

Operating Room Scheduling and Adaptive Control Using a Priority First Fit Decreasing Heuristic

Wei Li [†]
Victoria L. Mitchell [‡]
Barrie R. Nault [‡]
Denise Brind ^{*}

wei.mike.li@uky.edu
victoria.mitchell@haskayne.ucalgary.ca
nault@ucalgary.ca
denise.brind@albertahealthservices.ca

[†] Department of Mechanical Engineering, University of Kentucky

[‡] Haskayne School of Business, University of Calgary

^{*} Foothills Medical Centre, Alberta Health Services

June 14, 2015

Copyright © Wei Li, Victoria L. Mitchell, Barrie R. Nault and Denise Brind, 2015.
All rights reserved.

Operating Room Scheduling and Adaptive Control Using a Priority First Fit Decreasing Heuristic

Abstract

Operating room (OR) scheduling is a critical factor affecting overall hospital performance. We examine OR scheduling from two perspectives. In the first perspective we propose a scheme for OR block scheduling that uses a heuristic developed for a three-machine flow shop where the three phases of the peri-operative process (pre-op, OR, and post-op) correspond to the three-machine flow shop. This approach facilitates a hospital-as-a-system perspective. The second perspective used to examine OR scheduling is adaptive control of the OR slate. Recognizing that there are many factors affecting OR throughput performance, especially preemptions from emergent and urgent cases, adaptive control of the OR slate is necessary. To realistically improve performance, adaptive control of the OR slate should incorporate constraints on how surgeries can be rescheduled. We examine the benefits from adaptive control of the OR slate that uses a priority first fit decreasing (PFFD) heuristic while incorporating constraints on OR slate rescheduling. The PFFD heuristic is a priority-driven variation of the classic FFD heuristic used in bin packing problems. We develop a scheme for OR block scheduling and our PFFD heuristic. We then demonstrate our PFFD heuristic in a simulation-based case study, and subsequently run a simulation using 1000 instances to test the performance of our PFFD heuristic in OR slate scheduling and OR slate adaptive control showing improvements in performance relative to the frequently used first-come-first-served rule.

Keywords: Operating Room Scheduling, Bin Packing, Heuristic, Flow Shop

1. Introduction

Hospital performance is tightly coupled with efficient operating room (OR) utilization. OR scheduling is a critical input to OR utilization, where resource deployment and throughput is organized by a series of OR timetables that culminate in a daily OR slate. Although OR scheduling has been the focus of research for decades, inefficiencies in OR utilization persist. These inefficiencies often arise from an inability to adapt to dynamic changes in OR demand, evidenced by lengthy OR wait lists. Static OR scheduling models tend to ignore pre- and post-operative influences that impact OR utilization and performance (Magerlein & Martin, 1978; Cardoen, Demeulemeester, & Belien, 2010). A systems engineering perspective enables a more fluid approach to OR scheduling that takes demand uncertainty and system interdependencies into consideration (e.g., Testi, Tanfani, & Torre, 2007; Vanberket et al., 2011).

In the context of systems engineering, OR procedures constitute one phase in a sequence of surgery-related activities called the Perioperative Process (peri-op). The peri-op process is divided into three distinct phases: preoperative, intraoperative, and postoperative periods. The preoperative (pre-op) phase encompasses all aspects of preparation for surgery – the necessary inputs to the surgery experience. The intraoperative (OR) phase entails the surgical experience from patient arrival in the operating room to arrival in postoperative (post-op) units such as the post-anesthesia care unit (PACU). Post-op activities comprise the third perioperative phase (Phillips, 2004). The three phases resemble tightly coupled subsystems whose adjustment in the face of uncertainty determines overall system success (Wierner, 1961).

Using a 3-machine flow shop as an analogy for the peri-op process (Li, Mitchell, & Nault, 2013), we apply the concept of adaptive control to the OR slate employing a priority first fit decreasing heuristic. In the 3-machine flow shop the first machine represents the pre-op phase of the peri-op process, the second machine represents the OR phase, and the third machine represents the post-op phase. Few studies have taken the entire peri-op process into consideration when constructing OR schedules (for a description of this challenge see Gupta & Denton, 2008), mainly because a hospital is a complex system and the 3-machine flow shop is *NP*-complete (Garey & Johnson, 1976). In other words, it is difficult to get optimal solutions to a general complicated *NP*-complete problem in a given time.

