# *Modeling the impact of changing patient transportation systems on peri-operative process performance in a large hospital: insights from a computer simulation study*



**Citation:** Segev, Danny, Levi Retsef, Peter F. Dunn, and Warren S. Sandberg. "Modeling the Impact of Changing Patient Transportation Systems on Peri-Operative Process Performance in a Large Hospital: Insights from a Computer Simulation Study." Health Care Management Science vol. 15, no. 2 February 2012, pp. 155–169.

**As Published:** http://dx.doi.org/10.1007/s10729-012-9191-1

**Publisher:** Springer-Verlag

**Persistent URL:** http://hdl.handle.net/1721.1/104855

**Version:** Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Health Care Manag Sci (2012) 15:155–169
DOI 10.1007/s10729-012-9191-1

# Modeling the impact of changing patient transportation systems on peri-operative process performance in a large hospital: insights from a computer simulation study

**Danny Segev · Retsef Levi · Peter F. Dunn · Warren S. Sandberg**

Received: 29 June 2011 / Accepted: 6 January 2012 / Published online: 14 February 2012


**Abstract** Transportation of patients is a key hospital operational activity. During a large construction project, our patient admission and prep area will relocate from immediately adjacent to the operating room suite to another floor of a different building. Transportation will require extra distance and elevator trips to deliver patients and recycle transporters (specifically: personnel who transport patients). Management intuition suggested that starting all 52 first cases simultaneously would require many of the 18 available elevators. To test this, we developed a data-driven simulation tool to allow decision makers to simultaneously address planning and evaluation questions about patient transportation. We coded a stochastic simulation tool for a generalized model treating all factors contributing to the process as JAVA objects. The model includes elevator steps, explicitly accounting for transporter speed and distance to be covered. We used the model for sensitivity analyses of the number of dedicated elevators, dedicated transporters, transporter speed and the planned process start time on lateness of OR starts and the number of cases with serious delays (i.e., more than 15 min). Allocating two of the 18 elevators and 7 transporters reduced lateness and the number of cases with serious delays. Additional elevators and/or transporters yielded little additional benefit. If the admission process produced ready-for-transport patients 20 min earlier, almost all delays would be eliminated. Modeling results contradicted clinical managers' intuition that starting all first cases on time requires many dedicated elevators. This is explained by the principle of decreasing marginal returns for increasing capacity when there are other limiting constraints in the system.

**Institution for Attribution** Department of Anesthesia, Critical Care and Pain Medicine, Massachusetts General Hospital

**Prior Presentations** None

**Summary Statement** A generalized model of patient transport indicated that relatively few transporters and elevators are needed for optimal performance for delivering patients, but that process redesign is important to provide enough time for transportation.

D. Segev
Department of Statistics, University of Haifa,
Haifa 31905, Israel

R. Levi
Sloan School of Management,
Massachusetts Institute of Technology,
E62-562,
Cambridge, MA 02139, USA

P. F. Dunn
Harvard Medical School Anesthetist, Department of Anesthesia and Critical Care and Executive Medical Director of the Operating Rooms, Massachusetts General Hospital,
Boston, MA 02114, USA

W. S. Sandberg (✉)
Department of Anesthesiology, Surgery & Biomedical Informatics,
Vanderbilt University School of Medicine,
1211 21st Avenue South, 701 MAB,
Nashville, TN 37212-1050, USA
e-mail: Warren.sandberg@vanderbilt.edu

**Keywords** Hospital transportation · Logistics · Operating room · Elevators · First case on time starts · Perioperative systems design

## 1 Introduction

Large hospitals with on-site clinics, operating rooms (ORs), intensive care units, diagnostic and procedure units, and inpatient beds routinely handle thousands of patient transports per day in multiple complex care environments. This makes physical patient flow a fundamental challenge. Efficient, timely and safe transport of patients throughout the hospital is critical to accomplish the hospital's medical and financial goals. For many patients – especially those having procedures – the surgical admission center is the first station in the patient flow. Movement of patients within and out of the surgical admission center is felt throughout the perioperative care processes: Patients are admitted, prepared for surgery, and then transported to the ORs. Any delays in this stage immediately translate into a significant (negative) cumulative effect on the ability to move patients smoothly throughout the perioperative care environments and to perform surgical operations according to schedule.

Through ongoing and incremental efforts, well-run admission centers establish processes that ensure a smooth, timely flow of patients into their downstream areas. New construction or relocation can fundamentally disrupt admission center function by obliterating the established systems. Perhaps the impact of such disruption can be minimized and addressed in advance by careful planning and modeling of the constraints and potentials of the new patient flow pathways and capacities.

Our hospital plans to expand its OR capacity (through construction of a new building) to 70 ORs, and to renovate its admission center to accommodate more patients. The current (and future) admission center is on the 3rd floor of the hospital immediately adjacent to the OR suite, which in its current form is entirely on the third floor and inter-connected. As part of the construction plan the surgical admission center will temporarily be moved to the 12th floor of a different building for a period of 18-24 months. There is no possibility to permanently (or even temporarily) close the existing ORs and to erect a completely new building that will accommodate all surgical facilities in a single location. Therefore, facilities must be moved around within the existing structure, while at the same time maintaining full patient flow through all 52 currently used ORs.

Planning for changed patterns of patient transport (including multiple changes over the course of the project) was expected to be a significant challenge. The new admission center location is distant from the ORs, and will require transporting patients through elevators. Eventually the admission center will be returned to its adjacent location on the 3rd floor. However, some of the new ORs will be on the floor above or below. Thus, elevators and loss of co-location will to some extent always exist in our transportation system going forward.

Patient flows involving elevators have been viewed as a significant source of complexity in our perioperative systems in the past, and this impression has been borne out in empirical studies [1]. There are 18 elevators that serve the 12th floor of the three connected buildings (named Ellison, Blake and Gray) that could potentially serve the admission center. However, these elevators currently serve patients, visitors and hospital staff in various functions located in these buildings. Thus, it is clear that only very few elevators could be solely dedicated to the use of the admission center. Data from the Hospital Physical Plant Department indicate that these elevators have very long waiting times during peak hours, so it has been decided to allocate only a small number, if any, for exclusive use by OR transporters. Hence, there have been serious concerns regarding the impact of the new location on the ability to start surgical operations on time, which is one of the central performance measures of the perioperative care system.

The proposed admission center relocation thus leads to several planning and evaluation questions:

1. What are the right number, type, and location of dedicated elevators? How many transporters are required? Also, for a predetermined configuration of elevators and transporters, what are the practical effects of limited coordination between personnel? These questions are interrelated, and must be addressed in a holistic way. Moreover, to answer the latter questions, one has to define a clear set of desired performance measures by which the system performance is evaluated. (This is discussed in detail in the Methods Section.)
2. Given the current limitations on the availability of physical and human resources, is it possible to design efficient and timely transportation of patients based only on dedicated elevators and transporters? Otherwise, what additional changes in the current plan are possible, and how will they affect the performance of the system? For example, is it beneficial to simply start the admission process sooner?
3. What are the expected effects of potential variability in some parameters of the underlying environment (e.g., traveling times and volume of patients)? In other words, how robust is a given configuration to reasonable deviations from the present system characteristics?

We sought to address these questions by developing a generalized transport modeling tool that begins with actual operational data from the current setting, and then allows decision makers to dynamically and simultaneously address all of the planning and evaluation questions enumerated above. More generally, this tool can be used to model similar situations involving various types of patient and equipment transportation systems.

The rest of the paper is structured as follows: In the next section we briefly review the literature on patient transportation in hospitals, so as to put our work in context. In Section 3, we describe the setting in which the modeling tool was developed, our simulation approach, our methods of collecting input data for the model, and we define the terms used to assess performance. In Section 4, we describe our results. In Section 5, we discuss the results in terms of the managerial insights the modeling tool provided and with respect to other available techniques.

## 2 Literature review

Operations research disciplines and analytical methods, such as queuing models and traditional discrete event and Monte Carlo simulation approaches, have been extensively used to study healthcare management problems. Patient appointment scheduling, patient flow, manpower staffing and resource allocation are just some representative examples. We refer the reader to several surveys and books [2–6], and to the references therein, for a comprehensive review of the literature.

