17 reports the number of weekly duty hours that would be needed in order for PGY1 residents to be able to cover all night call and weekend shifts. We considered varying weekday service, night call, and weekend shift lengths. From this figure, we see that with a weekday service shift length of 15 hours (and 13 hour night call and weekend shifts), a feasible night and weekend shift assignment can be made with PGY1 residents within the 80-hour mandate. However, this is only true if 19 or fewer weekday service shifts are desired for each resident. For each resident to be given 20 weekday service shifts, 1 additional duty hour would be needed each week, on average.

Figure 17 also reports results for weekday service shifts lasting 16 hours, with night call and weekend shifts varying from 13 to 16 hours in duration. Note that since PGY1 residents may not be assigned to consecutive shifts due to the 16-hour maximum on-duty period mandate, the fact that these day and night shifts may overlap by more than 1 hour is irrelevant with regards to the calculation of the hours needed to construct a feasible night and weekend schedule with only PGY1 residents. For the extreme case of 16 hour shifts for all shift types, we see that a large number of weekly hours are needed. However, if night call and weekend shifts last 13 or 14 hours, respectively, with only a few additional hours per week on average (above the 80-hour limit), residents may gain 1 or more weekday service shifts per month. These weekday service shifts provide a much richer educational experience than night and weekend call shifts.
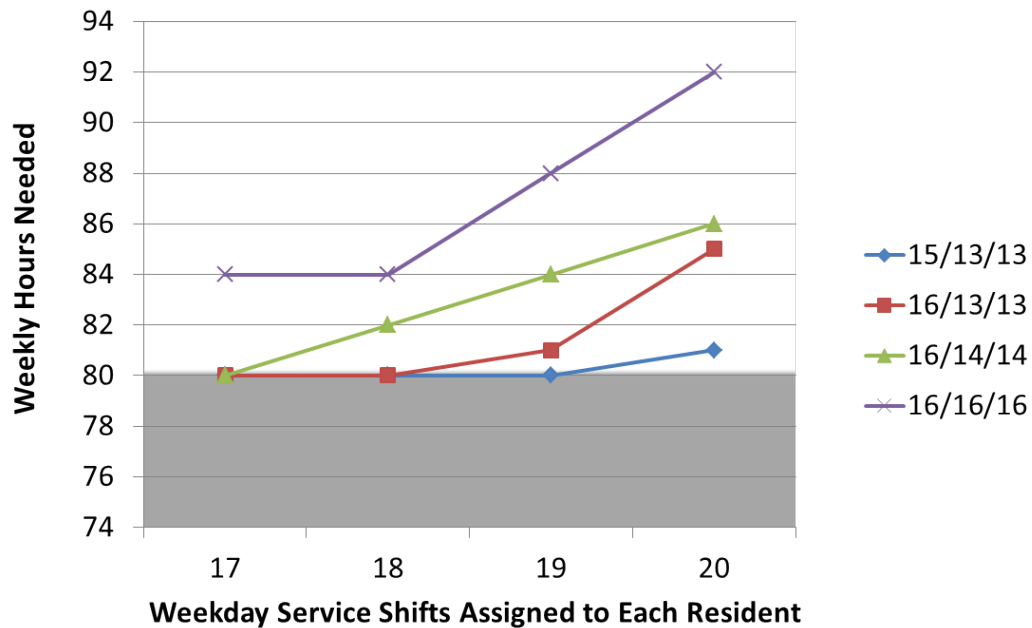
Figure 17: The weekly duty hours needed per resident for assignments to 17, 18, 19, and 20 weekday service shifts in a month, respectively. Duty hours are reported for 4 shift length variations listed as X/Y/Z to represent weekday service shifts, night shifts, and weekend shifts of durations X, Y, and Z hours, respectively. The ACGME mandates a maximum of 80 duty hours per week.

An alternative to attempting to extend the weekly duty hour limit is to reduce the weekday service shift workload for PGY1 residents, so that these residents may be assigned to as many weekday service shifts as preferred, while still being able to cover all night call and weekend shifts. One means of reducing this workload is to hire additional staff to manage many of the activities performed by these residents during a weekday service shift but which do not enrich their education. Alternatively, 2nd through 5th year residents can be assigned to night call and weekend shifts as well. This is likely the simplest solution, but requires a change in paradigm among supervising physicians in the Emory group accustomed to a day-to-day schedule with PGY1 residents covering all night and weekend shifts.

3.8 Conclusions

One of the greatest challenges in graduate medical education is navigating the complexities associated with resident scheduling. This is compounded by new and continuously evolving rules with respect to resident education, work hours, night and emergency call, and supervision. If a systematic, reproducible, widely applicable decision support tool/software can be designed to "solve" these complex manpower distribution puzzles, the process could be rapidly adopted across every residency training group in the country. It would simplify and streamline a process which can take several days worth of person-hours of time.

The efficient assignment of residents to services and effective scheduling while on those services could positively impact patient care [27], improve compliance with the duty-hour restrictions [44], improve resident education [20], and enhance patient care through increased supervision and reduced fatigue. The quality of patient care delivered by residents is affected by various factors. Continuity of care, resident fatigue, and proper resident supervision all play a role in patient health outcomes. By more efficiently assigning the appropriate PGY levels of residents to clinical services, it may be possible to improve both patient care and resident education on those services.

RRA-IP can be used in subsequent years to generate feasible schedules. Educational requirements do not change frequently, so the IP should be able to stay up to date with only the occasional small changes. Use of RRA-IP provides solutions which are better than manually generated schedules with respect to objectives, and can reduce the number of hours that the scheduler must work to generate a feasible schedule for the year, freeing up a staff member's time and energy to be spent on other valuable activities. Among residents, there is a perceived sense of fairness in having an objective model rather than a person create schedules for the year, and so residents may be less likely to complain in general. Finally, such a model can be

adapted in different hospital systems to help them create feasible schedules.

SRSS-IP highlights the impact of shift lengths and duty hour restrictions on physician scheduling, particularly with regards to resident education. Current scheduling practices which require that 1st-year residents be assigned to all night and weekend shifts are still feasible under new duty hour restrictions, but there may be negative impacts on resident education due to a reduction in the number of weekday shifts that can be worked by residents. Resident education can be improved by slightly relaxing ACGME duty hour restrictions, by reducing daily workloads, or by changing the preferred scheduling practices.

CHAPTER IV

ASSIGNING PHYSICIANS TO MULTIPLE TASKS TO MAXIMIZE SPACE UTILIZATION AND ACHIEVE EQUITY

In the final chapter of this thesis, we study the problem of assigning physicians to multiple tasks to maximize space utilization and ensure fairness. We outline the methodology we use to solve this problem as well as efforts towards implementation in a real-world setting. We discuss planned future work, and conclude with our contributions.

4.1 Introduction

When scheduling staff, regardless of the industry, there are often multiple objectives including minimizing cost, maximizing resource utilization or staff preferences, or maximizing fairness of assignments or quality of life. Balancing the necessary and sometimes conflicting objectives when manually creating a staff schedule can be a challenging and cumbersome task, particularly when attempting to optimize objectives within feasibility requirements. Scheduling can be further complicated by individual staff preferences and restrictions which create a non-homogeneous scheduling pool.

While many industries require staff to work in a single location, it is not uncommon to see instances where workers are expected to complete tasks in different locations on different days, particularly in hospital settings. Within many medical specialties, physicians provide services in operating rooms, the emergency room, as well as outpatient clinics, just to name a few, and physicians from one specialty must

often share these spaces with other physicians. Space capacities can restrict how assignments to various locations can be made, with or without sharing of space among various groups.

In this chapter, we present an integer programming model (IP) for assigning a group of heterogeneous physicians to multiple tasks with varying demand and space availability. Section 4.2 provides specific details of this problem faced in the Department of Gynecology and Obstetrics at Emory University Hospital (Emory OB/GYN) [15], and gives a discussion of previous relevant literature. Section 4.3 presents details of the IP we developed for solving this problem which considers multiple objectives. We used this model to construct a 6-month schedule for a group of physicians at Emory OB/GYN. In Section 4.4, we discuss our results and provide details of the efficiency of our model with alternative objective functions. We discuss efforts towards implementation and future work in Section 4.5, and provide conclusions in Section 4.6.

4.2 Problem Description and Literature Review

Emory OB/GYN has 39 faculty members (2011-2012), approximately 25% of whom are generalists who cover a selection of daytime activities including Labor and Delivery (L&D), the Emergency Room (ER), two outpatient clinics (which we will refer to as Clinic 1 and Clinic 2), and Surgery. Some of these activities, such as L&D and the ER, have fixed demand each day, while others (e.g., the clinics) should be staffed by available physicians if possible depending on physician and exam room availability, but assigning staff is not necessary. Some physicians in and outside this group of generalists have fixed time periods for which they provide coverage in the clinics, limiting the flexibility of further assignments to these clinics due to space availability. Not all generalists can be assigned to each activity. Thus, the problem of assigning physicians can be quite complex.

