REDSim: A SPATIAL AGENT-BASED SIMULATION FOR STUDYING EMERGENCY DEPARTMENTS

Ana Paula Centeno
Richard Martin

Computer Science Department
Rutgers University
110 Frelinghuysen Road
Piscataway, NJ 08854, USA

Robert Sweeney

Emergency Department
Jersey Shore University Medical Center
1945 State Route 22
Neptune, NJ 07754, USA

ABSTRACT

Faced with a mismatch between demand and resources, Emergency Department (ED) administrators and staff need to gauge the impacts of staff decision processes in lieu of increasing resource levels. In this paper we present REDSim, a spatial agent-based simulation framework for studying emergency departments. REDSim focuses on quantifying the impacts of staff decision processes, such as patient selection, on the length of stay, waiting and boarding times, and other variables. We show REDSim's effectiveness by comparing four patient selection strategies: longest waiting, shortest distance, random, and highest acuity. REDSim showed that although patient length of stay is not significantly reduced (1.4%) the throughput increases 17% when providers select the closest instead of the highest acuity patient for the next task.

1 Introduction

Over the past 20 years the demands on Emergency Departments (ED) have changed substantially, although the quantity and nature of the resources available has not kept pace. Today's ED is a primary evaluation and treatment site for not only trauma care and emergent medical care, but also acute psychiatric illness, domestic violence, sexual assault, drug and alcohol addictions, and other social and primary care issues. Indeed, as evidence of this increased role, visits to the ED have increased by 20% although the number of EDs has declined by 10% (Nawar et al. 2007, Reid et al. 2005).

The result is that many EDs are chronically overburdened. Symptoms of the mismatch between demands and resources include ambulance diversion, extended boarding times and high walk out rates. The negative impacts these problems cause have been quantified in numerous studies, as summarized in (Committee on the Future of Emergency Care in the United States Health System 2007), and involve serious health effects including patients leaving without being seen, increases in medical errors, and statistically significant correlations between overcrowding and increased patient mortality.

There are diverse strategies to reduce overcrowding in the ED, including staffing changes in the triage process, protocol changes during surges, electronic white boards systems, and systemic reforms to reduce or divert patient flow. One facet of the problem is that ED staff and administrators have little in the way of quantitative tools and techniques to judge the effectiveness of changes in staffing, space, procedures or technologies without resorting to costly pilots and trials.

Recent works have addressed this evaluation gap through simulation. (Zeltyn et al. 2011, Brenner et al. 2010, Duguay and Chetouane 2007). A key limitation of most of these works, however, is that they focus only on resources (number of staff, equipment, and rooms) as the dependent variables. The independent variables were often metrics such as length of stay (LOS), waiting times, and leave without being seen (LWBS) rates. Our simulation framework also focuses on a similar set of independent variables as metrics,

but has two key distinctions for the dependent variables. First, it models provider decision making, and second, it models the spatial layout of an ED. Both allow a greater range of studies on the impacts of how health-care providers make decisions, as well as the impacts of the ED's spatial reconfigurations.

Our agent-based ED simulation, called REDSim, thus allows evaluations of the impacts of provider task selection. In a modern ED, providers continuously face local scheduling decisions about what task should they do next. Anecdotal evidence suggests that health-care providers typically follow a few rules-of-thumb based on their situational awareness, resulting in a set of self-created guidelines. REDSim allows users to evaluate the impacts of these local rules-of-thumb on the operation of the ED.

Capturing spatial effects can be important because the physical layout of the ED can influence the relative impact of new technologies. For example, tracking technologies or tablets may allow the staff to accomplish tasks with less physical motion. Re-arraigning the layout of the ED may have similar impacts.

Using real data including patient arrivals as well as staff and equipment levels we show how REDSim can be used to evaluate questions in the ED. We demonstrate its use for evaluating a patient selection algorithm, traditional staffing decisions, as well as showing its use in a sensitivity study.

2 RELATED WORK

Extensive research using simulation has been conducted on EDs to help mitigate its excessive number of patients, long waiting times, patients being treated in hallways, ambulance diversions, and patients leaving without treatment (Paul, Reddy, and Deflitch 2010). Discrete-event simulation (DES) which offers few if any insights on human actions and interactions has been widely used (Jun, Jacobson, and Swisher 1999) whereas only on the past decade has agent-based simulation (ABS) been studied in the ED scope. ABS is highly attactive to model EDs not only because its entities are proactive, autonomous and intelligent but also because the simulation of the interactions of these entities create opportunities for people to better understand their nature (Chan, Son, and Macal 2010).

While ABS models have the means to study provider decision making, previous work has focused only on the impact of resource changes in the number of staff, equipment, and rooms. Cabrera et al. (2012) uses exhaustive search optimization to find the optimal staff configuration. Taboada et al. (2011) have studied the impact of staff experience (a senior member finishes its tasks faster than a junior member) over patient throughput and changes on the frequency of patient arrivals and staff levels. Zeltyn et al. (2011) combine simulation and analytical formulae to solve staff scheduling problems. REDSim can study both, the effects of *resource changes* and, in case of non-mobile resources (X-Ray) it can also analyze their *best placement into the floor plan*. REDSim enabled us to study the effects of different decisions made by the ED staff.

