Innovative Scheduling Solutions for Graduate Medical Education

TIFFANY N. S. BALLARD, MD
 TYLER R. GREENDA, MD
 AMY M. COHN, PhD
 MARK S. DASKIN, PhD
 F. JACOB SEAGULL, PhD
 RISHINDRA M. REDDY, MD, FACS

The emergency medicine chief resident sits at her desk, at first just staring at an empty grid. She then determinedly starts assigning shifts to the 15 residents rotating next month. Hours later, she isn't much further along than when she started, due to the different rules for each type of resident. Difficulty accounting for various scheduling conflicts, such as medicine continuity clinics on Thursday mornings and surgical grand rounds on Tuesday mornings, has caused her to start over several times. Elsewhere in the hospital, the cardiothoracic fellows have struggled to perform an adequate number of transplant cases within their standard 2-year residency, despite the appearance of sufficient case volume annually. Creating the optimal call schedule to maximize equitable distribution of these cases among the fellows has also become a challenge.

These manual scheduling issues have become an everyday part of graduate medical education but could be alleviated to some degree by adopting solutions other industries have used for years. We are all familiar with the fact that one snowstorm on the East Coast can devastate the schedules of every airport in the country, but more often than not, the airlines keep the majority of people flying.

So how do they do it? They use solutions from the field of operations research.

Operations research is the science of strategic decision making. It enables real world problems to be translated into mathematical models that can be analyzed and optimized to create actionable solutions. The field originated prior to World War II, when scientific teams were established to

study strategic and tactical dilemmas in military operations. Today, operations research is used to address complex personnel, logistical, and scheduling issues in many commercial industries, including air travel and health care, by using data to inform decisions. In health care, the earliest applications of operations research were related to nurse scheduling and operating room utilization,^{1,2} and now this research can be applied to solving numerous other logistical problems, including residency scheduling dilemmas.

Today's health care environment has led to residency and educational scheduling constraints for a variety of reasons. First, the institution of the 80-hour workweek in 2003 followed by the 16-hour maximum shift for interns in 2011 has limited the flexibility of scheduling. Within these restricted work hours, residents must comply with varying Accreditation Council for Graduate Medical Education (ACGME) operative, procedural, and educational requirements. When one adds in the unpredictable nature of patient or operative case exposure and evolving technologies, it becomes increasingly difficult for programs to ensure that all residents have the necessary experience to become competent. Within these constraints, balancing the educational goals with the service requirements of training further complicates the scheduling process.

Despite this increasing complexity, most programs rely on administrators or senior residents to create educational and service schedules. Often, they do so without any formal training in scheduling methods, using historical data and heuristics handed down from one chief to the next. Manual scheduling is time consuming, error prone, and it often results in poor resident satisfaction. In contrast, the University of Michigan Health System has conducted several collaborations with faculty and students from the Department of Industrial and Operations Engineering and the Center for Healthcare Engineering and Patient Safety, in which operations research techniques are used to build optimized schedules that satisfy all patient care and educational requirements.

One example is our pediatric emergency department, which has now been building such schedules collaboratively for several years. At the press of a button, a complete schedule is created, which can then be reviewed and revised

All authors are at the University of Michigan. **Tiffany N. S. Ballard, MD**, is PGY-5, Section of Plastic Surgery, Department of Surgery; **Tyler R. Grenda, MD**, is PGY-5, Section of General Surgery, Department of Surgery; **Amy M. Cohn, PhD**, is Arthur F. Thurnau Associate Professor of Operations Research, Department of Industrial and Operations Engineering, and Associate Director, Center for Healthcare Engineering and Patient Safety; **Mark S. Daskin, PhD**, is Clyde W. Johnson Collegiate Professor of Industrial and Operations Engineering and Chair of the Department of Industrial and Operations Engineering; **F. Jacob Seagull, PhD**, is Assistant Professor of Learning Health Sciences and Director of the Patient Safety and Quality Leadership Scholars Program; and **Rishindra M. Reddy, MD, FACS**, is Assistant Professor of Surgery, Section of Thoracic Surgery, Department of Surgery.

