Hospitals as Complex Social Systems: Agent-Based Simulations of Hospital-Acquired Infections

José M. Jiménez, Bryan Lewis, and Stephen Eubank

Network Dynamics and Simulation Science Laboratory, Virginia Bioinformatics Institute
1800 Pratt Drive, Research Building XV, Blacksburg, Virginia 24061
jjimenez@vbi.vt.edu

Abstract. The objective of this study was to develop a highly-detailed, agent-based simulation to compare medical treatments against healthcare-acquired infections (HAIs). A complex hospital model was built using patient information and healthcare worker data from two regional hospitals in Southwest Virginia. A specific HAI, *Clostridium difficile*, was chosen among other HAIs as the pathogen for the study due to its increased prevalence in the United States. The complex hospital simulation was created using the first principles of agent-based simulation. The simulation was then tested using a disease model with two different scenarios: a baseline with no medical treatment antimicrobials, and the use of an antimicrobial (fidaxomicin). The model successfully simulated over 164,000 personal contacts between patients and healthcare workers. Each medical treatment was evaluated one hundred times using one month of real hospital data. The mean case count was 2.66 for scenario 1 and 2.33 for scenario 2. The highest case count for scenario 1 was 21 cases whereas scenario 2 had a maximum of 11 cases. Understanding complex interactions between patients and hospital personnel could help hospitals understand transmission of infections while simultaneously reducing healthcare costs.

Keywords: Complex systems, healthcare systems, healthcare-acquired infections, *Clostridium difficile*, agent-based simulation.

1 Introduction

Hospitals are by definition complex systems containing multiple subsystems. The environment inside a hospital is a collection of interactions between numerous subsystems. Hospital subsystems include the internal and external environment, personnel, and technology. Although sometimes the subsystems may seem to work autonomously, there is an overarching objective that governs all of them. This goal aims to improve health of the patients by performing tasks unique to that subsystem. Take for example the interactions between the environmental services or janitorial department and the nursing staff within an intensive care unit (ICU). The objective of the environmental services department is to ensure that all surfaces in a patient room are cleaned and sanitized. Members of this department have very specific standard work that includes directives on

K. Glass et al. (Eds.): COMPLEX 2012, LNICST 126, pp. 165–178, 2013.

how to sanitize, what surfaces to clean, and cleaning times for each room and major area. Nurses, on the other hand, have a very different mission. They are tasked with patient care by following physician orders for treatment. If a patient's room needs to be cleaned, but at the same time the patient needs a particular treatment, there might be a conflict of objectives between the two subsystems. While at times it may appear that both subsystems are competing for the same space to conduct different activities, they share one goal. Their end goal is to guarantee that the patient regains health and that he is not harmed further during his stay in the hospital.

Healthcare-acquired infections (HAIs) occur within this complex system. Interactions between multiple subsystems such as hospital personnel, patients, and technology can create a difficult environment to identify and treat HAIs. In order to investigate this type of environment without reducing it to a very simplistic model, one can utilize highly-resolved simulations. Highly-resolved simulation refers to agent-based simulation models that can evaluate infection exposure, interventions and individual behavior change for populations in the hundreds of millions, while maintaining the resolution of the individual agent. This study is unique in its kind because it combines very specific factors. First, it makes use of multiple populations inside a hospital. These populations include patients, physicians, nurses, respiratory therapists, occupational therapists, speech therapists, physical therapists, and environmental services associates. Second, it utilizes actual patient data in the form of electronic medical records. Third, the study uses the technique of shadowing hospital workers to develop activity schedules from multiple disciplines to include in the simulation. Lastly, due to the amount of data being processed, it is clear that less advanced simulation software would not be able to calculate the interactions of thousands of agents in an efficient manner. The use of highly-resolved simulation was essential for the epidemiological study. For this reason it is necessary to combine the resolution of highly-resolved simulation with the power of high performance computing.

1.1 Objective of the Study

The objective of the study was to develop a highly-detailed, agent-based simulation to compare medical treatments against healthcare-acquired infections (HAIs). In order to achieve this objective the number of daily exposures and contacts between patients and healthcare workers was obtained through a hospital simulation. By determining these contacts, their durations, and locations it was possible to suggest to hospital staff better treatment measures against HAIs. This paper represents the initial stages of the study and therefore looks at the practicability of conducting highly-resolved simulations for hospitals.