Many optimization techniques do not have the capacity to handle the complexities of medical systems, therefore requiring too many unrealistic assumptions, which may render solutions invalid. Such assumptions can be logical (i.e., over-simplifying the way a given system behaves in reality), structural (e.g., ignoring certain physical characteristics that cannot be sensibly handled by existing mathematical models), and quantitative (e.g., probabilistic assumptions such as sticking to particular distributions, independence, etc.). There has been some work on more complex environments and systems, in which multiple factors interact and give rise to more involved and hard-to-model dynamics [7].

To our knowledge, there are relatively few studies of hospital transportation systems[8–12]; these approaches seem to be rather limited to the specific applications in question [12]. The focus of most operations research literature in healthcare is investigating systems with relatively simple structure, such as appointments in an outpatient practice [13], desired level of staffing in a medical department [14–16], or surgery scheduling [17, 18]. In these examples, one can employ rather elementary modeling and optimization methods (e.g., queuing theory and linear programming).

Previous work on hospital transportation systems has focused on generalized transportation systems of the entire hospital [7, 19, 20]. In one study, transportation was optimized by creation of a centralized dispatch system [20]. The problem being addressed was one of organization. Elevators and modes of transport were not specifically considered. There are many instances, such as transport of patients, instruments or materiel from one location to another, as is common in operating room or procedure suites, where the organizational task is simple, but the timing and reliability are important. Moreover, little of the prior work considers transportation modeling as a planning tool to guide process changes in the setting of new construction.

Simulation of perioperative function, capacity and staffing has focused on the recovery room or staffing to cover emergency cases. There have been very few effective applications of these methods to large academic hospitals and to transportation systems of the kind addressed in our project. We consider all of the features of the intake transport process all at once, including transporters, mode of transport, elevators and the timing of process steps. Additionally, our contribution has value as a case study because of the managerial insight it provided, as will be elucidated in the Discussion. Similar to what other researchers inferred [21, 22], we find simulation methodology to be particularly suitable for modeling the complexities of health care facilities.

## 3 Methods

The use of retrospective OR process data to inform this simulation study was reviewed and approved by the Human Research Committee of Partners Healthcare, Boston, MA.

### 3.1 Setting description

This study was performed at Massachusetts General Hospital, where 52 operating rooms receive approximately 150 patients per day (mixed same-day-admission and day surgery patients) from an admission center with approximately 30 beds. The surgical admission center and all ORs are currently co-located on the third floor of several interconnected buildings, making the task of transporting patients rather straightforward. Specifically, transportation distances are short, and there is no need to use elevators. OR managers are also able to “manually” coordinate clinical staff and transportation services with ease by telephone or overhead paging. Intuitively, greater distances between units imply significantly less coordination between most parties involved.

The proposed new transportation process is schematically drawn in Fig. 1. We did not attempt to model the preoperative preparation process itself,[1] but rather treated the admission center as a black box that generates a sequence of ready-for-transport patients. In other words, the starting point of the process under consideration is not the time when a patient arrives to the hospital. Instead, it is the moment in time where all preoperative procedures have been completed and the patient can be transported to the corresponding OR. We refer to this time as Patient Ready for Transport (PRT) time. [23],[2]

As soon as a patient is ready, (s)he is assigned a transporter, whose only purpose from that point on is to complete the transportation task without any unnecessary delays. In the event of all transporters being momentarily unavailable (due to transporting prior patients), subsequent patients are placed in a queue; the rules for picking the next patient to be transported from this queue will be described below. Our fundamental assumption, which matches the present and future status of this system, is that there is a transportation team of limited size, dedicated solely to patient transportation, without engaging in any other activity. Due to financial constraints, this team is not planned to be augmented beyond the current 10 full-time transporters, although additional members could be considered if the number of transporters appears to be a significant constraint on process flow.

Once a transporter becomes available,[3] the next patient can be taken to one of an unspecified (but small – no more than 2 or 3 due to constraints imposed by the physical plant management) number of dedicated elevators. There are a total of 18 possible elevators, arranged in groups of 6, one group in each of three buildings (the Ellison, Blake or Gray buildings; see Fig. 1). All elevators have the same capacity: two stretchers (or gurneys) plus two transporters, or four wheelchairs, each with a transporter. Upon arrival to the chosen elevator group, the transporter may have to wait, due to current use of the elevator or due to a line of additional transporters that have already arrived and are currently waiting as well. Once the transporter enters the elevator (possibly with one or more other transporters, depending on elevator packing), they are all dispatched to the third floor, and from there to the corresponding OR. Transporters recycle back to the 12th floor on elevators.

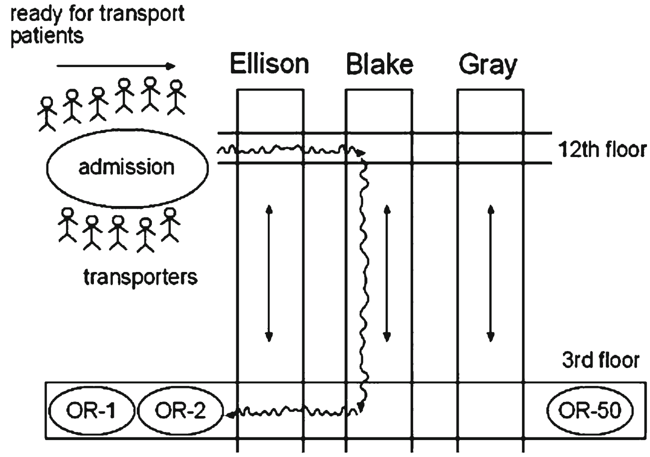


**Fig. 1** A schematic description of the proposed transportation system, including newly added elevator step. In the baseline condition, patients were transported from an adjacent holding area directly to one of 51 operating rooms. In the new system, patient admission and preoperative preparation would occur on a different floor, after which patients would be transported via elevator to the operating rooms. Elevators and transporters were potentially limited resources

## 3.2 Simulation approach

Rather than feeding the circumstantial environment of the proposed transport process into a high-level simulation software (such as Arena), we decided to create a JAVA-based transportation-oriented application from scratch, thereby gaining greater control and flexibility over all object characteristics, parameter setting, decision-making, and information extraction. To clarify the latter point, it is worth mentioning that the flexibility of general-purpose programming languages such as JAVA allows us to more accurately code the structural and logical properties of the studied process. Admittedly, we could have alternatively decided to make use of the SIMAN simulation language (in Arena), for instance, but the use of JAVA allowed us to develop a computer application that could be maintained and further developed by other researchers with general skills. We were especially concerned with making our implementation adaptable to additional transportation processes, putting a particular emphasis on providing a modular and extendable software, trying to anticipate patient flows of various natures, to incorporate environmental randomness, and to allow future users to integrate clinical and administrative preferences. The resulting simulation tool is easily accessible to non-experts, allowing essentially anyone to make interactive use of it for the workflow and spatial configuration described in this manuscript. Extending or changing the workflow and/or spatial configuration to be modeled would require basic JAVA skills. By this, we mean that the person has programming skills, and has basic working knowledge of JAVA. At present time, all source code is available online, along with extensive documentation, at: http://retsef.scripts.mit.edu/MGHTransport.zip

[1] This decision is reasonable because the first step in the process – the preadmission activities – is unchanged in the new configuration.

[2] Nonetheless, our analysis (to be discussed later) reveals interesting and fundamental relationships between the transportation component and the admission preoperative process. Specifically, it reveals delays that are inevitable unless one improves and expedites the admission process.

[3] To start 52 cases with 10 transporters, it is likely that all transporters are sometimes simultaneously occupied.

### 3.3 Data collection and interpretation

We eschewed evaluating the performance of various potential transport system configurations based on simulation with artificially created data, or based on standard probabilistic assumptions or on fitting some fixed distribution to already-observed data and sampling from that distribution. Instead, our approach is more data-driven, based on examining the patient transport process using real data collected from the current electronic OR information system, hospital Buildings and Grounds Department databases, transport services, and – on top of these sources – manually collected data from direct observation of travel times between locations.