Although a hospital is a complex system and 3-machine flow shop production is *NP*-complete, OR scheduling has to be integrated across the peri-op process. To take a step toward this integrated scheduling, we propose a scheme for OR block scheduling in which our state space with head, body and tail (SS-HBT) heuristic --

formerly used in the context of a 3-machine flow shop to sequence (Li et al., 2013) -- can be employed to set up a schedule of surgical block times, otherwise known as an OR block schedule. This is a long-term schedule, either semi-annual or annual, and essentially allocates hospital and surgeon resources. The 3-machine flow shop approach allows us to incorporate constraints from pre-op and post-op units.

Within an OR block schedule there is scheduling of the OR slate. The OR slate is a short-term schedule, daily or weekly, and essentially allocates surgical cases (patients) to time slots within the OR block schedule. OR slates face constant changes from surgeries that take longer or shorter than planned, and more critically from preemptive emergent and urgent cases. Consequently, the OR slate is rescheduled dynamically, and we propose a priority first fit decreasing (PFFD) heuristic for adaptive control of the OR slate. Our PFFD heuristic is developed to incorporate constraints based on hospital rescheduling policies, and is tested with a set of randomly generated case studies.

Our methodology begins by developing a scheme for OR block scheduling, and the mechanisms and constraints incorporated in our PFFD heuristic. We then demonstrate our PFFD heuristic in a simulation-based case study, showing both the OR slate scheduling phase and an example of the adaptive control of the OR slate, comparing these to the frequently used first-come-first-served (FCFS) rule. Subsequently, we run a simulation using 1000 instances of an OR slate in order to test the performance of our PFFD heuristic in OR slate scheduling and in OR slate adaptive control. This allows us to show improvements in different performance dimensions relative to the frequently used FCFS rule.

The remainder of this paper is organized as follows. Section 2 introduces our SS-HBT heuristic and our scheme for OR block scheduling and Section 3 introduces our PFFD heuristic for adaptive control of the OR slate. A case study of adaptive control of the OR slate is presented in Section 4 together with our 1000 instance simulation, and finally we draw conclusions and suggest future work in Section 5.

2. A Scheme for OR Block Scheduling

Both minimization of the maximum completion time, $\min(C_{max})$, or of the total completion time, $\min(\sum C_j)$, have the potential to minimize patient waiting times. Based on the literature in flow shop production scheduling, we know that we can get optimal solutions to $\min(C_{max})$ or $\min(\sum C_j)$ for single machine flow shop production. For a 2-machine flow shop production setting, we can get optimal solutions to maximum completion time minimization, although total completion time minimization is *NP*-complete (Hoogeveen & Kawaguchi, 1999). For 3-machine flow shop production, even the maximum completion time minimization is *NP*-complete (Garey & Johnson, 1976). Moreover, Li et al. (2013) have shown that these two objectives, $\min(C_{max})$ and $\min(\sum C_i)$, are inconsistent even in the case of 2-machine flow shop production. That is, minimization of one objective can result in an increase in the other. The general reason for this is that in an *m*-machine flow shop for machines $i = 2, \dots, m$ completion times are evaluated on the last machine *m*, and the inconsistency of $\min(C_{max})$ and $\min(\sum C_j)$, exists between every pair of adjacent machines. If we model OR scheduling as a 3-machine flow shop, then inconsistencies exist between input units and ORs (i.e., pre-op and OR), and between ORs and output units (i.e., OR and post-op), which complicates the wait time minimization. Indeed, this partially explains difficulties in applying Industrial Engineering technologies to healthcare systems.

Taking into consideration the complexities in healthcare systems and in production scheduling, we propose the following scheme in Figure 1 for overall OR scheduling based on expected performance in the long and short term. Expected performance can be any metric used in OR management over a period of time, such as utilization of OR blocks in a year or patient flows over a week, etc., which are based on known inputs, such as the distribution of surgery type, average surgery times, variation in surgery times, patterns of emergencies, etc.

In Figure 1, the loop in thin lines represents the computer-based scheduling phase, and the loop in thick lines represents the actual schedule execution and adaptive control in all pre-op units, ORs, and post-op units. The scheme we use for OR block scheduling, and OR slate scheduling with adaptive control is broken into the following steps.