The objective in assigning physicians to these daytime activities has many parts, including maximizing space utilization within the clinics while balancing assignments of physicians to those clinics. Further, we seek to construct a fair schedule with regards to the number of day and night L&D calls, weekends, holidays, and ER assignments given to each physician.

Previous work in scheduling daily hospital operations is extensive [5, 17], including problems in nurse scheduling [9, 10] and assignment of medical residents to shifts [11, 13, 41, 45, 46]. Mansdorf (1975) develops a mathematical model with the objective of determining the most fair and efficient allocation of staff among a group of outpatient clinics [31]. Hodgson et al. (1977) present an integer programming approach to the problem of assigning physicians from different specialties to cover multiple outpatient clinics which share space [22]. Similar to our problem, they consider other commitments such as operating room schedules of the physicians. Isken et al. (1999) present a simulation framework for outpatient obstetrical clinics, and consider the assignment of exam rooms to specific physicians [23]. Other applications such as academic course scheduling also consider utilization of available space when making assignments [14].

Assignments to various unique tasks are not uncommon in hospital settings. For education purposes, medical residents are often assigned to different medical services during their training, most commonly for a one-month period at a time, and multiple approaches have been taken to solve this problem. Franz and Miller (1993) use a rounding procedure to assign medical residents to such services over a one-year period [19], and Day et al. (2006) use integer programming to assign medical residents and fellows to services [13]. Belien and Demeulemeester (2007) construct a trainee schedule at a hospital using column generation, and compare two decomposition methods based on (1) physicians available for scheduling and (2) activities trainees can be assigned to [6]. Javeri (2011) develops an integer programming model for making

96

mothly service assignments while considering fairness among residents of the same training level [24].

Rostering problems in hospitals often must incorporate multiple objectives, and rarely is the scheduling pool of physicians homogeneous. Maenhout and Vanhoucke (2010) develop a branch and price approach for assigning heterogeneous nurses to shifts while considering multiple objectives including equity and cost [29]. Li et al. (2012) use goal programming and meta-heuristic search to create a nursing roster given nine conflicting objectives [25].

The main characteristics of the problem of scheduling physicians for Emory OB/GYN include:

- Assignments to multiple activities

- Variations in clinic space available each day

- Heterogeneous physicians in scheduling pool

- Multiple objectives

In combination, these characteristics create a complex scheduling environment which has yet to be addressed in the literature to the best of our knowledge. Mathematical modelling can provide a tool for automating schedule generation while considering the unique characteristics of this problem. This chapter presents an integer programming model, the Physician Scheduling model (PS-IP), which assigns physicians to multiple tasks over the schedule horizon while considering space availability, differentiates between unique characteristics of each physician in the scheduling pool, and incorporates multiple objectives. PS-IP is highly efficient, and can greatly improve the process of schedule generation. While the model formulation is specific to the problem faced by Emory OB/GYN, the methodology could be applied to other units and institutions with slight modifications to accommodate varying physician and institutional preferences.

4.3　An Integer Programming Model

Table 25 displays the possible daytime assignments that can be given to each physician in the scheduling pool for Emory OB/GYN. If a physician is assigned to ER/Clinic/Surgery, then there are several distinct possible subassignments that can be given, shown in Table 26.

Table 25: Emory OB/GYN Daytime Assignments

	Assignment
1	Labor and Delivery
2	ER/Clinic/Surgery (weekdays and non-holidays only)
3	Off Duty

Table 26: Emory OB/GYN ER, Clinic, and Surgery Weekday Assignments

	Assignment
4	Emergency Room and Clinic
5	Clinic 1 - Full Day
5	Clinic 1 - Morning
7	Clinic 1 - Afternoon
8	Clinic 2 - Full Day
9	Clinic 2 - Morning
10	Clinic 2 - Afternoon
11	Surgery

The sets, parameters, and primary decision variables of PS-IP are presented in Tables 27, 28, and 29. Additional dependent variables, which are used mainly to form the objective function, are given in Table 30.

4.3.1　Model Constraints

Each physician must be given exactly one daytime assignment each weekday and holiday (constraints (1)). Possible assignments are L&D day call, ER/Clinic/Surgery, and Off Duty (see Table 25). On weekends (non-holidays), physicians can either be

98

Table 27: Emory OB/GYN Physician Scheduling Model - Sets and descriptions

Set	Description
P_t	Set of physicians available for assigning to task t, $t \in \{1, 2, ..., 12\}$
D	Set of days in schedule horizon
C_{pt}	Set of days in schedule horizon for which physician p can be assigned to task t
D_w, D_e, D_h	Set of weekdays (excluding holidays), weekends (excluding holidays), and holidays in schedule horizon, respectively
D_m	Set of mondays in schedule horizon
D_i	Set of days in month i in schedule horizon, $i \in \{1, 2, ..., 6\}$
D_s	Set of days in the 2nd week of each month in the schedule horizon
H_i	Set of holidays in week i of the schedule horizon, $i \in \{1, 2, ..., \lvert D \rvert /7\}$
A	Set of possible daytime assignments (i.e., $A = \{1, 2, 3\}$)
T	Set of ER/Clinic/Surgery assignments (i.e., $T = \{4, 5, ..., 11\}$)
L_{ij}	Set of possible morning (i=1) and afternoon (i=2) assignments at clinic j (inclusive of full-day assignments), $j \in \{1, 2\}$
J	Set of availability periods, defined as the disjoint time periods for which no physicians leave or join the group during the period
F_j	Set of days in availability period j, $j \in J$

given a daytime assignment or L&D night call (constraints (2)). If a physician is assigned to night call, then she must be off duty the next day (constraints (3)). No physician can be assigned to night call if they are off duty (constraints (4)).

$$(1) \qquad \sum_{t \in A: p \in P_t, d \in C_{pt}} X_{ptd} \;=\; 1 \quad \forall p \in P, d \in D_w \cup D_h$$

$$(2) \quad \sum_{t \in A: p \in P_t, d \in C_{pt}} X_{ptd} + Y_{pd} \;=\; 1 \quad \forall p \in P, d \in C_{p(12)} \cap D_e : d \notin D_h$$

$$(3) \qquad X_{p3(d+1)} - Y_{pd} \;\geq\; 0 \quad \forall p \in P, d \in \{1, 2, ..., \lvert D \rvert - 1\} \cap C_{p(12)}$$

$$(4) \qquad X_{p3d} + Y_{pd} \;\leq\; 1 \quad \forall p \in P, d \in C_{p(12)}$$

Table 28: Emory OB/GYN Physician Scheduling Model - Parameters and descriptions

Parameter	Description
R_{pd}	1 if physician p requests assignments to all weekend Labor and Delivery night call shifts in the week beginning with day d; 0 otherwise. $p \in P$, $d \in D_m$
U_{ijd}	Preferred additional number of physicians that can staff clinic j during time period i on day d, $j \in \{1,2\}$, $i \in \{1,2\}$, $d \in D_w$
V_{ijd}	Maximum additional number of physicians that can staff clinic j during time period i on day d, $j \in \{1,2\}$, $i \in \{1,2\}$, $d \in D_w$
λ_k	The weight placed on objective k

Table 29: Emory OB/GYN Physician Scheduling Model - Primary Decision Variables and descriptions

Decision Variable	Description
X_{ptd}	1 if physician p is given assignment t on day d; 0 otherwise. $p \in P_t$, $d \in C_{pt}$, $t \in \{1, 2, ..., 11\}$
Y_{pd}	1 if physician p is assigned to Labor and Delivery night call on day d; 0 otherwise. $p \in P_{12}$, $d \in C_{p(12)}$

If a physician is assigned to ER/Clinic/Surgery, then she must be assigned to one subassignment within that category (see Table 26) (constraints (5)).

$$(5) \quad \sum_{t \in T: d \in C_{pt}} X_{ptd} - X_{p2d} \quad = \quad 0 \quad \forall p \in P, \, d \in C_{p2}$$

L&D day and night call require exactly one physician per day (constraints (6) and (7)). ER requires exactly one physician per weekday (constraints (8)).

$$(6) \quad \sum_{p \in P_1: d \in C_{p1}} X_{p1d} \quad = \quad 1 \quad \forall d \in D$$

$$(7) \quad \sum_{p \in P_{12}: d \in C_{p(12)}} Y_{pd} \quad = \quad 1 \quad \forall d \in D$$

$$(8) \quad \sum_{p \in P_4: d \in C_{p4}} X_{p4d} \quad = \quad 1 \quad \forall d \in D_w$$

L&D day call and ER are assigned on a weekly basis. In other words, one physician is assigned to each weekday L&D day call shift in a given week, and another physician

Table 30: Emory OB/GYN Physician Scheduling Model - Additional Variables and descriptions