Our data set is comprised of the entire population of patients who checked in for any surgical procedure at the surgical admission center during 2008, and comprises 17682 different cases. Each case record consists of:

1. Medical record number, age, date of service and time-stamps of process steps: These properties were kept confidential for the purpose of our study. We used these identifiers to assure that each case was unique and then assigned new, unique identifiers to each case, removing the medical record number from the primary data.
2. Patient Ready for Transport time: In principle, this time stamp should indicate the exact moment at which all preoperative procedures have been completed and the patient can be transported to the corresponding OR. However, repeated independent observations by our research assistants indicate that nurses mark a patient record as being ready about 5 min prior to the actual occurrence of this event due to workflow considerations (i.e., when it is convenient to log on to the computer for purposes of documentation). Consequently, we have shifted Patient Ready for Transport times forward accordingly.
3. Arrival Deadline: Theoretically, this time stamp should indicate the exact point in time by which the transport process is planned to be completed, meaning that the patient should arrive or already be present at their assigned OR. To estimate Arrival Deadlines, we obtained the Start Time[4] of each surgery, and set the Arrival Deadline by subtracting 30 min from the scheduled start. At our institution, the Start Time is the time the patient is expected to enter the OR. Our institution does not have a large holding area for preoperative patient preparation. Instead, each OR has a small patient holding area adjacent to it. This is where the patient is greeted by OR nursing and anesthesia personnel. Final documentation checks, day-of-surgery anesthesia consent, IV access and regional anesthesia are all performed in this holding area. The Arrival Deadline is the time a patient should be present in this holding area to begin these activities. The "correct" shifting from the Start Time to create the Arrival Deadline is very much case-dependent, and due to the obvious inability to electronically infer it in retrospect, 30 min was suggested by the medical staff as a conservative estimate. In the modeling software we developed, this parameter is a variable that can be set by the user. For first cases, the scheduled Start Time should equal the actual Start Time, since there are no prior cases in the OR to potentially delay subsequent starts.
4. Destination: OR to which the patient should be transported. This OR matches the one that was actually used in the dataset.
5. Transportation mode: In our current process, all patients are transported from the pre-surgical admission center to their corresponding ORs on a hospital stretcher. However, as elevator capacity may be a major limiting factor, we designed the software so that we could examine other means of transportation, such as wheelchairs.

We consulted the Buildings and Grounds Department to obtain accurate descriptive data about the elevators, such as dimensions, weight capacity, and speed. We used these data to decide how waiting patients can be packed into a single elevator ride (depending on transportation modes), and to estimate one-way and round-trip times for each elevator. These properties were also measured independently by research assistants who timed multiple transport/elevator trips by pushing stretchers along the proposed routes throughout the day. The research assistants also directly measured the extra time incurred by loading actual patients onto the elevators. We conducted measurements of actual traveling times using research assistants following transporters moving patients between different locations. Finally, we evaluated the following: Walking time with a stretcher from the location proposed for the temporary admission center (12th floor) to every prospective elevator (i.e., can be allocated as a dedicated one), and from each of these elevators to every OR. There is minimal variability in observed traveling times, so these were initially treated in our simulation as being deterministic. However, travel times are intuitively expected to vary. Therefore, we set up the simulation so that travel average times can be adjusted, and we used the software to examine the effects of deviations from the actual values observed.

[4] Here it is particularly important to mention that by "scheduled start time" we mean the actual point in time where the surgery took place (when looking at historical data). For instance, if a given surgery should have started at 11:00, but due to various delays actually started at 14:00, then the right number (i.e., 14:00) was used. This way, we simulate and test exactly what happens in reality, instead of what was *supposed* to happen.

### 3.4 Simulation construction

To create the simulation, we duplicated the patient transport environment, including all human and physical entities previously discussed in this section, as JAVA objects. The properties of each such object are dictated by those observed in real-life. For instance, patient properties consist of an unique identification number, destination, transportation method, Patient Ready for Transport time, Arrival Deadline, and various other time stamps obtained throughout their transport process; elevators have characteristic patient-dependent capacity, round-trip time, and loading time; corridors have their corresponding travel times; so on and so forth, noting that these lists are by no means exhaustive.

Our simulation was created to allow managers to assess the likely effect of interventions intended to minimize lateness in arrival to the OR. In other words, we seek to assure that as many patients as possible are at the OR by the Arrival Deadline. It remains to be explained how momentary patient lateness is estimated, noting that (as the name indicates) our computation is not an exact one, but rather an educated approximation. For this purpose, given a patient that currently resides in the admission center, we break the transportation time, going through a specific elevator, into five non-overlapping parts: (1) The time a ready-for-transport patient waits in the admission center to be picked up by a transporter; (2) Travel time from admission to the elevator; (3) waiting time prior to entering the elevator; (4) one-way elevator travel time; and (5) travel time from elevator to destination. As previously mentioned, each of these parameters (except the dynamic time periods, which are those of waiting for a transporter (item 1 above) and waiting for the elevator (item 3)) has been deterministically fixed in advance based on simple measurements of current performance. As a side note, it is worth mentioning that the way to evaluate item 3 will be explained later on, while item 1 will be easy to evaluate, since we will compute momentary lateness only when a patient could be transported right away (so there is no further waiting time for a transporter). Lateness is estimated as the degree to which the sum of the 5 intervals above, when added to the Patient Ready for Transport timestamp, exceeds the Arrival Deadline. In other words, $Lateness_{(minutes)}=\{Patient\ Ready\ for\ Transport\ Time+(sum\ of\ 5\ transport\ intervals)\}-Arrival\ Deadline$; negative values are set to zero.

Having characterized all objects in a precise way, the remaining issue to consider is how these entities interact with each other. We proceed by describing the type of decisions that need to be made, coupled with two solution approaches, one that requires very little coordination (henceforth called the simple algorithm), and one that involves more coordination (improved algorithm). The decisions are:

1. *Next patient to transport*: At any point in time, it is quite possible that multiple patients have already completed their preoperative preparation, and are ready to be transported to the appropriate ORs. The question arises: Who is the next patient to be transported? In the simple algorithm, we pick the patient whose preparation has been completed the earliest, i.e., with earliest Ready-for-Transport time. This is basically a first-in-first-out queue. In the improved algorithm, for each patient waiting, we estimate the individual lateness incurred, assuming that: (a) his transport starts right away, and (b) we pick an elevator based on the selection rule stated in item 3 below. We pick the patient with maximal estimated lateness as the next patient to transport, so as to minimize lateness.
2. *Transporter assignment*: Once the next patient identity has been determined, we move on to pick a transporter from the pool of currently available ones (i.e., not in the midst of an ongoing patient delivery or return trip). Due to lack of other policies that may be considered in practice, both algorithms make a completely random assignment, meaning that any available transporter is equally likely to be picked. (The underlying assumption is that all transporters are "identical", although this can be easily relaxed in the simulation to create a population of transporters with different speeds, for example.)
3. *Elevator dispatching*: As soon as a transporter is assigned to the next patient, the next immediate question is which route to pick, or equivalently, which elevator should be used? In the simple algorithm, we pick an elevator at random, where different elevators have identical chances to be chosen. In the improved algorithm, we pick the elevator that minimizes the estimated individual lateness, assuming that the current transport begins right away, going through the elevator in question.[5]
4. *Packing patients*: It might be the case that, at some point in time, multiple transporters have been dispatched to the same elevator, which is currently in use by previously arriving transporters. Once the elevator door "opens", i.e., arrives back to the admission floor, which transporters should be let in first, given a limited elevator capacity? Both algorithms use a first-in-first-out policy for the

[5] Elevator dispatching in the improved model: For this purpose, since all traveling times from one location to the other are known (i.e., estimated in advance), the only dynamic ingredient of this estimate is resolving the question: "how many transporters are currently waiting in line for some given elevator?". This could be easily answered by asking the last transporter that was dispatched to that elevator – via mobile phone, or video monitoring of the elevator lobbies.

queue waiting for the elevator. In other words, transporters and their patients are packed into the elevator one after the other, in order of arrival, until the residual elevator capacity is insufficient to pack the next transporter-patient in line. Practically speaking, this means two patients and two transporters, as patients are customarily transported on stretchers, and our elevators can accommodate two stretchers with transporters.

### 3.5 Performance measures – individual performance

Individual patients can be transported to their respective destinations as soon as all preoperative procedures have been completed, which is indicated by each patient's ready-for-transport time. However, due to transporter availability, elevator congestion, large traveling distances, delays due to problems with the preoperative process, and numerous other reasons, he may or may not be delivered on time (that is, prior to his OR Arrival Deadline). Consequently, *individual lateness* is defined as the difference (in minutes) between the actual arrival time and the corresponding deadline, when this quantity is positive (i.e., when the patient is late); otherwise, lateness is defined as 0. For instance, for a patient who arrives to the OR at 9:13, but whose Arrival Deadline is 9:00, lateness will be defined as 13 min; however, if the Arrival Deadline is 9:20, lateness is 0 min.

