Resource Scheduling through Resource-Aware Simulation of Emergency Departments

Seung Yeob Shin[UM], Hari Balasubramanian[⋆], Yuriy Brun[UM], Philip L. Henneman[†], and Leon J. Osterweil[UM]
[UM]School of Computer Science, University of Massachusetts, Amherst, MA, USA
[⋆]Department of Industrial Engineering, University of Massachusetts, Amherst, MA, USA
[†]Department of Emergency Medicine, Tufts-Baystate Medical Center, Springfield, MA, USA
{shin@cs, hbalasubraman@ecs, brun@cs, ljo@cs}.umass.edu, philip.henneman@baystatehealth.org

***Abstract*—This paper proposes using resource-aware, discrete-event simulation to measure the effects of resource scheduling in hospital emergency departments. Determining staffing and resource allocation is a complex constraint-optimization problem that has significant impact on hospital costs and patient care quality. We developed detailed models of the emergency department process of caring for patients, the resources available to support that process, and the scheduling constraints on the deployment of those resources. We then ran a battery of discrete-event simulations of this process, varying details of process, resource mixes, and scheduling constraints, to analyze the effects of resource availability (e.g., staffing patterns) on patient length of stay. Our simulation approach proved to be particularly adept at supporting the systematic investigation of two issues of particular interest to domain experts: (1) an excessive focus on minimizing the average length of stay (the objective most typically used for optimizing emergency department staffing) can have undesirable, previously unappreciated effects, (2) too strong a focus on one particular kind of resource as the preferred vehicle for decreasing patient length of stay can tend to obscure the value of considering other kinds of resources. The unexpected nature of some of our results raises open questions about how to validate the results of complex simulations.**

I. Introduction

The set of available resources and the scheduling of those resources can play an important role in determining the efficiency and cost of many systems. However, determining an optimal set of resources and scheduling those resources are NP-hard [19]. This creates questions about how to improve the state-of-the-practice of resource allocation and scheduling.

Hospital emergency departments are examples of systems whose performance is heavily dependent on resource availability. Hospital administrators typically determine and schedule resources using informal heuristics, lower-bound calculations [11], [12], existing national averages, and extensive historical data at similar size and type institutions. Even just estimating the effects of a particular resource allocation and schedule — the number of doctors, nurses, beds, x-ray rooms, etc., and when those nurses and doctors are made available for work — is costly and imprecise in practice. Recent advances in discrete-event simulation have reduced the cost of such estimation [1], [7], [8], [21], but the estimates are still imprecise because they are based on models that do not account for variability in patient arrival rates and severity types, certain details of patient care, variability in the characteristics of the involved resources (especially the human resources), and the complex interplays between all of these different dimensions of variability.

We address the resource allocation and scheduling problems by creating discrete-event simulations based on detailed models of system processes, and detailed models of resource characteristics and constraints. This requires access to domain knowledge, extending existing discrete-event simulation capabilities, and careful validation. In this paper, we present the results of developing these detailed models and running simulations to study the effects of various approaches to resource allocation on such measures of emergency department quality as the average patient length of stay (LoS). Our simulations have provided insights and perspectives that domain experts have found to be provocative. Consequently, we have also explored the issue of simulation validation.

Two provocative observations we have based upon our simulations are:

1) While hospital administrators, and previous related simulation research, focus on minimizing the average patient LoS, we find that this optimization goal results in staffing patterns having undesirable consequences, such as overstaffing during low-demand patient periods (which, at a low cost in resources, lowers the average LoS), and understaffing during high-demand, and peak patient periods (when adding extra resources has less of an effect on the high LoS). Thus, optimizing for the average LoS creates staffing patterns that result in undesirably high variability in individual patient LoS.
2) While the approach of adding extra beds is often used in attempting to achieve significant reductions in LoS, we find that scarcity of other resources, such as x-ray rooms, may at other times be the key bottleneck, and a more effective way to reduce LoS.

The rest of this paper is structured as follows: Section II summarizes the relevant background and related research in resource allocation and discrete-event simulation. Section III describes our approach of resource-aware simulation. Section IV details our experiments and findings. Section V discusses open questions in simulation validation. And finally, Section VI concludes with a summary of our contributions.

II. Background and Related Work

Resource allocation and scheduling are NP-complete problems [19], [22]. The number of ways in which resources can be allocated and scheduled to optimize some objective function is exponential in the number of the resources, and no efficient ways exist to efficiently search through this space to find optimal solutions. The objective functions, and constraints that may restrict the search space are often domain-specific.