Sensing technology could be used for staff to make decisions based on patients location. REDSim is not the first study to simulate sensing technology on the medical field to track staff, resources, and patients. Liu, Wang, and Cheng (2010) used data collected by ambulance GPS system to develop a simulation model to study opportunities to minimize ambulance response time. Fry and Lenert (2005) propose a system location monitoring system to enhance management of resources during catastrophic events. Miller et al. (2006) used RFID tags to collect simulation data and demonstrate its successful use. Lee et al. (2006) propose a system that uses location and context awareness to infer notifications for reminding physicians and nurses. None of these studies has delved into the impact of localization awareness in the provider decision process. Furthermore, all previous studies failed to address the amount of time it takes to gather resources for a task. According to our interviews, ED staff members spend a considerable amount of time looking for misplaced resources which contributes to their frustration. REDSim models not only movement and resource utilization but also human behavior.

3 REDSIM FRAMEWORK

In this section we describe our REDSim framework. Our goal was to devise a *general and flexible* model to simulate a variety of patient flows, study resource allocation and spatial disposition as well as to capture human movements and behavior.

3.1 Overview

REDSim uses a *agent-based simulation (ABS) model with a discrete-event infrastructure*. State variables change at separate instants of time at which agents can initiate actions, communicate with other agents and make decisions of their own. The terms entities and agents are used interchangeably throughout the text.

The process is modeled using a workflow approach which can be seen as a series of connected tasks involving one or more entities. The workflow represents the patient moving through the facility from the initial *arrival* task, then probabilistically transitioning from one task to the next until its final *departure* task.

Each task is divided into four-phases to model the patient waiting for a provider (**wait**), the coordinating provider gathering resources for the task (**gather**), the task being executed (**action**), and the replacement of all gathered resources during the gather phase (**replace**). Once all four phases are completed the task is retired and the patient moves to its next task.

Using the four-phase task the framework ensures that every patient in the system is part of a task and is accounted for at all times. After the patient completes a task it is immediately transferred to the **wait** phase of the next task.

The next sections describe in detail the independent parts of the REDSim framework: the *input components* that model the system architecture (e.g. patient flow and the interaction between entities), the *hooks* that models provider behavior, and the simulation execution.

3.2 Input Components

REDSim models the patient ED visit using the workflow approach where each task (nodes) represents an activity that takes place during the visit. The input components in Figure 1a specify the entities, their activities or tasks and their relationship with other entities for the ED workflow.

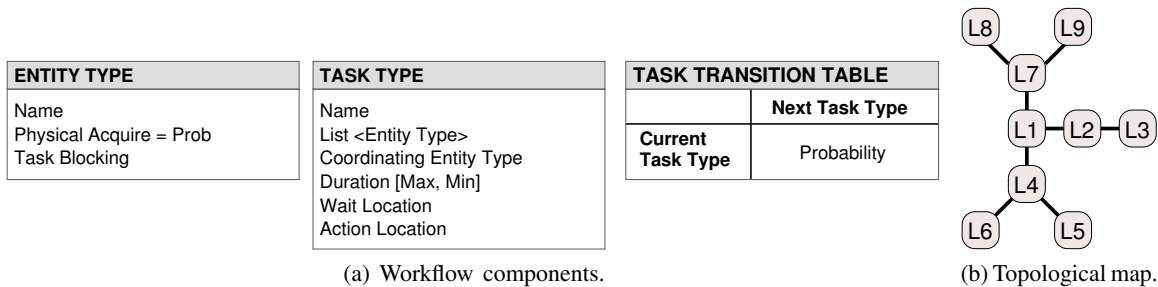
(a) Workflow components.   (b) Topological map.

Figure 1: Input components.

Entities such as patients, staff, equipment, requests and tests are specified by the **entity type** component. Each entity on the system is an instance of an *entity type*. An *entity type* with the *task blocking* attribute models entities that need to go through an activity before the patient's current task moves from the *wait*-phase to the *gather*-phase. For example, if the patient needs an EKG test, it cannot be performed until the physician enters its request into the system, therefore request entity is blocking the EKG test task. Even though the patient is at the EKG test task the task stays at the *wait*-phase until the physician enters the request. The *task blocking* attribute forces the dependency between tasks, the pharmacist cannot fill a prescription until a physician enters its request into the system.

An *entity type* with the *physical acquire* attribute models stationary resources (e.g. stretchers) and their likelihood of being returned to their original location during the task *replace*-phase. For example, at the *gather*-phase the task coordinating entity must walk to the stretcher storage location to check if it is there or not. If available it is acquired for the task and will be returned to its storage location with some probability during the task *replace*-phase. If the stretcher is not into place at the *gather*-phase the coordinating entity must walk around the facility until it finds it. The *physical acquire* attribute models the frequent situations when equipment is not returned to its proper location.

The **task type** specifies the workflow task nodes. Each task represents an activity that takes place or is related to an ED visit. It includes a set of entities (defined by the *entity type* list) that must be at the the *action location* for the activity to happen. The *task type* also specifies the task *coordinating entity type* which is responsible for gathering entities and the replacement of *physical acquire* entities to its storage position. It also specifies the range of time to be allotted for the task *action*-phase, the location where the patient is waiting at, and the location where the activity takes place. Using the triage task as an example, the nurse (coordinating entity) must conduct the patient from the waiting room (wait location) to the triage room (action location) in order for the triage activity to happen.