1.2 Healthcare-Acquired Infections

Healthcare-acquired infections are those infections defined as being transmitted and acquired inside a healthcare facility. Healthcare facilities include hospitals, clinics, outpatient facilities, and nursing homes. There are multiple types of pathogens that can infect a patient or be transmitted within a healthcare facility. The transmission, control and prevention, and treatment are very different for each type of infection.

Clostridium difficile is a normal occurring bacteria in the intestinal flora, however certain strains of the bacteria can cause *Clostridium difficile*-associated disease (CDAD). CDAD can produce watery diarrhea with the mildest of cases, but can also produce severe colitis requiring surgery in the harshest cases. The severe form of the infection can also lead to death. CDAD has been linked to risk factors such as being 65 years-old or older, suffering a severe underlying illness, going through a nasogastric intubation, taking antiulcer medications, having a prolonged hospital stay, and receiving treatment with certain antibiotics [1, 2]. New treatments are currently being developed to fight CDAD to include new antibiotics, fecal transplants, and a new vaccine. *Clostridium difficile* was chosen as a model infection for this study due to its prevalence in the hospital environment and chain of infection.

A recent review of data from the Healthcare Cost and Utilization Project (HCUP) showed that the average cost of *Clostridium difficile* treatment could be as high as $24,400 and the aggregate costs of all *Clostridium difficile* treatments in the United States was $8,238,458,700 in 2009. Of the total $8 billion in *Clostridium difficile* infections costs 67.9 percent was covered by Medicare, 9.1 percent by Medicaid, and 18.8 percent was covered by private insurance [3]. Several studies have estimated the cost of individual treatment between $3,000 and $32,000 [4-7].

1.3 A Hospital as a Complex System

Hospitals are by virtue of their structure complex systems. Each hospital is a set balanced connections of multiple subsystems and thousands of individual agents. The personnel subsystem and the technological subsystem are two of the most important subsystems in the hospital. The personnel subsystem is composed of all the individuals that interact within the hospital. This subsystem includes patients, hospital workers and hospital visitors. The technological subsystem is sometimes seen as synonymous with machines or computers, but it also includes the knowledge shared by the hospital personnel, the level of automation of its internal processes, and as mentioned before the equipment that is utilized. Furthermore, the organizational structure of a hospital can be daunting. A complex organizational structure can follow a large number of procedures, rules, and guidelines. Nurses, as an example, can have supervisors from different reporting structures. For example, a nurse can report to a nurse manager, a resident doctor, and a chief of staff all at the same time. With nested structures such as these, standard operating procedures can become complex and with additional stress added to a person, errors are more prevalent [8].

In addition to the multiple subsystems there are thousands of individuals interacting constantly within the walls of a hospital. For this study, one of the regional hospitals had over 37,000 patients (non-recurring admissions) and over 4,000 hospital workers. The same hospital had over 1,000 locations. The combination of personnel and locations can produce millions of activities during a year inside a hospital.

Complex systems with so many individual parts interacting together are difficult to study without reducing to very simplistic models. Most of the regularly used studies do not have the capacity to analyze the number of agents that have been described above. It is also impractical and costly to conduct complex experiments on interventions in a hospital. The size of the overall population and the interactions of the multiple subsystems make this task very difficult to achieve. In the past, researchers have looked at different types of study designs to test their hypotheses such as case-control studies or retrospective studies. One of the weaknesses of a case-control or a retrospective study is that experiments of this fashion would not be able to encompass the detailed dynamics within a hospital. The use of highly-resolved, agent-based simulation can help overcome these weaknesses. Highly-resolved simulation uses data obtained from the population data sources to create a realistic synthetic (*in-silico*) population. The *in-silico* representation of the population includes information such as location, daily activities, age, and other important demographics [9]. Even though all this information is incorporated into the simulation, the identity of the patients and hospital workers are protected through anonymity. Highly-resolved simulation has been used in the past to analyze possible interventions in response to hypothetical disease outbreaks such as influenza [10, 11].

1.4 Highly-Resolved Simulations

In the past, simulations have been used on multiple occasions to help explain the complexities of health systems. Simulations are representations of complex systems, but they can never be a perfect representation of reality. Additionally, one cannot say that any kind of simulation is better than the other. Simulations that are more graphical and user-friendly have the advantage of a short learning curve and ease of use. People that unfamiliar with simulation can learn to use the software in a matter of hours. However, user-friendly simulations lack the resolution needed for analyzing large number of individual agents. On the other hand, one can find simulations that are computer code-intensive and that are not intuitively user-friendly. Written directly in language such as C, C++, or Java language, these type of simulations can provide better resolution for large amounts of data, but require training and users that are computer-code savy [12]. The selection of the right simulation software depends on the type of problem that a researcher seeks to answer.