Variable	Description
α_{dij}	Number of morning (i=1) and afternoon (i=2) assignments to clinic j on day d above the preferred number, $j \in \{1, 2\}$, $d \in D_w$
$\bar{\beta}$	Largest number of weeks without a Surgery assignment for all physicians
$\bar{M}_j$	Largest number of Labor and Delivery night call shifts worked by any physician in availability period j, $j \in J$
$\bar{W}_j$	Largest number of weekend Labor and Delivery day call shifts worked by any physician in availability period j, $j \in J$
Q	Largest number of holidays worked by any physician
M_{ij}	Largest number of Labor and Delivery night call shifts worked by any physician in intersection of month i and availability period j if nonempty, $i \in M$, $j \in J$
W_{ij}	Largest number of weekend Labor and Delivery day call shifts worked by any physician in intersection of month i and availability period j if nonempty, $i \in M$, $j \in J$
B_d	Largest number of assignments to either clinic on day d, $d \in D_w$
Θ_j	Largest number of ER assignments worked by any physician in availability period j, $j \in J$
Δ_j	Largest number of Labor and Delivery day call shifts worked by any physician in availability period j, $j \in J$
Ψ_d	1 if some physician works both L&D night call and Surgery on day d; 0 otherwise. $d \in D_w$

is assigned to each ER shift (constraints (9)).

$$(9) \quad \sum_{r \in \{0,1,...,4\}:d+r \in C_{pt}} X_{pt(d+r)} \geq 5 \cdot X_{pt(d+\hat{r})} - |H_i| \quad \forall p \in P_t, \, d \in D_m$$

$$\hat{r} \in \{0, 1, ..., 4\} :$$

$$d + \hat{r} \notin D_h \,\, \& \,\, d + \hat{r} \in C_{pt}$$

$$t \in \{1, 4\}, \, i = (d + 6)\,/7$$

Any physician assigned to L&D day call or ER during the week should not be assigned to L&D night call any of those days (excluding holidays) (constraints (10)).

$$(10) \quad Y_{pd} + X_{ptd} \leq 1 \quad \forall p \in P_t, \, d \in C_{pt} \cap D_w \cap C_{p(12)} : d \notin D_h, \, t \in \{1, 4\}$$

The same physician should be assigned to L&D day call on both Saturday and Sunday (excluding holidays) (constraints (11)). This should not be the same physician assigned to L&D day call during the week (constraints (12)).

$$(11) \quad X_{p1(d+5)} - X_{p1(d+6)} = 0 \quad \forall p \in P_1, \, d \in D_m : d + r \in C_{p1},$$

$$d + r \notin D_h \text{ for } r \in \{5, 6\}$$

$$(12) \quad \sum_{r \in \{0,1,...,6\}:d+r \in C_{p1}, d+r \notin D_h} X_{p1(d+r)} + \leq 5 \quad \forall p \in P_1, \, d \in D_m$$

$$\left| H_i \cap \left(\bigcup_{\bar{r} \in \{0,1,...,4\}} \{d + \bar{r}\} \right) \right|$$

If a physician is assigned to L&D day call on Saturday and Sunday, then they must be off duty on Monday (constraints (13)).

$$(13) \quad X_{p1(d+5)} + X_{p1(d+6)} - X_{p3(d+7)} \leq 1 \quad \forall p \in P_1, \, d \in \{1, 8, ..., |D| - 13\} :$$

$$d + r \in C_{p1} \text{ for } r \in \{5, 6\}$$

The physician assigned to weekend L&D night call should work night call on Friday and Sunday (constraints (14)). Saturday L&D night call can be assigned to this physician as well if requested (constraints (15) and (16)).

$$(14) \qquad Y_{p(d+4)} - Y_{p(d+6)} \;=\; 0 \quad \forall p \in P_{12},\, d \in D_m : d + r \in C_{p(12)},\, d+r \notin D_h \text{ for } r \in \{4,6\}$$

$$(15) \qquad \sum_{r \in \{4,5,6\}: d+r \in C_{p(12)},\, d+r \notin D_h} Y_{p(d+r)} - R_{pd} \;\leq\; 2 \quad \forall p \in P_{12},\, d \in D_m$$

$$(16) \qquad \sum_{r \in \{4,5,6\}: d+r \in C_{p(12)},\, d+r \notin D_h} Y_{p(d+r)} - 3 \cdot R_{pd} \;\geq\; 0 \quad \forall p \in P_{12},\, d \in D_m$$

Physicians should not be assigned to L&D day call and/or ER in consecutive weeks (constraints (17)). We set the right-hand side of this inequality to 7 because there are at most 2 holidays in any week in the year. Therefore, this allows someone to be assigned to L&D (or ER) in one week, and then given 2 holiday L&D assignments the previous or following week.

$$(17) \qquad \sum_{t \in \{1,4\}} \sum_{r \in \{0,1,\dots,4\} \cup \{7,8,\dots,11\}: d+r \in C_{pt}} X_{pt(d+r)} \;\leq\; 7 \quad \forall p \in P$$
$$d \in \{1, 8, \dots, |D| - 13\}$$

There is a maximum number of physicians that can staff each clinic (constraints (18)). Each week, some physicians have a fixed schedule of clinic assignments, and these fixed assignments impact the number of additional physicians that may be scheduled (see Tables 31 and 32). Physicians listed in these tables with numbers make up the pool of physicians we are tasked with scheduling. Some of these physicians already have fixed clinic schedules, and are not given any further clinic assignments. The remaining physicians, represented by letters, help provide coverage in the clinics during predetermined time periods.

$$(18) \qquad \sum_{p \in P_t} \sum_{t \in L_{ij}: d+r \in C_{pt}} X_{pt(d+r)} \;\leq\; V_{ij(d+r)} \quad \forall d \in D_m,\, r \in \{0,1,\dots,4\},$$
$$i \in \{1,2\},\, j \in \{1,2\}$$

Each physician may be assigned to Surgery at most once in a week (constraints (19)). At most two physicians may be assigned to Surgery each day (constraints (20)).

Table 31: Emory OB/GYN Weekly Clinic 1 Fixed Assignments

Clinic 1	Monday		Tuesday		Wednesday		Thursday		Friday	
	am	pm	am	pm	am	pm	am	pm	am	pm
Physician 1	Fluctuates with schedule									
Physician 2	Fluctuates with schedule									
Physician 3	Fluctuates with schedule									
Physician 4	Fluctuates with schedule									
Physician 5	Fluctuates with schedule									
Additional physicians covering clinic										
Physician A	X	X	X	X	X	X	X	X	X	X
Physician B			X	X					X	X
Physician C			X	X						

$$\text{(19)} \quad \sum_{r \in \{0,1,\ldots,4\}:d+r \in C_{p(11)}} X_{p(11)(d+r)} \ \leq \ 1 \quad \forall p \in P_{11},\, d \in D_m$$

$$\text{(20)} \quad \sum_{p \in P_{11}:d \in C_{p(11)}} X_{ptd} \ \leq \ 2 \quad \forall d \in D_w$$

Physicians should not be assigned to L&D night call and clinic duties in the same day (constraints (21)). Physicians should be assigned to L&D night call at most twice during the week (excluding weekends) (constraints (22)).

$$\text{(21)} \quad Y_{pd} + \sum_{t \in \bigcup_{i,j \in \{1,2\}} L_{ij}:d \in C_{pt} \cap C_{p(12)}} X_{ptd} \ \leq \ 1 \quad \forall p \in P_t \cap P_{12},\, d \in D_w$$

$$\text{(22)} \quad \sum_{r \in \{0,1,\ldots,4\}:d+r \in C_{p(12)}} Y_{p(d+r)} \ \leq \ 2 \quad \forall p \in P_{12},\, d \in D_m$$

On holidays, the physician assigned to L&D day call also works L&D night call (constraints (23)).

$$\text{(23)} \quad X_{p1d} - Y_{pd} \ = \ 0 \quad \forall p \in P_1 \cap P_{12},\, d \in D_h : d \in C_{p1} \cap C_{p(12)}$$

There are additional constraints when scheduling physicians for Emory OB/GYN, but these constraints are taken into account in the definition of the decision variables (i.e., if certain assignments are not possible, then the variables corresponding to those assignments are not created). Such constraints include the following:

- No assignments to ER/Clinic/Surgery are given on weekends or holidays.

Table 32: Emory OB/GYN Weekly Clinic 2 Fixed Assignments

Clinic 2	Monday		Tuesday		Wednesday		Thursday		Friday	
	am	pm	am	pm	am	pm	am	pm	am	pm
Physician 6	X	X				X				X
Physician 5	Fluctuates with schedule									
Physician 7	Fluctuates with schedule									
Physician 8	Fluctuates with schedule									
Physician 10	X	X	X	X						
Physician 11					X	X	X	X		
Additional physicians covering clinic										
Physician D	X	X	X	X	X	X	X	X	X	X
Physician E			X							
Physician F							X			
Physician G							X			
Physician H			X	X					X	X
Physician I	Schedule varies by week									
Physician J	Schedule varies by week									
Physician K					X	X	X	X		
Physician L		X					X	X		
Physician M									X	X
Physician N										X
Physician O	X									

- No one is assigned to clinic duties on Wednesday mornings due to necessary meetings.