### 3.6 System-level performance

The system performance was evaluated on a daily basis, per each day of the week. For a given day, we are interested in several performance measures:

1. Average Lateness: Lateness taken over all cases in that day.
2. First Case Lateness: Similar to 1 above, but limited to first cases of the day (i.e., the first case in each of the ORs).
3. Serious Delays: The percentage of cases seriously delayed with respect to the Arrival Deadline among all the cases in the day. For this report, any arrival delay longer than 15 min was considered serious by the clinical leaders of the perioperative system. The definition of serious delay is a user definable variable in the simulation.
4. First Case Serious Delays: Similar to 3 above but limited to first cases.

From the overall operational perspective, identifying and developing systems that routinely function well is the most important objective. For our overall system, clinical leaders decided that minimizing serious delays (i.e., cases that missed the Arrival Deadline by more than 15 min) was the most important objective. Thus, we focused on serious delays by day of the week (both for first cases and for the overall day), and divided model-predicted performance into 'good days' and 'bad days'. A day is considered "bad" if the percentage of seriously-delayed cases exceeds 10%. We then focus on each day of the week and consider various statistics of the distribution of the average daily delay and percentage of seriously-delayed surgical operations. (For example, consider all Mondays). In addition, we consider the percentage of bad days predicted by a given set of model constraints.

As a validation step, we tested the simulation software on current operations. Specifically, we used current travel times, transporter numbers, arrival deadlines and definitions of lateness to simulate the current process. We used the simulation to assess the fraction of bad days for each day of the week, and compared these results to current actual performance using the same definitions. In all cases, the simulation was run at least 3 times.

## 4 Results

### 4.1 Simulation of current state

In the simulation of the current state, we assessed the fraction of bad days for each day of the week using the measured values for the model inputs. We compared the simulation results (using the improved algorithm) to actual results for first cases. These results are shown in Table 1, and demonstrate reasonable agreement between the observed and predicted fractions of lateness by day of week. OR managers agreed that the simulation provided sufficient precision to monitor the effects of process improvement efforts. In other words, process improvement schemes that would be judged successful would tend to yield changes in performance that were larger than the observed variations between simulation and actual performance in Table 1, and so the model was sufficiently sensitive.

### 4.2 Tested scenarios

The data-driven simulation tool that we developed was used to test several scenarios in light of the managerial and

**Table 1** Comparison of actual performance with simulations of performance of the current system

| WEEKDAY | Fraction of Days with ≥10% of Cases Delayed > 15 min (%) | |
|---|---|---|
| | ACTUAL | SIMULATION |
| Monday | 31.9 | 44.7 |
| Tuesday | 32.7 | 32.7 |
| Wednesday | 43.4 | 37.7 |
| Thursday | 24.0 | 14.0 |
| Friday | 41.2 | 43.1 |

operational challenges described above. In particular, we tried to identify how many elevators should be dedicated to the new location of the admission center, and how the system would perform under various scenarios of elevators, transporters, and transport time. We used the improved algorithm throughout.

As a benchmark, we first simulated an *ideal* scenario, in which one assumes that there are infinitely many elevators and transporters available. This is not a realistic or even a feasible scenario, but it provides a good benchmark to evaluate more realistic scenarios. Moreover, this scenario reveals "built-in" delays that are not due to the elevators or transporters, and in some sense provides a lower bound on delays (or, conversely, an upper bound on how well the system can work) when elevators and transporters are the only variables. These results are given in Table 2. Table 2 indicates that moving the admission center to a more distant location incurs a significant fraction of 'bad days', both for first cases and all cases, even with an infinite number of available elevators and transporters.

Table 2 follows the format of all of the subsequent tables in the manuscript, and so a detailed description of these tables is given below. For each day of the week, the tables specify the total number of simulated days (second column); the average, standard deviation and 80, 90, and 95 quantiles of the *average daily lateness* in minutes (columns 3-7); the average, standard deviation and quantiles of the *daily percentage of seriously delayed cases* (columns 7-11) and the *percentage of 'bad days' – the days in which the fraction of seriously delayed cases exceeds 10%.* For example, for the ideal scenario – First Cases, there were 47 Mondays that were simulated; the average (over all the Mondays) of the *average delay* was 3.8 min, the standard deviation was 2.7 min and the 80th, 90th, and 95th quantiles were 5.2, 6.5 and 8.6 min, respectively. On 80% of the Mondays the average delay was below 5.2 min. The *average percentage of seriously delayed cases* was 10.4% and the 80th, 90th, and 95th quantiles were 15.1, 20.2 and 27.9%, respectively. In particular, 80% of the Mondays had percentage of seriously delayed cases lower than 15.1%. Finally, 44.7% of the Mondays were 'bad', (i.e., the fraction of seriously delayed cases exceeded 10%).

If this is the best possible performance that can be obtained without changing something else in the system (see below), then what practical allocation of elevator and transporter resources comes close to this ideal scenario? To answer this question we carried out simulations for all days of the week. Of the 17,682 cases in the data set, 7,283 were first cases. These were evenly distributed across the weekdays, and there were 1,404 Monday first cases. In the results that follow, we focused – for ease of presentation – on first case Mondays (where the proportion of Mondays predicted to be 'bad' with respect to first case arrivals from Table 2 is 44.7%). Thus, for Monday first cases, we used the simulation tool to vary the number of transporters and elevators between 1 (unrealistically parsimonious) to infinity. The results are shown in Figs. 2 and 3.

Figures 2 and 3 demonstrate that seven transporters and two elevators nearly match the performance of infinite resources, with no benefit gained by increasing the number of transporters beyond the 10 already employed. Therefore, we fixed the number of transporters at 10 (the existing number) and tested scenarios in detail with limited but dedicated elevators. Specifically, we tested one or two elevators (Tables 3 and 4), simulating the elevator resources likely available in reality. We stopped increasing the number of elevators once we observed that the performance is very close to the ideal scenario (i.e., 2 elevators).

**Table 2** Ideal scenario with no limitation on number of elevators or transporters

| Day of Week | Number of Days | Mean Daily Lateness (min) | SD (min) | Quantiles (min) | | | Fraction Seriously Delayed (%) | SD (%) | Quantiles (%) | | | Fraction of Days with≥10% of Cases Seriously Delayed (%) |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | 0.80 | 0.90 | 0.95 | | | 0.8 | 0.9 | 0.95 | |
| First Cases | | | | | | | | | | | | |
| Monday | 47 | 3.84 | 2.69 | 5.23 | 6.46 | 8.57 | 10.4 | 9.7 | 15.1 | 20.2 | 27.9 | 44.7 |
| Tuesday | 52 | 3.04 | 2.10 | 3.89 | 5.16 | 6.50 | 7.7 | 6.4 | 12.5 | 15.1 | 18.0 | 28.9 |
| Wednesday | 53 | 3.11 | 1.75 | 4.88 | 5.40 | 5.93 | 8.3 | 6.1 | 14.4 | 17.6 | 18.4 | 37.7 |
| Thursday | 50 | 1.49 | 1.74 | 2.32 | 3.49 | 4.01 | 4.1 | 4.9 | 6.3 | 10.5 | 14.2 | 14.0 |
| Friday | 51 | 3.35 | 1.62 | 4.93 | 5.70 | 6.03 | 9.5 | 6.0 | 15.6 | 17.2 | 18.6 | 43.1 |
| All Cases | | | | | | | | | | | | |
| Monday | 47 | 2.78 | 1.50 | 3.70 | 4.08 | 5.01 | 6.9 | 5.0 | 11.1 | 12.4 | 15.6 | 25.5 |
| Tuesday | 52 | 2.24 | 1.17 | 3.31 | 3.71 | 4.15 | 5.5 | 3.3 | 8.2 | 10.1 | 10.7 | 13.5 |
| Wednesday | 53 | 2.33 | 1.08 | 3.10 | 3.92 | 4.33 | 5.8 | 3.0 | 8.7 | 9.9 | 10.9 | 9.4 |
| Thursday | 50 | 1.40 | 0.88 | 2.02 | 2.24 | 2.97 | 3.3 | 2.2 | 5.1 | 5.9 | 7.5 | 0.0 |
| Friday | 51 | 2.52 | 1.23 | 3.26 | 3.93 | 4.18 | 6.6 | 3.3 | 9.2 | 10.5 | 11.9 | 15.7 |