Other theoretical work has focused on efficiently approximating optimal solutions [14]. These approaches are relevant in practice, but are general and do not benefit from domain-specific constraints. Additionally, in practice, and in some domains, it may be worth expending additional computation to find a solution closer to the optimum, as the costs of using suboptimal solutions may far exceed the cost of extra planning (as is often the case in healthcare).

In hospital emergency departments, resource allocation and scheduling are especially critical because suboptimal staffing can result in poor patient care, loss of life, and less importantly, cost inefficiency. In this domain, there are many types of resources (e.g., doctors, nurses, beds, x-ray rooms, etc.), and also many factors that affect the quality of care, including variable patient arrival rates and variable severity of patient injuries and needs.

While much previous work in the space of resource allocation and scheduling has focused on the computational and theoretical aspects of these problems [9], [14], [15], we are interested in addressing practical instances of the problems having the added complexity of domain-specific constraints, complex resource specifications, intricate allocation policies, and objective functions that compose potentially inconsistent goals.

In the domain of hospital emergency departments, discrete-event simulation has proven fruitful in exploring effects of crowding, varying patient arrival rates, resource allocation, and scheduling [1], [7], [8], [10], [13], [21]. Our approach focuses on detailed models of the involved resources, especially human resources, to further improve the models and produce results that are closer to the real world in terms of uses and scheduling of those resources.

In the closest work to our own, Beck uses the Arena simulation software [18] to provide a simulation-based, iterative method for resource allocation while allowing for heterogeneous resource types [1]. This method assumes that patient arrival rates are fixed. In contrast, our approach allows for patient arrival rates that vary in the course of a typical 24-hour day, and for more-detailed specifications of both the patient-handling process and the involved resources. In theory, Beck's approach can be extended to handle the details of the patient-handling process, and some, but not all, of the resource details. Our view, however, is that our simulation architecture and process and resource specification approaches make it easier to specify these details and vary them as driven the the needs of our research.

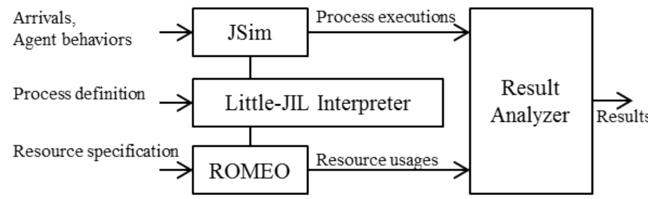

Fig. 1. The architecture of our resource-aware simulator separates the simulation engine (JSim [17]), process definition interpreter (Little-JIL interpreter), and resource manager (ROMEO).

III. Resource-Aware Simulation

We have built a resource-aware simulator to address the resource allocation problem. Figure 1 shows our resource-aware simulator architecture. Our approach uses a discrete event simulator, JSim [17], to simulate the 24-hour operation of an emergency department. JSim's extension for application-specific simulation result analyzers greatly facilitated the construction of our specific emergency department simulator. The simulation uses the Little-JIL process definition language [23] to specify an emergency department process. The rich semantics of Little-JIL make it easy to capture the complex nature of an emergency department process, and is especially effective in supporting the clear specification of what kinds of resources are needed by each activity in a process. Management of the resources available for allocation to the activities in the process specification is separated into the ROMEO [16] component, which facilitates specifying the characteristics of different resources, and constraints on those resources' allocation and availability. This separation makes it easy to keep track of resource allocations, utilizations, waiting times, and other properties. In addition, resource constraints can be changed in a flexible manner. This separation also facilitates clear visibility into resource utilization levels, leading to new insights into emergency department resource management.

A. Emergency Department Characteristics

We developed detailed models of the emergency department based on advice from one of our coauthors, a domain expert with extensive experience as a doctor and a Director of the Emergency Department at the Baystate Medical Center, in Springfield, MA, USA. Due to space constraints, we focus only on a few of the interesting characteristics of the process that, while ostensibly intricate to specify, nevertheless, proved to be relatively straightforward for us to implement by virtue of the resource specification architecture and specification language we used.

- **Six acuity levels**: The patients are classified into six acuity levels, based on the severity of their ailments. After a patient is placed in a bed, the process the patient goes through varies based on the acuity level. A level-six patient is the sickest patient, and an MD (medical doctor), an RN (registered nurse), a bed, and an x-ray room, and CT room resources are allocated to such a patient with the highest priority. In contrast, simulation of a level-one patient entails

