Proceedings of the 2015 Winter Simulation Conference
L. Yilmaz, W. K V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, eds.

A SCALABLE DISCRETE EVENT STOCHASTIC AGENT-BASED MODEL OF INFECTIOUS DISEASE PROPAGATION

Paul J. Sanchez
Susan M. Sanchez

Department of Operations Research, GL-239
Naval Postgraduate School
1411 Cunningham Rd
Monterey, CA 93943, USA

ABSTRACT

We propose a new stochastic model of infectious disease propagation. This model tracks individual outcomes, but does so without needing to create connectivity graphs for all members of the population. This makes the model scalable to much larger populations than traditional agent-based models have been able to cope with, while preserving the impact of variability during the critical early stages of an outbreak. This contrasts favorably with aggregate deterministic models, which ignore variability, and negates the requirement to assume "convenient" but potentially unrealistic distribution choices which aggregate stochastic models need in order to be analytically tractable. Initial explorations with our new model show behaviors similar to the observed course of Ebola outbreaks over the past 30+ years—while many outbreaks will fizzle out relatively quickly, some appear to reach a critical mass threshold and can turn into widespread epidemics.

1 INTRODUCTION

Ebola is a virulent disease, with an extremely high mortality rate. It is endemic in the regions near the Ebola river in the Democratic Republic of the Congo, where the first outbreak noted by the outside world occurred in 1976. It is found in the local animal populations, including bats and chimpanzees, and sporadically crosses over into the human population. Prior to 2013, there were 25 outbreaks in the human population that would flare up in hot-spots and then die out. The current outbreak started in 2013 in Guinea (WHO Ebola Response Team. 2014) and grabbed world-wide headlines by becoming a pandemic, spanning multiple countries and infecting a conservatively estimated 26 thousand people (World Health Organization 2015).

The level of concern has declined given the failure of the disease to spread to the rest of the world. As of this writing, a few isolated cases have been identified in the US and in Spain, but were quickly identified, isolated, and cared for without triggering further widespread outbreaks. However, if the disease were to evolve into a pneumonic form or one with a longer onset between exposure and becoming symptomatic, or if it were to get a significant foothold in a densely populated city outside of Africa, our ability to identify, track, and isolate exposures would quickly be overwhelmed.

2 MODELING BACKGROUND

Analytical models for epidemiology date back to the early 20th century, pre-dating modern computing (Kermack and McKendrick 1927). These so-called "compartmental" models categorize individuals comprising the population into subsets that reflect their current status. Various categorizations can be used, but a common one for diseases such as Ebola, where survivors acquire immunity, is the so-called SEIR model. The population is broken into the following classifications: *Susceptible* (S) is the set of people who have

not contracted the disease, *Exposed* (E) is the set of people who have contracted the disease but cannot yet transmit it, *Infectious* (I) is the set of people who are capable of transmitting the disease to susceptibles, and *Removed* (R) is the set of people who are removed from the process by death or recovery. The categories describe the progression of the disease for each individual, but the models describe the transfer of portions of the population from category to category using rate descriptors and differential equations as a continuous approximation to what is in reality a discrete process happening to individuals. The resulting systems of equations could be solved analytically and thus represented a major step forward in understanding disease propagation. One major drawback of these models is that the aggregate approximation is questionable with smaller populations. Additionally, as we study lower levels of infection, variation can play a significant role in the progression of the disease. In short, there is a need for stochastic models.

Differential equation based models have been extended to include stochastic behaviors, although doing so requires Markovian assumptions to preserve analytic tractability. A very recent example is the general epidemic model EPIMOD developed by Washburn (2015), that motivated our work.

Both deterministic and stochastic variations of the aggregated compartmental models can be and have been implemented computationally using time-stepped models. At every delta-t, the rates describing aggregate movement from one compartment to another are crossed with the compartment populations to calculate the number of people who get moved between each of the compartments. Note that this is a discretized approximation of a continuous model of a discrete process. The amount of error introduced by these approximations can be reduced by using smaller values of delta-t, but some error is unavoidable, and error propagation is cumulative. An example of a stochastic SEIR-like model appears in Broeck et al. (2011), who include a graphical user interface for setting up, executing, and visualizing the results of simulations.