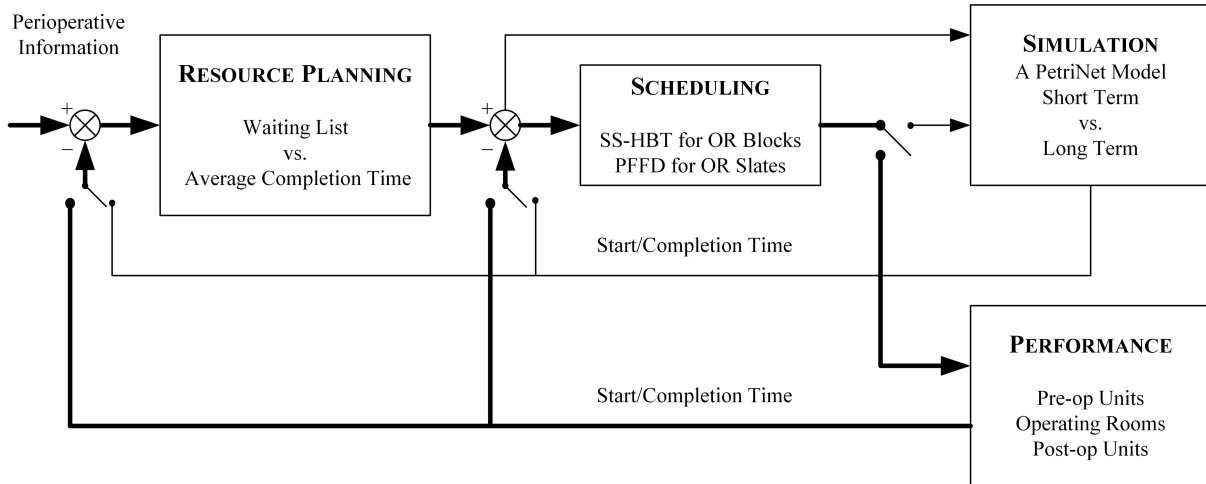


Figure 1. A Scheme for OR Scheduling

1. Resource planning. Based on the existing waiting list of surgery cases and historical OR data from past surgery cases, we can generate the necessary information for planning, including the number and types of surgery, average surgery times and variation in surgery times for types of surgeries by surgeon as well as processing times in pre-op and post-op units for each types of surgery. Accordingly, we can generate a resource matrix including the number of ORs and surgeons for each type of surgery etc., which is used to estimate an expected average completion time, or a maximum completion time for all of n surgeries in the waiting list. The number of ORs is typically fixed over the relevant time horizon and each OR has a range of surgery types that can be performed. Each OR may, for example, operate for 10 hours per day, six days a week, giving 3000 hours of OR time. Time in each OR is usually allocated by surgery type (e.g., cardiac, orthopedic, etc.), and demand for each surgery type is known from the existing waiting list and historical OR data. Moreover, OR time required for each surgery type is also known from medical standards and historical OR data, and these times may vary by surgeon.
2. Scheduling. With the necessary information and the resource matrix from the resource planning step, our SS-HBT heuristic can generate a long-term OR block schedule for a hospital or group of hospitals. The OR block schedule is essentially allocating OR surgical blocks (i.e., a day in an OR) by surgery type of even surgeon, taking into account the capacity constraints from pre-op and post-op units, with long-term being an annual or semi-annual OR block schedule. Breaking down the long-term block schedule for all n surgeries, we can set up necessary constraints, e.g., block times, and subsets of n individual surgical cases for our PFFD heuristic to generate short-term schedules known as the OR slate, where short-term is either daily or weekly. These short-term schedules for the OR slate incorporate the constraints set by the OR block schedule. Moreover, they can be adaptively controlled in the face of the variations in surgery times as well as pre-emptive emergent and urgent cases, while also accounting for hospital rescheduling policies.
3. Simulation. Long-term OR block schedules and short-term OR slates can be simulated in a Petri Nets-based model. This Petri Nets-based model allows us to determine inconsistencies across the peri-op process, and the interaction of performance in the long and short term. Human schedulers can adjust resources for pre-op, OR and post-op units to finalize the OR block schedules and the short-term OR slates. In this simulation phase, when it comes to dealing with disturbances the computation time and the flexibility of our heuristics are critical, and the resulting changes in schedules from this adaptive control improves peri-op performance. For example, a human scheduler can change the expected frequency of surgery cancellations in a week, which in turn affects utilization and resource allocation not only in ORs, but also in pre-op and post-op units.
4. Adaptive control. Our adaptive control applies to the OR slate whereby both the flexibility and computation time of our heuristics are critical for generating solutions to disturbances in real-time. After the desired