- Physicians only work in designated clinic locations.

- Physicians should be assigned to Off Duty during predetermined vacation days.

- Physician 6 is not assigned to L&D day or night call on weekdays (except holidays). Rather, her schedule will be entered manually.

- Physicians 6, 10, and 11 should not be assigned to Surgery or ER.

- Physician 1 is not assigned to clinic duties on Thursday mornings.

- Physician 8 is not assigned to L&D day call during the 2nd week of each month.
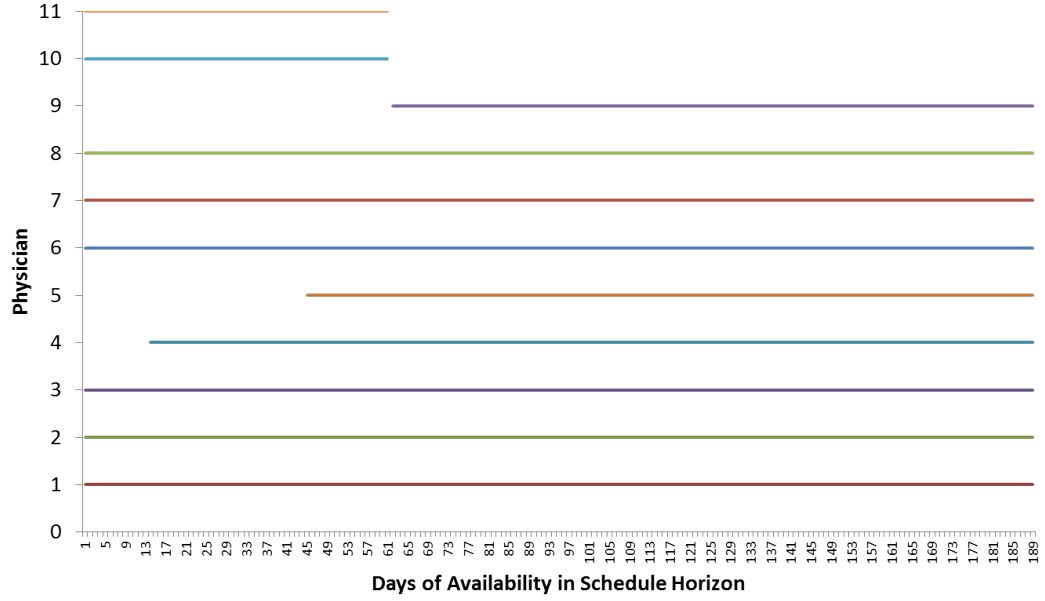
Figure 18: Emory OB/GYN Physician Availability Across Schedule Horizon

- Physicians 6, 10, and 11 have fixed clinic schedules, and therefore, we assume that no additional assignments to clinic duties are allowed for these physicians.

- Physicians should be assigned to Off Duty during periods of unavailability (see physician availability across schedule horizon in Figure 18).

4.3.2 Objective

There are multiple objectives. Fairness is important with regards to the following:

1. monthly L&D night call assignments (O.1)

2. total L&D night call assignments over schedule horizon (O.2)

3. total L&D day call assignments over schedule horizon (O.3)

4. monthly weekend L&D day call assignments (O.4)

5. total weekend L&D day call assignments over schedule horizon (O.5)

6. total holiday assignments (O.6)

106

7. total ER assignments (O.7)

To ensure that fairness for each of these items is considered with respect to the time periods that individual physicians are available, soft constraints (24) through (30) consider only physicians that are available, and during the appropriate windows of availability.

$$(24) \quad \sum_{d \in C_{p(12)} \cap D_i \cap F_j} Y_{pd} \leq M_{ij} \quad \forall p \in P_{12}, \, i \in \{1, 2, ..., 6\}, \, j \in J$$

$$(25) \quad \sum_{d \in C_{p(12)} \cap F_j} Y_{pd} \leq \bar{M}_j \quad \forall p \in P_{12}, \, j \in J$$

$$(26) \quad \sum_{d \in C_{p1} \cap F_j} X_{p1d} \leq \Delta_j \quad \forall p \in P_1, \, j \in J$$

$$(27) \quad \sum_{d \in C_{p1} \cap D_i \cap D_e \cap F_j} X_{pld} \leq W_{ij} \quad \forall p \in P_1, \, i \in \{1, 2, ..., 6\}, \, j \in J$$

$$(28) \quad \sum_{d \in C_{p1} \cap D_e \cap F_j} X_{p1d} \leq \bar{W}_j \quad \forall p \in P_1, \, j \in J$$

$$(29) \quad \sum_{d \in C_{p1} \cap D_h} X_{p1d} \leq Q \quad \forall p \in P_1$$

$$(30) \quad \sum_{d \in C_{p4} \cap D_w \cap F_j} X_{p4d} \leq \Theta_j \quad \forall p \in P_4, \, j \in J$$

Additional objectives include the following:

8. Maximize the total number of assignments to clinics (O.8).

9. Each physician should be assigned to Surgery once per week if possible. Therefore, we seek to minimize the maximum number of weeks for which any physician is not assigned a Surgery day (O.9). The maximum number of weeks for which any physician is not assigned a Surgery day is determined by constraints (31).

$$(31) \quad \sum_{d \in D_m} \left(1 - \sum_{r \in \{0,1,...,4\}: d+r \in C_{p(11)}} X_{p(11)(d+r)}\right) \leq \bar{\beta} \quad \forall p \in P_{11}$$

10. If a physician is assigned to night call during the week (excluding weekends and holidays), he or she may be assigned to Surgery one of those days. However, while this is acceptable, it is not ideal and thus the total number of assignments to both Surgery and night call in the same day should be minimized (O.10). Soft constraints (32) determine if a physician is assigned to both Surgery and night call in the same day, for each day in the schedule horizon.

$$(32) \quad Y_{pd} + X_{p(11)d} - 1 \leq \Psi_d \quad \forall d \in D_w \cap C_{p(11)} \cap C_{p(12)}, \, p \in P_{11} \cap P_{12}$$

11. We seek to minimize assignments to clinics above the preferred number (O.11). Soft constraints (33) determine the number of assignments above the preferred number in a given day.

$$(33) \quad \sum_{p \in P_t} \sum_{t \in L_{ij}:d+r \in C_{pt}} X_{pt(d+r)} \leq U_{ij(d+r)} + \alpha_{(d+r)ij} \quad \forall d \in D_m$$
$$r \in \{0, 1, ..., 4\}$$
$$i \in \{1, 2\}, \, j \in \{1, 2\}$$

12. Among those physicians for which the clinic schedule is flexible, the number of physicians assigned to each clinic location should be balanced as much as possible (O.12). Constraints (34) and (35) determine the minimum number of assignments at each clinic location.

$$(34) \quad \sum_{p \in P_t} \sum_{t \in \{5,6,7\}:d \in C_{pt}, d \notin D_h} X_{ptd} \geq B_d \quad \forall d \in D_w$$
$$(35) \quad \sum_{p \in P_t} \sum_{t \in \{8,9,10\}:d \in C_{pt}, d \notin D_h} X_{ptd} \geq B_d \quad \forall d \in D_w$$

So the objective with weights λ_k is the following:

Objective Function 1:

Minimize

$$(\text{O.1}) \quad \lambda_1 \sum_{j \in J} \sum_{i \in \{1,2,\dots,6\}} M_{ij}$$

$$(\text{O.2}) \quad + \lambda_2 \sum_{j \in J} \bar{M}_j$$

$$(\text{O.3}) \quad + \lambda_3 \sum_{j \in J} \Delta_j$$

$$(\text{O.4}) \quad + \lambda_4 \sum_{j \in J} \sum_{i \in \{1,2,\dots,6\}} W_{ij}$$

$$(\text{O.5}) \quad + \lambda_5 \sum_{j \in J} \bar{W}_j$$

$$(\text{O.6}) \quad + \lambda_6 Q$$

$$(\text{O.7}) \quad + \lambda_7 \sum_{j \in J} \Theta_j$$

$$(\text{O.8}) \quad - \lambda_8 \sum_{d \in C_{pt} \cap D_w} \sum_{p \in P_t} \sum_{t \in \{5,6,\dots,10\}} X_{ptd}$$

$$(\text{O.9}) \quad + \lambda_9 \bar{\beta}$$

$$(\text{O.10}) \quad + \lambda_{10} \sum_{d \in D_w} \Psi_d$$

$$(\text{O.11}) \quad + \lambda_{11} \sum_{d \in D_w} \sum_{i \in \{1,2\}} \sum_{j \in \{1,2\}} \alpha_{dij}$$

$$(\text{O.12}) \quad - \lambda_{12} \sum_{d \in D_w} B_d$$

4.4 Results and Discussion

We attempted to use PS-IP to construct a physician schedule for July - December, 2012, but no feasible solution was found due to violated constraints. However, one simple modification to physician availability can allow for a feasible schedule to be produced; i.e., we remove the restriction that physicians 6, 10 and 11 not be assigned to ER or Surgery. This is clearly not a feasible option, but we see that a solution is possible with greater physician availability. With this adjustment, the model finds an optimal solution (optimality gap = 0.00%) in approximately 2 minutes. Thus, PS-IP could efficiently solve the problem faced at Emory OB/GYN if more physicians were available for assignments to ER and Surgery. This is not the case, however, so we considered an alternative means of finding a feasible schedule.