Agent-based models have also been used to study disease propagation by tracking the progression of the disease in individuals. All members of the population other than a small pool of initially infected people start off as susceptibles. The disease then progresses through the various stages for each infected person as time advances. Upon becoming infectious, the agent can potentially expose any susceptible agents amongst their contact pool to the disease, until such time they qualify for "removal" by death or recovery. Adding randomness is fairly easy with this approach. The times spent in each state can be random, as can the likelihood and timing of transmitting the infection to each of the potential contacts. Drawbacks to this approach include (1) the need to specify either a connectivity graph between the various agents or use the dynamics of mutual discovery in a spatial representation; (2) the fact that most agent-based modeling platforms are time-stepped, which has been demonstrated to introduce modeling artifacts and potential errors (Al Rowaei, Buss, and Lieberman 2011); and (3) the consequent need to track and evaluate the status of every agent at every time step. Because of these issues, agents are often used to represent groups rather than individuals to offset the otherwise rapid growth of run times.

One example of agent-based modeling appears in Andradóttir et al. (2011), who develop a detailed simulation of influenza spread in Hamilton (Ontario) Canada, a city of roughly 650,000 people. They focus on the average overall illness attack rates for a variety of vaccine intervention conditions. They generate detailed attributes for every person in the population, including three or four contact groups to which the person belongs: household; neighborhood (approximately 500 people); community (approximately 2000 people); and daycare, school, or workplace. The agents go through an latent/incubation stage, and one of two potential infectious stages (symptomatic or asymptomatic), before being removed from the model. Each day, susceptible agents have the opportunity to become infected in their contact groups, where the daily probability of infection within each group is based on the number of infectious contacts and their per-contact influenza infection transmissions. Note that because state changes and transmission probabilities are evaluated on a daily basis, this is a time-stepped agent-based model. Although there is a large total population, each agent's potential contacts are much more limited in number. A similar approach is taken by Grefenstette et al. (2013) in their Framework for Reconstructing Epidemic Dynamics (FRED) platform. The authors use a fixed simulation step of one day, but suggest that this does not appear to be a

limitation except for diseases with extremely short latency and infectious periods, or extremely short-period intervention strategies.

Other agent-based models have been proposed. For example, Parker and Epstein (2011) develop a Global-Scale Agent Model (GSAM) software platform for agent-based epidemic modeling. It is a "highly parallel distributed model" that evaluates agents and their travel itineraries to track epidemic spread. Written in Java, its U.S. submodel with 300 million agents executes in "roughly twenty minutes." The FluTE simulation of Chao et al. (2010) models the spread of influenza, using two time steps per simulated day to allow for differences in daytime and nighttime social interactions, and restricting the community sizes within which agents interact to less than 3000 agents to improve model scalability.

Discrete event formulations also exist. For example, Perumalla and Seal (2012) present a discrete event model of a reaction-diffusion simulation model of epidemic propagation. They use optimistic scheduling with rollback to achieve massive parallelization. Their reaction-diffusion model uses an abstract geographic structure, where the disease propagates as individuals travel between locations within a region, or between geographic regions. They observe that systems characterized by many locations containing smaller numbers of individuals are computationally more efficient than systems characterized by fewer locations with larger populations. Their focus is on model scalability in terms of speed and memory, rather than the dynamics of disease propagation.

The papers mentioned above involve creating a structure for modeling epidemics. Papers that seek to model specific epidemics are also available, such as the Ebola models of Merler et al. (2015) and Siettos et al. (2015). Others have focused on fitting models to cumulative incidence data, or estimating model parameters as incidence data become available. King et al. (2015) discuss the potentially large errors in parameter estimates that can be masked when fitting deterministic models to cumulative incidence data. They go on to recommend against the use of aggregated models, in favor of checking the effects of violated model assumptions such as independence, and to state that stochastic models are generally preferable to deterministic ones.

