Proceedings of the 2017 Winter Simulation Conference
W. K. V. Chan, A. D'Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page, eds.

DATA-DRIVEN GENERIC DISCRETE EVENT SIMULATION MODEL OF HOSPITAL PATIENT FLOW CONSIDERING SURGE

Carolyn R Busby
Michael W Carter

University of Toronto
5 King's College Road
Toronto, ON M5S 3G8, CANADA

ABSTRACT

Many Canadian hospitals run at or near capacity, frequently experiencing congestion due to surges in demand. Hospital-wide bed management strategies, called "surge protocols", that formally define when and what kind of operational steps can be taken to alleviate congestion are routinely in use. Decisions across the hospital, regarding bed capacity and allocation, staffing levels, and the master surgical schedule influence the frequency and severity of congestion, which in turn manifests in high bed occupancy, delayed admissions, a crowded emergency department (ED), surgical cancellations and increased use of surge protocols. A generic, data-driven, discrete event simulation (DES) is developed to help hospitals assess the impact of hospital-wide decisions, including surge policies, on congestion. The model has been developed in cooperation with two hospitals, and has been validated at two additional hospitals.

1 INTRODUCTION

Canadian hospitals often run at or near capacity, sometimes causing ED crowding, cancelled surgeries, delays in admitting patients, and added effort to find solutions to the congestion. To ease and coordinate this effort, most hospitals have formal surge policy procedures in place. The surge protocols specify how hospital behavior changes when in a surge situation. The definition of "being in surge", the degree or "level" of surge, the responses, and the hospital units involved, is defined by each individual hospital. Hospitals make use of surge policies to adapt when full, avoiding cost and wasted resources associated with maintaining bed capacities for peak periods full time. However, despite the common use of surge procedures, and the corresponding change in hospital-wide structures and behaviors to respond to the surge, models fully incorporating these protocols are not found in the literature.

There are two types of surge protocols: those for catastrophic events such as natural disasters or terrorist attacks and for normal demand variability called "daily surge" (Jenkins, O'Connor, and Cone 2006). This model focuses on the latter. Daily surge responses include: a) increased effort to discharge eligible patients; b) opening extra beds; c) delaying transfers from other institutions, d) adding pressure to downstream providers such as home care and long-term care to allow discharge of non-acute patients; and e) moving patients to lower level of care within the hospital, where possible, to make room for more acute patients.

This work also builds upon past success of a widely implemented generic perioperative model for surgical patient flow (Sniekers et al. 2017). A generic DES of the perioperative system was constructed, using Simul8 software, that is transportable from hospital to hospital without altering the simulation code. All customization is done by populating an Excel input file with hospital specific data. Upon importing the file to the simulation model, the structure and rules of the model adjusts based on the inputs. Although many papers claim that models are generic and can be easily used at other hospitals, they rarely design the model explicitly for this purpose or test its transportability. The perioperative model did both. Currently the model has been implemented at 12 hospitals in two Canadian provinces. Some examples of uses at the various hospitals include: reallocation of surgery beds; testing proposed new block schedules; benchmarking operating room (OR) efficiency; meeting new surgical volume goals; planning accommodation of urgent surgeries; and smoothing use of surgical beds.

While implementing the above model, medical patient use of surgical beds was often a concern leading to possible surgical cancellations and delays. At the same time, hospitals were focused on reducing overcrowding in the ED under "Pay for Results" or "P4R" funding. P4R is part of the Ontario government's Wait Time Strategy, that offers funding for ED wait time improvement and additional variable funding for meeting wait time goals. Therefore, capturing these interactions across the hospital, and understanding their effect on congestion, required the incorporation of the ED, complete inpatient units, and surge protocols. The current model is built with the same generic data-driven structure and reuses some of the code from the perioperative model for the surgery scheduling and operating room flow.

Continuing to build the simulation in Simul8 as a flexible, generic, data-driven model, designed for reuse, makes sense given: a) the ability to reuse some of the existing code; b) the proven success of applying the perioperative model across various hospitals; c) the size and complexity of the model, making it expensive to build a model from scratch each time; d) interest from multiple hospitals to use the new model.

Elements that distinguish this work from previously published literature include:

1. Inclusion of full surge protocols for more accurate representation of congested hospitals
2. Inclusion of ED, inpatient units and ORs to capture interactions between each area of the hospital.
3. Designed and tested for reuse across several hospitals without altering simulation code

2 DISCUSSION AND LITERATURE REVIEW

Simulation has been used extensively in health care, and there are many reviews of this work. Recent reviews include (Mielczarek 2016; Katsaliaki and Mustafee 2011; Gunal and Pidd 2010). The use of generic models within hospitals is reviewed in Fletcher and Worthington (2009). They define various levels of genericity and review papers that fall into the category of "generic" and "specific" including seven papers focused on modelling multi-departmental and whole system hospital flows. The review by Gunal and Pidd (2010) focuses specifically on DES in hospitals but also offers a useful overview of healthcare simulation reviews from 1978 to 2003. Their review of DES in hospitals also concludes that most models are unit specific and facility specific, with very few hospital-wide generic models and even fewer that include all inpatients rather than specific services.