EpiSimdemics is a novel simulation software that looks at simulation in large social networks. EpiSimdemics is a parallel scalable algorithm used to simulate the spread of infectious diseases over extremely large contact networks. This algorithm can also simulate other population factors such as fear and behaviors. The Network Dynamics and Simulation Science Laboratory (NDSSL) at Virginia Tech developed EpiSimdemics with the objective of studying the effects of pharmaceutical and non-pharmaceutical interventions. EpiSimdemics is an agent-based simulation in which every person is an individual agent. EpiSimdemics is based on social network theory, but can overcome its limitations of use for only small groups of people. Prior studies have utilized this tool for analysis of large populations at a scale of 100,000,000 people. In the simulation, agents and locations are identified as nodes in a network graph (or a sociogram) and the edges between the nodes are considered contacts between agents or visits of an agent to a specific location. Every activity for a particular agent is defined in an activity schedule. The schedule determines the location of the activity, the activity performed, and the duration of the activity. Interactions between people are calculated using a stochastic model. If one or more individuals are in the same location as an infected individual then they might or might not become infected based on probability calculations. These probabilities are defined in the simulation in the disease model, which we explain later on. Other simulation models do not have the same level of resolution that EpiSimdemics has due to the ability to utilize extremely large populations through social networks. Other advantages of EpiSimdemics are the ability to combine policy interventions as well as individual behaviors in the same model. NDSSL has used EpiSimdemics for several studies of large populations for multiple government agencies [13-15].

Furthermore, "first principles" or guidelines have also been developed by NDSSL[16]. The first step in developing the simulation is the creation of a synthetic population or proto-population. In other models that describe epidemics or pandemics in large areas, the use of massive databases was needed to replicate the population of a city [10], a region, or an entire country [16, 17]. In the case of the hospital simulation, the study uses data collected directly from patient records and shadowing of hospital workers to develop the proto-population. The second step in developing highly-resolved simulations is to create activity schedules for the entire proto-population. The third step is to create a disease model that will be used to "infect" agents as they come in contact with other agents. All these steps will be explained further in the next section.

EpiSimdemics has the distinct ability to quantify not only the number of agents that become infected with a specific disease in the simulation but it can also identify the contacts between individual agents in the simulation. These contacts occur for different times depending on the activity and location of the person-agent. The software can count the number of contacts and the time that each contact lasts. Figure 1 below, shows an example representation of contacts in an ICU's social network. Every time that there is an agent-to-agent contact the software keeps a record of its location, time, and of the agents who interacted. This record is used to create the sociogram or network diagram at each time step.

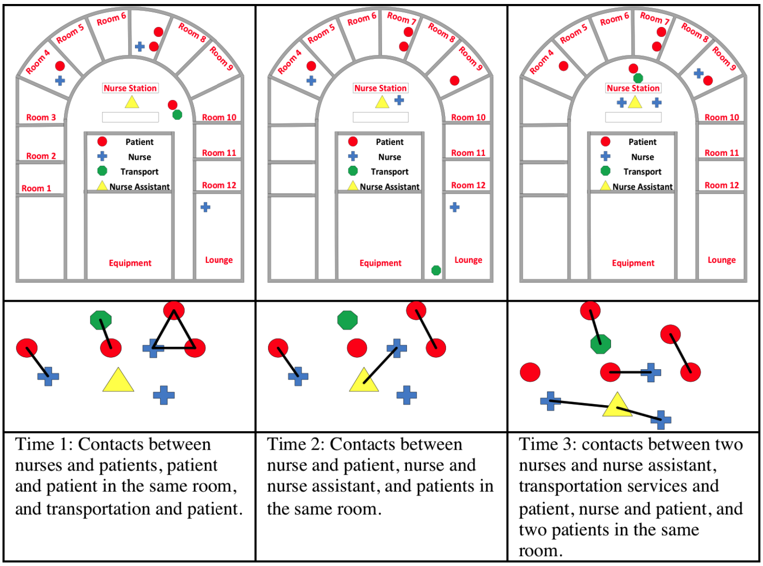

Fig. 1. Contacts at different times in an ICU social network. In this representation of the simulation different agents come into contact with each other causing a potential transmission of the infection. With more contacts there is a higher probability of transmission. The *red circles* represent patients, the *blue crosses* represent nurses, the *green hexagons* represent transportation services, and the *yellow triangle* represents the nurse assistant.