One such approach is to allow consecutive assignments to L&D night call. This

is not an ideal scenario, but it is acceptable if necessary to satisfy feasibility requirements and physician preferences such as vacation requests. Constraints (4) state that no physician can be assigned to L&D night call if they are off duty, and a physician is forced to be off duty in a given day if they were assigned to L&D night call the previous day (constraints (3)) or are on vacation or unavailable. If a physician is assigned to L&D night call, it is important that they not be assigned to daytime duties the following day. Therefore, rather than removing constraints (3) to allow consecutive night assignments, we make constraints (4) soft constraints with the addition of penalty variables Γ_{pd} (see new constraints ($\bar{4}$)).

$$(\bar{4}) \quad X_{p3d} + Y_{pd} \quad \leq \quad 1 + \Gamma_{pd} \quad \forall d \in C_{p3} \cap C_{p(12)},\ p \in P_3 \cap P_{12}$$

This adjustment allows for physicians to be scheduled to L&D night call even if they are considered to be off duty due to a night call assignment the previous day, but we include a penalty and seek to minimize these penalties in the objective function (see Objective Function 2 below). So the new additional objective is:

13. We seek to minimize the number of times physicians are given consecutive assignments to L&D night call.

Objective Function 2:

Minimize

Objective Function 1

$$(O.13) \quad +\lambda_{13} \sum_{p \in P_3 \cap P_{12}} \sum_{d \in C_{p3} \cap C_{p(12)}} \Gamma_{pd}$$

A complete list of the objectives is given in Table 33.

With this objective and the modified constraints ($\bar{4}$), a good feasible schedule (with optimality gap $< 1\%$) can be found in approximately 3 minutes (optimality gap $< 0.01\%$ achieved in less than 1 hour). This schedule includes 9 instances of physicians given consecutive assignments to L&D night call.

Table 33: Emory OB/GYN - Physician Scheduling Problem Objectives

Objective	Description
1	Balance monthly L&D night call assignments
2	Balance total L&D night call assignments over schedule horizon
3	Balance total L&D day call assignments over schedule horizon
4	Balance monthly weekend L&D day call assignments
5	Balance total weekend L&D day call assignments over schedule horizon
6	Balance total holiday assignments
7	Balance total ER assignments
8	Maximize total clinic assignments
9	Minimize the maximum number of weeks for which any physician is not assigned a Surgery day
10	Minimize the total number of assignments to both Surgery and L&D night call in the same day
11	Minimize assignments to clinics above the preferred number
12	Balance assignments to clinics across both locations
13	Minimize the number of times physicians are given consecutive assignments to L&D night call

The results reported above are found using an objective weight vector λ which places high value on objectives 3, 5, 7, and 13 (objectives which balance assignments to weekdays and weekends on L&D day call, ER assignments, and minimize consecutive assignments to L&D night call, respectively), placing the highest weight on objective 3. The remaining objectives are given a weight of 1. With a unit weight vector, we see a great increase in solution time, and the large number of objectives results in largely unbalanced L&D day call and ER assignments in a solution produced after running PS-IP for 2 hours (optimality gap 8.62%). In this solution, physicians available for short time periods in the schedule horizon are given a large number of L&D day call and ER assignments during those periods as compared to other physicians available at the same time, which is an undesirable result. Thus, we prioritize the objectives as described.

PS-IP is highly efficient for solving this instance of the physician scheduling problem faced by Emory OB/GYN. In addition, we tested various objective functions (reported in Table 34) to understand their impact on the solutions. Note that an optimality gap $< 1\%$ is achieved for each objective function tested in under 5 minutes.

From Table 34, we see from Test 2 that removing objective 11 has very little impact on the values of the objective function components, and the solution time is about the same. This is not surprising, as this objective seeks to minimize the daily number of assignments to clinics above the preferred number, and PS-IP gives no great advantage to assigning more than this preferred number. Test 3, which excludes objective 12 (i.e., balanced clinic assignments across locations) as well as objective 11, has a similar result as well, but assignments to clinics are less balanced, as expected (i.e., there is a larger gap between the number of assignments to each clinic). If this objective loses importance, we see that the objective function used in Test 3 could provide an acceptable schedule. While each objective function tested achieves an optimality gap $< 1\%$ in 5 minutes or less, Tests 1 and 2 require up to 1 and 2 hours, respectively, to achieve an optimality gap $< 0.01\%$. Test 3 reaches an optimality gap $< 0.01\%$ in less than 12 minutes. Thus, the objective function used in Test 3 may be an acceptable alternative if an optimality gap $< 0.01\%$ is strongly desired, and balanced assignments to clinics become less important.

Test 4 produces a solution which also achieves very similar values of the objective function components compared to that found with the full objective. We expect less balanced assignments to clinics given the objectives omitted by this test, but we would also expect to see more assignments to Surgery and L&D night call in the same day. This is not the case, however. Thus, we conclude that some combination of model constraints and objectives prevent an improvement in the number of assignments to Surgery and L&D night call in the tests which include objective 10.

Table 34: Emory OB/GYN - Objective Function Analysis

Test	Objectives													Comparison of Solution to Model with Full Objective (i.e. Test 1)
	1	2	3	4	5	6	7	8	9	10	11	12	13	
1	X	X	X	X	X	X	X	X	X	X	X	X	X	-
2	X	X	X	X	X	X	X	X	X	X		X	X	Similar values of objective function components, slightly more clinic assignments
3	X	X	X	X	X	X	X	X	X	X			X	Similar values of objective function components, less balanced clinic assignments across locations
4	X	X	X	X	X	X	X	X	X				X	Similar values of objective function components, less balanced clinic assignments across locations, identical number of assignments to Surgery and L&D night call in same day, faster solution time
5	X	X	X	X	X	X	X	X					X	Similar values of objective function components, less balanced clinic assignments across locations, identical number of assignments to Surgery, as well as Surgery and L&D night call assignments in same day, faster solution time
6	X	X	X	X	X	X	X	X						Large (411%) increase in number of consecutive assignments to L&D night call, 0 assignments to Surgery over schedule horizon, less balanced clinic assignments across locations, slower solution time

The result of Test 5 is very similar to that of Test 4. Tests 4 and 5 reach an optimality gap $< 0.01\%$ in under 5 minutes. The solution found by Test 6 varies greatly from that of Test 1 (i.e., PS-IP with full objective). Not only is there a slower solution time (to an optimality gap $< 0.01\%$), but there is also a large increase in consecutive assignments to L&D night call, and no assignments to Surgery are given over the schedule horizon. Each of these issues prevent this objective function from being useful in generating an acceptable schedule.

From this analysis, we conclude that the objective function used by Test 2 would be an acceptable alternative to use of the full objective function, but with no improvement in solution time, there is really no benefit in using this objective function instead. If balanced assignments to clinics across locations is not important, the objective function used by Test 3 would be useful, particularly as it produces a similar result with a faster solution time. The objective functions used in Tests 4 and 5 may also provide a reasonable substitute, with even faster solution times, but further tests on varying instances are needed for comparison to determine if the tests produce acceptable schedules under different circumstances. The objective function used by Test 6 is much too limited to provide a solution acceptable for implementation.

In light of this comparison, we see that, for this instance where balanced clinic assignments are important, the full objective function (i.e., incorporating the 13 objectives) is most appropriate. Given this, we investigated whether or not we could improve on the 3 minute solution time required to reach an optimality gap $< 1\%$. To accomplish this, we limited our objective function to our high priority objectives (i.e., 3, 5, 7, and 13). An optimal solution (optimality gap $< 0.01\%$) to PS-IP with this objective function is found in seconds. Fixing the L&D day call assignments from this solution, and then using the full objective and solving the model again, an optimal solution (optimality gap $< 0.01\%$) can be found in less than 1.25 minutes. Table 35 reports the results from this analysis. We refer to this new two-part test as "Test 7".