Two recent simulation papers (Devapriya et al. 2015; Demir, Gunal, and Southern 2016) describe models of multiple hospital departments. Devapriya et al. (2015) include most inpatient beds (excludes mental health, pediatrics, and obstetrics) with arrivals from the ED, direct admissions, hospital transfers and surgery, but does not fully model the ORs or the ED. The intention was to apply the model at multiple sites within the Geisinger Hospital System. Some customization of the simulation code was required between hospitals and expansion outside the original hospital system was not reported. The model presented in Demir, Gunal, and Southern (2016) incorporates patient flow through the ED, outpatient and inpatient beds. Bed requirements by service are predicted based on forecasted demand by service.

Outside of healthcare, data-driven generic models are particularly prevalent in defense applications. Brown (2010) provides a literature review of generic data-driven models including literature raising concerns about their usage. Monks, Robinson, and Kotiadis (2009) also raise concerns over model reuse, such as lack of critical involvement by stakeholders in the model development phase.

It has been our experience while implementing this model as well as a previous generic OR model (Sniekers et al. 2017) across multiple hospitals, that when hospital administrators have a problem they have limited financial resources, and can rarely wait the time it would take to develop a detailed simulation from scratch. In addition, a model that has had success at other hospitals is more readily adopted by hospital stakeholders. Concerns expressed by (Monks, Robinson, and Kotiadis 2009) on insufficient involvement from hospital stakeholders if using a generic model, have not materialized in our experience. The stakeholders continue to be involved in validating the data inputs to the model and development of what-if scenarios. That may be because a) the "data" required to run the model also include selection and definition of process rules that can be adjusted in several dimensions (such as surge policy rules) b) the input is sufficiently flexible to allow options concerning the level of detail to be included. Because of this there remains a necessity to walk stakeholders through the options and require them to think in a structured manner about how they operate. We have often found the validation is as enlightening for stakeholders as the what-if portion of the model despite it being generic.

No papers have been found that study hospital-wide, daily surge protocols. Miller, Shahi, and Ferrin (2009) study the effect of ED surge on ED boarding at various inpatient occupancy levels. They also consider the impact on ambulance diversion levels. A few papers also study the impact of a single congestion relief technique on a particular area of the hospital. These include impact of ambulance diversion (Lin, Kao, and Huang 2012; Ramirez, Fowler, Wu 2009), accurately capturing the practice of expediting movement between levels of care (Mallor and Azcarate 2014; Barado et al. 2012), the effect of early discharge policies on ED boarding (Crawford et al. 2014), and inclusion of transfers to other hospitals (Devapriya et al. 2015). In addition, there are papers that look at surge capacity for a catastrophic event such as a terrorist attack (Hirshberg, Stein, and Walden 1999) or a natural disaster (Yi et al. 2010; Paul and Hariharan 2007). Watson, Rudge and Coker (2013) provide a review of disaster surge.

Finally, the inclusion of housekeeping in the model appears to be unique. Housekeeping was included in anticipation of scenarios investigating the effects of a morning discharge policy. (Chaiyachati, et al. 2016; Powell et al. 2012). Such a policy aims to free beds up earlier in the day thereby reducing ED wait times. The concern was, that without changes to housekeeping resources, the gains made by discharging patients earlier in the day could be lost as a backlog of beds sit empty longer waiting to be cleaned. Without modeling these resources all beds would be cleaned simultaneously overestimating the impact of the policy.

3 MODEL DETAILS

The generalized simulation model includes patient flow through the ED, ORs and inpatient units as shown in Figure 1. Actions taken during surge that alter the patient flow between these units are also modelled. Surge actions are shown in the shaded ovals in Figure 1.