3 AN EVENT-BASED SEIR FORMULATION

All these considerations motivated us to pursue development of a new epidemic modeling paradigm. We wanted to avoid aggregation, include stochastic behaviors, avoid unnecessarily restrictive modeling assumptions, and eliminate the artifacts associated with time-stepped modeling. At the same time, we wanted to retain the simplicity and extensibility of the SIR/SEIR class of epidemic models, and have a model that is scalable. Our view of scalability focuses on the size of the population that can be studied, rather than the degree of parallelization that can be achieved.

We protoyped our model in the Ruby programming language (Flanagan and Matsumoto 2008, Ruby Community 2015) using discrete-event scheduling. We describe the model, which has quite simple logic, with an event graph based on the notation of Schruben (1983). We present our model in Figure 1.

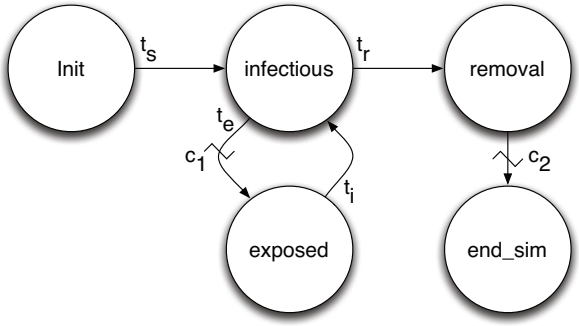

Figure 1: An event graph representation of the Ebola model's event scheduling relationships.

The events, and their associated state transitions and scheduling behaviors follow. Events other than **`Init`** and **`end_sim`** are scheduled on a per-agent basis.

`Init` : Set all state variables to initial values and schedule when each of the initially infected members of the population will become infectious.
- For each of the initially exposed agents, determine how long until that agent becomes infectious and schedule an **`infectious`** event with a delay of t_s.

`infectious` : This agent can now infect susceptibles in their pool of contacts.
- Determine t_r, the time until this agent is removed from the infectious state by either death or recovery.
- Schedule a **`removal`** event to occur in t_r time units.
- Determine the pool size of potential susceptible contacts for this agent.
- For each potential contact calculate t_e, the amount of time until that contact would be exposed.
- Determine Boolean condition c_1: $t_e < t_r$.
- If c_1 is true, this contact will be exposed to the current agent while that agent remains infectious. Schedule an **`exposed`** event for this contact.

`exposed` : A new agent is exposed to the infection. The new agent is now a carrier of the disease, but is not yet infectious.
- Schedule an **`infectious`** event for this agent t_i time units after their exposure.

`removal` : This agent is no longer infectious.
- Decrement the number of **`currently_infected`** agents.
- Determine Boolean condition c_2: **`currently_infected`** == 0.
- If condition c_2 is true, schedule the **`end_sim`** event.

`end_sim` : Print final reports and terminate the simulation.

Interested readers can find more information about implementing event graph models in Sanchez (2006).

4 DISCUSSION

This model differs from prior models in several important respects. First, this model tracks the state of individual agents, but does so using a discrete-event scheduling framework. This means that it avoids the order of execution issues which time-stepped models are subject to, and which Al Rowaei, Buss, and Lieberman (2011) demonstrated can alter model behaviors in significant ways. Our dynamic event-oriented approach also creates significant runtime efficiencies relative to time-stepped models, since evaluating the number of state transitions for active agents can be orders of magnitude less than evaluating the pairwise interactions of all agents in the population at every delta-*t* increment of the simulation clock.

Second, susceptibles are maintained implicitly by the model. Agents are not instantiated until such time as they become exposed to the disease, and are disposed of by removal from the pool of infecteds. This coordinates with the event-scheduling aspects of the model to greatly reduce both the number of events to be tracked and the memory footprint associated with the model state. This is made possible by doing dynamic instantiation of each agent's contact pool when the agent becomes infectious. Upon becoming infectious, we generate a pool size for the set of contacts. We have used a Poisson distribution for the pool size. As the model progresses, the Poisson rate is scaled by the proportion of the population which remains susceptible. The effect of this is to stochastically reduce the typical pool size for agents who become infected later in the simulation because many other agents have already been eliminated as potential infectees due to death or prior infection.