Table 35: Emory OB/GYN - Comparison of Model with Full Objective to Two-Part Model

Test	Solution Time	Comparison to Model with Full Objective
1	3 Minutes*	-
7	< 1.5 Minutes+	Almost identical values of objective function components, slightly less balanced holiday assignments, monthly L&D night call assignments, and monthly weekend L&D day call assignments, total assignments to Clinics 1 and 2 vary by 7 (less) and 1 (more) over schedule horizon, respectively

* Optimality gap < 1%; optimality gap < 0.01% achieved in 1 hour
+ Optimality gap < 0.01%

This two-part test (i.e., Test 7) constructs an optimal schedule (gap < 0.01%) with a much smaller solution time than the original model with the full objective. Whether or not the solution provided by this test is acceptable for implementation, given the slightly less balanced holiday, night call, and weekend assignments, can be determined by decision-makers at Emory OB/GYN. We do see that the solution time can be greatly reduced by limiting the objective in this way.

4.5   Implementation and Future Work

During model development, we have been in direct contact with decision-makers at Emory OB/GYN, with the goal of providing them with a tool for constructing physician schedules on a regular basis. Communications included sample schedules generated by PS-IP and feedback regarding issues which usually revolved around miscommunicated scheduling constraints.

Initial discussions regarding the physician scheduling problem faced at Emory OB/GYN resulted in the inclusion of objectives 1-2, 4-6, 8-9, and 11-12 in the objective function. However, schedules produced by our model with these objectives made it clear that the additional objectives 3, 7, and 10 were needed due to a lack of

balance among L&D day call and ER assignments over the schedule horizon, as well as an overwhelming number of assignments to both Surgery and L&D night call in the same day.

We evaluated and used as a benchmark a manually constructed schedule for July-December, 2012, which included assignments of physicians to L&D day call, L&D night call, and ER. We were also given vacation requests for that time period. We used PS-IP to generate a schedule for the same time period, considering vacation requests. The manual schedule had 27 back-to-back night call assignments, compared to only 9 for the IP-generated schedule. The maximum number of night call assignments per physician was higher in the manual schedule than for the IP-generated schedule, for each availability period (the same is true for each month in the schedule horizon). The maximum number of weekdays on ER was higher for the manual schedule for 3 of the 4 availability periods. The maximum numbers of weekend L&D day call assignments given to any physician were identical for 4 of the 6 months in the schedule horizon. The maximum number of holiday L&D day call assignments given to any physician were also identical.

One limitation in this comparison is that no assignments of physician 6 to L&D day or night call were made on weekdays (excluding holidays), as requested by decision makers at Emory OB/GYN. The manual schedule includes L&D assignments of this physician. However, despite this limitation, we conclude that automated schedule generation outperforms manual schedule construction with respect to solution time and the scheduling objectives.

Feedback has been very positive regarding the ability of our model to generate a schedule which meets all scheduling requirements, as well as the usefulness of an automated tool which schedulers can use at the operational level. Future work includes the development of a stand-alone tool for decision-makers to generate schedules at the operational level using a freely available solver such as GLPK [30] or OpenSolver

[32] and incorporating the integer programming model presented in this chapter. We will develop a user-friendly interface which allows physician specific restrictions to be entered as parameters into the model to create greater flexibility in schedule construction for years to come.

4.6   Conclusions

The physician scheduling problem faced by Emory OB/GYN can be efficiently solved with the use of an integer programming model. While solution times are considerably low, they can be further reduced depending on the importance of some of the scheduling objectives. The solution approach presented here allows for creating feasible schedules while considering multiple objectives, a task which can be cumbersome to accomplish manually. The model considers the varying availability of clinic space each day, and allows for scheduling of physicians with varying restrictions.

While the model formulation we present is specific to the problem faced by Emory OB/GYN, a similar methodology could be applied in other settings with various types of availability and demand constraints.

CHAPTER V

CONCLUSIONS AND FUTURE RESEARCH DIRECTIONS

This thesis presents three applications of optimization methods to problems in physician scheduling, and provides methods for improving schedule construction over manual methods. More efficient scheduling of physicians, both with regards to rotations and day-to-day shifts, could positively impact patient care [27]. With methods for measuring the continuity of a schedule, handoff efficiency can be improved through mathematical modeling. Optimization models can further inform decisions with regards to shift scheduling, particularly as increased duty hour restrictions combined with physician preferences may negatively impact resident education.

As discussed in Chapter 2, there are many factors which impact the efficiency of a handoff, including but not limited to bed occupancy levels, new admissions through the day, disease acuity of current patients, and fatigue of physicians both starting and ending a duty period. ACGME duty hour restrictions were created in part to reduce physician fatigue. However, fatigue cannot be completely eliminated. One future research direction is to investigate the impact of fatique on handoff efficiency. An expansion of the HCS which includes a factor for measuring fatigue of physicians at handoff may provide a tool for understanding this impact, in combination with careful analysis of the communication that occurs during those handoffs.

An additional research direction includes attempting to better understand the impact that too much familiarity among oncoming physicians may have on the handoff process. If increased familiarity reduces a physician's level of attention to communication shared at handoff, this is an undesired result. Thus, further study is needed

to identify the true impact of an HCS-maximized schedule on handoff efficiency and quality.

While the heuristic we developed, which incorporates the Children's PICU Physician Scheduling MIP (CPPS-MIP), efficiently solves the physician scheduling problem faced at Children's, results are not as attractive when removing the institutional preference for service blocks. We believe the alternative model we developed has the potential to improve on the efficiency of the heuristic and CPPS-MIP on problems with a more general structure (i.e., without service blocks). Constraints which prevent assignments to overlapping shift sequences could tighten the alternative model formulation. Further, column generation techniques could be employed to reduce the size of the problem and improve efficiency.

In Chapter 3, we discussed two physician scheduling problems faced by the Department of Surgery at Emory University School of Medicine (EUSOM): (1) resident rotation assignment, and (2) day-to-day shift scheduling. We developed an efficient model for rotation assignment, and using a simple integer program, provided insights into the impact of current practices for day-to-day shift scheduling. We developed a decision support tool for construction of rotation assignment schedules. Future work includes the combining of the excel-based decision support tool we developed for entering problem specific parameters with a freely available solver such as GLPK or OpenSolver. With enough flexibility to handle various types of demand and supply constraints, such a tool could be easily adopted by similar programs to that of EUSOM.

In Chapter 4, we show that a physician scheduling problem which includes a non-homogeneous physician pool, multiple objectives, multiple tasks, and space considerations, can be efficiently solved using integer programming. We solved a specific instance of such a problem faced by the Department of Gynecology and Obstetrics at Emory University Hospital. Future work includes the combining of the integer

119

programming model we developed with a user-friendly decision support tool, using a freely available solver so that schedules can be easily constructed at the operational level.

APPENDIX A

CHILDREN'S PICU PHYSICIAN SCHEDULING MIP (CPPS-MIP) - ADDITIONAL CONSTRAINT DEFINITIONS

Additional constraints were added to CPPS-MIP during 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, are presented below. Note in the set definitions that we refer to a physician M. This physician routinely requested to be assigned to specific night call shifts which 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.

Table 36: Children's PICU Physician Scheduling MIP (CPPS-MIP) - Additional Sets, Parameters, and Decision Variables

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, $p \in \{1, 2, ..., 8\}$
Decision Variable	**Description**
Φ_{pij}	Binary variable, $=$ penalty p incurred for physician i in week beginning with day j. $i \in A, j \in \bar{J}, p \in \{1, 2, ..., 8\}$

We modified constraints (2a) to be soft constraints, to allow the possibility that attendings not be assigned to 7 consecutive days if necessary to satisfy requests and/or fellowship requirements.

$$(\bar{2}a) \quad Y_{ijn} - \sum_{k \in K_1} X_{isk} \quad \leq \quad \Phi_{1is} \quad \forall i \in A,\ s \in \{j, j+1, ..., j+n-1\}$$
$$j \in \{1, 2, ..., N - n + 1\},\ n \in B_p,\ P \in \wp$$

Attendings should not be assigned to night call two days in a row. This is a modification of original constraints (2e).

$$(\bar{2}e) \quad \sum_{k \in K_2} \left(X_{ijk} + X_{i(j+1)k} \right) \quad \leq \quad 1 \quad \forall i \in A,\ j \in \{1, 2, ..., N-1\}$$

Each attending should be assigned to no more than 3 night call shifts per week.

$$(4a) \quad \sum_{k \in K_2} \sum_{r \in \{0,1,...,6\}} X_{i(j+r)k} \quad \leq \quad 3 \quad \forall i \in A,\ j \in \bar{J}$$

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

$$(4b) \quad \sum_{k \in K_2} \left(X_{i(j+1)k} + X_{i(j+3)k} + X_{i(j+5)k} \right) \quad \leq \quad 2 \quad \forall i \in A,\ j \in \bar{J}$$
$$(4c) \quad \sum_{k \in K_2} \left(X_{i(j+1)k} + X_{i(j+3)k} + X_{i(j+6)k} \right) \quad \leq \quad 2 \quad \forall i \in A,\ j \in \bar{J}$$

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 they work the service shift on Monday.

$$(4d) \quad \sum_{k \in K_2} X_{ijk} - \sum_{k \in K_1} X_{ijk} \quad \leq \quad \Phi_{2ij} \quad \forall i \in A,\ j \in \bar{J}$$
$$(4e) \quad \sum_{k \in K_2} X_{i(j+3)k} - \sum_{k \in K_1} X_{ijk} \quad \leq \quad \Phi_{3ij} \quad \forall i \in A,\ j \in \bar{J}$$

If a physician is one of the "on-service" physicians for the week, they should be assigned to at least one weekend night call shift.

