Agent-Based Simulation of Emergency Departments with Patient Diversion

Marek Laskowski[1] and Shamir Mukhi[2]

[1] Internet Innovation Centre,
Dept. Electrical and Computer Engineering, University of Manitoba
R3T 5V6, Canada
mlaskows@ee.umanitoba.ca
[2] The Canadian Network for Public Health Intelligence (CNPHI)
1015 Arlington Street,
Winnipeg, MB R3E3R2, Canada

Abstract. We describe an agent based model of an emergency department and its utility for evaluating workflow and assessing patient diversion policies. The overall goal of the research is to develop tools to better understand and manage emergency departments. There are several modes in which agent based modeling tools may be of benefit. In a self contained manner the operation of an emergency department can be modeled. In this mode, policies such as staffing could be changed and the effect on parameters such as waiting times and throughput could be quantified. In an extended version, multiple emergency departments can be modeled and would allow for the evaluation of ambulance or patient redirection policies. In either case we also suggest an effective means of augmenting the simulation with empirical data collected using a proximity location and tracking system within an emergency department. Our agent based model allows for a simulation of a number of emergency departments and introduces a method of extracting real time patient data from emergency departments throughout a city allowing for the evaluation of patient diversion policies.

Keywords: Agent based modeling, waiting time reduction, ambulance redeployment, emergency department, and bio-inference.

1 Introduction

Hospitals represent one of the most promising areas where agent based modeling may be seen as an effective tool in evaluating policy and improving efficiencies. In many cases, the operations of an emergency department are over taxed and may not be guided by optimal policies. Although policies evolve over time and efforts are made to reduce wait times etc. often there is little quantitative analysis or feedback in the process. Using agent based modeling an emergency department can be modeled and in this manner better use of resources can be made possible through identification of anomalies and bottlenecks which may be difficult to detect otherwise.

In addition, technologies are emerging that will be leveraged by Hospitals to improve patient care. Two of the more obvious technologies and applications include

D. Weerasinghe (Ed.): eHealth 2008, LNICST 1, pp. 25–37, 2009.

inter Hospital tracking and internetworking. These technologies will also allow for a more distributed approach to managing a number of interacting emergency departments. Along these lines the emergency department simulator can be used to evaluate ambulance redirection or diversion policies.

The research presented here provides a specific emergency room example application and an extension to a wide area Hospital/ED/Ambulance and patient diversion scenario. Section 3 discusses work completed to date on an open source visualization, simulation, and wait time forecasting simulation suite. Section 4 presents simulation results for the scenarios discussed above, for use in evaluating workflow within an ED, and for use in patient diversion policies.

2 Applications

The applications discussed below are of increasing interest here and elsewhere. Although agent based modeling of an emergency department is of utility on its own [1], it is of considerably more utility when combined with tracking technologies and networking capabilities. Specifically our agent based model is a distributed model across a number of regional hospitals with emphasis on utilizing data collected in real time. Another novelty is the use of congestion avoidance algorithms from Telecommunications Engineering redeployed as a model for evaluating patient diversion policies. The following applications are included as context for which the emergency department simulator is well suited.

2.1 Local Area Patient Tracking

An application we are working on is directed towards monitoring or estimating the location of patients as they enter an emergency room until treated and released. This is an extremely active area of research and concern for both the general public as well as practitioners. In the following scenario, emergency department simulation is enhanced with empirical data collection technologies [2].

In one scenario a patient is tagged with an RFID tag when registering at emergency, an RFID reader at the desk records the event. As the patient enters the waiting area another reader reads this event. The event and reader location are recorded at an Hhost (Hospital host subsystem). Data logged by the Hhost (timestamp, tag and reader ID) would be uploaded via the nearest Wi-Fi access point to be relayed to CORE (Central Observation and Reporting Environment). Once stored at the CORE it can be made available for augmenting the agent based simulator, mining or made available over the Internet to other stakeholders. In certain circumstances the information can be made available to the public such that they may be better informed as to which emergency department or clinic they may choose to attend. A specific example would be information concerning wait times at facilities for injuries such as broken bones and whether or not a physician was on staff to set a break at that time.

Figure 1 illustrates an emergency department with RFID readers deployed strategically throughout the area capable of reading and logging patient tags for subsequent uploading to a more central network service.