Once a task is complete the patient transitions to the next task according to the probabilistic **task transition table**. The table specifies the numerous patient flows available during a visit. The different flows are due to every patient having its own path of treatment, one can have only blood work done while another will have multiple exams.

To model entities' movement over the ED floor plan a **topological map** is used (Figure 1b). The map represents the connectivity of the environment in a graph structure, where vertices are distinctive locations on the floor plan and edges represent a direct path between them. To move an entity from one location to another the entity must follow the edges passing through each vertex between the two locations. To simulate an entity moving from a bed at location L4 to the pharmacy at location L9, the entity's location is updated at separate instants of time to follow the edges passing the locations L1 and L7 to arrive at the pharmacy.

3.3 Runtime Components

The REDSim framework models the operation of the ED as a discrete sequence of events which are processed in a timely order. The two types of events either change an entity's location or its status. Figure 2a has the three main runtime components. A **task** is an instance of one of the input *task type* components, it maintains its entities and events, and keeps track of which phase it is at. The **movement** event moves the entity from one location to another by taking one adjacent edge on the topological map. The **action** event models change its entity's status to busy.
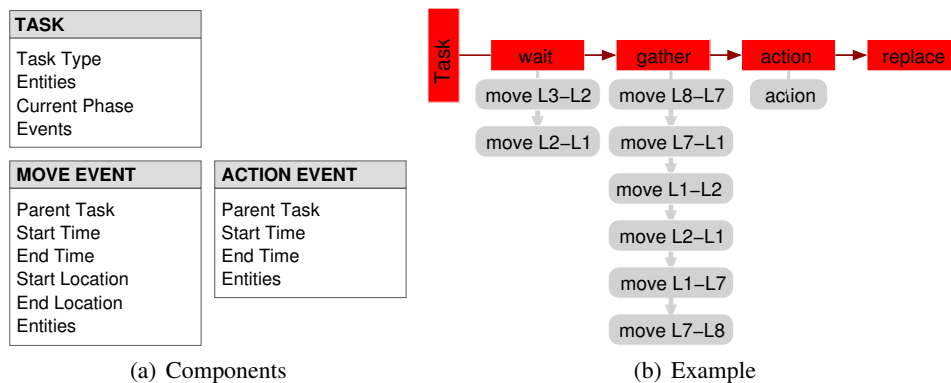

(a) Components   (b) Example

Figure 2: Runtime.

Figure 2b has the example of a triage task execution (at triage the patient waits in the main waiting aerea until a nurse calls him/her into the triage room). The execution proceeds as follows: once the triage task is created it immediately goes into the *wait*-phase. To move the patient from its current L3 location to the waiting L2 location one move event is created **(move L3-L2)**.

When a nurse (coordinating entity) becomes available it will choose one of the waiting tasks. At this point the chosen task transitions into the *gather*-phase and the nurse tries to acquire all remaining entities (patient at waiting area) and move into the action location (triage room). To acquire the patient the nurse moves from the triage room (L8) to the patient location (L2); for that 3 move events are created **(move L8-L7)**, **(move L7-L1)**, and **(move L1-L2)**. The nurse then acquires the patient and both entities move to the triage room (action location) which requires 3 more move events **(move L2-L1)**, **(move L1-L7)**, and **(move L7-L8)**. Once all entities are gathered at the action location the task transitions into its *action*-phase and the **action** event representing the triage being done is created. The triage task does not have the *replace*-phase which starts when the action event ends and the coordinator replaces all used resources to its original location. After that the task is retired.

3.4 Behavior Components

Providers (nurses, physicians and technicians) are constantly faced with decisions about what to do next, for example, which patient should be seen next.

Provider behavior is modeled by functionalities/plugins added to pre-defined hooks. At the *end task* hook it must choose what to do next and at the *pick next task* hook which order the tasks are executed.

Once a provider becomes available it chooses what to do next by selecting one of the *end task* hook functionalities. It can choose to *take a break* or to *select a task to attend to*. If the latter is chosen, one of the functionalities of the *pick next task* hook is used to choose among the available waiting tasks.

4 CASE STUDY

In this section we describe how we applied the REDSim framework to model the ED at the Jersey Shore Medical University Center (JSUMC).

Figure 3 shows a map of the ED at the Jersey Shore University Medical Center (JSUMC). The area is over 350 ft. x 200 ft. It is divided by functionally: (1) Trauma, (2) Pediatrics, (3) Fast Track/Minor Care, (4) Waiting, (5) Triage, (6) Radiology, (7) Behavioral/Crisis, (8) General exam rooms/Urgent Care, and (9) Administration. In addition, within area 8 there are further zones; the central area is for clinical staff, and the outer areas are exam rooms.
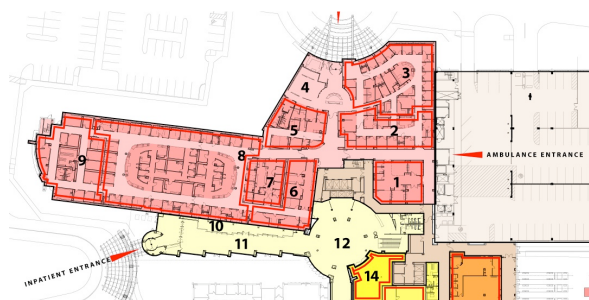

Figure 3: JSUMC Floor Plan.