Corresponding author: Tiffany N. S. Ballard, MD, University of Michigan, Taubman Center, SPC 5340, 1500 E Medical Center Drive, Ann Arbor, MI 48109, 734.936.5895, tballard@med.umich.edu

DOI: <http://dx.doi.org/10.4300/JGME-D-14-00581.1>

by the chief resident or administrator. This has led to a number of benefits, including improved efficiency in schedule creation, elimination of favoritism in the creation of the schedule, and the ability to quickly recover from last-minute changes in the schedule. Perhaps even more importantly, these schedules can prioritize equity across residents in terms of the number of night shifts, the frequency of day-to-night shift transitions, personal requests, and so on.

A second example pertains to our procedural-based specialties, where the number of cases required for training presents another scheduling conundrum. The unpredictable arrival of emergent cases, such as cardiothoracic transplants, and the fact that the experience of any given resident can vary greatly as a matter of being in the right place at the right time, adds complexity to this issue. In another collaboration between engineering and clinical faculty and students, we were able to evaluate our institutional volume for cardiothoracic transplants and the ability to train fellows in light of the United Network for Organ Sharing certification requirements.³ Using historical transplant data, this team built a model to simulate the number and timing of transplants that may occur within a given time period. The model allowed the simulation of 10 000 resident years using different call schedules and varied case volumes. This showed that the minimum number of annual operations needed to train our residents is much greater than expected. Intuitively, 4 residents, who each need 15 cases, would require a minimum of 60 cases at an institutional level. Based on the inherent variability of transplant case arrival, the minimum number needed to train 4 residents using ACGME-compliant call schedules was closer to 90 cases. Improving the flexibility and responsiveness of residency scheduling will allow for more efficient training and reduced institutional case volume by promoting equity.

Given the individual benefits of operations research-based tools that can schedule residents and predict the annual volume necessary to adequately train residents, synergy can be found in taking an integrated approach to addressing both problems. Conceptually, computer applications could build schedules programmed to address case requirements, incorporate work hour regulations and service-specific needs, and account for the number and

varying levels of trainees. Utilizing historical case data for each training program would enable a tailored schedule to be created for a particular institution. With such a system, unexpected changes in the scheduling of a particular clinic to a different day might not necessitate the rebuilding of an entire schedule. Rather, the modification could be addressed by altering the inputs to the model and then rerunning the algorithm, which automatically accounts for all other previously determined requirements. Furthermore, if a training program expands its faculty with an associated increase in procedural case volume or adds another resident, the system could account for the impact of these changes.

To incorporate these scheduling strategies at readers' own institutions, we recommend contacting the nearest industrial engineering department at a local university. In fact, our initial contacts were identified through a quick Internet search. Many industrial engineering departments will have faculty who are well versed in optimization and interested in collaboration with clinical practitioners, as was our experience. Alternatively, there are commercial entities that specialize in physician scheduling, although university collaboration may be more mutually beneficial, leading to greater customization and ongoing joint efforts. If you encounter difficulty finding a collaborator experienced in optimization and scheduling, please contact the Center for Healthcare Engineering and Patient Safety at the University of Michigan.

To continue to meet the demands of today's changing medical education environment, training programs will need to employ more advanced approaches to solving scheduling dilemmas such as the examples provided. Operations research, with its established record of success in other fields, offers innovative approaches to enable graduate medical education programs to efficiently address this constantly changing field.

References

- 1 Bennett TR, Duckett SJ. Operations research and nurse staffing. *Int J Biomed Comput.* 1981;12(5):433–438.
- 2 Kuo PC, Schroeder RA, Mahaffey S, Bollinger RR. Optimization of operating room allocation using linear programming techniques. *J Am Coll Surg.* 2003;197(6):889–895.
- 3 Chung J, Obi AT, Chen R, Lin W, Sun S, Chen Z, et al. Estimating minimum program volume needed to train surgeons: when 4×15 really equals 90. *J Surg Educ.* 2014;72(1):61–67.