performance of the hospital as a system is finalized during the OR block scheduling phase and setting of the subsequent OR slates, the two schedules can be sent to the peri-op units. As seen in Figure 1, it is at this point the thin line loop is switched to the thick-line loop. In the real execution phase, if disturbances happen to any pre-op or post-op units, or in an OR, the thick-line loop may be switched back to the thin-line loop for adaptive control. Human schedulers may adjust the resource matrix and related input data to our heuristics accordingly, such as the number of surgeries, processing times in each peri-op unit, etc. The changes to the OR slate in response to the disturbance should be evaluated based on feasibility, interaction of long-term and short-term performances of a hospital and surgical suites. For example, consider a scheduled surgery that has to be cancelled at a specific time. OR slate rescheduling in real-time (adaptive control) is based on the short-term availability of patients, surgeons, nurses, and necessary equipment in ORs, pre-op and post-op units. Performance of this rescheduling should be aggregated to the OR block schedule to capture whether there are impacts on OR utilization and patient flow in the long-term as well. If desired results can be reached in this adaptive control phase, then the change in OR slate in response to the disturbance can be sent to the relevant peri-op units.

3. Adaptive Control of the OR Slate

3.1 Description of the PFFD Heuristic

For OR slate scheduling and adaptive control of the OR slate we propose our priority first fit decreasing (PFFD) heuristic based on research in bin packing (Coffman Jr., Garey & Johnson, 1996; van Houdenhoven et al., 2007). In general, in order to pack (or schedule) n items, where the volume of each item is less than or equal to 1, the bin packing problem is effectively how to use the minimum number of unit-capacity bins to pack all n items. The analogy between bin packing and OR slate scheduling is that OR block times can be regarded as bin capacities. The surgery times are less than or equal to the block times, and OR slate scheduling uses the minimum number of OR blocks to finish all n surgical cases. This short-term adaptive control of the OR slate is consistent with the long-term objective of minimizing the maximum completion time. Different from the first fit decreasing (FFD) heuristic in the bin packing literature, our PFFD heuristic does the following. First, it groups the surgical cases into priority classes, and sorts the surgical cases in each priority class group into a non-increasing order of surgery times. Then it applies FFD to put surgical cases into each block, and calculates the capacity left not only by from mean surgery times but also incorporates the variation in surgery times by adding the standard deviation assuming surgical cases are independent. This means slack time -- the combination of standard deviations of surgery times -- is also taken into consideration for OR slate scheduling, and this slack time affects OR utilization. The traditional FFD heuristic does not consider the priority classes or slack times.

3.2 OR Slate Scheduling Case Study

As illustrated in our scheme, there are two phases: an OR slate scheduling phase and an adaptive control of OR slate scheduling phase. In the OR slate scheduling phase the PFFD heuristic generates daily or weekly OR slates. After being generated in the scheduling phase, an OR slate will be executed in real-time in a hospital.

In the adaptive control of the OR slate scheduling phase, if any disturbance happens along the peri-op process as the OR slate is being executed, then adaptive control is necessary to generate solutions that reschedule surgical cases in the OR slate. We provide case studies for these two phases.

3.3 OR Slate Scheduling Phase

Assuming the waiting list is long and consequently the number of surgical cases n is large, we generate the annual OR block schedule using our SS-HBT heuristic. Taking the OR block schedule as given, for our case study of OR slate scheduling we set the daily block times at 10 hours. For a general type of surgery (e.g., orthopedic), we randomly select 40 surgical cases for our PFFD heuristic to generate daily OR slates, the surgery times of which are normally distributed in a range of [30, 240] minutes and with a variation in processing time (standard deviation) of 10%. Moreover, each surgical case has a patient ID and a priority.

Usually, a hospital that takes surgical case priorities into consideration schedule surgical cases with higher priorities earlier in the slate, and uses the FCFS rule for OR slate scheduling. This means surgical cases with the same priority are scheduled by the FCFS rule (Magerlein & Martin, 1978; Cardoen et al., 2010). Table 1 shows the outcomes of OR slate scheduling using the FCFS rule for the 40 surgeries, which are sorted first by priority and then scheduled by the FCFS rule.