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

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.

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

(4h) $\qquad\qquad\qquad \sum_{k \in K_2} \left(X_{i(j+3)k} - X_{i(j+5)k} \right) \quad \leq \quad \Phi_{7ij} \qquad\qquad \forall i \in A,\, j \in \hat{J}$

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.

(4i) $\quad \sum_{k \in K_2} \left(X_{i(j+3)k} - X_{i(j+6)k} \right) \quad \leq \quad \Phi_{8ij} \quad \forall i \in \bar{A},\, j \in \tilde{J}$

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

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

(4k) $\qquad\qquad\qquad \sum_{j \in J} \sum_{k \in K_2} X_{ijk} \quad \geq \quad 0.9 \cdot \Psi_{2i} \quad \forall i \in A$

(4l) $\quad \sum_{j \in \bar{J}} \sum_{r \in \{5,6\}} \sum_{k \in K_1} X_{i(j+r)k} \quad \geq \quad 0.9 \cdot \Psi_{3i} \quad \forall i \in A$

The objective function then becomes:

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

Note that requests are given precedence over the service block structure, as well as the preferred call shift structure.

APPENDIX B

RESIDENT ROTATION ASSIGNMENT MODEL - SUPPLEMENTARY TABLE AND DATA

Table 37 reports the demand from each service staffed by surgical residents from Emory University School of Medicine. An asterisk next to a resident level means: (i) if contained in an "or" statement, then the asterisk means that this level is preferred over the other, (ii) if not contained in an "or" statement, then the asterisk means that the level is preferred, but not demanded, (iii) if next to a number, then the asterisk means that that number of a particular resident level is preferred, but not demanded.

The following services represent equivalent experiences in general surgery.

- GMH General Surgery A

- GMH General Surgery B

- EUH General Surgery A

- EUH General Surgery B

- EUH ACS

- EUH Surgery Oncology

- EMH General Surgery

- PGH General Surgery

- VAH General Surgery

Table 37: Services staffed by surgical residents from Emory University School of Medicine, with monthly demand requirements. GMH = Grady Memorial Hospital, EUH = Emory University Hospital, EMH = Emory University Hospital Midtown, VAH = Veterans Affairs Hospital, PGH = Piedmont General Hospital, HEH = Emory University Hospital at Egleston.

Rotation	Service Name	Demand
1	GMH General Surgery A	1 PGY-1, 1 PGY-3* or PGY-4, and 1 PGY-5
1	GMH General Surgery B	1 PGY-1, 1 PGY-3* or PGY-4, and 1 PGY-5
1	EUH General Surgery A	1 PGY-1, 1 PGY-2 or PGY-3*, and 1 PGY-4 or PGY-5*
1	EUH General Surgery B	1 PGY-1 or PGY-2*, and 1 PGY-4 or PGY-5*
1	EUH ACS	1 PGY-1 or PGY-2*, and 1 PGY-4
1	EUH Surgery Oncology	1 PGY-1, 1 PGY-3, 1 PGY-4, and 1 PGY-5
1	EMH General Surgery	1 PGY-1, 1 PGY-2, and 1 PGY-4
1	PGH General Surgery	1 PGY-1*, 1 PGY-2*, 1 PGY-3*, 1 PGY-4*, and 1 PGY-5*
1	VAH General Surgery	1 PGY-2 and 1 PGY-4* or PGY-5
2	EUH Selective	1 PGY-3* and 1 PGY-5*
3	GMH Trauma Day	1 PGY-1
4	GMH Trauma Night	1 PGY-3, and 1 PGY-4
5	EUH SICU	1 PGY-2* or PGY-3
6	GMH SICU	1 PGY-1 and 2 or 3* PGY-2's
7	EUH Vascular Surgery	1 PGY-1, 1 PGY-2, and 1 PGY-4* or PGY-5
8	VAH Vascular Surgery	1 PGY-2, and 1 PGY-3* or PGY-4
9	HEH Pediatric Surgery Day	3 PGY-1's and 1 PGY-4
10	GMH Burns	1 PGY-1, 1 PGY-2, and 1 PGY-3
11	GMH Plastic Surgery	1 PGY-1
12	GMH Urology	1 PGY-1
13	EUH Transplant	1 PGY-1* or PGY-2 and 1 PGY-3 or PGY-4*
14	EUH Cardiothoracic	1 PGY-2
15	EUH Night	1 PGY-1 and 1 PGY-3
16	HEH Pediatric Surgery Night	1 PGY-1
17	VAH Cardiothoracic	1 PGY-2
18	GMH Quad Night	1 PGY-1
19	EUH Colorectal	1 PGY-1, 1 PGY-2, and 1 PGY-5
20	GMH General Surgery C	1 PGY2 and 1 PGY4

REFERENCES

[1] ACGME, *Accreditation Council for Graduate Medical Education - Common Program Requirements.* 2011. Accessed December 1, 2011, http://www.acgme.org/acWebsite/home/Common_Program_Requirements_ 07012011.pdf.

[2] ACGME, *Resident duty hours in the learning and working environment; comparison of 2003 and 2011 standards.* 2011. Accessed July 5, 2012, http://www.acgme.org/acWebsite/dutyHours/dh-ComparisonTable2003v2011.pdf.

[3] ARORA, V., MANJARREZ, E., DRESSLER, D., BASAVIAH, P., HALASYAMANI, L., and KRIPALANI, S., "Hospitalist handoffs: A systematic review and task force recommendations," *J Hosp Med*, vol. 4, no. 7, pp. 433–440, 2009.

[4] BEAULIEU, H., FERLAND, J., GENDRON, B., and MICHELON, P., "A mathematical programming approach for scheduling physicians in the emergency room.," *Health Care Management Science*, vol. 3, no. 3, pp. 193–200, 2000.

[5] BEKKAN, C., *Quantitative methods of physician scheduling at hospitals.* Masters dissertation, University of Agder, Department of Industrial and Information Management, August 2010.

[6] BELIEN, J. and DEMEULEMEESTER, E., "On the trade-off between staff-decomposed and activity-decomposed column generation for a staff scheduling problem," *Ann Oper Res*, vol. 155, pp. 143–166, 2007.

[7] BRASEL, K., PIERRE, A., and WEIGELT, J., "Resident work hours - what they are really doing," *Arch Surg.*, vol. 139, pp. 490–494, 2004.

[8] BRUNNER, J., BARD, J., and KOLISCH, R., "Midterm scheduling of physicians with flexible shifts using branch and price.," *IIE Transactions*, vol. 43, pp. 84–109, 2010.

[9] BURKE, E., DE CAUSMAECKER, P., BERGHE, G., and VAN LANDEGHEM, H., "The state of the art of nurse rostering.," *Journal of Scheduling*, vol. 7, pp. 441–499, 2004.

[10] CHEANG, B., LI, H., LIM, A., and RODRIGUES, B., "Nurse rostering problems - a bibliographic survey.," *European Journal of Operational Research*, vol. 151, pp. 447–460, 2003.

[11] COHN, A., ROOT, S., KYMISSIS, C., ESSES, J., and WESTMORELAND, N., "Scheduling medical residents at boston university school of medicine.," *Interfaces*, vol. 39, no. 3, pp. 186–195, 2009.

[12] CONNORS, R., DOTY, J., BULL, D., MAY, H., FULLERTON, D., and ROBBINS, R., "Effect of work-hour restriction on operative experience in cardiothoracic surgical residency training.," *J Thorac Cardiovasc Surg.*, vol. 137, pp. 710–713, 2009.

[13] DAY, T., NAPOLI, J., and KUO, P., "Scheduling the resident 80-hour work week: An operations research algorithm.," *Current Surgery*, vol. 63, no. 2, pp. 136–141, 2006.

[14] DINKEL, J. J., MOTE, J., and VENKATARAMANAN, M. A., "An efficient decision support system for academic course scheduling," *Operations Research*, vol. 37, no. 6, pp. 853–864, 1989.

[15] EMORY UNIVERSITY SCHOOL OF MEDICINE, *Department of Gynecology and Obstetrics.* 2012. Accessed July 9, 2012, http://www.gynob.emory.edu/.