We define an activity as a process in the patients' visit to the ED. Figure 4 is a simplified representation of the process and is the workflow used during evaluation. Each node in the graph is an activity, and each edge is a dependency.

When a patient arrives he is seen by a greeter, who takes the patients' name. The patient then waits to be interviewed by a triage nurse in a triage room. During triage the nurse uses the Emergency Severity
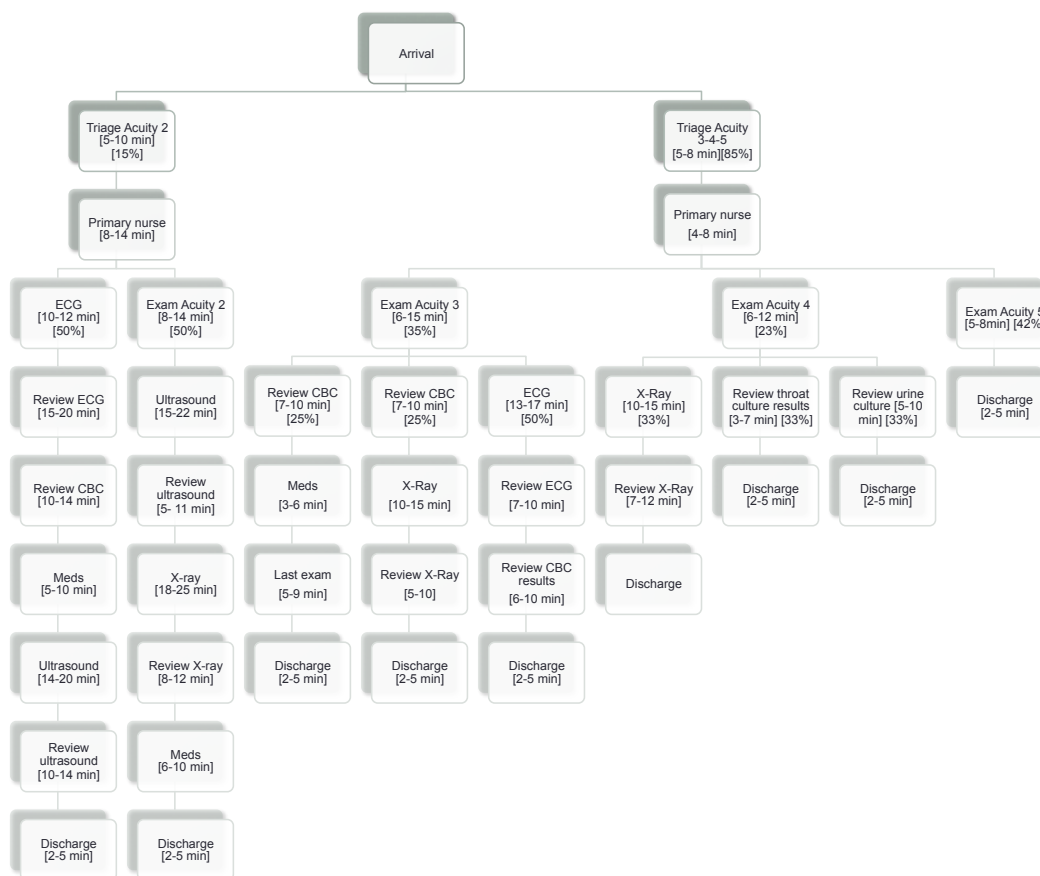
Figure 4: JSUMC Patient Flow.

Index (ESI) (Gilboy et al. 2011), which is a five-level triage algorithm that categorizes ED patients by evaluating both patient acuity and resource needs. First, the triage nurse assesses the acuity level. If the patient does not meet a high acuity level (ESI level 1 or 2), the nurse then evaluates expected resource needs to determine if the patient is ESI level 3, 4, or 5. The algorithm depicted in Figure 5 has four decision points that aid the nurse in assigning a ESI level. Patients assigned ESI level-1 have a high risk of death and require immediate physician involvement. ESI level-2 patients are high risk but stable and the primary nurse can initiate care through protocols without a physician immediately at the bedside. ESI levels 3, 4, and 5 have a lower death risk and can wait in the waiting room for the next available bed. These patients are assigned an ESI level with an estimate of how many resources the they are going to consume in order for the physician reach a disposition decision.
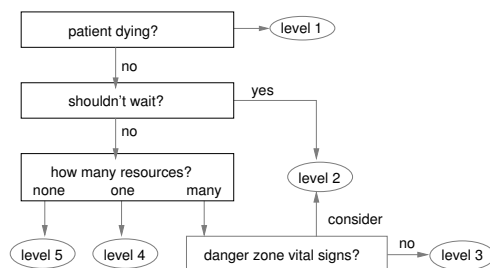

Figure 5: Emergency Severity Index Algorithm.

Once the patient is placed on a bed, vital signals are taken by a primary nurse followed by the consultation where the patient is seen by a doctor accompanied by the primary nurse. The next set of tasks depends on the patient condition as determined by the physicians' assessment. Tests are typically blood, urine and X-ray tests. When all the tests are completed, the ED physician makes a decision if the patient is well enough to be discharged, or must be admitted to the hospital.

When a patient is discharged, non-ambulatory patients require transport. Ambulatory patients can simply leave the ED with instructions for follow up care, if necessary. Admitted patients go through one of two possible paths: either their own physician is contacted for admission orders, i.e., the house doctor, or a resident is contacted. In either case, admission orders are required, which results in a hospital bed assignment. The patient hence remains boarded in the ED until the hospital makes the bed assignment.