2 Methodology

The methodology of the study consisted of four phases: electronic data collection, shadowing of hospital workers, disease modeling and simulation of the hospital. In order to create a more realistic simulation the study gathered data that represented reality as close as possible. For this reason, electronic medical records were requested directly from the hospital. Additionally, the study included shadowing of different hospital disciplines while hospital workers performed their daily activities. It is difficult to obtain 100 percent validation on simulation models, however when the agents in the model are based on accurate data directly from patient records and activities from direct observation, the fidelity of a large portion of the simulation can be verified.

2.1 Data Collection

As explained before, the first step in creating a highly-resolved simulation is to develop a synthetic or proto-population that resembles the real population with high resolution and high fidelity. In order to achieve this detail of resolution it was necessary to obtain actual data of the multiple populations that interact within a hospital. Two regional hospitals from Southwest Virginia provided data for the study. The first type of data obtained was in the form of electronic medical records. The patient records were de-identified to protect patient information by the hospital before submitting them to the research team. The electronic records included one year of patients' locations and activities throughout the hospital. The records included over 37,000 patients and over 400,000 single activities by the patients. The anonymous data was stored on a SQL database and the information was protected by numerous measures such as deidentification and restricted access only to those researchers participating in the study.

2.2 Shadowing of Hospital Workers

The next phase of data collection was the shadowing of hospital workers. Hospital workers are not typically tracked throughout the healthcare facility as they conduct their daily routines. It was important for the study however, to obtain a realistic representation of those daily activities in the form of vignettes. These vignettes are schedules that summarize the most common activities that a worker would perform throughout the day. The vignettes include locations, activities, contacts with other people or technology, and durations of those activities and contacts. Table 1 below shows the different healthcare disciplines that were observed during the study at the two hospitals. Even though data exists on most departments for the hospitals, not all the disciplines have been included currently in the simulation. Further work is needed to develop specialty departments within the hospitals such as the emergency department, the operating rooms, the neonatal unit, and the cafeterias.

Table 1. List of hospital disciplines shadowed during the study

Intensive Care Unit (ICU) Nurses	Infection Preventionists	Dieticians
Vascular Intensive Care Unit Nurses	Infection Control Physicians	Social Workers
Progressive Care Unit Nurses	Resident Physicians	Case Managers
Ostomy Nurses	Respiratory Therapists	Imaging Specialists (radiography/ultrasound)
Neonatal Intensive Care Unit (NICU) Nurses	NICU Respiratory Therapists	Facility Managers
Dialysis Nurses	Physical Therapists	Phlebotomists
Emergency Department Nurses	Occupational Therapists	Laboratory Technicians (General)
Nurse Assistants	Speech Therapists	Laboratory Technicians (specializing in *Clostridium difficile*)
	Environmental (Janitorial) Services	

Figure 2 below shows an example of a vignette. As shown on the diagram a day shift Intensive Care Unit (ICU) nurse arrives to the hospital in the morning and participates in a shift meeting. Then, the nurse conducts a preliminary assessment of his two patients. After that the nurse could continue to monitor patients, relieve another nurse, or goes to the cafeteria to eat lunch. This information was coded into a computer program that create multiple schedules for a large number of nurses that had similar jobs. Stochasticity was also added into the program so that all nurses would not be conducting the same activity at the same time and location.