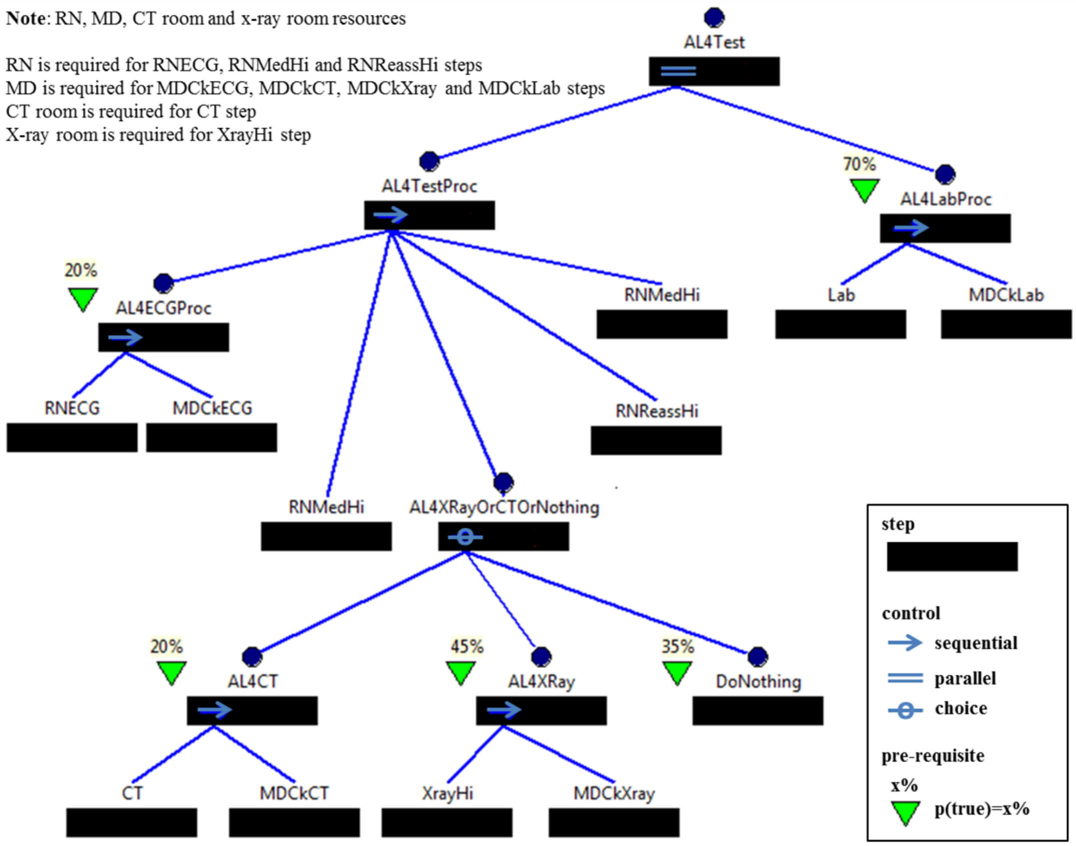

Fig. 2. The Little-JIL [23] definition of the patient testing process, which is part of the care an acuity-level-four patient undergoes in an emergency department.

fewer procedures, such as x-rays, each of which will have a lower priority for the acquisition of resources.

- **Staffing**: MD and RN resources work on a shift system; the numbers of available MD and RN resources vary over 24 hours. Typically, an MD or an RN will work one of three different 8-hour shifts, although our simulations have suggested that greater flexibility in the start times and durations of shifts could lead to improved LoS results.
- **Same MD-RN constraints**: A patient assigned to a specific bed is cared for by the same MD and RN throughout the patient's stay, with changes only due to shift changes.
- **Fast & Main tracks**: The emergency department operates two separate tracks. The *fast track* cares for low-acuity patients (levels 1 to 3), and the main track treats high-acuity patients (levels 4 to 6). The two tracks have their own bed, MD, and RN resources. At night, however, the fast track closes and fast-track patients are transferred to the main track. During this transfer, fast-track resources are deallocated and appropriate main-track resources are allocated for the patients. One interesting emergent property of the simulation is that it is possible that no main-track resources are available when the fast track closes. In such a case, the treatment of the patient must continue with a fast-track bed, but main-track MD and RN resources.

B. *Emergency Department Model*

Both to improve the state of the art in resource scheduling in the ED domain, and also to provide a stringent test of the specification power of our approach, we sought to developed very precise and very detailed models of both ED processes and ED resources. This required understanding low level details of both, and required access to an expert in ED operations and resources. Our domain expert from the Baystate Medical Center provided those details, which did prove to be quite intricate and challenging to model. Indeed, some of the details (e.g. how the two tracks of an ED relate to each other) would have been difficult for a non-expert to invent or infer, and provided good examples of the challenges that consideration of the real world presents. Figure 2 illustrates one small piece of this very detailed process definition: the patient testing process of an acuity-level-four patient.