Agent-based models are not the only ones that suffer from lack of scalability. Washburn (2015) develops a mixed epidemic model in the form of a a continuous time Markov chain. He says "in spite of considerable previous work, there is still no practical method of computing the exact, complete SIR case count distribution

for large [values of the initial number of susceptibles]." He shows that the computational effort required to determine the exact distribution for the lower tail of the total number infected is bounded. Throughout his paper, but particularly for large populations, Washburn (2015) emphasizes "estimating the average case count, since it is anticipated that minimizing that count will be the goal of any countermeasures taken."

Another characteristic of our approach is that we can incorporate stochastic behavior into any event, state transition, and scheduling behavior. For example, we currently use a Poisson process to generate variation in the number of potential contacts associated with each agent. This can yield very different results than assuming a fixed number of exposures for each infected person.

5 EXAMPLES

Our objective is to demonstrate the variety of epidemic behaviors that can be obtained from a single parameterization of our model, rather than attempt to calibrate it to a specific disease or outbreak. We illustrate our model using a case with a population of 100,000 agents with one initially infected person. We ran the model for 5,000 replications at three different transmission rates: 0.8, 1.4, and 2.0 average new exposures per infectious agent. Histograms and summary statistics are given in Figure 2. Note that the scales for the histograms vary by orders of magnitude. The 90^{th} percentile for the lowest transmission rate is 8 people—less than 0.01% of the population. In contrast, with transmission rates of 1.4 and 2.0, the 90^{th} percentiles are approximately 34% and 69% of the population infected, respectively—indicating that the outcomes are potentially devastating.

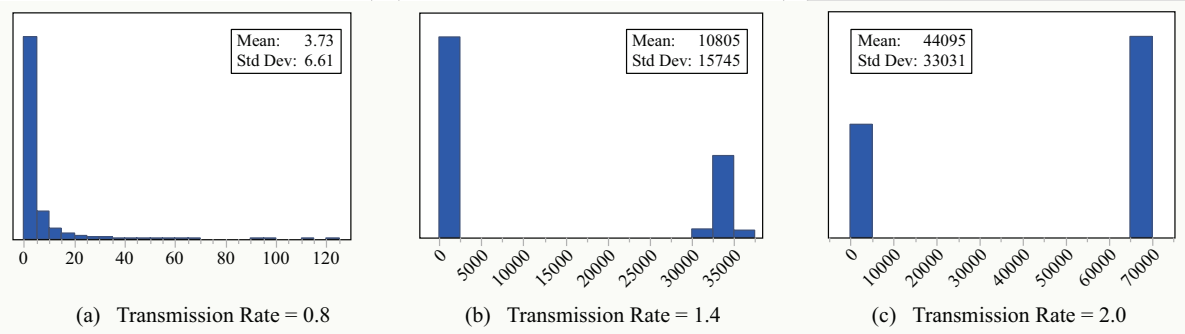

(a) Transmission Rate = 0.8 (b) Transmission Rate = 1.4 (c) Transmission Rate = 2.0

Figure 2: Epidemic distributions for transmission rates of 0.8, 1.4, and 2.0, in a population of 100,000 with one initial infection.

Deterministic aggregate models report only one outcome for a given transmission rate. With stochastic models, we see there is some potential good news—many potential epidemics fizzle. With a transmission ratio of 0.8 the epidemic fizzled 100% of the time. The results are strikingly bimodal for the higher transmission rates: the epidemic fizzled 68% of the time with a transmission rate of 1.4, and 36% of the time with a rate of 2.0. In contrast, a deterministic SEIR model will always flare with a transmission rate greater than 1.0. Given the bi-modal nature of many of these configurations, it is clear that means-based analysis tells a very misleading story about what might happen. Note that the sample means for the two highest transmission rates in Figure 2 are far away from either mode, making them singularly unsuitable for discussing likely cases.

Although there have been a number of smaller Ebola outbreaks in sub-Saharan Africa since 1976, most of them run their course infecting anywhere from a handful to a few hundred cases. However, the 2013 epidemic flared into a pandemic (Corum 2014). Although we have not attempted to calibrate our model to the Ebola epidemics, our model's qualitative behavior corresponds to the pattern observed in the real world. Relatively few outbreaks actually flare to pandemic levels—and even for situations with quite high transmission rates, pandemics are not a given.