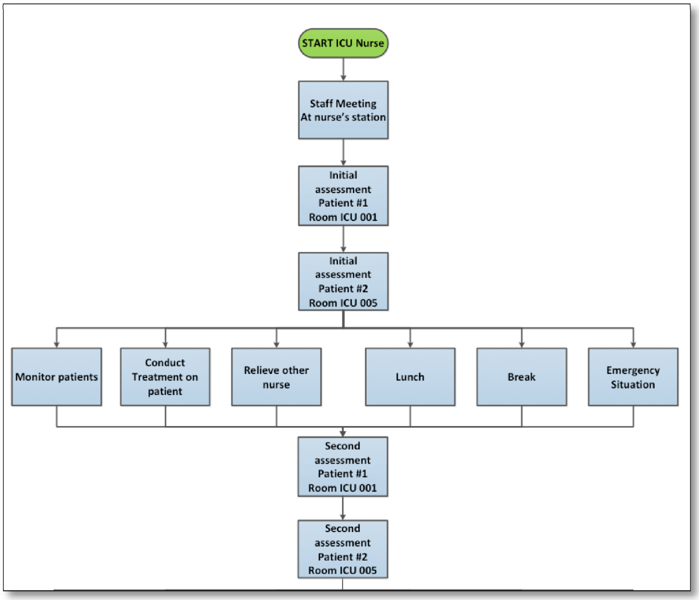

Fig. 2. Example of a vignette for an ICU nurse

2.3 Disease Models

The next step in highly-resolved simulation is designing a disease model based on the medical literature from the infection that will be used to "infect" the proto-population. Information gathered from the literature or from the facility's infection prevention professionals to identify the problem served as a base to build the disease model. The disease model is a probabilistic finite-state machine (FSM), a state-transition model that is used on the entire proto-population. Probabilities and distributions are fed into the FSM in order to develop a realistic model of how the disease spreads inside the hospital. Figure 3 shows an example of a disease model for *Clostridium difficile*.

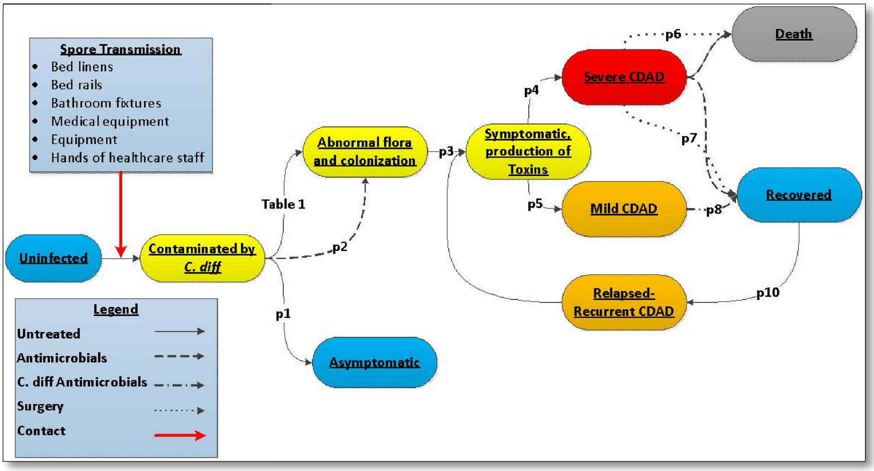

Fig. 3. Disease model for *Clostridium difficile*

The disease model represents the different health states that any agent in the simulation can be in at any particular time. For this study two different scenarios were used to compare several treatments. Both scenarios had nine health states: uninfected, colonized, not colonized, infected, asymptomatic, severe CDAD, mild CDAD, death, and recovered. The first scenario was a control or baseline therefore No medical or preventive treatment was utilized. The second scenario incorporated *Clostridium difficile* antimicrobials for mild and severe CDAD as well as a vaccine to avoid infection. Specific antimicrobials have been identified as a risk factor for *Clostridium difficile*. For this reason, use of antimicrobials is the first health state (uninfected) as a link towards the next step, colonization, if the agent was exposed to the pathogen. Table 2 and 3 below shows the parameters utilized for the different disease models.

Table 2. Scenarios used in the simulation

	Disease Model 1	Disease Model 2
First Disease Link	Exposure to *Clostridium difficile,* use of antimicrobials to trigger infection	Exposure to *Clostridium difficile*, use of antimicrobials to trigger infection
Treatments Tested	No preventive or medical treatments	*Clostridium difficile* antimicrobials for mild CDAD *Clostridium difficile* vaccine *Clostridium difficile* antimicrobials for severe CDAD

Table 3. Parameters for the two scenarios

Disease State	Disease Link	Probability	Next Disease State
Uninfected	Untreated	0.95	Colonized
	Untreated	0.05	Not Colonized
	Antimicrobials	0.96	Colonized
	Antimicrobials	0.04	Not Colonized
	Vaccine (Scenario 2)	0.25	Colonized
	Vaccine (Scenario 2)	0.75	Recovered
Colonized	Untreated	0.50	Infected
	Untreated	0.50	Asymptomatic
Infected	Untreated	0.1	Severe CDAD
	Untreated	0.9	Mild CDAD
Not Colonized	None	None	None
Asymptomatic	None	None	None
Severe CDAD	Untreated	1.0	Death
	Clostridium difficile antimicrobials (Scenario 2)	0.68	Recovered
	Clostridium difficile antimicrobials (Scenario 2)	0.32	Death
Mild CDAD	*Clostridium difficile* antimicrobials (Scenario 2)	0.85	Recovered
	Clostridium difficile antimicrobials (Scenario 2)	0.15	Severe CDAD
	Untreated	1.0	Severe CDAD
Death	None	None	None
Recovered	None	None	None