A Little-JIL specification is a graphical, hierarchical decomposition of activities (called steps), with each step represented graphically by a black bar. Steps are connected by edges to parents (above) and children (below), with edges also specifying the flow of arguments between parents and children. Parent steps both define scopes, and also specify the flow of control between children. The legend in Figure 2 indicates three different control flow possibilities: sequential (children performed in left-to-right order), parallel (children performed in any order, possibly concurrently), and choice (only one of the children selected for performance). Each step also incorporates a specification (not shown) of needed resources (e.g. doctor, nurse, x-ray machine) to be allocated at run time. It is useful to note that these specifications can set up contentions that can further constrain execution order, for example by making concurrent performance either possible or impossible. Thus, Figure 2 specifies that `AL4Test` is a parallel step, which means that a lab test process, `AL4LabProc`, can be performed in parallel with the other tests, although contention for needed resources (in this case the patient) may make concurrency impossible. As noted in the legend of Figure 2, steps may have prerequisites that may be used by our simulations to specify the relative frequency with which exceptions should be thrown, or alternatives in a choice steps should be selected. Thus, the pre-requisite on `AL4LabProc` means that 70% of acuity-level-four patients require the lab test. For the other tests, a nurse checks a patient's `ECG` first, `RNECG`, and then a doctor checks the ECG result, `MDCkECG`, because `AL4ECGProc` controls its child steps sequentially. After the ECG test, a nurse gives a medication to the patient, `RNMedHi`, and then the patient will be transferred to the CT or x-ray room. This behavior is represented by the `AL4XrayOrCTOrNothing` choice step which means only one of its child steps will be executed.

While Figure 2 describes only 19 steps, a small part of the overall process, in total, our detailed process definition contains 164 steps, each with detailed resource requirements, resource assignment policies, execution probabilities, and completion times drawn from a distribution, all developed with the help of our Baystate Medical Center domain expert.

Our process definition allows for patients to arrive in the emergency department in two ways: (a) critical patients are assumed to always arrive by ambulance, while (b) other patients are assumed to always arrive by their own transportation. Critical patients are the sickest (acuity level six), while the others are categorized into the remaining five acuity levels. Patient arrival rates over a 24 hour period are specified by a Poisson distribution. A critical patient is placed in bed immediately; however, other patients are first triaged by a TrRN (triage nurse) and then cared for according to their acuity levels. The times for performing each of the activities in the process for each acuity level are specified by triangular distributions. These estimates dictate that the amount of utilization of each of the resources is determined by the amount of time required by each of the activities.

Figure 3 specifies the emergency department resource allocation policy and availability for the resources that are available to support the execution of the emergency department process. Bed, MD and RN resources are managed separately by each of the two different tracks, while, TrRN, Clerk, x-ray, and CT resources are shared by the two tracks.

Resource	Track	Allocation policy	Available
bed	main	sickest first	24 hours
bed	fast	sickest first	9AM–1AM
MD	main	sickest first	shift hours
MD	fast	sickest first	shift hours
RN	main	sickest first	shift hours
RN	fast	sickest first	shift hours
TrRN	shared	first-come-first-serve	24 hours
Clerk	shared	first-come-first-serve	24 hours
x-ray	shared	sickest first	24 hours
CT	shared	sickest first	24 hours

Fig. 3. Emergency department resource allocation policy and availability specification.

IV. Evaluation

An emergency department has potentially conflicting goals of decreasing patients' LoS and increasing the net revenue derived from its operations. An emergency department can decrease patients' LoS by hiring more staff and making more facilities available; however, these actions increase the cost of operating the emergency department and reduce the net revenue. The work described in this paper demonstrates how our resource-aware simulation can be used to support decisions and to suggest staffing patterns and numbers of facilities that result in obtaining a desirable balance between reducing LoS and sustaining desired levels of net revenue. Further, it can be used to explore the effects of changing the resources on LoS, revenue, and other factors. In the simulations described in this paper, we experiment with adjusting levels of staff resources and numbers of x-ray rooms and beds, but assume fixed distributions of patient arrivals, patient acuity levels, and other statistical parameters.

A. *Average LoS*

In this section, we describe how we used our simulator to analyze the effects of RN staffing on patient LoS and RN utilization.

We first determined that the theoretical minimum LoS (achieved when there are no resource limits placed on the simulation) was 118 minutes, for the given patient arrival and acuity rate functions. In finding reasonable RN staffing patterns, based on our domain expert's advice, we established a realistic target of an average LoS between 136 and 153 minutes (between 115% and 130% of the minimal LoS). As an initial suggestion, we hypothesized nurse staffing levels based on patient arrival rates since staffing needs appear highly dependent on incoming volumes of patients over time. This initial pattern supported an average patient LoS of 143 minutes, which satisfied our domain expert's specified, although estimated constraint.