4.1 Scheduling Policies

To study providers' behavior we have four different selection algorithms. These algorithms are added to the *pick next task* hook described in section 3.4 at which the provider chooses its next task.

In the first algorithm, Highest Acuity (HA), the provider gives preference to the patient with the highest acuity (lowest ESI level index). If there are more than one patient with the same ESI level the provider chooses the one that has been waiting the longest. In the second algorithm, Longest Waiting (LW), the providers select the next patient based on waiting time. The idea here is to maximize fairness of patients' time. In the third algorithm, Shortest Distance (SD), providers select the patients who are the physically closest, thus reducing the amount of walking during its shift. In the fourth algorithm, Random (R), patients are selected at random.

5   EVALUATION

In this section we present preliminary results of the JSUMC ED simulations. Figure 6 is a screenshot of REDSim during an execution. We measure the performance of the different decision making policies described in section 4.1 in terms of the length of stay (LOS) of patients. Lower LOS translates into higher patient throughput, thus alleviating overcrowding and enhancing the overall patient care experience. In addition, increased patient throughput using the same resources in the ED directly translates into reduced costs per patient.
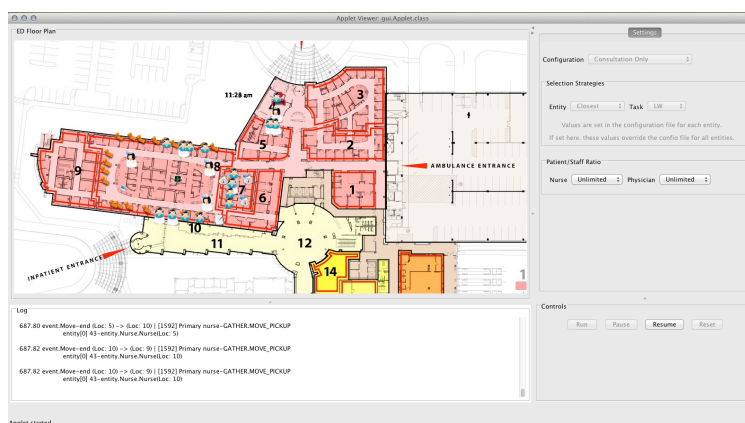

Figure 6: A REDSim screenshot during execution.

The scheduling policy Highest Acuity (HA) gives preference to tasks related to patients with the highest acuity (lowest ESI level index), Longest Waiting (LW) gives preferences to tasks that have been waiting the longest while in Shortest Distance (SD) tasks that are closest to the provider are preferred, and in Random

(R) tasks are selected randomly. We simulate patients with ESI level 2, 3, 4, and 5. ESI level 2 patients have preference over others because they are considered high risk patients.

The simulations mirror the *real values used in the JSUMC ED* for frequency of patient arrivals, staffing and equipment levels. Table 1 has the staff levels for each shift.The equipment levels are as following: 30 beds, 2 mobile EKG, 1 X-Ray, 1 ultrasound, and 4 workstations. Workstations are used to enter test requests, review test results and choose the next patient to be seen. Except for physicians, all personnel leave the ED at 7 am, at 11 am and at 11 pm. Physician shifts overlap in order to have higher number of attending physicians during peak hours (11 am to 9 pm). At 7 am there are two physicians, at 11 am there are three and at 4 pm there are five attending physicians.

Table 1: Staff levels

Entity Type	Count							
	7am-11am	7am-5pm	11am-9pm	11am-11pm	2pm-12am	3pm-1am	9pm-7am	11pm-7am
Triage Nurse	2	-	-	3	-	-	-	2
Nurse	6	-	-	9	-	-	-	6
Physician	-	2	1	-	1	1	1	1
PCA	4	-	-	6	-	-	-	4
Pharmacist	1	-	-	1	-	-	-	1

Figure 7 shows the frequency of patient arrivals. From midnight to 1 am, eleven patients arrive, from 1 am to 2 am one patient arrives, from 2 am to 3 five patients and so forth. Patients arrive at random times within the hour period to simulate the unpredictable nature of the ED arrivals. The time of arrival for each patient is the same throughout executions. ESI level 2, 3, 4 and 5 constitute 15%, 30%, 20%, and 35% of ED patients (Gilboy, Tanabe, Travers, and Rosenau 2011). As of now we do not simulate ESI level 1 patients that are approximately 1% of ED patients. Instead we simulate these patients as ESI level 2.
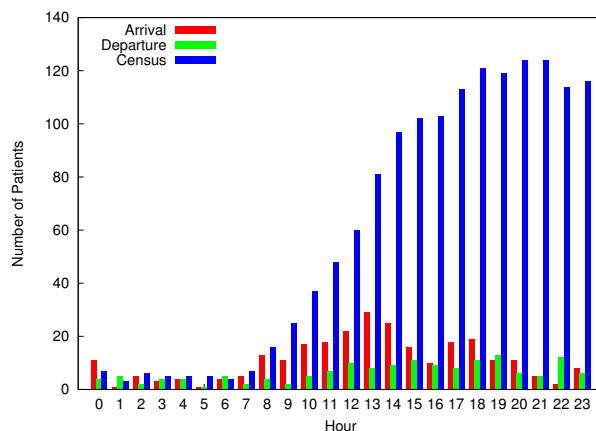
Figure 7: Number of arrivals, departures and the total number of patients currently in the ED.