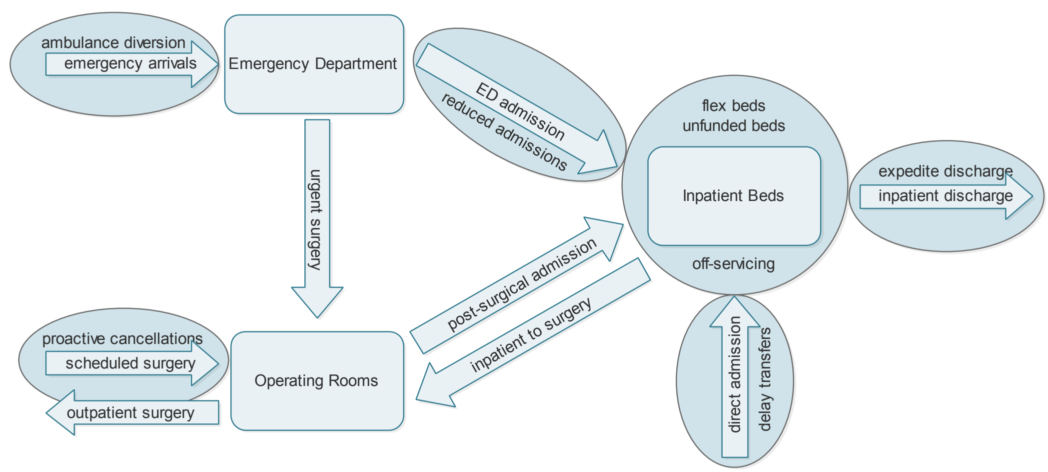

Figure 1: Model coverage with surge protocol responses.

3.1 Data-Driven Generic Simulation Model

While there is a large amount of variation between hospitals in terms of size, specialty, location, and patient mix, most hospitals have many structural and procedural elements in common. These common procedures are used to build a data-driven generic simulation model. By using the excel input file to store all the hospital specific data, the model itself can remain generic. The hospital specific model is created by importing the excel file to the simulation model with no additional changes to the simulation code required. Data stored in the input file is discussed throughout Section 3.

3.2 Flexibility

The complexity of the model leads to large data requirements, so there are built-in options to reduce the detail modelled in case data is not available or less detail is desired. For example, bed turnover time can be modelled explicitly based on housekeeping shifts to detect time-of-day bottlenecks when bed turnover volume is high or other responsibilities increase, or it can be modelled implicitly as a duration distribution. In addition, the OR and ED can be modeled independently, be omitted, or be replaced by input distributions. All hospitals tested so far have chosen to include all inpatient units, the full OR module and the ED from the point of ED disposition.

3.3 Patients

Historical patient data listed below is stored in the Excel input file, each record representing a patient visit to the hospital. At least a year's worth of patient data is preferred. Instead of using distributions to choose patient attributes throughout the simulation, the arriving entity is assigned a patient record and adopts all the listed attributes of that patient. This is done for two reasons: a) patient characteristics are correlated so multiple conditional probabilities and distributions would need to be formulated to capture these dependencies; and b) this approach allows hospitals to more easily understand, update and continue to use the model after our initial consultation. The attributes contained in the patient file include:

1. Arrival Type – ED, Surgical Direct Admit, Medical Direct Admit
2. Arrival day-of-week – for urgent arrivals
3. Arrival time-of-day – for urgent arrivals
4. Booked surgery duration – time used when scheduling the surgery
5. Actual Surgery duration – total time the patient actual spent in the OR
6. Surgical service
7. Surgeon
8. Surgery Type – urgent or elective
9. Surgery priority – priority within type
10. Surgical Patient Category – can be used for any categorization for scheduling purposes, or omitted
11. Number of inpatient stay phases
12. Inpatient Category for Phase i – typically inpatient service and/or level (e.g. acute psychiatry, geriatric, obstetrics) in phase i.
13. Length of stay for Phase i – time in the hospital for inpatient phase i.

If entries are not applicable to the patient, they are left blank. For example, if a patient does not have surgery, the fields related to surgery are left blank. Inpatient stays can be broken down into as many phases as required. These phases represent a change in patient status that requires a change in bed or service. For example, a patient who spends two days as a general medicine patient and then moves to rehab for five days would have two phases. Inpatient Category and Length of Stay are entered for each phase of the inpatient stay. In the previous example, the inpatient category would be general medicine for phase 1 and rehab for phase 2, and the length of stay would be two for phase 1 and five for phase 2. While Inpatient Category typically refers to their inpatient clinical service, the model flexibility allows the categories to be defined in any way. Patient categories can be any group of patients that will be directed through the model in a similar manner. For example, a patient could simply be categorized as "medical" and be sent to an inpatient bed unit representing the medical ward, or could be categorized as "medical with isolation requirements" and be sent to a set of inpatient beds that represent single rooms on the medical ward. If a patient requires surgery during their inpatient stay this is also indicated in the patient category. Since this categorization typically relates to clinical service (e.g. medicine, maternity, cardiology, orthopedics, gastroenterology, etc.) "patient service" is used interchangeably with "patient category" in this paper.

3.4 Arrivals

Entities arrive in the model: a) through the ED; b) as elective surgical patients; c) as urgent surgical patients, not entering through the ED; or d) as direct admission not captured in the other three arrival types. This last category typically includes maternity patients, and transfers from other facilities or clinics.