3 Results

3.1 Simulation Runs and Parameters

This study simulated the entire patient population of the hospital for an entire year. Particular importance was given to the ninth floor of the simulated hospital for the purposes of studying the interactions of the multiple populations of hospital workers. Only hospital workers of the ninth floor and their activities were simulated due to the complexity of creating additional departments for the hospital such as the emergency department and the operating room. These departments will be included in further simulations. The study was conducted for thirty days of hospital activities and it included agents acting as patients, physicians, nurses, respiratory therapists, occupational therapists, speech therapists, physical therapists, and environmental services associates. Each of the agents had individual activity schedules that identified each location, activity type, and duration. Each simulation of thirty days ran through 100 iterations for each scenario.

3.2 Assumptions and Simplifications

In order to run the simulation in a fast and efficient manner some assumptions were made regarding the model. The model does not distinguish between the ages of the agents. In reality, the age of the patient is an important risk factor as people over 65 years old are more likely to be infected with HAIs than younger people. Additionally, the medical condition of the agents did not play a factor in the simulation with the exception of those individuals that are under an antimicrobial regiment. Infections other than CDAD were not taken into account. The next assumption was that every agent in the *in-silico* hospital population could be infected. Similarly, all agents could progress through the disease models equally with no difference on whether they were patients or hospital workers. Finally, the study assumed a "barrier" around the hospital for agents. This means that once a patient-agent left the premises of the hospital, the infection and the disease model did no longer affect him.

3.3 Exposures and Contacts

A very interesting point of the simulation is to be able to observe the different contacts that occur between *in-silico* agents. The average number of activities for the entire month was 8,234 activities. The total number of agent-to-agent contacts during the 30 days was 164,176. These contacts are represented on figure 5 as a sociogram. Every node in the sociogram represents an agent-person. When two agent-persons are in the same location at the same time they are considered in contact with each other. An edge or line represents their contact.

3.4 Scenario Summary

Two graphs were obtained from the multiple iterations of the simulation. Each graph includes the cumulative case count of infected agents for each of the two scenarios as explained in section 3.1. Each scenario was run 100 times with different probability seeds for each run. The stochasticity of the models based on the seeds explains the differences in the number of cases from one curve to another. During one of the iterations for scenario 1 the number of cases grew to 21, whereas in several other iterations there were no cases at all. It is important to understand that the purpose of the simulation is not to replicate the exact results from the two regional hospitals but to evaluate the possible contacts between the different populations. These contacts and exposures can later be analyzed for improving systemic interventions such as hand washing, education programs, or better disinfection in rooms. The cumulative case curves are presented to show the disease transmission in a manner that is consistent, but not exactly the same as an actual hospital. Scenario 1 had an average number of cases of 2.66, compared to 2.33 from scenario 2. The highest number of cases for scenario 2 was only 11 compared to 21 on scenario 1. Figures 5 and 6 show the cumulative cases for scenarios 1 and 2.