Helped in part by our architectural separation of resource management concerns into a distinct architectural component, we were then able to study the variation over time of such factors as RN utilization levels and patient LoS. The solid lines in Figure 4 show this initial (original) hourly RN staffing pattern (Figure 4(a)), the RN utilization (Figure 4(b)), and the patient average LoS (Figure 4(c)), over the course of 24 hours. Note that RN utilization is a high 91% at 2AM and that the LoS during the night can be as long as 198 minutes. On the other hand, in late afternoon, the RN utilization and LoS are both considerably lower, when the staffing levels are dramatically higher.

The size of the variation in utilization and LoS levels over the 24-hour period were surprising to the domain expert. More

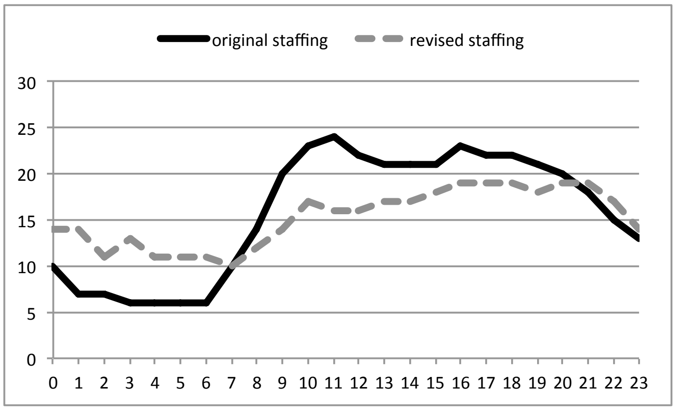

(a) #RN

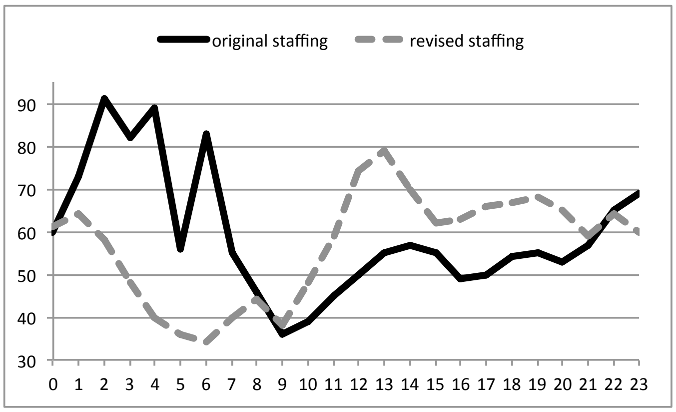

(b) RN utilization (%)

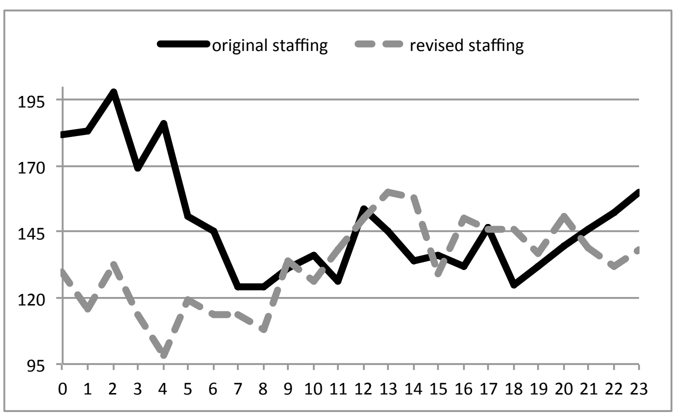

(c) average LoS (minutes)

Fig. 4. The 24-hour, by-the-hour number of RNs, RN utilization, and patient average LoS for two RN staffing patterns. Both staffing patters aim to achieve an LoS between 136 and 153 minutes (between 115% and 130% of the minimal possible LoS), but while the original staffing imposes no other constraints, the revised staffing also attempts to decrease the variance in the RN utilization.

consistent utilization and LoS levels were considered to be more desirable, suggesting the need to identify staffing patterns that reduce these variations.