There are two alternate ED arrival points coded into the model: triage, or disposition (time of decision to admit to an inpatient bed). If arrival at triage is chosen in the input file all ED patients, including those that are eventually discharged home with no hospital admission are included in the model as well as their processes through the emergency room. If the arrival point is chosen to occur at the point of admission to the hospital, then just the time spent in the ED waiting for an inpatient bed is modelled. Participating hospitals so far have felt that patients waiting in the ED for inpatient beds were a significant cause of ED crowding, and it was sufficient to start at ED disposition, thereby reducing the amount of data required. Therefore, the model has only been validated from the point of ED disposition. As such, this paper will focus on that option.

The daily volume of entities arriving to the ED is based on a Poisson distribution that is day-of-week dependent. As each entity arrives, a set of patient characteristics is selected from the subset of patient records, described in the previous section, that indicate an arrival through the ED on the corresponding day of the week. Once chosen, the entity adopts all characteristics from this patient record. Entities arrive to the model each day at midnight and then enter the ED at the time-of-day specified in their patient record. This method allows day-of-week patterns to be maintained while still having random arrivals throughout the day. Patient arrival time of day is considered part of the patient characteristic that may be correlated to other characteristics such as length of stay and service. For example, it is expected that there is a different patient mix on a Saturday evening than at midday on a Tuesday. The Poisson distribution is coded into the model so only the mean is required in the input file. Although the Poisson distribution may not always be the best fit to the data this distribution is chosen for the following reasons: a) Poisson distribution represents random independent arrivals to a process; b) there may not be sufficient data to improve on this distribution given the day-of-the week specification; and c) use of a standard distribution that is reasonably representative, with easily obtained parameters, allows hospitals to easily continue using the model.

Direct Admission patients also arrive to the model each midnight according to a Poisson distribution by day-of-week. A patient record is selected by the arriving entity from a subset of records with a direct admission arrival type and current day-of-the-week arrival. Patients then move to a direct admission arrival queue, at the arrival time-of-day indicated in their patient record, and wait to be assigned a bed.

Each elective surgical service has a unique arrival object in the model for patients arriving for surgery by direct admission. These patients arrive at a rate based on a service specific exponential distribution. Since the surgical service is therefore known on arrival to the model, the entity selects a patient record from the subset of patient records with a matching elective surgical service. The patient is then sent to their surgeon's waiting list (surgeon is specified in the patient file) to wait to be scheduled for surgery. The patient is not physically waiting in this case, but is on a waiting list for surgery. Entry to the hospital occurs on the day the patient is scheduled to have surgery.

Urgent surgery patients entering as direct admissions are likely arriving from a surgeon's clinic or a transfer from another hospital. On a weekly basis, entities arrive to the model according to an exponential distribution by surgical service. A patient record is assigned to the arriving entity from a subset of patient records that have a direct surgical admission arrival, urgent patient type, and matching surgical service. They move to the urgent waiting list at the specified arrival day and time given in their patient record.

The parameters for each arrival distribution are entered in the input file.

3.5 Emergency Department

Once the decision to admit has been made in the ED, a check is done to see if a bed is available. Each patient service is associated with a set of inpatient units that can serve that patient group. Units associated with each patient service are either considered "on-service" or "off-service". On service units are the ideal unit for this category of patients, while off-service units are better suited to a different category of patients, but still acceptable if on-service beds are not available. Therefore, the model will first look for available on-service beds, and if none are found, a check for off-service beds is done. Details on when off-service beds can be used are further provided in section 3.7.2. If no beds are available upon admission, the patient will remain in the ED until one is ready.

If the patient requires urgent surgery, they are placed on an urgent surgical list. While waiting, they remain in the ED or are moved to an inpatient unit associated with their surgical service, if one is available. This is achieved in the model by splitting the entity in two parts, one part waiting on the urgent surgery list while the other physically occupies a space in the ED or inpatient bed. When the OR is ready, the two entities are rejoined and the patient is moved to the OR for surgery.

The number of admitted patients waiting in the ED typically influences the surge level in the hospital. The effect on the ED by changes made throughout the hospital, such as OR block schedules, inpatient bed allocations, prioritization rules, surge policies etc., must therefore captured.

3.6 Operating Rooms

ORs are fully modeled, based on a master surgical schedule, to accurately capture surge and bed occupancy, and to allow what-if scenarios that consider the impact of the OR schedule on the rest of the hospital.

Each day the block schedule, in the input file, dictates the surgical service in each OR. Each block is then assigned a surgeon from the appropriate surgical service. Patients are selected from the scheduled surgeon's waiting list while making sure that all constraints are met. These include: unscheduled time remaining in the block compared to the expected duration of their surgery ("Booked time" in patient record); and max/min constraints, specified in the input file. These constraints can be applied, to the day, the OR, or the surgical service being scheduled, and are based on a user-defined patient categorization specified in the patient file in the field "Surgical Patient Category". For example, an orthopedic block may include a constraint of "at least three total knee replacements".