Each execution simulates the ED over a 24 hour period. At the end of the 24th hour the execution is terminated even if not all patients have departed.

Figure 8 shows two simulations, in **Starving tasks** the next task is chosen according to the scheduling policies in place, there are no exceptions. On the other hand, **No starving tasks** aims to leave no task waiting longer than 2 hours even if it has a related high acuity patient. For example, if amongst the waiting tasks there are two X-Ray requests, one for a patient of ESI level 5 waiting for 2 hours and the other for a patient of ESI level 2 waiting for 1 hour the former is executed first.

(a) Patient Length of Stay (LOS) (b) Discharged Patients (c) Time to see a doctor
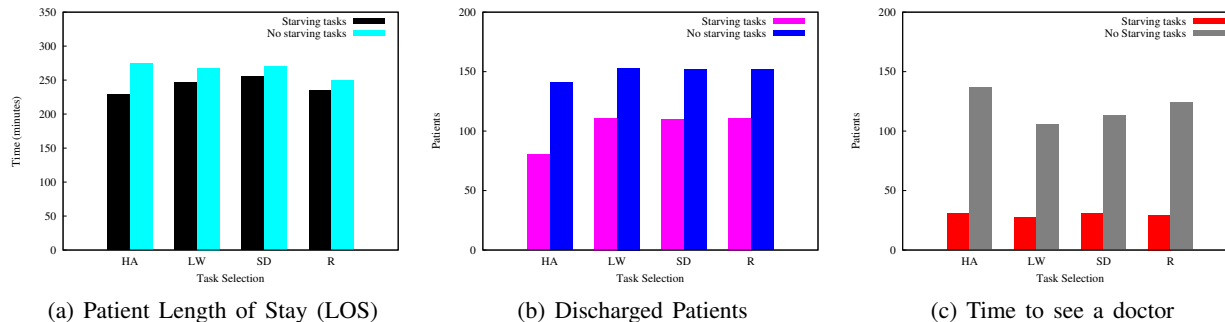
Figure 8: Trade-off of favoring tasks that are waiting for more than 2 hours.

Figure 8a shows that HA policy has the lowest LOS for the *Starving tasks* simulation. Patients stay on average 26 minutes (10%) less when compared to SD. But giving priority to high acuity patients leaves others starving, decreasing the number of patients discharged. HA discharges 40 patients (25%) less than SD which can drastically increase the number of patients leaving the ED without being seen (LWBS rate). The number of patients leaving without being seen is one of the metrics used by ED to measure its performance. Therefore, lower LOS does not translate into higher patient throughput and in fact it if not careful it can increase mortality.

While SD is not the ideal policy to reduce LOS or increase patient throughput it may help reduce nurse burn out and increase patient satisfaction. Patient satisfaction levels are lower in hospitals with nurses that are dissatisfied or burned out; improving their working conditions may improve quality of care (McHugh, Kutney-Lee, and Cimiotti 2011). According to our simulations, nurses walk on average 300 more steps than physicians throughout their shift. Making small improvements such as reducing the amount of walking may contribute to an increase in job satisfaction.

When preference is given to tasks that are waiting longer than 2 hours the LOS increases as seen Figure 8a. In the worst case (HA) LOS increases 20% while at the same time the number of discharged patients increases by 76%, which in some situations may be a good compromise. Figure 8c shows that the higher LOS is not the only downside of not letting tasks starve, the time to see a doctor for ESI level 2 patients increases by 105 minutes (70%). This increase may translate into higher mortality rates since high acuity patients have to wait much longer to see a doctor.

Although LOS is not significantly reduced (1.4%) the throughput increases 17% when comparing SD with HA. Therefore, on top of helping reduce nurse burn out it increases throughput.

In EDs things change very rapidly, while in one minute a scheduling policy may be the best choice in the next a high acuity patient arrives and the policy may become the worst option. Emergency departments need a system to analyze the trade-offs given a set of patients and conditions to determine which is the next task to be executed in order to improve the outcome.

5.1 Sensitivity of Results

To better understand the relationship between input and output variables we conducted simple variations on the input. We use the **No starving tasks** execution as our base case for comparison. We varied staff levels and the time a task action takes. For the first variation we reduced 1 nurse from the 7am - 11am and 11pm - 7am and 2 nurses from 11am - 11pm. There is little output variation, as seen in Figure 9. What happens is that most of the time patients are waiting for tests (Lab work, X-Ray, Ultrasound) to be done and results analyzed.

For the second variation we increased in 5 minutes the time it takes for test results to be analyzed. Three of the scheduling policies had an increase in LOS while all of them had the throughput reduced (Figure 10). The longer the tests take the longer the patient occupies a bed the lower is the patient throughput.
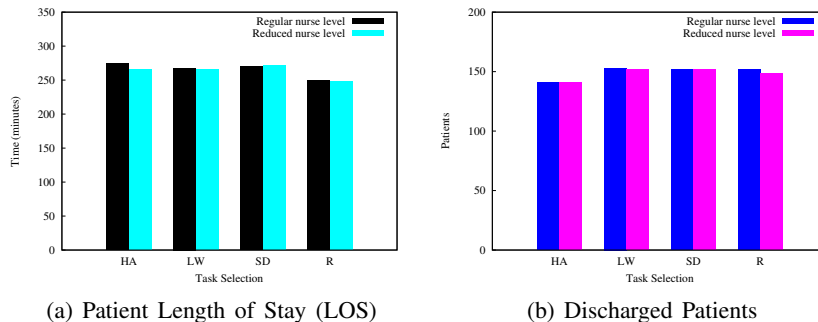
(a) Patient Length of Stay (LOS)   (b) Discharged Patients

Figure 9: Impact of reducing nurses level.