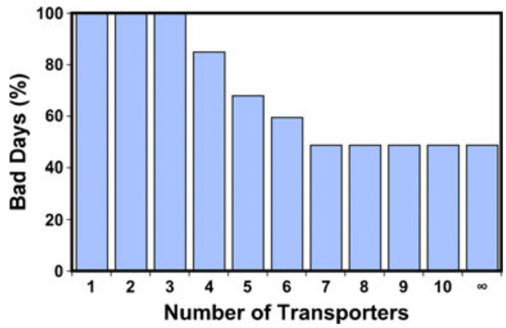


**Fig. 2** Modeled effect of changing the number of transporters (in the setting of infinite available elevators) on first case on time start performance. The bars represent the percentage of bad days for first case starts on Mondays. A 'bad day' is defined as one in which more than 10% of first cases start more than 15 min late

We then sought the impact of making changes elsewhere in the system to improve performance. Specifically, we started with ample (10) transporters and plateau (2) elevators, and modeled the impact of progressively rolling back the time that patients would be made ready for transport, shown in Fig. 4 using the percentage of cases with serious delays for Mondays. For modeling purposes, the admission center process (assumed to be a black box that functions independently of the rest of the hospital) was started progressively earlier. Performance as assessed by the Monday first-case bad day criterion improves as the start time and ready for transport time is rolled back, and approaches zero (2.1%) 'bad days' if the process is started 20 min earlier. We then tested the performance of the system under the assumption that the ready-to-transport times of all patients are shifted 20 min earlier over all days (Table 5 below); in this scenario we assumed the use of 2 dedicated elevators and 10 transporters.

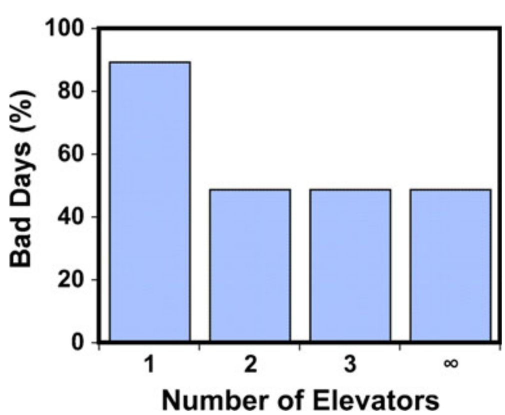


**Fig. 3** Modeled effect of changing the number of elevators (in the setting of infinite available transporters) on first case on time start performance. The bars represent the percentage of bad days for first case starts on Mondays. A 'bad day' is defined as one in which more than 10% of first cases start more than 15 min late

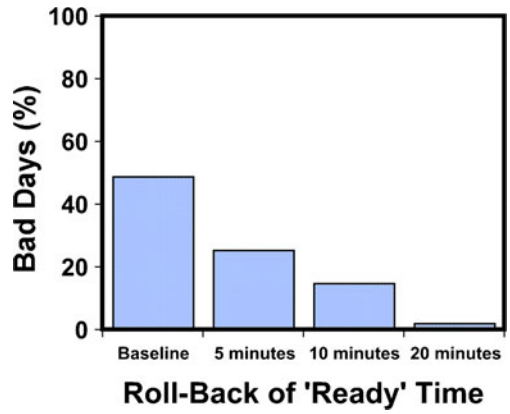


**Fig. 4** Modeled effect of changing the start time for the pre-operative process (i.e., beginning patient preparation in the admission center earlier) on first case on time start performance. Modeling is conducted with two elevators and 10 transporters. The bars represent the percentage of bad days for first case starts on Mondays. A 'bad day' is defined as one in which more than 10% of first cases start more than 15 min late

If rolling the admission center start time back by 20 min (using 2 dedicated elevators and 10 transporters) could virtually eliminate bad days, then it might be possible to revisit the question of whether two dedicated elevators are required. Thus, we performed a nested simulation analysis to see whether the transport system could function with only one dedicated elevator, or how well it could function with one or two elevators and decreasing numbers of transporters. Figure 5 shows the results of these simulations, again focused on first cases, Monday mornings. With only one elevator (Fig. 5a), all days are bad with only one or two transporters. Addition of a fourth transporter yields a large performance improvement but the fraction of bad days only drops to 6% when 10 transporters are used. Thus, we conclude that one dedicated elevator would be sufficient if 6% of days being bad is acceptable. Referring back to Table 1 indicates that under current conditions, 24% to 41% of days are bad, so perhaps a single dedicated elevator, an early start, and ample transporters is a reasonable tradeoff. On the other hand, with 2 dedicated elevators and an early start, the process could be run with as few as four transporters, though optimum performance requires at least six. In practice, given the low cost of transport personnel, maintaining a buffer would be cost effective.

Finally, we tested the robustness of the system to variability and inaccurate estimations of the travel times. In other words, we were interested in the impact of slower than expected transport times. In Fig. 6, we progressively increased the average travel time, again using the Monday first-case bad day proportion as the test case. We then assumed that the actual travel times are 50% higher than the original estimations for all days and cases (Table 6 below); again we assumed 2

**Table 3** One elevator and 10 transporters

| Day of Week | Number of Days | Mean Daily Lateness (min) | SD (min) | Quantiles (min) | | | Fraction Seriously Delayed (%) | SD (%) | Quantiles (%) | | | Fraction of Days with≥10% of Cases Seriously Delayed (%) |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | 0.80 | 0.90 | 0.95 | | | 0.8 | 0.9 | 0.95 | |
| First Cases | | | | | | | | | | | | |
| Monday | 47 | 6.68 | 3.25 | 9.06 | 10.50 | 11.22 | 19.5 | 9.2 | 26.4 | 29.2 | 30.7 | 89.4 |
| Tuesday | 52 | 6.39 | 2.82 | 8.08 | 9.00 | 11.22 | 19.4 | 8.5 | 25.0 | 29.7 | 32.8 | 88.5 |
| Wednesday | 53 | 6.58 | 2.75 | 9.04 | 10.14 | 10.37 | 19.6 | 8.7 | 27.7 | 29.7 | 31.3 | 81.1 |
| Thursday | 50 | 1.58 | 1.79 | 2.60 | 3.49 | 4.01 | 4.5 | 5.1 | 9.5 | 10.6 | 14.8 | 18.0 |
| Friday | 51 | 6.31 | 3.00 | 8.59 | 9.87 | 11.23 | 4.5 | 9.1 | 25.0 | 30.6 | 32.4 | 82.4 |
| All Cases | | | | | | | | | | | | |
| Monday | 47 | 4.29 | 1.85 | 5.58 | 6.18 | 7.20 | 11.7 | 4.9 | 15.5 | 17.2 | 17.9 | 66.0 |
| Tuesday | 52 | 3.96 | 1.56 | 5.41 | 5.69 | 6.30 | 11.3 | 4.4 | 14.5 | 15.6 | 17.6 | 63.5 |
| Wednesday | 53 | 4.25 | 1.70 | 5.29 | 5.61 | 7.13 | 11.8 | 5.0 | 15.6 | 17.2 | 18.6 | 64.2 |
| Thursday | 50 | 1.79 | 1.02 | 2.47 | 3.25 | 3.38 | 4.4 | 2.8 | 6.6 | 7.9 | 9.0 | 2.0 |
| Friday | 51 | 4.07 | 1.60 | 5.27 | 6.24 | 6.81 | 11.3 | 4.2 | 14.5 | 17.1 | 19.5 | 58.8 |

dedicated elevators and 10 transporters. Inspection of the figure and of Table 6 (relative to Table 4) indicates a modest but potentially meaningful decline in on-time performance if the travel by the transporters was increased by 50%.