Urgent surgeries are completed within target times based on acuity level and availability of ORs. For example, in Ontario there are four priority levels for urgent cases with 2-hour, 8-hour, 48-hour, and 7-day deadlines. However, the model allows the user to enter any number of priority levels with user-defined target times for each priority level and surgical service. There are several ways that urgent patients can be worked into the schedule to meet these deadlines during scheduled operating hours. This includes: a) scheduling into the elective block; b) scheduling into urgent blocks that may be included in the block schedule for urgent surgeries; or c) adding them to the end of an elective block if time remains once elective patients have been completed. If this has not occurred by the time patients reach their deadline, they can then bump a scheduled elective patient during operating hours, or be completed outside of operating hours (overnight and on weekends) if necessary.

Once a patient enters the OR they remain there for the time specified in their patient record ("Actual Surgery Duration"). The completed surgical patient then moves to the surgical recovery room, an inpatient bed, or leaves the hospital. If the patient is meant to go to an inpatient bed that is not yet ready, the patient will wait in the surgical recovery room. If the surgical recovery room is full, the patient will remain in the OR blocking the next patient from beginning surgery.

After the surgery is completed, and the patient has left the OR, a check is done to see if the next scheduled surgery can proceed. Because surgery time is booked based on the "Booked Surgery Duration" value in the patient record, there may be a discrepancy between the two times, leading to surgeons being ahead or behind schedule when the next patient is ready to go. A check is done to ensure that surgery will be completed within the operating hours of the OR, plus a user-defined overtime allowance. In addition, a final bed check is done, before the patient can proceed for surgery. If there is not enough time to complete the surgery, or a bed cannot be found the surgery is cancelled. If the patient is cleared to go ahead, they enter the OR after the OR turnover time has been completed. Turnover time is based on an average time by surgical service to clean up and prepare the room between surgeries. If the patient is emergent they will proceed with surgery and wait in the PACU after surgery until a suitable bed is found.

The surgical block schedule, surgeon assignments, surgical priority levels and deadlines, overtime allowance and turnover time are defined by the user in the input file.

3.7 Inpatient Beds

3.7.1 Capacity

Modelling inpatient beds is more complex than simply setting a capacity for the beds and allowing patients to flow into them from other areas of the hospital. Inpatient beds are divided into units, typically aligned with clinical service and/or level of care. Beds in each unit are then further divided into three categories: funded, flex, and unfunded. Funded beds are included in the hospital budget with planned, consistent nursing levels and set nurse-to-bed ratios. Flex beds are extra physical beds that can be temporarily used without adding nursing staff. They are intended to be short term and therefore nurses can “flex up” the nurse-to-bed ratio. Finally, unfunded beds are physical beds that can be opened when necessary by bringing in extra nursing staff. These flex and unfunded beds are routinely used in busy hospitals as part of their surge protocols. The number of inpatient units and the number of each type of bed available (funded, flex, unfunded) in each unit are defined in the model input file.

3.7.2 “On-service” and “off-service”

As discussed above, there are also “on-service” and “off-service” inpatient units associated with each patient category. There is no limit to how many inpatient units can be associated with each patient service and one unit can be associated with more than one patient service. There are four sets of rules, contained in the model input file, that govern inpatient unit assignments:

1. List of on-service and off-service units - for each patient category
2. Surge level required for off-servicing to be allowed - for each patient category and associated unit
3. Probability that a patient will be suitable for off-servicing - for each patient category and unit
4. Number of beds in potential off-service unit protected from off-servicing - for each patient category

3.7.3 Bed Prioritization

When a bed becomes available, there may be several patients waiting for that bed. For example, when a bed becomes available in an orthopedic ward there may be patients waiting in the surgical recovery unit, in the ICU, in another ward where they are off-service or blocked, in the ED, or on the direct entry list. In addition, patients in these areas may be on-service or off-service for the receiving unit. If multiple patients are waiting for the bed, prioritization rules are used to decide who will get the bed. For each inpatient unit, all possible feeder locations and on-service/off-service combinations are ranked in order of preference in the input file.

3.7.4 Subdivided Inpatient Stays

If a patient completes the current phase of their length of stay, they will either exit the model if all phases of their inpatient stay have been completed, or will attempt to move to another inpatient unit associated with the next phase of their inpatient stay. If this bed is not available, the patient will remain in their current unit. For example, if a patient is a general medicine patient on admission, and then changes to a complex care patient, but all units accommodating complex care are full, the patient will remain in general medicine.

If the next phase of the inpatient stay is to have surgery, the patient will remain in their current unit until their surgery is scheduled to begin. If the same bed is appropriate for their post-surgical care, their existing bed will be held until they return from surgery.