Accordingly, we designed a new staffing pattern by considering the need to keep the average LoS constrained between 136 and 153 minutes, but also the need to decrease the variance in nurse utilization over the 24-hour period. The dashed lines in Figure 4 show this revised staffing, along with the resulting utilization and LoS measurements. In comparing the original and the revised staffing, we note that the revised staffing has more nurses at night and fewer nurses in the daytime. The average LoS for the revised staffing is 137 minutes, which is within the required range, and indeed is an improvement over the LoS for the original staffing. Moreover, the standard deviation of nurse utilization levels over 24 hours for the revised staffing is 12, which is less than the standard deviation of 15 in the original staffing. We observed a similar reduction in the variation of LoS, with a standard deviation of 16 in the revised staffing and 21 in the original staffing.

Even though the revised staffing pattern in Figure 4 is not the optimal solution in terms of minimizing the variance of nurse utilization, it provides a good example of how consideration of more constraints can provide better solutions. In particular, it indicates the clear danger of using only one objective (e.g., LoS) as the basis for the evaluation of resource allocations and staffing patterns. As a consequence, together with the domain expert, we are now investigating the objective functions that should be used to identify superior staffing patterns, considering more factors than only the average LoS.

B. Resource Bottlenecks

In addition to adjusting human-resource scheduling, an emergency department can also balance its conflicting goals of decreasing LoS and increasing net revenue by modifying the numbers of non-human resources, such as x-ray rooms, beds and other facilities. Our observation is that the bed resource is most commonly increased in order to effect desired reductions in LoS, because LoS appears to be most negatively correlated with full, or nearly full beds. However, our work hypothesizes that other resources can also be serious bottlenecks.

To study this hypothesis, we ran a battery of simulations that assumed access to infinite quantities of MDs, RNs, and other key resources, but controlled the number of x-ray rooms and beds. In this way, we were able to focus on these two resources that were the only possible sources of delay and increase in LoS for these simulations. Figure 5(a) shows how LoS decreases with the growth in the number of x-ray rooms from 2 to 5. Meanwhile Figure 5(b) shows how LoS decreases with growth in the number of beds from 18 to 25. Adding a third x-ray room can substantially decrease LoS, whereas adding subsequent x-ray rooms has little effect. Similarly adding a modest number of beds can have a similarly beneficial effect, but adding a large number of new beds may be less effective. Our approach allows finding the point at which adding additional resources of a particular type is of limited benefit, and help balance how much of each resource to invest in.

This work has reinforced our view that a powerful discrete-event simulation capability can be helpful in suggesting what mixes of non-human resources are most likely most cost effective. The complex interplay among the various kinds of resources during different phases of the healthcare processes for different patients requires using a precise and detailed model of

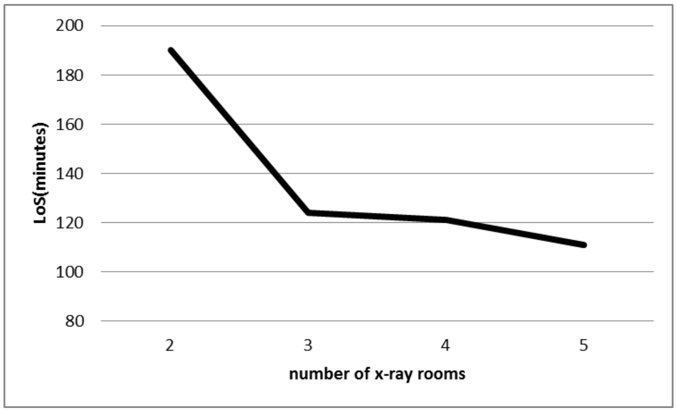

(a) LoS(minutes), as a function of the number of x-ray rooms

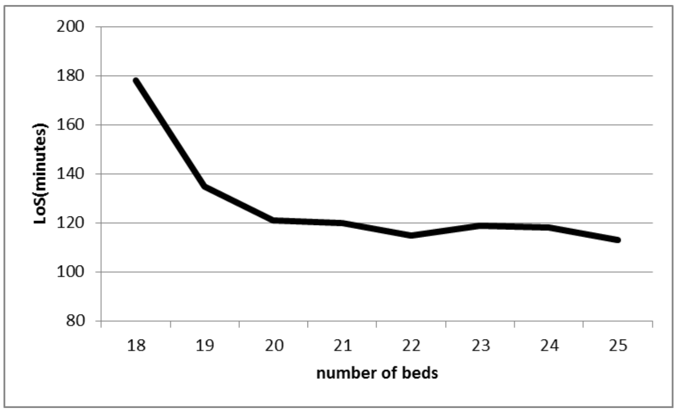

(b) LoS(minutes), as a function of the number of beds

Fig. 5. Independently varying the number of x-ray rooms and beds affects the patient average LoS.

processes and resources as the basis for such simulations. The use of a clearly separated component for the management of resources facilitates the wide-ranging simulation studies that have suggested the value of investigating considering various combinations of new resources in deciding which are likely to comprise the most effective investments.