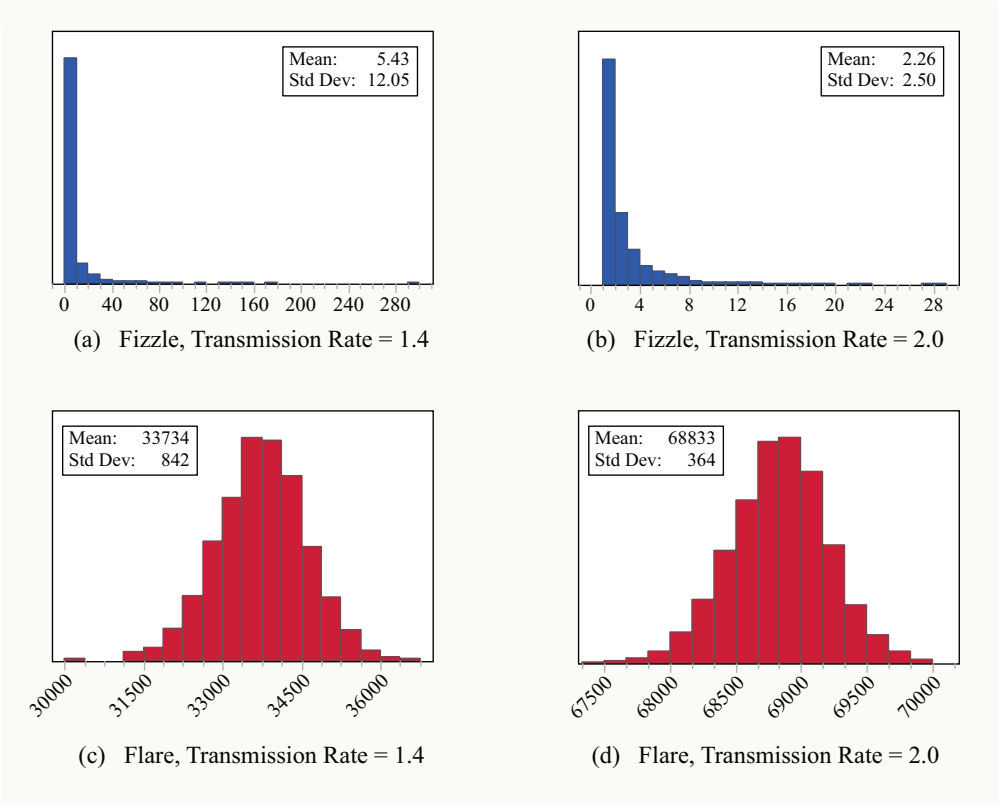

Figure 3: Distributions of fizzled epidemic behavior (in blue, at most 10% of the population infected) and flared epidemic behavior (in red, over 10% of the population infected) for transmission rates of 1.4 and 2.0, in a population of 100,000 with one initial infection.

Now let's take a closer look at the fizzle and flare distributions in Figure 3, where we have separated the original distributions into distinct histograms based on whether or not 10% of the population became infected. Once again, the scales differ. Note that all of the fizzle distributions are skewed right, indicating that low infected counts are most likely to occur. A transmission rate of 0.8 was sufficiently low that no flares happened in the 5,000 replications. By contrast, with a higher transmission rate, once a critical mass of infected agents is achieved, a major epidemic is a virtual certainty. This corresponds more closely to the predictions based on aggregate models. Still, the histograms provide additional information about the variability that can occur in the end result.

6 CONCLUDING REMARKS

In a recent paper, Lucas et al. (2015) assert that it is time for simulation, rather than a stylized analytical model, to be considered a first resort for modeling complex phenomena. Specifically, they point out that (1) analytical models often make oversimplifying assumptions for modeling convenience, while simulation models need not; (2) simulation models are computationally tractable, in the sense that simulation results, such as graphs of relationships, can be made as precise as desired; and (3) simulation models lend themselves to large-scale computational experiments.