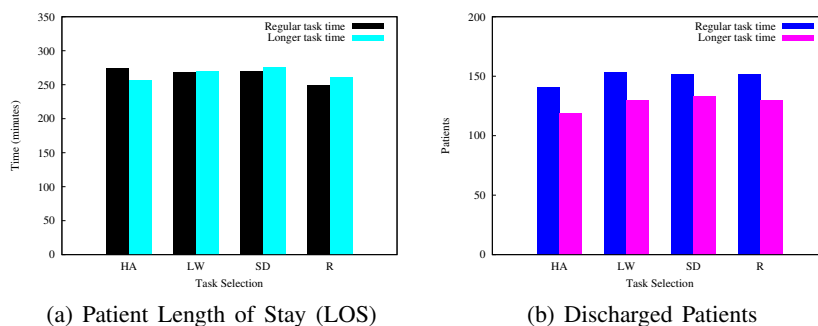(a) Patient Length of Stay (LOS)   (b) Discharged Patients

Figure 10: Impact of increasing task time.

5.2 Validation

To validate REDSim we compare the mean service time for a 24 hour period of simulation against the real ED. Figure 7 shows the number of arrivals, departures and the number of patients in the system for a time interval T, and it can be interpreted as a queuing system. Let $a(t)$, $d(t)$ and $c(t)$ be the number of arrivals, departures and patients in the system at time t. The arrival rate $\lambda(T)$ is the total number of arrivals $A(T)$ divided by T, and the departure rate $\delta(T)$ is the total number of departures $D(T)$ divided by T.

In equation 1, we use Little's law (Jain 1991) to compute the mean service time. Little's law only applies if the number of arrivals equals the number of departures. To compute the time at which all patients have left ($TTotal$) we must drain the remaining patients in the ED at the end of the experiment. Therefore, $TTotal = \frac{remaing\,patients}{\delta(T)} + T$. We can compute the mean number of patients dividing the total number of patients in the system ($C(T)$) by $TTotal$.

$$Mean\,service\,time = \frac{Mean\,number\,of\,patients}{\lambda(T)} \qquad (1)$$

While the simulation LOS for the LW policy is around 57 minutes the real ED LOS is 138 minutes. We did not expect these values to match since we are only simulating the main ED while the hospital uses a fast track for ESI level 4 and 5 patients and an additional track for children that are not included in the simulation. Instead we use the entire patient frequency of arrivals for the main ED. We also do not have the information of how patients are distributed among the ESI levels, we used an estimate from (Gilboy, Tanabe, Travers, and Rosenau 2011).

JSUMC uses the time from arrival to triage, to bed assignment, to be seen by a doctor as well as LOS as performance metrics. We use these values to aid in the simulator validation as well.

Running experiments with the complexity of visits and accurate patient distribution among ESI levels to reflect the JSUMC ED operation is part of the future work.

6   CONCLUSION AND FUTURE WORK

In this work we presented REDSim, a simulation framework focused on investigations of the decisions made by ED staff members. We showed it can also analyze more traditional scenarios that examine resource optimization problems, such as staffing levels, as well.

Initial experiments show that the order in which physicians execute tasks has a direct impact on LOS and that a lower LOS does not necessarily increase patient throughput. We also note that there is no scheduling policy that works for all scenarios. We conclude that the ED would benefit from a system to provide heuristics that the staff can apply when the ED is in different situations.

An open question is if fine-grained agent-based spatially-aware simulations are needed to accurately model decision processes of ED staff. While our simulation approach is not overly difficult, building such a simulator requires significant programming effort that might not be available in many environments. Future work will parameterize the tasks, staffing level, and spatial layout of the simulation to allow for testing different scenarios, and possibly running simulation on different EDs.

A next step is also implementing a trace-driven approach to validation. In this strategy, the outputs of the simulator should match the measured times from a trace of the actual ED. That is, given a real set of patient arrivals, discharges, staffing levels and ED layouts, the simulated LOS and waiting times should closely track the actual times in the ED. Fortunately, many EDs have electronic systems where patient arrivals and departures are manually entered, thus making an trace-driven validation tractable.

Finally, REDSim would allow us to test the impact of various technologies before deployment. For example, we could examine if tablets improve LOS and reduce waiting times because staff move through the ED less. The find-grained spatial component of RED could also allow us to quantify the amount of movement that was reduced. We are also planning an experiment to see if technology that flags tasks that are taking excessive time can improve waiting time and patient throughput. For example, tracking technology could identify if lab samples and results are held up, allowing staff to take action. If such delays could be reduced, it may greatly improve average waiting times.

REFERENCES

Brenner, S., Z. Zeng, Y. Liu, J. Wang, J. Li, and P. K. Howard. 2010, July. "Modeling and Analysis of the Emergency Department at University of Kentucky Chandler Hospital Using Simulations". *YMEN* 36 (4): 303–310.