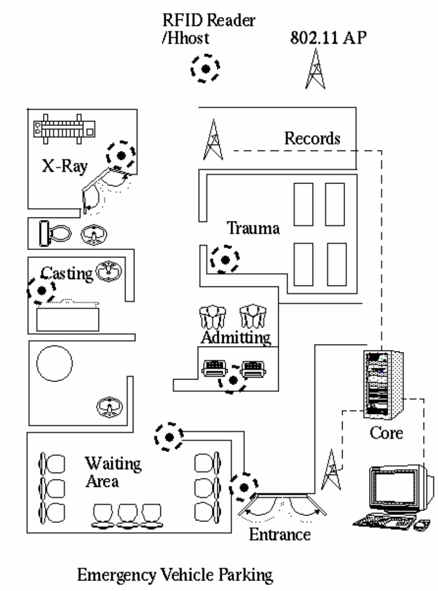

Fig. 1. Schematic of tagging locations in a Hospital Emergency Facility

In the future it is also conceivable that a patient's tag could be updated during treatment such that a more complete record of a where bottlenecks occur could be extracted. For example, upon receiving an X-ray the RFID tag could be programmed to store this information and subsequently uploaded to CORE when read by the next reader encountered. The tag could also be programmed at admitting indicating the general type of aliment or complaint for more complete although anonymous data collection and analysis.

2.2 Wide Area ED Scenario

In a wide area hospital/ED/ambulance scenario the basic CORE functionality would scale and encompass alternative communication modalities. At present, our prototype CORE is designed to support wide area networking having evolved from a geographically dispersed event management system [3]. Also as CORE is IP centric, security issues are largely resolved using VPNs and encryption standards associated with IP. The following figure illustrates the basic components and agents in the extended framework.

Figure 2 illustrates a wide areas scenario incorporating participation hospitals and emergency service vehicles.

In the wide area scenario (Figure 2) each hospital emergency department would be equipped as in Figure 1 with data extracted from RFID proximity location systems augmenting the agent based modeling.

More proactive modeling extensions include the ability to notify and receive information from ambulances and other emergency vehicles. The actual communication services likely would be over GSM or similar communication infrastructures. Here

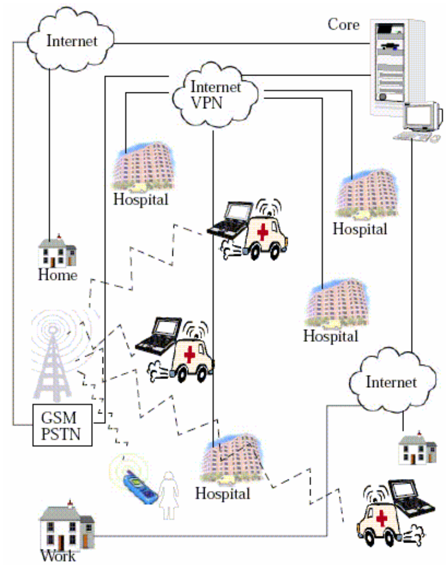

Fig. 2. Wide area deployment of the framework illustrating the major stakeholders or agents

these services can be modeled as messages between agents and the platform would be used to assist in optimizing ambulance diversion policies. Other types of considerations required would be in the estimation of travel time as these factors would significant in an effective model. Although not addressed here, with the proliferation of GPS and mapping technologies, these estimates can again become empirical inputs to the multiple emergency department simulation.

At present ambulance diversion is principally based on best effort reporting and operating in good faith based on regional guidelines, an example of which can be found at [4]. We suggest that in addition to these heuristics, emergency department modeling can benefit from algorithms more commonly associated with the Internet and congestion avoidance schemes that deal with overcrowding of routers. As an example, we have adapted the Random Early Detection (RED) algorithm [5] as a candidate for consideration when attempting to optimize ambulance or patient redirection. This is an ideal initial algorithm for adaptation as it has many of the attributes well suited to improving system throughput. RED accommodates limited bursts and can be effective even when there is limited sharing of information between emergency departments. In our case, we would be using RED as a model for redirection based on emergency department congestion information being made to potential patients/ ambulances to model the redirection.