Our discrete-event, stochastic, agent-based simulation model makes fewer assumptions than current analytical models: it does not require exponential transition times between different disease states. It reflects the wide variety of trajectories that an epidemic may follow—a single starting configuration yields a distribution of end states. In contrast, deterministic models yield single trajectories—depending on the

starting conditions, the epidemic will always fizzle or always flare. Historically-accepted models often fail to capture even the qualitative behavior of outbreaks. Given that deterministic models often predict disaster, which then fails to materialize, they have fostered an attitude of skepticism towards well-intended public health warnings. Stochastic models such as ours allow for risk assessment.

Second, even for the deterministic SEIR models, exact distributional results of epidemic spread are not possible to obtain computationally in a reasonable amount of time. Consequently, despite all the simplifying assumptions, the results must be approximated. Our simulation model scales more readily than time-stepped models, and so it is more computationally tractable. Thus we can explicitly assess the probabilities of different epidemic sizes, and by replicating the experiment we can make these results as precise as desired.

As a result of the differences highlighted above and in Section 4, the new model is not subject to the inaccuracies introduced as modeling artifacts in time-stepped models. The model prototype, implemented in a singly-threaded scripting language, is capable of simulating an epidemic in a population of one million in less than 40 seconds on a modern laptop, i.e., it is extremely scalable. We further anticipate that the model will run one to two orders of magnitude faster when implemented in a compiled language, and a well-implemented parallelized version might achieve even greater speed-ups.

With simulation we are not constrained in our distributional choices. We can construct the model using realistic distributions when data are available, and use data farming techniques to explore the impact of various distributional choices and parameterizations when data are unavailable. Last, but certainly not least, we can look at a much richer and more informative set of output measures, such as quantiles of the distributions of case counts or epidemic duration. Given the bi-modal nature of many scenarios, it is clear that means-based analysis gives a very incomplete and misleading representation of what is happening.

Although we have not done so in this paper, our model can easily be expanded to incorporate additional compartments that represent disease states or subpopulation groupings, and interventions that change the dynamics of the system. potential contacts for each agent. These are areas of ongoing exploration.

REFERENCES

Al Rowaei, A., A. H. Buss, and S. Lieberman. 2011. "The Effects of Time Advance Mechanism on Simple Agent Behaviors in Combat Simulations". In *Proceedings of the 2011 Winter Simulation Conference*, 2431–2442. Piscataway, NJ: Winter Simulation Conference: IEEE.

Andradóttir, S., W. Chiu, D. Goldsman, M. L. Lee, K.-L. Tsui, B. Sander, D. N. Fisman, and A. Nizam. 2011. "Reactive Strategies for Containing Developing Outbreaks of Pandemic Influenza". *BMC Public Health* 11(Suppl 1) (S1): 1–15.

Broeck, W. V. d., C. Gioannini, B. Gonçalves, M. Quaggiatto, V. Colizza, and A. Vespignanni. 2011. "The GLEaMviz Computational Tool, a Publicly Available Software to Explore Realistic Epidemic Spreading Scenarios at the Global Scale". *BMC Infectious Diseases* 37 (11). Available via http://www.biomedcentral.com/1471-2334/11/37 [accessed 16 July 2015].

Chao, D. L., M. E. Halloran, V. J. Obenchain, and I. M. Longini Jr.. 2010. "FluTE, a Publicly Available Stochastic Influenza Epidemic Simulation Model". *PLoS Computational Biolology* 6 (1): e1000656.

Corum, J. 2014, 29 December. "A History of Ebola in 24 Outbreaks.". New York Times.

Flanagan, D., and Y. Matsumoto. 2008. *The Ruby Programming Language*. First ed. O'Reilly.

Grefenstette, J. J., S. T. Brown, R. Rosenfeld, J. DePasse, N. T. B. Stone, P. C. Cooley, W. D. Wheaton, A. Fyshe, D. D. Galloway, A. Sriram, H. Guclu, T. Abraham, and D. S. Burke. 2013. "FRED (A Framework for Reconstructing Epidemic Dynamics): An Open-Source Software System for Modeling Infectious Diseases and Control Strategies Using Census-Based Populations". *BMC Public Health* 13 (Article 940): 1–14.