## 5 Discussion

### 5.1 Solution approach

The main methodology used in this paper is based on data-driven simulation modeling. Our primary objective was to develop a comprehensive data-driven simulation tool in order to illuminate the relationship between key system parameters, such as number of dedicated elevators and staff, transportation means (beds, stretchers, wheelchairs, etc.), and coordination methods on the one hand, and the magnitude of patient lateness and significantly delayed surgical operations on the other. The resulting simulation tool differs from standard discrete event Monte Carlo simulations, in that it originates from the actual performance data of the hospital. We have examined and evaluated various system configurations and policies simulated with real data (obtained from clinical and administrative information systems at MGH), as well as with manually collected data to model added process steps. This feature allowed us to accurately simulate the system being studied. However, the model software design allows it to be tailored to any proposed change in patient transportation system in any setting

**Table 4** Two elevators and 10 transporters

| Day of Week | Number of Days | Mean Daily Lateness (min) | SD (min) | Quantiles (min) | | | Fraction Seriously Delayed (%) | SD (%) | Quantiles (%) | | | Fraction of Days with≥10% of Cases Seriously Delayed (%) |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | 0.80 | 0.90 | 0.95 | | | 0.8 | 0.9 | 0.95 | |
| First Cases | | | | | | | | | | | | |
| Monday | 47 | 4.35 | 2.94 | 5.87 | 7.08 | 9.21 | 11.5 | 9.6 | 16.6 | 22.9 | 27.9 | 48.9 |
| Tuesday | 52 | 3.51 | 2.19 | 4.61 | 5.58 | 7.16 | 9.3 | 6.8 | 14.0 | 18.1 | 20.2 | 42.3 |
| Wednesday | 53 | 3.59 | 1.98 | 5.62 | 6.18 | 6.70 | 9.7 | 6.3 | 15.5 | 18.1 | 19.8 | 39.6 |
| Thursday | 50 | 1.55 | 1.77 | 2.41 | 3.49 | 4.01 | 4.5 | 5.1 | 9.5 | 10.6 | 14.8 | 18.0 |
| Friday | 51 | 3.82 | 1.81 | 5.31 | 6.41 | 6.58 | 11.1 | 6.1 | 16.2 | 18.8 | 20.9 | 56.9 |
| All Cases | | | | | | | | | | | | |
| Monday | 47 | 3.04 | 1.62 | 4.00 | 4.42 | 5.41 | 7.4 | 5.1 | 11.2 | 12.7 | 16.1 | 27.7 |
| Tuesday | 52 | 2.49 | 1.23 | 3.75 | 4.13 | 4.48 | 6.4 | 3.4 | 8.4 | 11.3 | 12.4 | 17.3 |
| Wednesday | 53 | 2.59 | 1.17 | 3.58 | 4.27 | 4.62 | 6.5 | 3.3 | 9.4 | 11.0 | 12.7 | 17.0 |
| Thursday | 50 | 1.48 | 0.90 | 2.14 | 2.44 | 3.07 | 3.5 | 2.3 | 5.2 | 6.5 | 8.0 | 0.0 |
| Friday | 51 | 2.77 | 1.29 | 3.75 | 4.11 | 4.43 | 7.4 | 3.2 | 10.1 | 11.7 | 12.1 | 23.5 |

(i.e., any hospital). In our specific example, this flexibility allowed us to "see" how the system would have worked today if the new setting of distant admission center would be implemented. In addition, it enables us to answer "what will happen if…" questions. Specifically, we can now address the following issues:

1. Understand the interactions between resource availability, coordination mechanisms (policies), and different means of patient transportation.
2. Search the space of practical policies and system configurations, and to sensibly compare between them from clinical, physical, and financial perspectives.
3. Examine how potential changes in the environment, and in particular unpredictable variability, will affect the system dynamics.

The main findings of this simulation study can now be interpreted in terms of basic questions raised in the introduction.

*Question 1. Configuration of elevators and transporters, including coordination* There is a significant difference in the performance of the system comparing the use of one versus two dedicated elevators (see Tables 3, 4, 5 and 6). When only one elevator is allocated, the percentage of 'bad days', either throughout the day or for first cases is very high. For most days of the week, 60% of days are bad, and more than 80% of the mornings are expected to be bad. However, when a second elevator is allocated the percentage of bad days decreases by at least 50%.

On the other hand, the scenario of 2 elevators seems to perform almost as well as the 'best case' scenario (unlimited elevators) with maximum differences of 2-7% in the percentage of bad days comparing 2 elevators to infinite elevators. Consequently, there is no real reason to increase the number of elevators beyond two. Similar analysis yields minimal benefit to increasing the number of transporters beyond 7. Because there is little further performance improvement with two elevators, there is little motivation to use other means to carry patients (e.g., wheel chairs instead of stretchers). Specifically, 4 patients in wheel chairs could be packed into a single elevator, and two such elevators could carry 8 patients. However, this is equivalent to 4 elevators with two patients each on stretchers, and 4 such elevators yields no better performance than 2 elevators with 2 patients each. These results are consistent with the fundamental phenomenon of decreasing marginal returns of additional capacity in systems with capacity constraints elsewhere in the system. [24] The other insight is that adding capacity to a component of the system (elevators and transporters in this example) might not be as effective in improving the overall throughput or wait times if there exist other bottlenecks.

*Question 2. Feasibility of the current plan* The results of the simulation study suggest that the transportation aspects are not likely to become a bottleneck provided that two dedicated elevators will be used. That said, the distances between the new location of the admission center and the ORs will lead to relatively high percentage of seriously delayed patient arrivals at the ORs, particularly during the mornings. These delays cannot be eliminated by adding additional transportation resources. Even in the ideal scenario with infinitely many elevators and transporters these delays persist, so the current plan of simply moving the admission center and adding transport resources is infeasible. Thus, the initial hypothesis shared by decision makers (that transportation resources would be the limiting factor) turns out to be incorrect as judged by the simulation results. Instead, the bottleneck is the ability to make

**Table 5** Start 20 min earlier, with 2 elevators and 10 transporters

| Day of Week | Number of Days | Mean Daily Lateness (min) | SD (min) | Quantiles (min) | | | Fraction Seriously Delayed (%) | SD (%) | Quantiles (%) | | | Fraction of Days with≥10% of Cases Seriously Delayed (%) |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | 0.80 | 0.90 | 0.95 | | | 0.8 | 0.9 | 0.95 | |
| First Cases | | | | | | | | | | | | |
| Monday | 47 | 1.05 | 1.25 | 1.69 | 2.24 | 2.80 | 2.8 | 3.9 | 5.4 | 6.2 | 6.8 | 2.1 |
| Tuesday | 52 | 0.70 | 1.04 | 1.08 | 1.48 | 2.78 | 1.3 | 2.2 | 3.0 | 3.8 | 6.1 | 0.0 |
| Wednesday | 53 | 0.90 | 0.96 | 1.52 | 2.34 | 2.82 | 2.0 | 2.3 | 3.7 | 5.4 | 6.7 | 0.0 |
| Thursday | 50 | 0.34 | 0.93 | 0.41 | 1.05 | 1.55 | 0.7 | 1.6 | 0.0 | 3.8 | 4.7 | 0.0 |
| Friday | 51 | 0.81 | 0.89 | 1.36 | 1.79 | 2.48 | 1.8 | 2.7 | 3.7 | 5.3 | 6.3 | 2.0 |
| All Cases | | | | | | | | | | | | |
| Monday | 47 | 1.02 | 0.78 | 1.60 | 1.71 | 2.28 | 2.4 | 2.2 | 3.9 | 4.7 | 5.6 | 2.1 |
| Tuesday | 52 | 0.71 | 0.67 | 1.31 | 1.52 | 1.82 | 1.5 | 1.5 | 3.1 | 3.6 | 3.8 | 0.0 |
| Wednesday | 53 | 0.84 | 0.67 | 1.39 | 1.64 | 2.26 | 1.9 | 1.2 | 2.9 | 3.7 | 4.3 | 0.0 |
| Thursday | 50 | 0.56 | 0.53 | 0.95 | 1.09 | 1.39 | 1.3 | 1.2 | 2.5 | 2.8 | 3.3 | 0.0 |
| Friday | 51 | 0.87 | 0.81 | 1.39 | 1.75 | 2.17 | 2.0 | 2.2 | 2.9 | 3.2 | 5.8 | 2.0 |