3 Visual Simulation Suite

As mentioned in section 2, we are developing an object oriented (OO), open-source, visual simulator which could be used with data gathered from the previously discussed architecture, and may be used to analyze and forecast emergency department waiting times. The simulator was written using C++ and makes use of the Qt4 API for cross platform windowed applications [6]. By virtue of Qt4 itself being open-source as well as our code, once the source is released in Beta stage, our project will benefit from other researchers customizing and extending our code. Neither of these things would be possible had we used an off-the-shelf proprietary solution instead of an open-source paradigm. Qt4 also allows us to deploy the simulator on Windows, Mac, or Linux. A screenshot of the simulator window is shown in Figure 3. The spatial aspect of the visualization reflects the spatial nature of the underlying data we will collect.

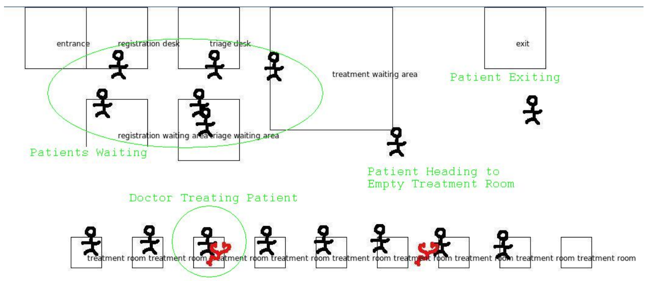

Fig. 3. Screen Capture of the Basic Simulator

3.1 Emergency Department Model

The aspects of the emergency department treatment process which we are interested in modeling are depicted in Figure 4. Patients arrive either by ambulance or walk in. Patients in need of immediate care are sent straight through to the treatment area. All ambulance arrivals are considered to be in need of immediate care, as well as some small fraction of walk-in patients.

Walk-ins that do not require immediate care proceed to the registration desk. If the registration desk is busy with an earlier arrival, the arriving patient then waits in a queue. Once the registration process is complete, the patient proceeds to the triage

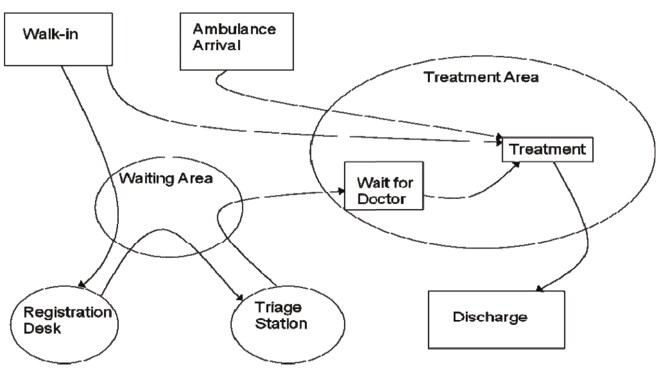

Fig. 4. Model of Emergency Department Patient Service

station. Again, if the triage station is busy, the patient waits in a queue. The nurses at the triage station will assign the patient a priority based on the severity of their condition. The arriving patient then waits with other patients in what is effectively a priority queue to be assigned a treatment room.

Once assigned a treatment room the patient waits for a doctor to come around in order to receive treatment. It is assumed that doctors will treat the patients in order of urgency, then in order of arrival. Upon completion of the treatment, the patient leaves the system, and both the treatment room and the doctor become available for another patient.

3.2 Basic Architecture

The simulator maintains an instance of the erWorld class which keeps references to all the simulated objects described in this section. Each erWorld represents one emergency department. To implement the Wide Area Scenario to simulate and investigate ambulance diversion policies suggested in section 2.2, and further discussed in section 4, we would have numerous instances of erWorld. The flexibility and reuse of code made possible by the OO architecture makes this possible by subclassing or extending existing classes to allow communication between instances of erWorld. As mentioned, each erWorld maintains a collection of patient generators, agents representing nursing stations (registration, triage), patient agents, and doctor agents which represent corresponding aspects of the model discussed earlier. A special erController agent is used to mediate patient flow through the emergency department process. Creating subclasses of erController is necessary to be able to handle variations on the basic emergency department processes in order to reflect different policies for individual emergency departments we are trying to model. For example, one would create a subclass of erController for an emergency department that allows for bedside registration for all patients versus an emergency department that requires most patients to register at a desk (as per current implementation). The classes that represent doctors and the nursing stations may also be subclassed to reflect procedures that vary between emergency departments. Patients too can be subclassed should the need arise.