V. Simulation Validation

In the previous section, we presented experiments that show the effectiveness of our approach in providing support for exploring the effects of resources on properties of the emergency department. These experiments led to preliminary conclusions that were surprising to a domain expert who is well familiar with previous discrete-even simulation work on emergency departments. This suggests one area of this field that needs further exploration is the validation of these simulations, in both prior work and our own work. We now outline a few open research questions that our work has suggested in the space of simulation validation:

> *Q1: How do we know a simulation is accurate and that the results have real-world implications?*

In general, we use the simulation to predict results of alternative design decision choices. In this scenario, there is no oracle that can identify if the simulation results are correct. Simulating cases with design decisions present in the real world and for which real-world measurements have been made, and verifying those results against the real-world measurements can begin to address this issue; however, deeper validation techniques that do not require an oracle are also necessary.

> *Q2: What properties and behavior, and types of properties and behavior, need to be validated, and how can they be validated?*

There are many kinds of properties that should be validated between the simulation and the real world. Some properties are domain specific to the application domain, whereas others are general. Dynamic testing and assertion approaches can handle some of these properties, whereas others are better suited for static reasoning (e.g., finite state verification, model checking, and fault-tree analysis [5], [6], [20]). Identifying both the properties and the validation techniques that are suitable for supporting the study of these properties remains open research. We have begun using a dynamic model-inference approach [2], [3], [4] to explore which simulation properties can be verified automatically during, as well as after, the simulation execution. We are interested in the degree to which this approach can be complemented by static verification of these properties. Our view is that the growing arsenal of increasingly powerful software testing and analysis tools and technologies should be evaluated against the increasingly challenging needs of our increasingly complex discrete-event simulation systems. This is a key future direction of our research.

> *Q3: How well has the architectural separation of the resource management concern actually worked out?*

Our work appears to have produced evidence that the clear separation of resource management concerns greatly facilitates the identification of effective staffing patterns, as it creates a clear method for modifying resource specifications, resource attributes, resource availability specifications, and other resource constraints. At present, the user interface through which these modifications are made is quite primitive and idiosyncratic, suggesting the need for more focus on that part of our approach. The work has also suggested the need for stronger tools to support reasoning about the consistency of different resource specifications, and indeed reasoning about the consistency between resource specifications and other process specifications, which are concentrated in other architectural components.

VI. Contributions

Staffing is an important problem for the emergency department. The staffing impacts the quality of patient care, efficiency of resource use, ability to treat large volumes of patients, and hospital revenue. Discrete-event simulation can help determine the necessary resources, and schedule those resources to achieve certain objectives. We have demonstrated that separating the concern of resource modeling into a distinct component in the simulation allows modeling domain-specific constraints, complex resource specifications, and intricate allocation policies.

Further, such models allow simulations to consider objective functions that compose potentially inconsistent goals.

In a case study, with the help of a domain expert, we developed a detailed model of the process an emergency department undergoes in caring for a patient. Using this model, we evaluated various staffing and resource allocation patterns, which lead us to two observations: (1) Aiming only to minimize the patients' average LoS may result in high variability in the staff (e.g., nurse) utilization and in the LoS itself across patients. As such variation is undesirable, the average LoS is an insufficient objective function on its own. (2) While increasing resources (e.g., beds and x-ray rooms) reduces LoS, different staffing patterns result in different resource bottlenecks. Discrete-event simulation can help identify when adding a bed is the most effective way to reduce LoS, and when adding other resources is more effective.

Our preliminary experiments show promise that discrete-event simulation that focuses on detailed, precise models of the involved resources can help hospital administrators in evaluating staffing and resource allocation patterns, potentially improving efficiency and quality of patient care.

Acknowledgments

We wish to thank our co-author Dr. Philip L. Henneman for his expertise in designing the detailed model of the emergency department process of handling patients. We also wish to thank Alexander Wise for his help in developing and supporting various Little-JIL tools.

This material is based upon work supported by the National Science Foundation under grants IIS-1239334, CNS-1258588, and IIS-0705772.

References

[1] Ekkehard Beck. A discrete event simulation approach to resource management, process changes and task prioritization in emergency departments. Master's thesis, Department of Mechanical and Industrial Engineering, University of Massachusetts, Amherst, MA, USA, 2009.

[2] Ivan Beschastnikh, Jenny Abrahamson, Yuriy Brun, and Michael D. Ernst. Synoptic: Studying logged behavior with inferred models. In *Proceedings of the 8th Joint Meeting of the European Software Engineering Conference and ACM SIGSOFT Symposium on the Foundations of Software Engineering Tool Demonstration Track (ESEC/FSE11)*, pages 448–451, Szeged, Hungary, 2011. DOI: 10.1145/2025113.2025188.