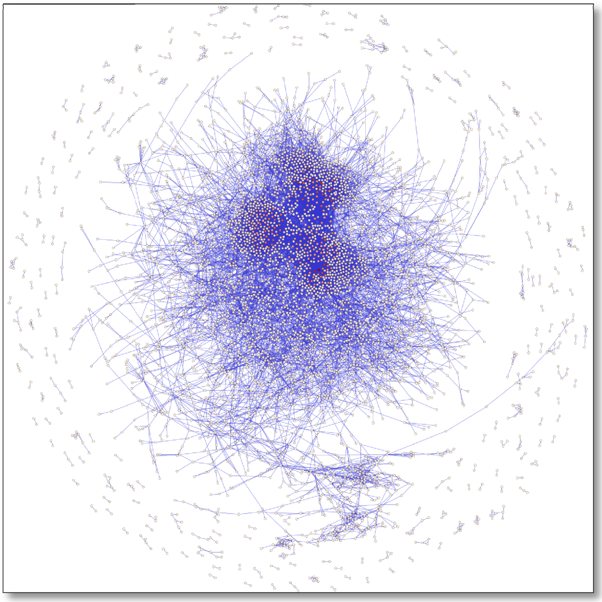

Fig. 4. Network diagram of the *in-silico* hospital population during one month

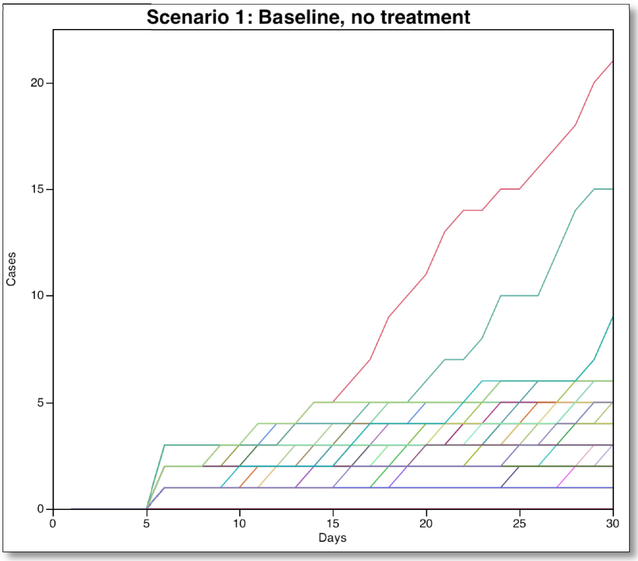

Fig. 5. Cumulative case curves for scenario 1. Scenario 1 is the baseline scenario and did not include any medical or preventive treatment against the infection.

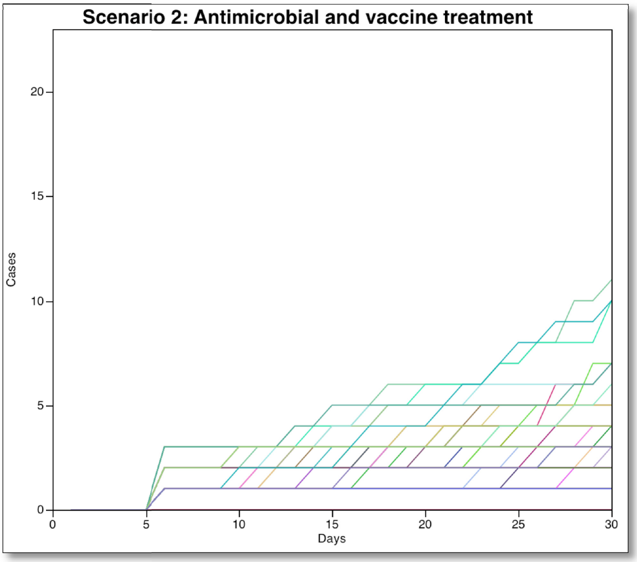

Fig. 6. Cumulative case curves for scenario 2. Scenario 2 included medical treatment of the infection through the use of a vaccine and a new antibiotic called fidaxomicin.

4 Conclusions and Future Work

A hospital is a complex system composed of multiple subsystems and thousands of individuals that work for the health of the patients. It is difficult and impracticable to perform complex experiments taking into consideration multiple populations within a hospital. The use of highly-detailed simulation for the study of healthcare-acquired infections can be of great use not only for researchers but also for healthcare professionals. In an era where "every penny counts" for healthcare system managers a simulation can represent an easy and affordable tool to understand infection control. This initial paper represents the first steps in our study of simulations of HAIs. The study provided evidence that simulation of a large complex system with high resolution and high fidelity is practicable. One of the advantages of this type of simulation is that real data from hospitals, such as electronic medical records, can be modified to create *in-silico* representations of people. The simulation was able to produce interesting results that should be studied further. One of the important results is that for over 164,000 contacts between agents, the highest number of cases was only 21. This is significant because it could explain that the current prevention practices in the hospitals are effective. Further study is needed in the area of simulation of prevention and control of HAIs.

Additional work in the study of HAIs in hospitals is already being performed at NDSSL in diverse areas to improve the simulation performance. First, to ensure the completeness of the hospital system it is necessary to add additional *in-silico* populations. These populations should include the entirety of the hospital workers and the visitors to the hospital. Visitors can be a very important factor in the transmission of infections. Visitors are susceptible to infections, especially if their demographics carry any of the risk factors for a specific disease. Hospital visitors are also more mobile than patients and could potentially transmit infections to different locations inside and outside of the hospital. The demographics of the population should also be included to future iterations of the study. Risk factors linked to the demographics of the agents could give further insight into effective treatment. The current simulation model did not include all the floors of the hospital due to the complexity of certain specialty departments. Further models should include every floor and department of the hospital to include the operating room and the emergency department.