3.8 Surge Policies

One of the key contributions of this model is the explicit modelling of hospital surge policies. Levels of surge correspond to the degree of congestion and therefore have different trigger thresholds and associated degrees of response. Personnel engaged also varies by surge level. For example, a level 1 surge may be called when there are four or more patients in the ED with no bed coming open for them today, whereas a level 4 surge may be called when the hospital is entering a second day of surge and has ten or more patients with no bed. Responses at level 1 may include allowing use of flex beds and limited off-servicing with no physician or management involvement, whereas level 4 could involve opening unfunded beds and coordinating with out-of-hospital downstream care providers requiring upper management involvement.

Triggers included in the model are:

1. Number of admitted patients in the ED
2. Number of beds free or coming free today minus number incoming patients
3. Previous days in surge
4. Number of scheduled surgeries with no reserved bed
5. Number of flex or unfunded beds currently open

Triggers can be used independently or in combination, and can specify the units included. There can be multiple ways that a given surge level can be called. For example, a hospital-wide surge level is declared when admitted patients in the ED exceeds a threshold, or an ICU-specific surge level is declared when ICU beds required exceeds available ICU beds. All triggers, combinations and units affected are specified in the input file. Responses included in the model are:

1. Expediting discharge – e.g. speeding up pending tests, physician rounds, and transfers out
2. Expediting move to lower level of care – e.g. move most stable ICU patient to Ward if feasible
3. Allowing increased off-servicing
4. Opening flex beds
5. Opening unfunded beds
6. Delaying outside transfers into the hospital
7. Cancellation of elective surgery
8. Transfers to other hospital from ED
9. Ambulance diversion or consideration

The degree of response (e.g. how early a patient can be considered for expedited discharge), the order in which the responses are executed, and the surge level required to invoke each response is entered in the input file by the user. This can also be differentiated by unit or groups of units, allowing for a great deal of variation in surge protocols represented in the model.

Although elective surgery could be proactively cancelled to alleviate congestion, it is more typical that surgeries are cancelled only when it's apparent there is no bed available. In other words, it is reactive rather than proactive. Including it as a surge response implies a proactive decision.

Because the model has only been validated from the point of ED disposition, ambulance diversion has not been validated. It is believed that the effect of ambulance diversion is small given it affects the total number of patients arriving to the ED, and so only indirectly the number of patients that may be admitted. The restrictions on ambulance diversion are such that only low risk patients will be considered for diversion, further reducing the likelihood that patients being admitted to inpatient beds will be significantly affected.

3.9 Housekeeping

In many cases the time to changeover a bed between patients can safely be represented with a fixed distribution. However, when focusing on surge, and therefore congestion, it can be useful to model the resources directly to better detect when bottlenecks may occur. Bottlenecks in housekeeping could have a significant effect on ED wait time when the hospital or unit is at or close to full capacity.

Shifts for housekeeping resources, assigned inpatient units, and time to clean a bed by unit are defined in the input file. When a patient leaves a bed, a housekeeping resource will be called to clean the bed. If no resource is available, the bed will sit empty while waiting for the housekeeping resource to arrive. Once the resource arrives, bed cleaning time begins. When cleaning is complete, the bed is made available to a new patient and the housekeeping resource looks for another bed to clean. Resource availability can also be restricted due to other duties such as lunch service, or breaks. Modeling this detail will catch when there is a mismatch between demand and supply of housekeeping resources leading to longer periods of beds sitting empty. If this level of detail is not desired, the fixed time to clean a bed will start as soon as the bed is vacated and there will be no limit on the number of beds that can be turned over simultaneously.

4 MODEL IMPLEMENTATION

Characteristics of the Hospitals involved are outlined in Table 1. Hospitals A and B were part of the design phase of the model. Hospital C and D were used to validate the transportability of the completed model.

Table 1: Pilot Hospitals' Characteristics.

Hospital	Location	Type	Inpatient Beds	ED visits	Model Role
A	Suburban	Regional Centre	~400	~90,000	Basis
B	Downtown Urban	Trauma Centre Teaching Hospital	~535	~75,000	Basis and Validation
C	Suburban	Community Hospital	~270	47,000	Validation and Scenarios
D	Small Town	Small Community Hospital	~60	~21,000	Validation and Scenarios

Hospital A was the first hospital studied to understand the patient flow, structure and rules that govern surge policy and bed management. Hospitals B was the second hospital studied to inform the construction of the generalized model and input file. For this hospital, data was collected, and the input file populated to test and validate the model. Hospital C and D used the completed model from the outset, testing transportability. The input file was populated with new data to represent each hospital. Populating the input file with the correct data and validating the output took approximately six months for both hospitals. It is expected that future application of the model at new hospitals will have a similar timeline. The two largest segments of time are for data collection from the hospital (due in part to high demand on the hospital IT departments) and manipulation of historical patient data into the input file format. This is expected to be typical as patient data must be pulled from multiple hospital databases, and then combined. The structure of hospital data varies from hospital to hospital, making data manipulation difficult to automate. However, when compared to the time taken to model the entire surgical process, emergency arrivals, inpatient bed management and surge impacts from scratch, this time frame is much more appealing to hospital administrators.