Table 1 shows that these 40 surgical cases are scheduled in 9 blocks, the numbers under each block are patient ID numbers from 1 to 40. If there is only one OR assigned to handle these 40 surgical cases, then it would take 9

days (or blocks) to complete them. Otherwise, the number of days to finish these surgical cases equals the number of blocks divided by the number of ORs assigned to these surgical cases. In this particular case study, we assume there is only one OR assigned to these 40 surgical cases, as Daily Schedule indicated in the first column

Table 1. OR Slate Scheduling by the FCFS Rule

	BLK:1	BLK:2	BLK:3	BLK:4	BLK:5	BLK:6	BLK:7	BLK:8	BLK:9
	27	25	31	18	24	17	4	40	16
	3	28	39	20	33	34	26	8	
Daily Schedule	5	30	6	23	2	35	29	9	
	13	1	32	38	12	36	7	11	
	14	10	15		37	21			
		19							
		22							
Slack Time	19.82	15.84	15.59	13.53	18.95	19.31	15.26	16.16	8.00
Surgery Time	574	555	577	581	573	530	519	505	121
Utilization (%)	99.00	95.17	98.83	99.17	98.67	91.67	89.17	87.00	21.50

and the first row in Table 1. As defined earlier, Slack Time means the standard deviation of variation in processing times for surgeries scheduled in a block. Surgery time is the sum of surgery times for surgeries in a block. Utilization equals the roundup of the slack time plus the surgery time divided by 600 minutes, stated in percentages.

If we use our PFFD heuristic to schedule these 40 surgical cases instead of the FCFS rule, then we use only 8 blocks instead of 9 blocks. Table 2 shows the schedule generated by our PFFD heuristic.

Obviously, the PFFD heuristic achieves a higher OR utilization than the FCFS rule, with no violation of 10-hour or 600-minute constraint set by the OR block schedule (potentially generated by the SS-HBT heuristic).

To test the computation time, we generated a case study that consisted of 10000 surgical cases. On a standard IBM laptop with 4 GB RAM and 2.30 GHz Dual Core it took only 6.0684 seconds for our PFFD heuristic to generate an OR slate. In practice, this is fast enough for adaptive control of the OR slate.

Table 2. OR Slate Scheduling by the PFFD Heuristic

	BLK:1	BLK:2	BLK:3	BLK:4	BLK:5	BLK:6	BLK:7	BLK:8
	27	28	31	20	18	2	4	9
	3	30	1	24	33	34	26	11
Daily Schedule	5	25	10	23	17	36	40	7
	13	22	19	32	35	12	21	15
	14	39	6		37	29	8	16
			38					
Slack Time	19.82	15.26	16.12	22.49	13.45	15.87	18.79	13.49
Surgery Time	574	565	577	567	574	580	570	528
Utilization (%)	99.00	96.83	99.00	98.33	98.00	99.33	98.17	90.33

3.4 Adaptive Control of the OR Slate Example

To examine adaptive control of the OR slate, we generate an example where we assume that the OR slate is determined by our PFFD heuristic as in Table 2, and the schedules for the first two days have been carried out. We randomly chose surgical case number 25 to be canceled on day 2 or Block 2. The surgery could be cancelled due to emergent cases, a patient failing to arrive, etc. Surgical case number 25 must be rescheduled on a later day, possibly from day 3 to day 8.

Table 3. Adaptive Control of the OR Slate by the PFFD Heuristic

	BLK:3	BLK:4	BLK:5	BLK:6	BLK:7	BLK:8
	31	20	18	2	4	9
	1	23	33	34	26	11
Daily Schedule	10	25	32	36	40	7
	19	24	17	12	21	15
	6		35	29	8	16
	38					37
Slack Time	16.12	22.49	13.45	15.87	18.79	13.53
Surgery Time	577	575	584	580	570	561
Utilization (%)	99.00	99.67	99.67	99.33	98.17	95.83

Hospitals often have constraints from rescheduling policies such as patient matches with their surgeon, surgeon availability, the inflexibility of patients that are from out of town, etc. In our example, we randomly generated a set of constraints from hospital rescheduling policies: all surgeries on days 3 and 6 have to be preserved, which means no change is allowed to the OR slate moving surgical cases from day 3 to day 8. In addition, surgical case numbers 20 and 23 on day 4, surgical case number 18 on day 5, and surgical case numbers 4 on day 7 must also be preserved. Taking these constraints into consideration, our PFFD heuristic adaptively generated a solution as shown in Table 3.