[16] EPSTEIN, K., JUAREZ, E., EPSTEIN, A., LOYA, K., and SINGER, A., "The impact of fragmentation of hospitalist care on length of stay," *J Hosp Med*, vol. 5, no. 6, pp. 335–338, 2010.

[17] ERNST, A., JIANG, H., KRISHNAMOORTHY, M., OWENS, B., and SIER, D., "An annotated bibliography of personnel scheduling and rostering.," *Annals of Operations Research*, vol. 127, pp. 21–144, 2004.

[18] FORTUNA, R., PALFREY, J., SHELOV, S., and SAMUELS, R., "Paediatric experiences with work-hour limitations," *J Eval Clin Pract*, vol. 15, pp. 116–120, 2009.

[19] FRANZ, L. and MILLER, J., "Scheduling medical residents to rotations: solving the large-scale multiperiod staff assignment problem," *Operations Research*, vol. 41, no. 2, pp. 269–279, 1993.

[20] FREIBURG, C., JAMES, T., ASHIKAGA, T., MOALEM, J., and CHERR, G., "Strategies to accommodate resident work-hour restrictions: Impact on surgical education.," *J Surg Educ.*, vol. 68, no. 5, pp. 387–392, 2011.

[21] GOPALAKRISHNAN, B. and JOHNSON, E., "Airline crew scheduling: state-of-the-art.," *Annals of Operations Research*, vol. 140, pp. 305–337, 2005.

[22] HODGSON, T. J., KILPATRICK, K. E., and LONGINI, I. M., "Scheduling multispecialty clinics using integer programming," *AIIE Transactions*, vol. 9, no. 1, pp. 69–74, 1977.

[23] ISKEN, M., WARD, T., and MCKEE, T., "Simulating outpatient obstetrical clinics," in *Simulation Conference Proceedings, 1999 Winter*, no. 2, pp. 1557–1563, 1999.

[24] JAVERI, M., *Rotation planning and scheduling for medical resident education.* Masters dissertation, Purdue University, Department of Industrial Engineering, August 2011.

[25] JINGPENG, L., BURKE, E. K., CURTOIS, T., PETROVIC, S., and QU, R., "The falling tide algorithm: A new multi-objective approach for complex workforce scheduling," *Omega*, vol. 40, no. 3, pp. 283–293, 2012.

[26] KOHL, N. and KARISCH, S., "Airline crew rostering: problem types, modeling, and optimization.," *Annals of Operations Research*, vol. 127, no. 1-4, pp. 223–257, 2004.

[27] KORT, K., PAVONE, L., JENSEN, E., HAQUE, E., NEWMAN, N., and KITTUR, D., "Resident perceptions of the impact of work-hour restrictions on health care delivery and surgical education: Time for transformational change.," *Surgery.*, vol. 136, pp. 861–871, 2004.

[28] LANDRIGAN, C. P., "The handoff: a critical point of vulnerability," *CRICO/RMF Forum*, vol. 25, pp. 8–9, 2007.

[29] MAENHOUT, B. and VANHOUCKE, M., "Branching strategies in a branch-and-price approach for a multiple objective nurse scheduling problem," *Journal of Scheduling*, vol. 13, no. 1, pp. 77–93, 2010.

[30] MAKHORIN, A., *GLPK (Gnu Linear Programming Kit).* 2012. Accessed July 9, 2012, http://www.gnu.org/software/glpk/.

[31] MANSDORF, B., "Allocation of resources for ambulatory care - a staffing model for outpatient clinics," *Public Health Reports*, vol. 90, no. 5, pp. 393–401, 1975.

[32] MASON, A. and DUNNING, I., *OpenSolver for Excel*. 2012. Accessed August 17, 2012, http://opensolver.org/.

[33] MCCOY, C. P., HALVORSEN, A. J., LOFTUS, C. G., MCDONALD, F. S., and OXENTENKO, A. S., "Effect of 16-hour duty periods on patient care and resident education," *May Clin Proc.*, vol. 86, no. 3, p. 192, 2011.

[34] MITCHELL, C., ASHLEY, S., ZINNER, M., and MOORE JR., F., "Predicting future staffing needs at teaching hospitals - use of an analytical program with multiple variables.," *Arch Surg.*, vol. 142, pp. 329–334, 2007.

[35] OVCHINNIKOV, A. and MILNER, J., "Spreadsheet model helps to assign medical residents at the university of vermont's college of medicine.," *Interfaces*, vol. 38, no. 4, pp. 311–323, 2008.

[36] PICKERING, B., HURLEY, K., and MARSH, B., "Identification of patient information corruption in the intensive care unit: Using a scoring tool to direct quality improvements in handover," *Crit Care Med*, vol. 37, pp. 2905–2912, 2009.

[37] RIESENBURG, L., LEITZSCH, J., MASSUCCI, J., JAEGER, J., ROSENFELD, J., PATOW, C., PADMORE, J., and KARPOVICH, K., "Residents' and attending physicians' handoffs: A systematic review of the literature," *Acad Med*, vol. 84, pp. 1775–1787, 2009.

[38] RODRIGUEZ, H., TURNER, J., SPEICHER, P., DASKIN, M., and DAROSA, D., "A model for evaluating resident education with a focus on continuity of care and educational quality," *Journal of Surgical Education*, vol. doi: 10.1016/j.jsurg.2010.09.004, 2010.

[39] Rousseau, L., Pesant, G., and Gendreau, M., "A general approach to the physician rostering problem.," *Annals of Operations Research*, vol. 115, no. 1, pp. 193–205, 2002.

[40] Saultz, J. and Albedaiwi, W., "Interpersonal continuity of care and patient satisfaction: A critical review," *Ann Fam Med*, vol. 2, no. 5, pp. 445–451, 2004.

[41] Sherali, H., Ramahi, M., and Saifee, Q., "Hospital resident scheduling problem.," *Production Planning and Control*, vol. 13, pp. 220–233, 2002.

[42] Smalley, H., Keskinocak, P., and Vats, A., "Development of a handoff continuity score to improve pediatric icu physician schedule design for physician and patient continuity," *Crit Care*, vol. 15, no. R246, 2011.

[43] Solet, D., Norvell, J., Rutan, G., and Frankel, R., "Lost in translation: challenges and opportunities in physician-to-physician communication during patient handoffs," *Acad Med*, vol. 80, pp. 1094–1099, 2005.

[44] Tabrizian, P., Rajhbeharrysingh, U., Khaitov, S., and Divino, C., "Persistent noncompliance with the work-hour regulation," *Arch Surg.*, vol. 146, no. 2, pp. 175–178, 2011.

[45] Topaloglu, S., "A multi-objective programming model for scheduling emergency medicine residents.," *Computers and Industrial Engineering*, vol. 51, pp. 375–388, 2006.

[46] Topaloglu, S., "A shift scheduling model for employees with different seniority levels and an application in healthcare.," *European Journal of Operational Research*, vol. 198, pp. 943–957, 2009.

[47] TOPALOGLU, S. and OZKARAHAN, I., "A constraint programming-based solution approach for medical resident scheduling problems.," *Computers and Operations Research*, vol. 38, pp. 246–255, 2011.

[48] TURNER, J., *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, June 2011.

[49] TURNER, J., RODRIGUEZ, H., DASKIN, M., MEHROTRA, S., SPEICHER, P., and DARosa, D., "Overcoming obstacles to resident-patient continuity of care," *Annals of Surgery*, vol. 0, no. 0, pp. 1–5, 2012.

[50] VAN EATON, E., TARPLEY, J., SOLORZANO, C., CHO, C., WEBER, S., and TERMUHLEN, P., "Resident education in 2011: three key challenges on the road ahead.," *Surgery.*, vol. 149, pp. 465–473, 2011.

[51] VPS, *MyVPS Online.* 2012. Accessed July 6, 2012, https://portal.myvps.org/Default.aspx.

[52] WILLIS, R., COVERDILL, J., MELLINGER, J., COLLINS, J., POTTS III, J., and DENT, D., "Views of surgery program directors on the current acgme and proposed iom duty-hour standards.," *J Surg Educ.*, vol. 66, no. 4, pp. 216–221, 2009.

VITA

Hannah Smalley was born in Cheraw, South Carolina. She received Bachelor of Science and Master of Science degrees in Mathematics from Winthrop University in 2004 and 2005, respectively. Her undergraduate studies included a semester at Deakin University in Geelong, Australia. Following graduation from Winthrop University in 2005, she began pursuing her Ph.D. in Optimization in the H. Milton Stewart School of Industrial and Systems Engineering at Georgia Institute of Technology. In 2008, she received a Master of Science degree in Operations Research from Georgia Institute of Technology. Her research experience and interests are in the fields of health care operations, public health (including immunization scheduling and disease spread modeling), and emergency response logistics.