Cabrera, E., E. Luque, M. Taboada, F. Epelde, and M. L. Iglesias. 2012. "ABMS optimization for emergency departments". In *Proceedings of the 2012 Winter Simulation Conference*, edited by C. Laroque, J. Himmelspach, R. Pasupathy, O. Rose, and A. M. Uhrmacher, 1–12.

Chan, W., Y.-J. Son, and C. Macal. 2010. "Agent-based simulation tutorial - simulation of emergent behavior and differences between agent-based simulation and discrete-event simulation". In *Proceedings of the 2010 Winter Simulation Conference*, edited by B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, and E. Yucesan, 135–150.

Committee on the Future of Emergency Care in the United States Health System 2007. *Hospital-Based Emergency Care: At the Breaking Point*. National Academies Press.

Duguay, C., and F. Chetouane. 2007, April. "Modeling and Improving Emergency Department Systems using Discrete Event Simulation". *Simulation* 83 (4).

Fry, E. A., and L. A. Lenert. 2005. "MASCAL: RFID tracking of patients, staff and equipment to enhance hospital response to mass casualty events.". In *Proceedings of the 2005 AMIA Annual Symposium*, 261–265.

Gilboy, N., P. Tanabe, D. Travers, and A. M. Rosenau. 2011. *Emergency Severity Index (ESI): A Triage Tool for Emergency Department Care*. Agency for Healthcare Research and Quality.

Jain, R. 1991. *The Art of Computer Systems Performance Analysis*. John Wiley and Sons. Parts IV and V.

Jun, J. B., S. H. Jacobson, and J. R. Swisher. 1999, February. "Application of Discrete-Event Simulation in Health Care Clinics: a Survey". In *Journal of the Operational Research Society*, 109–123.

Lee, S.-w., S.-y. Cheng, J. Y.-j. Hsu, P. Huang, and C.-w. You. 2006. "Emergency Care Management with Location-Aware Services". In *Proceedings of the 2006 Pervasive Health Conference and Workshops*, 1–6.

Liu, J., X. Wang, and E. Cheng. 2010. "Simulation modeling and analysis on asset planning for Emergency Medical System (EMS)". In *Proceedings of the 2010 IEEE International Conference on Industrial Engineering and Engineering Management (IEEM)*, 1353–1357.

McHugh, M., A. Kutney-Lee, and J. Cimiotti. 2011. "Nurses' widespread job dissatisfaction, burnout, and frustration with health benefits signal problems for patient care". *Health Affairs* 30 (2): 202–210.

Miller, M., D. Ferrin, T. Flynn, M. Ashby, K. White, and M. Mauer. 2006. "Using RFID Technologies to Capture Simulation Data in a Hospital Emergency Department". In *Proceedings of the 2006 Winter Simulation Conference*, edited by L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, 1365–1371.

Nawar, E. W., R. W. Niska, and J. Xu. 2007, June. "National Hospital Ambulatory Medical Care Survey-2005 Emergency Department Survey; Advance Date for Vital Health Statistics". (386).

Paul, S. A., M. C. Reddy, and C. J. Deflitch. 2010, August. "A Systematic Review of Simulation Studies Investigating Emergency Department Overcrowding". *Simulation* 86 (8-9).

Reid, P. P., W. D. Compton, J. H. Grossman, and G. Fanjiang. 2005. *Building a Better Delivery System: A New Engineering/Health Care Partnership*. 500 Fifth Street, N.W, Washington DC 20001: National Academies Press, also available www.nap.edu/catalog/11378.html.

Taboada, M., E. Cabrera, M. L. Iglesias, and F. Epelde. 2011. "An Agent-Based Decision Support System for Hospitals Emergency Departments". *Proceedings of the 2001 International Conference on Computational Science* 4:1870–1879.

Zeltyn, S., Y. N. Marmor, A. Mandelbaum, B. Carmeli, O. Greenshpan, Y. Mesika, S. Wasserkrug, P. Vortman, A. Shtub, T. Lauterman, D. Schwartz, K. Moskovitch, S. Tzafrir, and F. Basis. 2011, August. "Simulation-Based Models of Emergency Eepartments: Operational, Tactical, and Strategic Staffing". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 21 (4).

AUTHOR BIOGRAPHIES

ANA PAULA CENTENO holds a B.S. in computer science from Catholic University of Pelotas in Brazil, a M.S. in computer science from Sunny at Stony Brook and is working toward her Ph.D. degree at Rutgers University. Her research interests include operating systems, distributed and parallel systems and sensor networks. Her email address is anapaula@cs.rutgers.edu.

RICHARD P. MARTIN is an associate professor of computer science at Rutgers University and a member of the Rutgers Wireless Network Information Laboratory (WINLAB). His current research interests include wireless sensing and device localization, and advanced scheduling. He received a B.A. from Rutgers University and a, M.S. and Ph.D. in computer science the University of California at Berkeley. His email address is rmartin@cs.rutgers.edu.

ROBERT SWEENEY is the Chairman of Emergency Medicine at Jersey Shore University Medical Center and the Regional Medical Director of EmCare, Inc. He is a graduate of Villanova Univesity, B.S. Biology, CMDNY-School of Osteopathic Medicine, and Monmouth University MBA program. His email address is drrsweeney@meridianhealth.com.