Our PFFD heuristic is flexible enough to take multiple constraints into consideration such as detailed above, and the OR utilization achieved by adaptive control of the OR slate is high -- no lower than 95% for any block.

3.5 1000 Instance Simulation

To further test the performance of our PFFD heuristic, we generate 1000 random instances, of which there are 40 surgical cases in each instance and each instance is scheduled as an OR slate. The details of the parameters setting for the random data generation are the same as in the earlier section. We can see from Table 1 that the utilization of the last day or block L is low in the initial OR slate schedule because the number of surgical cases is fixed. In our example it is 40. Therefore we use the average utilization from day 1 to the second last day $L-1$ to compare performance of the PFFD heuristic and the FCFS rule in the OR slate scheduling phase. Moreover, we also compare the number of blocks that resulted in the OR slate generated from the PFFD heuristic and from the FCFS rule.

In the adaptive control of the OR slate phase, we randomly chose that the last surgical case on day 2 is canceled for some reason and has to be rescheduled into a later day from day 3 to the last day L ; surgical cases on days 3 and 6 have to be preserved as scheduled, and the first surgical case has to be preserved on days 4 and 5. For adaptive control with the above constraints, we ran the PFFD heuristic and the FCFS rule, and then compared the average utilization from day 3 to the second last day $L-1$ respectively. The results are summarized in Table 4.

Table 4. Average OR Utilization in Scheduling and Control Phases

	Util. of day 1~L-1 (%)		Util. of day 3~L-1 (%)		Improvement (%)	
	In scheduling phase		In control phase		PFFD over FCFS	
	FCFS	PFFD	FCFS	PFFD	Scheduling	Control
Average	95.60	96.75	95.13	96.27	1.22	1.23
Maximum	99.04	98.92	99.53	99.25	7.33	10.18
Minimum	88.71	89.35	87.29	89.17	-5.77	-4.59
Deviation	1.60	1.19	1.96	1.57	1.75	1.99

From Table 4 we can see that in the OR slate scheduling phases our PFFD heuristic can achieve an average utilization of 96.75%, which is better than that of 95.60% obtained using the FCFS rule. On average, the PFFD heuristic has a 1.22% improvement in utilization rates compared to the FCFS rule for OR slate scheduling, where the maximum improvement is 7.33%. A t-test of the difference in utilization rates between our PFFD heuristic and the PCPS rule for OR slate scheduling is significant at $p=7.71E-89$.

In adaptive control of the OR slate, the PFFD heuristic also performs better than the FCFS rule for utilization rates, with an average utilization of 96.27% for the PFFD heuristic compared to that of 95.13% for the FCFS rule. The average improvement of PFFD over FCFS is 1.23% with a maximum improvement of 10.18%. However, for some OR slates, the PFFD heuristic performs worse than the FCFS rule. That is why we have a minimum improvement of -5.77% in the scheduling phase, and that of -4.59% in the adaptive control phase. Nonetheless, t-test for the difference in utilization rates between our PFFD heuristic and the PCPS rule for adaptive control of the OR slate is $p=4.57E-4$.

In terms of the most critical measure of performance -- the number of OR blocks needed to complete the n surgical cases -- the PFFD heuristic uses one less OR block than the FCFS rule for 107 out of the 1000 instances in both OR slate scheduling and adaptive control of the OR slate. The FCFS rule uses one less block than the PFFD heuristic for 19 instances. Thus, over time the net increase in efficiency as measured by the number of OR blocks needed to complete a set of surgical cases using the PFFD heuristic instead of the FCFS rule is almost 9%. This finding shows that the small improvements in OR utilization result in much greater efficiencies in the number of OR blocks needed.

4. Conclusion and future work

Waiting list minimization is a broadly defined objective in OR scheduling. There are two specific objectives in flow shop production scheduling that are related to minimizing the waiting list. One is to minimize maximum completion time, $\min(C_{max})$, and the other is to minimize total completion time, $\min(\sum C_j)$. Seemingly, the maximum completion time is involved in the total completion time.

In this work we drew the analogy between a 3-machine flow shop and the hospital peri-op process made up of pre-op, OR and post-op units. This analogy has the potential to transfer useful results and methods from manufacturing to healthcare. We also described the process of setting up a long-term (annual, semi-annual) OR block schedule using an SS-HBT heuristic, and demonstrated through a case study and simulation how a short-term (weekly, daily) OR slate can be scheduled using our PFFD heuristic and how it can be used to adaptively reschedule the OR slate in reaction to disturbances.