As shown in Figure 3, we have the ability to place functional areas of the emergency department at arbitrary locations. However, it is our goal to eventually overlay the locations on the actual floor plans of the emergency departments we wish to simulate.

At every simulated time step all relevant agents are refreshed. For example, patients move between nursing, waiting, and treatment areas, as well as track the time spent in each activity. Doctors move to occupied treatment rooms and treat the patient within. Nursing stations count down the time required to process patients for the relevant activity. Patient generators model a Poisson arrival process for each patient class (i.e. classified by urgency of care required). At each time step each decides whether to introduce a new patient. Currently, since we have yet to collect data, we base our arrival rates and service times on estimates obtained by other researchers [7]. However, it is our goal to have the simulator driven by real world data, collected in real time, such as that proposed in section 2 where RFID proximity location system data would be used to augment the simulator.

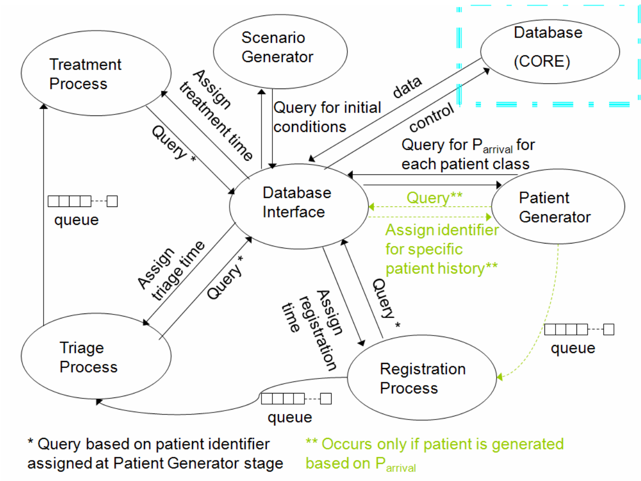

Fig. 5. Data Flow for Simulation

3.3 Data Driven Simulation

Once collected, and after some additional processing, data can be used to drive the simulator by employing an intermediary database interface class, which insulates the simulator code from the specific database implementation used. Processing may include sanitizing the data to ensure patient anonymity, and determining the arrival rate for the appropriate time of day and time of year. Figure 5 illustrates how the gathered data may be used during the course of the simulation.

When the simulation is started, the state of the simulated emergency department is seeded with the current state of the actual emergency department as inferred from the most recent data sent to CORE from the MHosts. The Patient Generator periodically queries for current arrival rates for each simulated patient class as illustrated in Figure 5. Those rates are used at subsequent time steps to determine whether a patient of a particular class arrives at that time step. If a patient arrives the agent is assigned a random patient history of the appropriate patient class upon creation and that history is used to determine service and treatment times for that patient. Note that times spent in queues are determined by the current state of the simulator, i.e. the number of patients waiting ahead of that patient, rather than the patient history.

In order to forecast future patient wait times, the simulation can be run into the future a number of times, keeping track of the wait times experienced by patients arriving at future times – until some reasonable level of confidence is reached. During this process the visualization can be disabled in order to speed multiple trials.

Prediction based on modeling and simulation is extremely difficult and potentially error prone. Confidence can be enhanced as the system is in operation and predictions tracked. Our conjecture is that the model of individual or interacting emergency departments augmented with whatever available empirical data is available would still be better than open loop policy decision making.

4 Simulations

In this section we describe two experimental results that suggest the utility of the discussed simulator for making informed policy decisions. The first scenario we investigate, while simple, illustrates the effects a policy decision has, such as changing staffing levels, using multiple performance metrics. In the second scenario, despite not having an actual RFID data collection system in place, we attempt to model what impact it would have if the infrastructure were in place to gather and disseminate ED utilization in real time. This would have the effect of informing the Ambulance Service as well as individual citizens (perhaps through a web portal) of the real-time status of EDs in the city in order to make better informed decisions about which ED to visit (based on current and projected wait-times).

4.1 Staffing Change Scenario

In this scenario we simulate the basic ED scenario mentioned above, and with Triage Classes, Service Times, and Patient Arrival rates based on [7]. We compare 3 different staffing scenarios of 2, 3, and 4 doctors working in the ED. The simulation was allowed a "warm-up" period of 24 hours, then observations were made during the following 24 hours. Ten independent trials were run, treatment queue length and doctor utilization were averaged, while individual patient waiting times are shown unaggregated. These results are presented in figures 6, 7, and 8, respectively.