Multiple what-if scenarios were also run for Hospital C. For example, construction of a new tower, containing a new OR suite and inpatient units, was nearing completion so the model was used to test OR schedules and inform new bed allocations. The impact across the hospital was assessed when making decisions. The model is also being used to explore differing base and flex staffing configurations by season.

In all cases a 13-week warm-up period was used to ensure that the hospital was full when results collection began. Most patients spend less than a week in hospital, so thirteen weeks of warm up allows plenty of time for several cycles of patient entry and discharge to take place before results collection begins. Results were collected for 26 weeks. Both can be adjusted through the input file if desired.

5 MODEL OUTPUT AND "WHAT-IF" SCENARIOS

The model provides detailed output on all areas of the hospital. Output available is listed below:

1. Average and maximum ED boarders (patients waiting for an inpatient bed) - total and by service
2. Average and maximum ED boarding time - total and by service
3. Percentage of time ED boarding time exceeds a given threshold – total and by service
4. Utilization - by OR and by surgical service
5. OR overtime and undertime – by OR and by service
6. Surgical throughput - overall, by service and by patient type
7. Surgical cancellation rates - by service and reason
8. Wait time for surgery - by surgeon and service
9. Midnight census - by inpatient unit
10. Average number of on-service, off-service and blocking patients per unit
11. Direct admission waiting times by service
12. Volume of patients discharged by service
13. Number of days each surge response is used.
14. Percentage of days in each surge level overall and by unit groupings

A secondary advantage of a data-driven model is that "what-if" scenarios are very easy to produce as any input used to create the model can be easily changed within the input file. Multiple simultaneous changes are also easy to test and preserve in a new input file. Five main areas have been, and are expected to be, the focus of what-if scenarios:

1. Bed capacity and allocation, including allocation by service and bed type (funded, flex, unfunded)
2. Impact of Surgical Schedule and rules on throughput, ED wait times and occupancy levels
3. Impact of future volume and patient mix changes
4. Effectiveness of various surge protocols
5. Testing of general operating principles

6 LIMITATIONS AND FUTURE WORK

This generic model speeds up the implementation time at hospitals by eliminating simulation construction time. However, initial consultation is still needed to interpret data, to determine the correct input for rules and processes, and to clarify any terminology differences. Once this initial consultation has occurred, hospital staff can be trained to run scenarios and/or update the data for future use, but out-of-the-box implementation without consultation is not currently feasible.

The ED module has been coded in the simulation from a triage point of entry, but has not been tested and refined with data. Next steps therefore include validating the pre-ED disposition part of the model with historical patient data at several hospitals. The impact of patients waiting in the ED prior to disposition, and the effect of ambulance diversion on overall ED crowding, can then be assessed.

Only formal surge policies have been included in this model. However, it is known anecdotally that there are additional informal actions that may be taken, when in surge. An example of this is the tendency to lean toward admitting a borderline patient when there are beds available, who might otherwise be discharged home if beds are tight. The degree to which this occurs is not captured and the practice is not part of a formal protocol, so has not yet been included in this study.

ACKNOWLEDGMENTS

This work has been funded by the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Ontario government.

REFERENCES

Barado, J., J. M. Guergue, L. Esparaza, C. Azcarate, F. Mallor, S. Ochoa. 2012. "A Mathematical Model for Simulating Daily Bed Occupancy in an Intensive Care Unit." *Critical Care Medicine* 40:1098-1104.

Brown, N. A. 2010. "Model Flexibility: Development of a Generic Data-Driven Simulation" In *Proceedings of the 2010 Winter Simulation Conference*, edited by B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, and E. Yücesan, 1366–1375. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Chaiyachati, K. H., A. N. Sofair, J. I. Schwartz, D. Chia. 2016. "Discharge Rounds: Implementation of a Targeted Intervention for Improving Patient Throughput on an Inpatient Medical Teaching Service." *Southern Medical Journal* 109:313-317.

Crawford, E. A., P. J. Parikh, N. Kong, C. V. Thakar. 2014. "Analyzing Discharge Strategies During Acute Care: A Discrete-Event Simulation Study." *Medical Decision Making : An International Journal of the Society for Medical Decision Making,* 34:231–41.

Demir, E., M. Gunal, D. Southern. 2016. "Demand and Capacity Modelling for Acute Services Using Discrete Event Simulation". *Health Systems* 5:1-8.

Devapriya, P., C. T. B. Stromblad, M. D. Bailey, S. Frazier, J. Bulger, S. T. Kemberling, K. E. Wood. 2015. "StratBAM: A Discrete-Event Simulation Model to Support Strategic Hospital Bed Capacity Decisions" *Journal of Medical Systems* 39:1-13.

Fletcher, A., D. Worthington. 2009. "What is a 'Generic' Hospital Model? -- A Comparison of 'Generic' and 'Specific' Hospital Models of Emergency Patient Flows." *Health Care Management Science* 12:374-391.