[3] Ivan Beschastnikh, Yuriy Brun, Michael D. Ernst, Arvind Krishnamurthy, and Thomas E. Anderson. Mining temporal invariants from partially ordered logs. *ACM SIGOPS Operating Systems Review*, 45(3):39–46, December 2011.

[4] Ivan Beschastnikh, Yuriy Brun, Sigurd Schneider, Michael Sloan, and Michael D. Ernst. Leveraging existing instrumentation to automatically infer invariant-constrained models. In *Proceedings of the 8th Joint Meeting of the European Software Engineering Conference and ACM SIGSOFT Symposium on the Foundations of Software Engineering (ESEC/FSE11)*, pages 267–277, Szeged, Hungary, 2011. DOI: 10.1145/2025113.2025151.

[5] Bin Chen, George S. Avrunin, Elizabeth A. Henneman, Lori A. Clarke, Leon J. Osterweil, and Philip L. Henneman. Analyzing medical processes. In *Proceedings of the ACM SIGSOFT/IEEE 30th International Conference on Software Engineering (ICSE08)*, Leipzig, Germany, May 2008.

[6] Lori A. Clarke, George S. Avrunin, and Leon J. Osterweil. Using software engineering technology to improve the quality of medical processes. In *Proceedings of the ACM SIGSOFT/IEEE 30th International Conference on Software Engineering (ICSE'08) companion volume*, Leipzig, Germany, May 2008.

[7] Lloyd G. Connelly and Aaron E. Bair. Discrete event simulation of emergency department activity: A platform for system-level operations research. *Academic Emergency Medicine*, 11(11):1177–1185, 2004.

[8] Christine Duguay and Fatah Chetouane. Modeling and improving emergency department systems using discrete event simulation. *Simulation*, 83(3):311–320, 2007.

[9] Naoki Fukuta and Takayuki Ito. Fine-grained efficient resource allocation using approximated combinatorial auctionsa parallel greedy winner approximation for large-scale problems. *Web Intelligence and Agent Systems*, 7(1):43–63, 2009.

[10] L. V. Green, P. J. Kolesar, and A. Svoronos. Improving the sipp approach for staffing service systems that have cyclic demands. *Operations Research*, 49, 2001.

[11] John D. C. Little. A proof for the queuing formula: $L = \lambda W$. *Operations Research*, 9(3):383–387, 1961.

[12] John D. C. Little. Little's law as viewed on its 50th anniversary. *Operations Research*, 59(3):536–549, 2011.

[13] M. L. McCarthy, R. Ding, J. M. Pines, and S. L. Zeger. Comparison of methods for measuring crowding and its effects on length of stay in the emergency department. *Academic Emergency Medicine*, 18(12):1269–1277, 2011.

[14] Rajeev Motwani and Prabhakar Raghavan. *Randomized Algorithms*. Cambridge University Press, New York, NY, USA, 1995.

[15] Ali ParandehGheibi, Atilla Eryilmaz, Asuman Ozdaglar, and Muriel Médard. On resource allocation in fading multiple-access channels — An efficient approximate projection approach. *IEEE Transactions on Information Theory*, 56(9):4417–4437, 2010.

[16] M. Raunak and L. Osterweil. Resource management for complex and dynamic environments. *IEEE Transactions on Software Engineering*, 99, 2012.

[17] M. S. Raunak, L. J. Osterweil, Alexander Wise, L. A. Clarke, and P. L. Henneman. Simulating patient flow through an emergency department using process-driven discrete event simulation. *Software Engineering in Health Care*, 2009.

[18] Rockwell Automation. Arena simulation software. http://www.arenasimulation.com/.

[19] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 3rd edition, 2009.

[20] Borislava I. Simidchieva, Sophie J. Engle, Michael Clifford, Alicia Clay Jones, Sean Peisert, Matt Bishop, Lori A. Clarke, and Leon J. Osterweil. Modeling and analyzing faults to improve election process robustness. In *Proceedings of the Electronic Voting Technology Workshop/Workshop on Trustworthy Elections (EVT/WOTE10)*, Washington, DC, USA, August 2010.

[21] David Sinreich and Yariv Marmor. Emergency department operations: The basis for developing a simulation tool. *IIE Transactions*, 37(3):233–245, 2005.

[22] Michael Sipser. *Introduction to the Theory of Computation*. PWS Publishing Company, 1997.

[23] Alexander Wise. Little-JIL 1.5 language report. Technical Report 2006–051, Department of Computer Science, University of Massachusetts, Amherst, 2006.