From these figures we see that intuitively they seem reasonable. Figure 6 shows the average number of patients waiting for treatment as a function of time (in seconds). For example, as seen in the case of an ED with two doctors the patient queue continually increases with time as the ED is clearly under-staffed. In an ED staffed with four doctors

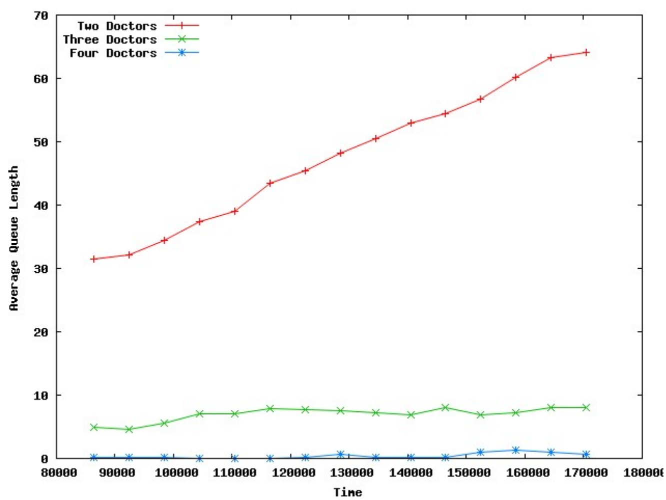

Fig. 6. Average Queue Lengths for Varying Number of ED Doctors on Duty

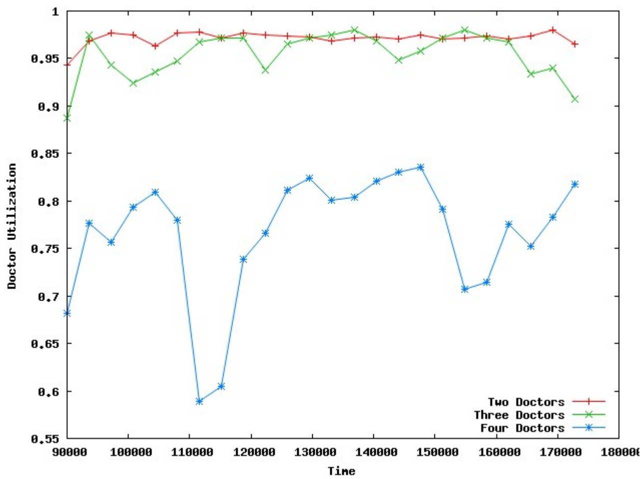

Fig. 7. Average Doctor Utilization for Varying Number of ED Doctors on Duty

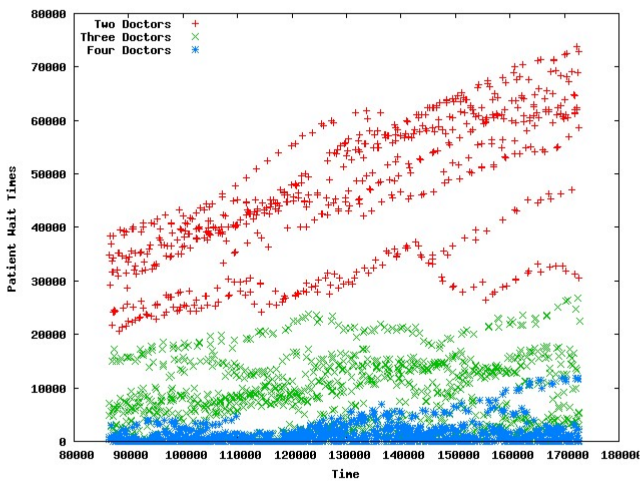

Fig. 8. Patient Waiting Times for Varying Number of ED Doctors on Duty

the patient queue is nearly zero however, Figure 7 illustrates that the doctors are under-utilized. This suggests that for this scenario, unless it is critical to have nobody waiting for service, resources would be better allocated elsewhere rather than adding a fourth doctor. It is interesting to note that in Figure 8 each data point is an individual patient, albeit simulated, passing through the ED - in which case outliers would perhaps correspond to patients which waited an unusually long or short time. For policy makers, it may be easier to empathize with individual patient cases than with a standard deviation or other such statistical measure.