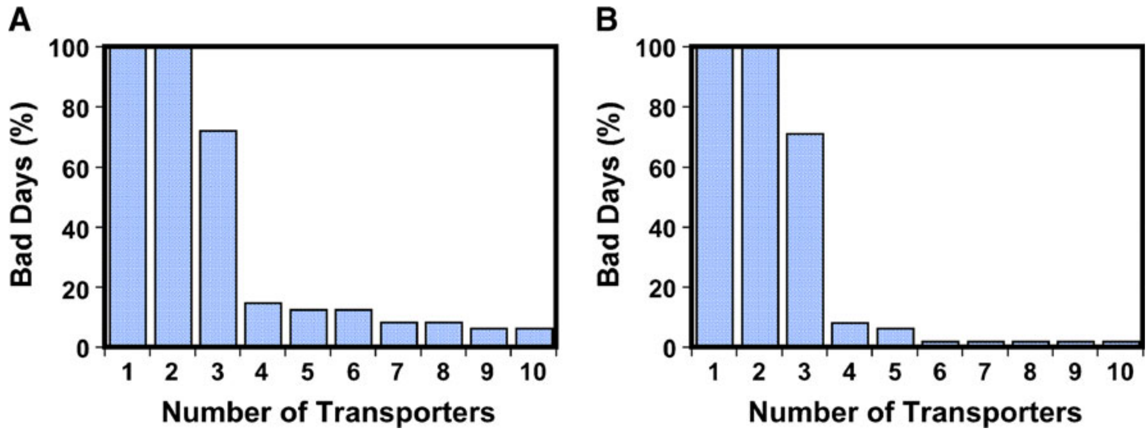


**Fig. 5** Modeled effect of starting the pre-operative process 20 min earlier (i.e., beginning patient preparation in the admission center earlier) and then systematically varying the number of elevators. Panel **a**: One elevator and varying the number of transporters from one to ten. Panel **b**: Two elevators and varying the number of transporters from one to ten. The bars represent the percentage of bad days for first case starts on Mondays. A 'bad day' is defined as one in which more than 10% of first cases start more than 15 min late

patients ready soon enough to be transported to the ORs. In this light, an initial plan to make the temporary admission center slightly smaller than the original is likely to exacerbate delays. However, Fig. 4 and Table 5 suggest that if all ready-for-transport times are made 20 min earlier, almost all the delays are eliminated. Moreover, when the start of the process is rolled back by 20 minutes for all patients, the system appears to require relatively few transporters and two elevators is ample (Fig. 5). However, moving all of the ready-for-transport times back by 20 min is not necessarily equivalent to scheduling all patients to arrive at the admission center 20 min earlier than today. This is because other linked processes in the hospital that are required to complete the patient readiness process may not necessarily have been pushed back, and so the patient may not become ready for transport earlier, despite arriving to the hospital earlier. For example, suppose the patients all arrive 20 min earlier, but the surgeons continue to arrive at their usual time to perform required site marking. In such a case, the effect of the patient's earlier arrival is negated by the interaction with other parts of the system whose function has not changed (and is potentially difficult to change). Our work sets up concrete goals for performance improvement in the admission process, and reveals the fundamental interactions between the admission aspects and the transportation aspects of the process.

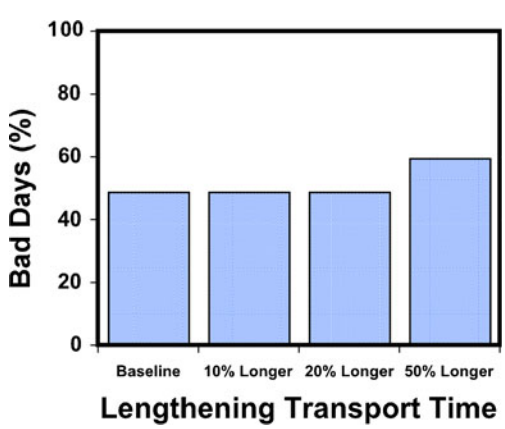


**Fig. 6** Modeled effect of changing the time each transporter takes to complete their circuit on first case on time start performance. Modeling is conducted with two elevators and 10 transporters. The bars represent the percentage of bad days for first case starts on Mondays. A 'bad day' is defined as one in which more than 10% of first cases start more than 15 min late

*Question 3. Expected effects of potential variability in environment parameters* Up until now, it may appear as if our conclusions could become invalid, should reasonable (or even predictable) deviations from the present system characteristics cause the overall configuration to behave badly. For example, a slow transporter might slow down the entire process. However, on the one hand, based on repeated measurements taken by research assistants, actual travel times exhibit small variability, and the use of dedicated elevators suggests that their related travel time can be predicted to a large degree of accuracy. This supports our modeling assumption that the traveling times are deterministic. On the other hand, to test the robustness of our tool to variability in transporter performance, we performed additional tests to study the impact of changes and fluctuations in the travel times compared to the initial estimates (see Fig. 6 and Table 6). The system performance is only weakly sensitive to these changes.

There are several underlying assumptions in our study. In particular, the following assumptions were used in the design of the model:

- During the transition, the volume of surgical operations will be similar to the volume in 2008. The data of our department indicate that the volume of surgical operations in the last couple of years is stable, and expected to stay stable in the next couple of years. That said, the tool that we developed can be used to test other assumptions.

**Table 6** Increase travel time by 50%, original start time, with 2 elevators and 10 transporters

| Day of Week | Number of Days | Mean Daily Lateness (min) | SD (min) | Quantiles (min) | | | Fraction Seriously Delayed (%) | SD (%) | Quantiles (%) | | | Fraction of Days with≥10% of Cases Seriously Delayed (%) |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | 0.80 | 0.90 | 0.95 | | | 0.8 | 0.9 | 0.95 | |
| First Cases | | | | | | | | | | | | |
| Monday | 47 | 4.88 | 3.14 | 6.53 | 7.93 | 10.00 | 13.1 | 9.5 | 19.9 | 22.3 | 27.9 | 59.6 |
| Tuesday | 52 | 4.08 | 2.33 | 5.18 | 6.39 | 7.98 | 11.3 | 7.0 | 15.6 | 17.9 | 23.3 | 55.8 |
| Wednesday | 53 | 4.13 | 2.09 | 6.13 | 6.99 | 7.22 | 11.0 | 7.0 | 16.5 | 20.0 | 22.4 | 49.1 |
| Thursday | 50 | 1.75 | 1.86 | 2.89 | 3.85 | 4.41 | 4.9 | 5.2 | 10.1 | 11.2 | 14.8 | 22.0 |
| Friday | 51 | 4.38 | 1.97 | 5.88 | 6.97 | 7.38 | 13.0 | 6.4 | 17.6 | 20.6 | 24.1 | 68.6 |
| All Cases | | | | | | | | | | | | |
| Monday | 47 | 3.35 | 1.72 | 4.44 | 4.80 | 5.94 | 8.5 | 5.2 | 12.1 | 13.7 | 15.6 | 38.3 |
| Tuesday | 52 | 2.84 | 1.32 | 4.24 | 4.72 | 4.88 | 7.4 | 3.6 | 10.1 | 11.6 | 15.0 | 25.0 |
| Wednesday | 53 | 2.95 | 1.31 | 3.90 | 4.77 | 5.08 | 7.6 | 4.3 | 11.7 | 13.7 | 14.4 | 22.6 |
| Thursday | 50 | 1.63 | 0.95 | 2.33 | 2.77 | 3.29 | 4.0 | 2.4 | 5.7 | 7.3 | 8.0 | 2.0 |
| Friday | 51 | 3.1 | 1.4 | 4.13 | 4.48 | 4.88 | 8.4 | 3.5 | 11.3 | 13.2 | 14.0 | 35.3 |

- The admission process and its capacity will stay the same as in the current location. Our experiments indicate that even with the same admission capacity, one should expect frequent serious delays in the start of time of surgical operations, especially at the beginning of the day. In fact, one of the important conclusions from simulation is that the admission process in the new location should be changed to produce patients ready for transport earlier than the current process, so as to compensate for the longer distance between the admission center and the ORs (see the discussion below).
- The admission center will be served by a few dedicated elevators that will solely serve the admission center, and dedicated transporters. We ignore the possibility to use, in parallel, the other elevators that will serve the functions within the 3 buildings, Ellison, Blake and Gray. The justification for that decision is that these elevators serve staff and patients in many other units in these building, so it is not practical or even medically safe to use them as a primary mean of transportation.
- Patients are transported on stretchers as they are today. Other assumptions – such as patients in wheel chairs - can be easily tested. Our analysis of current practices and future plans indicates there is no real motivation to change the way patients are transported.

## 5.2 Limitations of our approach

Despite the advantages of simulation tools, there are limitations to consider. In what follows, we attempt to highlight the main drawbacks of our methodologies, in addition to the assumptions above so that potential users would be able to apply or adapt this approach in a sensible way.