Working from the short-term perspective, we recognize that disturbances frequently affect OR schedules--patients or surgeons become unavailable, emergent cases require pre-emption, etc. As a consequence, adaptive control of the OR slate is necessary, and it must incorporate the constraints set up from the hospital perspective--in particular hospital policies. We proposed a method for OR slate scheduling and adaptive control using our PFFD heuristic, designed a case study to show how the PFFD heuristic is better than the current FCFS approach, and report results from a simulation comparing the two heuristics. The flexibility and computation time of our PFFD heuristic support this method of OR slate scheduling and adaptive control of the OR slate, and demonstrates improvement over the current FCFS approach.

In our method, the SS-HBT and PFFD heuristics partially fill the gap between the theoretical research on Industrial Engineering (IE) and the application of IE technologies to healthcare systems. As hospitals face dynamically occurring disturbances in real time, there is a need to carry out deep research on how adaptive

control can improve the performance of ORs and the peri-op process in short-term, and how this short-term improvement enhances the performance of hospitals in the long run. This work has the potential to provide robustness to OR scheduling and control.

Acknowledgments

We greatly appreciate the support from Alberta Health Services through the Foothills Medical Centre in Calgary, Alberta, and MITACS grant numbers IT01601/FR01; and support from the Natural Sciences and Engineering Council of Canada, and the Social Sciences and Humanities Research Council of Canada. We also thank Jeanette Burman for excellent editing assistance.

References

- Cardoen, B., E. Demeulemeester & J. Belien. (2009). Sequencing surgical cases in a day-care environment: An exact branch-and-price approach. *Computers & Operations Research*, 36(9), 2660-2669. <http://dx.doi.org/10.1016/j.cor.2008.11.012>
- Cardoen, B., E. Demeulemeester & J. Belien. (2010). Operating room planning and scheduling: a literature review. *European Journal of Operational Research*, 201(3), 921-932. <http://dx.doi.org/10.1016/j.ejor.2009.04.011>
- Coffman Jr., E.G., M.R. Garey & D.S. Johnson. (1996). Approximation algorithms for bin packing: a survey. In D. Hochbaum (Ed.), *Approximation Algorithms for NP-hard Problems*. Boston, MA: PWS Publishing Company.
- Garey, M.R., D.S. Johnson & R. Sethi. (1976). The complexity of flowshop and job shop scheduling. *Mathematics of Operations Research*, 1(2), 117-121.
- Gupta, D. & B. Denton. (2008). Appointment scheduling in health care: Challenges and opportunities. *IIE Transactions*, 40(9), 800-819. <http://dx.doi.org/10.1080/07408170802165880>
- Hoogeveen, J.A. & T. Kawaguchi. (1999). Minimizing total completion time in a two-machine flowshop: analysis of special cases. *Mathematics of Operations Research*, 24(4), 887-910. <http://dx.doi.org/10.1287/moor.24.4.887>
- Li, W., V.L. Mitchell & B.R. Nault (2013). Waiting-List minimization in operating room scheduling. Working paper.
- Magerlein, J.M. & J.B. Martin (1978). Surgical demand scheduling: a review. *Health Services Research*, 13(4), 418-433.
- Phillips, N. (2004). *Berry & Kohn's Operating Room Technique* (10th ed.). St. Louis, MO: Mosby.
- Testi, A., E. Tanfani & G. Torre (2007). A three-phase approach for operating theatre schedules. *Health Care Management Science*, 10(2), 163-172. <http://dx.doi.org/10.1007/s10729-007-9011-1>
- van Houdenhoven, M., J.M. van Oostrum, E.W. Hans, G. Wullink & G. Kazemier. (2007). Improving operating room efficiency by applying bin-packing and portfolio techniques to surgical case scheduling. *Anesthesia and Analgesia*, 105(3), 707-714. <http://dx.doi.org/10.1213/01.ane.0000277492.90805.0f>
- Vanberkel, P.T., R.J. Boucherie, E.W. Hans, J.L. Hurink, W.A.M. van Lent & W.H. van Harten. (2011). Accounting for inpatient wards when developing mater surgical schedules. *Anesthesia and Analgesia*, 112(6), 1472-1479. <http://dx.doi.org/10.1213/ANE.0b013e3182159c2f>
- Wiener, N. (1961). *Cybernetics* (2nd ed.). Cambridge, MA: MIT Press.