4.2 Data Infrastructure Scenario

In this section we attempt to model the impact of collecting real-time ED status using the RFID based framework discussed in section 2, then making available the real-time ED status to Ambulance operators and the public at large on a city-wide scale. Since this system is not yet in place, we borrow from computer network engineering the well established technique of Random Early Detection (RED) [5] to model this process. Random Early Detection is a method of network congestion management, whereby senders of data over the network (typically the Internet) are implicitly notified of network congestion by having their data packets (data over the Internet is divided into discrete chunks, called packets) probabilistically dropped from network queues. To avoid oscillation between intense bursts of traffic and choking off traffic entirely, the rate at which these packets are dropped is ramped up gently after a certain threshold in the queue length is reached. Similarly, in our ED model, we set a minimum threshold, below which ED queue lengths are considered acceptable by everyone and no dropping occurs. The rate at which patients are dropped increases linearly with queue length until some maximum threshold is reached, past which the drop rate remains constant. Since we consider "dropping", or turning away patients, as unacceptable our model instead considers a drop to be a patient being redirected to another ED. The mechanism for this is either self-redirection to another ED or an Ambulance being redirected by a central dispatcher. Two modes are considered, one where patients are redirected to a random ED with uniform probability, and one where patients are probabilistically redirected to an ED based on the ratio of doctors to patients waiting. In the latter case, this results in EDs that are less busy having a higher likelihood of patients being redirected there. This reflects an assumed patient preference for shorter waiting time, and also demonstrates the utility of having city-wide ED status information disseminated. This is contrasted with the former case, where patients are simply guessing as to which ED is a more preferable alternative without any guidance whatsoever.

To ground our simulations as much as possible in reality, we made use of a report on Emergency Department usage in Winnipeg released just prior to this writing by the Manitoba Centre for Health Policy [8]. There were 185,659 ED visits in Winnipeg among 6 Hospitals, the breakdown of which by CTAS [9] triage level roughly corresponds to the triage levels used in [7]. Since we don't have data on treatment times we thought it plausible to use the distribution of treatment times based on triage level from [7]. To date, we have no data on variation in patient arrival rate, based on time of day, day of the week, time of year, or variation between individual EDs so we assume that the rates are uniform for these variables. It should be noted that these types of variations can be easily incorporated in the simulator. With the information presented above, it was possible to estimate arrival rates of patients for each triage level at each simulated ED. An arbitrary but reasonable minimum threshold of 10 patients waiting in the queue was chosen for the RED model. The drop or redirection rate increases linearly to a maximum of 50% reached at a queue length of 20. It is not unreasonable to assume that staffing levels at each ED do not match demand. Because arrival rates are uniform among the simulated EDs, to make the simulation interesting, two EDs are staffed with 2 doctors, two EDs have 3 doctors, and two EDs have 4 doctors on staff during the simulation.

As in the previous section, three 24 hour scenarios were investigated with ten trials each, and a warm up time of 24 hours. For comparison, the first scenario, which we call No Redirection, assumes that there is no ED status information available and that patients are better off going to the nearest ED and remaining there no matter what the wait. The second is the mode where redirection occurs based on the discussed RED model, and the destination ED is chosen from a uniform probability, we refer to this as Random RED. The third is the form of RED redirection where EDs with lower expected waiting time are probabilistically chosen more often as the destination ED. We refer to this case as Guided RED.

Un-aggregated patient wait times are not shown for these scenarios. The reason is because of the disparity between ED conditions, patient wait times vary wildly. However,