In order to increase the realism of the simulation multiple hospitals could be added to the simulation. This network of hospitals could work similar to a state or regional health department evaluating an outbreak. Finally, the realism of the simulation could increase if other infections, not necessarily other HAIs, would be added into the activity files.

References

[1] Bignardi, G.E.: Risk factors for Clostridium difficile infection. J. Hosp. Infect. 40, 1–15 (1998)

[2] Sunenshine, R.H., McDonald, L.C.: Clostridium difficile-associated disease: new challenges from an established pathogen. Cleve. Clin. J. Med. 73, 187–197 (2006)

[3] Lucado, J., Gould, C.V., Elixhauser, A.: Clostridium difficile infections (CDI) in hospital stays, 2009: statistical brief #124 (2012)

[4] Scott, D.: The direct medical costs of healthcare-associated infections in US hospitals and the benefits of prevention. In: Centers for Disease Control and Prevention, N. C. f. P. Division of Healthcare Quality Promotion, Detection, and Control of Infectious Dieseases (2009)

[5] Dubberke, E.R., Reske, K.A., Olsen, M.A., McDonald, L.C., Fraser, V.J.: Short- and long-term attributable costs of Clostridium difficile-associated disease in nonsurgical inpatients. Clin. Infect. Dis. 46, 497–504 (2008)

[6] Kyne, L., Hamel, M.B., Polavaram, R., Kelly, C.P.: Health care costs and mortality associated with nosocomial diarrhea due to Clostridium difficile. Clin. Infect. Dis. 34, 346–353 (2002)

[7] Dubberke, E.R., Wertheimer, A.I.: Review of current literature on the economic burden of Clostridium difficile infection. Infect. Control Hosp. Epidemiol. 30, 57–66 (2009)

[8] Karsh, B.T., Brown, R.: Macroergonomics and patient safety: The impact of levels on theory, measurement, analysis and intervention in patient safety research. Appl. Ergon. 41, 674–681 (2010)

[9] Lewis, B.: In silico Public Health: The Essential Role of Highly Detailed Simulations in Support of Public Health Decision-Making. Doctor of Philosophy, Genetics, Bioinformatics, and Computational Biology, Virginia Tech, Blacksburg, VA (2011)

[10] Eubank, S., Guclu, H., Kumar, V.S., Marathe, M.V., Srinivasan, A., Toroczkai, Z., Wang, N.: Modelling disease outbreaks in realistic urban social networks. Nature 429, 180–184 (2004)

[11] Halloran, M.E., Ferguson, N.M., Eubank, S., Longini Jr., I.M., Cummings, D.A., Lewis, B., Xu, S., Fraser, C., Vullikanti, A., Germann, T.C., Wagener, D., Beckman, R., Kadau, K., Barrett, C., Macken, C.A., Burke, D.S., Cooley, P.: Modeling targeted layered containment of an influenza pandemic in the United States. Proc. Natl. Acad. Sci. U S A 105, 4639–4644 (2008)

[12] Keeling, M.J., Rohani, P.: Modeling infectious diseases in humans and animals. Princeton University Press, Princeton (2008)

[13] Barrett, C.L., Bisset, K.R., Eubank, S.G., Feng, X., Marathe, M.V.: EpiSimdemics: An efficient algorithm for simulating the spread of infectious disease over large realistic social networks. Presented at the Proceedings of the 2008 ACM/IEEE Conference on Supercomputing, Austin, Texas (2008)

[14] Bisset, K.R., Aji, A., Bohm, E., Kale, L.V., Kamal, T., Marathe, M., Yeom, J.-S.: Simulating the Spread of Infectious Disease over Large Realistic Social Networks using Charm++. Presented at the to appear at 17th International Workshop on High-Level Parallel Programming Models and Supportive Environments, HIPS (2012)

[15] Bisset, K.R., Feng, X., Marathe, M., Yardi, S.: Modeling interaction between individuals, social networks and public policy to support public health epidemiology. Presented at the Winter Simulation Conference, Austin, Texas (2009)

[16] Barrett, C., Beckman, R., Khan, M., Kumar, A., Marathe, M., Stretz, P., Dutta, T., Lewis, B.: Generation and Analysis of Large Synthetic Social Contact Networks. In: Winter Simulation Conference (2009)

[17] Chao, D.L., Halloran, M.E., Obenchain, V.J., Longini Jr., I.M.: FluTE, a publicly available stochastic influenza epidemic simulation model. PLoS Comput. Biol. 6, e1000656 (2010)