Günal, M. M., M. Pidd. 2010. Discrete Event Simulation for Performance Modelling in Health Care: A Review of the Literature. *Journal of Simulation* 4:42-51.

Hirshberg, A., M. Stein, R. Walden. 1999. "Surgical Resource Utilization in Urban Terrorist Bombing: A Computer Simulation." *Journal of Trauma - Injury Infection and Critical Care* 47:545-550.

Jenkins, J. L., R. E. O'Connor, D. C. Cone. 2006. "Differentiating Large-Scale Surge versus Daily Surge". *Academic Emergency Medicine.13*:1169-72.

Katsaliaki, K., N. Mustafee. 2011. "Applications of Simulation within the Healthcare Context." *The Journal of the Operational Research Society* 62:1431-1451.

Lin, C. H., C. Y. Kao, C. Y. Huang. 2012. "Managing Emergency Department Overcrowding via Ambulance Diversion: A Discrete Event Simulation Model." *Journal of the Formosan Medical Association* 114:64-71.

Mallor, F., C. Azcarate. 2014. "Combining Optimization with Simulation to Obtain Credible Models for Intensive Care Units." *Annals of Operations Research* 221:255-271.

Mielczarek, B. 2016. "Review of Modelling Approaches for Healthcare Simulation". *Operations Research and Decisions* 28:55-72.

Miller, M., N. Shahi, D. Ferrin. 2009. "Estimating Patient Surge Impact on Boarding Time in Several Regional Emergency Departments." In *Proceedings of the 2009 Winter Simulation Conference*, edited by M. D. Rossetti, R. R. Hill, B. Johansson, A. Dunkin and R. G. Ingalls, 1906–1915. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Monks, T., S. Robinson, K. Kotiadis. 2009. "Model Reuse versus Model Development: Effects on Credibility and Learning." In *Proceedings of the 2009 Winter Simulation Conference*, edited by M. D. Rossetti, R. R. Hill, B. Johansson, A. Dunkin and R. G. Ingalls, 759–770. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Ontario Ministry of Health and Long-Term Care 2009. "Ontario Targets Shorter ER Times: McGuinty Government Launches Public Reporting of Time Spent in the ER" Accessed Apr. 26, 2017 https://news.ontario.ca/mohltc/en/2009/02/ontario-targets-shorter-er-times.html

Paul, J. A., G. Hariharan, 2007. "Hospital Capacity Planning for Efficient Disaster Mitigation During a Bioterrorist Attack." In *Proceedings of the 2007 Winter Simulation Conference,* edited by S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. D. Tew, and R. R. Barton, 1139–1147. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Powell, E. S., R. K. Khare, A. K. Venkatesh, B. D. Van Roo, J. G. Adams, G. Reinhardt. 2012. "The Relationship Between Inpatient Discharge Timing and Emergency Department Boarding" *Journal of Emergency Medicine* 42:186-196.

Ramirez, A., J. W. Fowler, T. Wu. 2009. "Analysis of Ambulance Diversion Policies for a Large-Size Hospital." In *Proceedings of the 2009 Winter Simulation Conference,* edited by M. D. Rossetti, R. R. Hill, B. Johansson, A. Dunkin and R. G. Ingalls, 1875–1886. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Sniekers, D., C. R. Busby, M. W. Carter. 2017. "Implementation of Generic Perioperative Simulation at Multiple Hospitals" Working Paper, Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, Ontario, Canada.

Watson, S. K., J. W. Rudge, R. Coker. 2013. "Health Systems' 'Surge Capacity': State of the Art and Priorities for Future Research." *Milbank Quarterly* 91:78-122.

Yi, P., S. K. George, J. A. Paul, L. Lin, 2010. "Hospital Capacity Planning for Disaster Emergency Management." *Socio-Economic Planning Sciences* 44:151-160.

AUTHOR BIOGRAPHIES

CAROLYN R BUSBY is a PhD student in Mechanical and Industrial Engineering at the University of Toronto. Her research interest is operations research applied to healthcare, particularly simulation. Her email address is carolyn.busby@mail.utoronto.ca.

MICHAEL W CARTER is a Professor in Mechanical and Industrial Engineering at the University of Toronto since 1981. For the past 27 years, his research focus has been in the area of health care process improvement and policy modeling. He has supervised more than 260 engineering students in over 160 projects with healthcare institutions. Ten former students now work at CCO and most Toronto hospitals have a couple. He is on the editorial board for the journals "Health Care Management Science", "Operations Research for Healthcare", "Health Systems" and "IIE Transactions on Healthcare Systems". In 2012, he was inducted as a Fellow of the Canadian Academy of Engineering and in 2013, he became a Fellow of INFORMS. His email address is mike.carter@utoronto.ca.