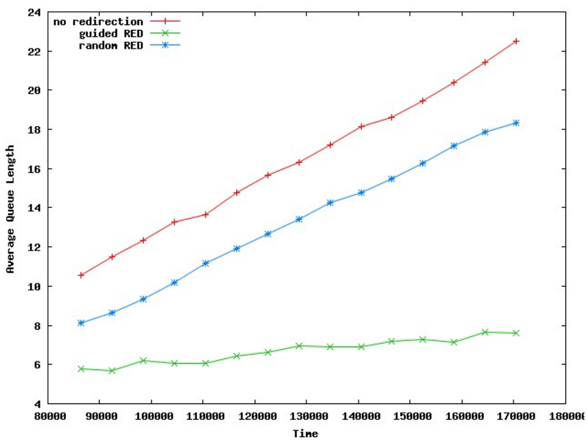

Fig. 9. Average Queue Lengths for Various Redirection Policies

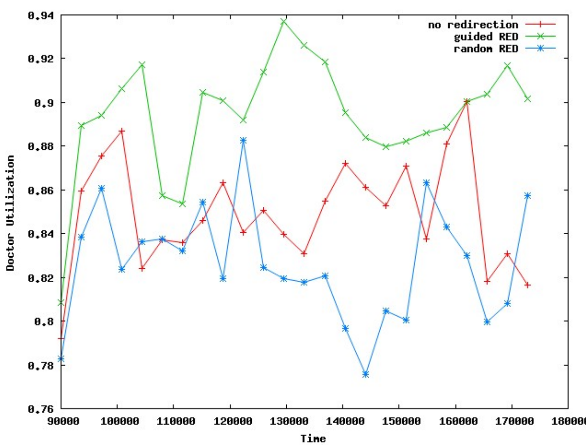

Fig. 10. Average Doctor Utilization for Various Redirection Policies

we can see from Figure 9 that average queue lengths among all hospitals are shortest for Guided RED. Also, in Figure 10 overall doctor utilization is highest in the Guided RED scenario. The improved doctor utilization results in significantly reduced queue lengths or equivalently reduced patient waiting times. It is interesting to note that a significant queue length reduction (waiting time) was achieved with only a modest increase in utilization and no additional resources.

5 Summary

This concept paper presented an agent based modeling system oriented to the simulation of emergency departments in either stand alone mode, multiple interacting emergency departments, as well as technologies well suited to enhance simulation with statistical empirical data collected in real time. Ideas from Telecommunications Engineering are introduced as a model for policy change regarding patient redirection. Simulation alone is not sufficient due to the difficulty in predicting highly uncertain patient arrival rates whereas the system presented here is oriented to augmenting simulation with empirical data when available. As such, the context of the work here also presented a means where emerging technologies such as RFID are suggested as data sources. These sources can be mined in a statistically significant manner and provide real world input for the simulation. The system we are developing is also open source and relies on open source components. It is extendable and can be ported or tailored to a variety of hospital IT applications several of which were identified here.

Acknowledgments. Many thanks go to individuals who have contributed to the platform presented here. These people include Professors B. McLeod and A. Alfa from the University of Manitoba. In addition, we acknowledge the collaboration with Mr. D. Sanders and Mr. B. Podaima from the Internet Innovation Centre, also at the University of Manitoba.

References

1. Kanagarajah, A.K., Lindsay, P.A., Miller, A.M., Parker, D.W.: An exploration into the uses of agent-based modeling to improve quality of health care. In: Y. Bar-Yam International Conference on Complex Systems, Boston, MA, June 25-30 (2006)
2. Sanders, D., Mukhi, S., Laskowski, M., Khan, M., Podaima, B.W., McLeod, R.D.: A Network-Enabled Platform for Reducing Hospital Emergency Room Waiting Times using an RFID Proximity Location System. In: 19th International Conference on Systems Engineering, Las Vegas (2008)
3. Geographical Intrusion Event Mapping System (GEMS) (June 2008), `http://gems.ee.umanitoba.ca`
4. County of Marin, Ambulance Diversion Guidelines (June 2008), `http://www.co.marin.ca.us/depts/HH/main/ems/documents/Policies/5400.pdf`
5. Floyd, S., Jacobson, V.: Random Early Detection Gateways for Congestion Avoidance. IEEE/ACM Transactions on Networking 1(4), 397–413 (1993)

6. Qt4 (June 2008), http://www.trolltech.com
7. Patvivatsiri, L.: A simulation model for bioterrorism preparedness in an emergency room (January 2008), http://www.informs-sim.org/wsc06papers/061.pdf
8. Doupe, M., Kozyrskyj, A., Soodeen, R., Derksen, S., Burchill, C., Huq, S.: An Initial Analysis of Emergency Departments and Urgent Care in Winnipeg (2008), http://mchp-appserv.cpe.umanitoba.ca/deliverablesList.html
9. Canadian Triage and Acuity Scale (This reference outlines Canadian Triage Scales illustrating the types of rules that can be embedded into the agent based model) (2008), http://www.caep.ca/template.asp?id=B795164082374289BBD9C1C2BF4B8D32