- *Data-driven simulation.* As previously mentioned, the patient-related input to our simulation model corresponds to the actual data that was observed and recorded, rather than being based on probabilistic assumptions or on distribution fitting, which is referred to in the literature as data-driven (or trace-driven) simulation. It is worth pointing out that this approach has its advantages, as well as its downsides [25]. On the one hand, trace-driven simulations are credible (i.e., easier to "market" to medical end-users than random inputs), preserve the correlation of events, and their input is deterministic (so there is less overall randomness). However, on the other hand, at least in the context of our model, such simulations require rather complex implementations, could be limited due to space considerations (when the input sequence is extremely long), and may be difficult to modify (for instance, increasing or decreasing the rate of arrivals).
- *Intensive computations*. What may be viewed as the main downside of our specific implementation is the rather massive computations required for processing data sets that accumulate over long period of times. These lengthy delays can be attributed to the large number of objects being simulated, as well as to the complicated interaction between them. As a consequence, in order that our original implementation would terminate within a reasonable amount of time, executing it on specialized hardware was a necessity. However, in the latest version we developed, this problem seems to be no longer a concern; in particular, realistic arrival sequences (consisting of thousands of patients) can be processed on inexpensive desktop computers within a time frame of several minutes.
- *Simplifying assumptions*. Clearly, the most frequent way of coping with the inherent complexity of real-life systems is that of introducing simplifying assumptions. This, in turn,

may create a false sense of security regarding the results obtained and their accuracy. It is worth noting that even the detailed system description in the materials and methods section cannot capture each and every influencing factor; some ingredients that have a concrete representation in reality could not have been implemented, mainly for sake of simplicity. Having that said, we believe that, as previously demonstrated in the sensitivity analysis, our model is robust enough, in the sense that small deviations have nothing but negligible effects on the performance guarantees.

- *No animation*. Another disadvantage of our simulation is the lack of corresponding animation, as the main focus has been on "crunching numbers". The ability to actually see the inner-workings of a given system configuration, albeit this happening on a computer screen, could potentially be a very useful media of presentation for a non-scientific audience, and serve as an additional mean of verification. We note that it would be relatively straightforward, albeit time consuming, to add these capabilities to the tool that we have developed.

### 5.3 Comparison and improvements over other prospective approaches

On top of the specific questions that were motivated by the anticipated patient transport setting at MGH, what may very well be the most important outcome of this project is a unique simulation tool for studying transportation systems within large academic hospitals. The JAVA code for the modeling program has been made freely available at the link given above in the manuscript, with sufficient documentation to allow it to be used by others. An important feature of our simulation is the ability to examine and evaluate an assortment of system configurations, not only on artificially-created data (such as that obtained by fitting a separate distribution for components like patient arrival sequences, destinations, deadlines, etc.), but also on real-life data collected from clinical and administrative information systems. Consequently, one is able to capture intricate dependencies and complexities that may be well-hidden within the actual process being considered, rather than implicitly ignoring these issues by gluing together components such as arrivals, destinations, and deadlines from independent distributions. In particular, we believe that the tool that has been developed in this work will be used in several other settings within the hospital, for example, the analysis of the transportation of surgical equipment. Furthermore, there is the previously mentioned new OR building under construction. After the building is functional, the ORs will be split into two distant locations (the old and the new buildings.) This will raise a new set of fundamental planning issues regarding the transportation of patients and equipment, which we believe could be addressed effectively, using the tool developed in this work.

**Financial Support** The work of the second author is supported in part by National Science Foundation grants DMS-0732175 and CMMI-0846554 (CAREER Award), an Air Force Office of Scientific Research (AFOSR) award FA9550-08-1-0369, a Singapore-MIT Alliance (SMA) grant and the Buschbaum Research Fund of Massachusetts Institute of Technology.

**Conflicts of interest** No conflicts of interest.

## References

1. Meyer MA, Seim AR, Fairbrother P, Egan MT, Sandberg WS (2008) Automatic time-motion study of a multistep preoperative process. Anesthesiology 108(6):1109–1116
2. Jun JB, Jacobson JB, Swisher JR (1999) Application of Discrete Event Simulation in Health Care Clinics: A Survey. J Oper Res Soc 50:109–123
3. Klein RW, Dittus RS, Roberts SD, Wilson JR (1993) Simulation modeling and health-care decision making. Med Decis Making 13 (4):347–354
4. Mahachek AR (1992) An introduction to patient flow simulation for health-care managers. J Soc Health Syst 3(3):73–81
5. Greene LV (2004) Capacity Planning & Management in Hospitals. In: Brandeau ML, Sainfort F, Pierskalla WP (eds) Operations Research in Healthcare: A Handbook of Methods and Applications. International Series in Operations Research & Management Science. Kluwer Academic Publishers, Boston, pp 15–42
6. Kachnal SK (2001) Industrial Engineering Applications in Health Care Systems. In: Salvendy G (ed) Handbook of Industrial Engineering: Technology and Operations Management, 3rd edn. John Wiley & Sons, New York, pp 737–750
7. Nickel S, Schmidt UA (2009) Process improvement in hospitals: a case study in a radiology department. Qual Manag Health Care 18 (4):326–338. doi:10.1097/QMH.0b013e3181bee127
8. Odegaard F, Chen L, Quee R, Puterman ML (2007) Improving the efficiency of hospital porter services, part 2: schedule optimization and simulation model. J Healthc Qual 29(1):12–18
9. Odegaard F, Chen L, Quee R, Puterman ML (2007) Improving the efficiency of hospital porter services, part 1: study objectives and results. J Healthc Qual 29(1):4–11
10. Dershin H, Schaik MS (1993) Quality improvement for a hospital patient transportation system. Hosp Health Serv Adm 38(1):111–119
11. Bryan W (1998) Rising to the challenge: portering services at the Queen Elizabeth II Health Sciences Centre. Int J Health Care Qual Assur Inc Leadersh Health Serv 11(4-5):i–v
12. McAleer WE, Turner JA, Lismore D, Naqvi IA (1995) Simulation of a hospital's theatre suite. J Manag Med 9(5):14–26
13. Zonderland ME, Boer F, Boucherie RJ, de Roode A, van Kleef JW (2009) Redesign of a university hospital preanesthesia evaluation clinic using a queuing theory approach. Anesth Analg 109 (5):1612–1621
14. van Oostrum JM, Van Houdenhoven M, Vrielink MM, Klein J, Hans EW, Klimek M, Wullink G, Steyerberg EW, Kazemier G (2008) A simulation model for determining the optimal size of emergency teams on call in the operating room at night. Anesth Analg 107(5):1655–1662. doi:10.1213/ane.0b013e318184e919
15. Marcon E, Kharraja S, Smolski N, Luquet B, Viale JP (2003) Determining the number of beds in the postanesthesia care unit: a computer simulation flow approach. Anesth Analg 96(5):1415–1423
16. Schoenmeyr T, Dunn PF, Gamarnik D, Levi R, Berger DL, Daily BJ, Levine WC, Sandberg WS (2009) A model for understanding

the impacts of demand and capacity on waiting time to enter a congested recovery room. Anesthesiology 110(6):1293–1304
17. Sokal SM, Craft DL, Chang Y, Sandberg WS, Berger DL (2006) Maximizing operating room and recovery room capacity in an era of constrained resources. Arch Surg 141(4):389–393, discussion 393-385
18. McManus ML, Long MC, Cooper A, Litvak E (2004) Queuing theory accurately models the need for critical care resources. Anesthesiology 100(5):1271–1276
19. Naesens K, Gelders L (2009) Reorganizing a Service Department: Central Patient Transportation. Production Planning & Control 20 (6):478–483
20. Hanne T, Melo T, Nickel S (2009) Bringing Robustness to Patient Flow Management Through Optimized Patient Transports in Hospitals. Interfaces 39(3):241–255
21. Davies R, Davies H (1994) Modelling Patient Flows and Resource Provision in Health Systems. Omega 22(2):123–131
22. Stafford EF Proceedings of the 1978 Summer Computer Simulation Conference. In, 1978. pp 153-159
23. Donham RT (1998) Defining measurable OR-PR scheduling, efficiency, and utilization data elements: the Association of Anesthesia Clinical Directors procedural times glossary. Int Anesthesiol Clin 36(1):15–29
24. Cachon G, Terwiesch C (2009) Batching and Other Flow Interruptions: Setup TImes and the Economic Order Quality Model. In: Matching Supply with Demand - An Introduction to Operations Management, 2nd edn. McGraw-Hill, Irwin, pp 118–121
25. Law AM, Kelton WD (2000) Building Valid, Credible, and appropriately Detailed Simulation Models. In: Simulation Modeling and Analysis, 3 rdth edn. McGraw Hill, Boston, pp 264–291