Kermack, W. O., and A. G. McKendrick. 1927. "A Contribution to the Mathematical Theory of Epidemics". *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 115 (772): 700–721.

King, A. A., M. D. de Cellès, F. M. G. Magpantay, and P. Rohani. 2015. "Avoidable Errors in the Modelling of Outbreaks of Emerging Pathogens, with Special Reference to Ebola". *Proceedings of the Royal Society of London B* 282 (20150347): 1–7.

Lucas, T. W., W. D. Kelton, P. J. Sánchez, S. M. Sanchez, and B. L. Anderson. 2015. "Changing the Paradigm: Simulation, Now a Method of First Resort". *Naval Research Logistics* 62 (4): 293–305.

Merler, S., M. Ajelii, L. Fumanelli, M. F. C. Gomes, A. P. y Piontti, L. Rossi, D. L. Chao, I. M. Longini, M. E. Halloran, and A. Vespignani. 2015. "Spatiotemporal Spread of the 2014 Outbreak of Ebola Virus Disease in Liberia and the Effectiveness of Non-Pharmaceutical Interventions: A Computational Modelling Analysis". *The Lancet Infectious Diseases* 15 (2): 204–211.

Parker, J., and J. A. Epstein. 2011. "A Distributed Platform for Global-Scale Agent-Based Models of Disease Transmission". *ACM Transactions on Modeling and Computer Simulation* 22 (1): Article 2, 2:1–2:25.

Perumalla, K. S., and S. K. Seal. 2012, July. "Discrete Event Modeling and Massively Parallel Execution of Epidemic Outbreak Phenomena". *Simulation* 88 (7): 768–783.

Ruby Community. 2015. "Ruby: A Programmer's Best Friend". https://www.ruby-lang.org/en/ [accessed 16 July 2015].

Sanchez, P. J. 2006. "As Simple As Possible, But No Simpler: A Gentle Introduction to Simulation Modeling". 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, 2–10. Piscataway, NJ: Winter Simulation Conference: IEEE.

Schruben, L. W. 1983. "Simulation Modeling with Event Graphs". *Communications of the ACM* 26 (11): 957–963.

Siettos, C., C. Anastassopoulou, L. Russo, C. Grigoras, and E. Mylonakis. 2015, 9 March. "Modeling the 2014 Ebola Virus Epidemic—Agent-Based Simulations, Temporal Analysis and Future Predictions for Liberia and Sierra Leone". *PLoS Currents Outbreaks* Edition 1, doi: 10.1371/currents.outbreaks.8d5984114855fc425e699e1a18cdc6c9.

Washburn, A. 2015. "GENMOD, a General Epidemic Model". Technical report, Naval Postgraduate School, Monterey, California.

WHO Ebola Response Team. 2014. "Ebola Virus Disease in West Africa — The First 9 Months of the Epidemic and Forward Projections". *New England Journal of Medicine* 371 (16): 1481–1495. PMID: 25244186.

World Health Organization 2015, May. "Situation Summary: Latest Available Situation Summary, 06 May 2015". http://apps.who.int/gho/data/view.ebola-sitrep.ebola-summary-latest?lang=en [accessed 16 July 2015].

AUTHOR BIOGRAPHIES

PAUL J. SANCHEZ is a faculty member in the Operations Research Department at the Naval Postgraduate School. His research interests span most of the simulation field, with a focus on large scale design of experiments. His email address is pjsanche@nps.edu.

SUSAN M. SANCHEZ is a Professor in Operations Research at the Naval Postgraduate School, and Co-Director of the Simulation Experiments & Efficient Designs (SEED) Center for Data Farming. She also holds a joint appointment in the Graduate School of Business & Public Policy. She has a B.S. in Industrial & Operations Engineering from the University of Michigan, and a Ph.D. in Operations Research from Cornell. She has long been active in within the simulation community, including the WSC Board of Directors. Her web page is http://faculty.nps.edu/smsanche/, and her email is ssanchez@nps.edu.