Proceedings of the 1999 Winter Simulation Conference
P. A. Farrington, H. B. Nembhard, D. T. Sturrock, and G. W. Evans, eds.

GERMS: AN EPIDEMIOLOGIC SIMULATION TOOL FOR STUDYING GEOGRAPHIC AND SOCIAL EFFECTS ON INFECTION TRANSMISSION

Andrew L. Adams
James S. Koopman

Department of Epidemiology, SPH-1
The University of Michigan
109 Observatory St.
Ann Arbor, MI 48109, U.S.A.

Stephen E. Chick
Peter J. Yu

Department of Industrial and Operations Engineering
The University of Michigan
1205 Beal Avenue
Ann Arbor, MI 48109, U.S.A.

ABSTRACT

The analysis, surveillance, and control of infectious diseases are important functions of public health organizations around the world. This article describes the design and implementation of simulation tools that include several innovations for modeling infectious disease transmission. These tools address several important issues for understanding the epidemiology of sexually transmitted infections. The model accounts for realistic infection transmission systems by explicitly modeling (i) heterogeneous populations of individuals with varying social and geographic characteristics, (ii) complex interaction between individuals to characterize opportunities for transmission, (iii) infection characteristics such as transmission probabilities and infection duration, and (iv) contact and infection histories. Since public health organizations collect and use information regarding infected individuals, including geographic location and partnership data, the tool is well equipped to help evaluate the effectiveness of interventions based on that data. We outline design decisions and present results of initial simulation analysis. We also discuss short-term goals for extending the simulation toolkit to address specific needs of the Centers for Disease Control.

1 INTRODUCTION

In addition to human suffering, the major sexually transmitted infections (excluding HIV) result in $10 billion in annual indirect and direct costs (Institute of Medicine 1997). There is therefore great potential benefit in developing tools to help design effective infection surveillance and control programs.

This article describes the design and implementation of GERMS (Geographic-Environmental Reinfection Modeling Simulator), a simulation toolkit that includes several innovations for modeling the transmission of infectious diseases. The development of GERMS is part of a three-year project to provide the CDC with analytical tools to augment the decision-making process for resource allocation for sexually transmitted infections.

The GERMS toolkit address several issues that are particularly important for understanding the epidemiology of sexually transmitted infections. In particular, the model accounts for realistic infection transmission systems by explicitly modeling (i) heterogeneous populations of individuals with varying social and geographic characteristics, (ii) complex interaction between individuals to characterize opportunities for transmission, (iii) infection characteristics such as transmission probabilities and infection duration, and (iv) contact and infection histories. Since public health organizations collect and use information regarding infected individuals, including geographic location and partnership data, this tool is well equipped to help evaluate the effectiveness of interventions based on those data.

Section 2 describes alternate model formulations in order to motivate the GERMS model formulation. Section 3 describes the implementation of the first phase of the project. Section 4 presents preliminary results, including a verification analysis that compares closed-form analytical results with the output of simulations of simplified populations. This verification helped both ensure the correctness of the computer implementation of the model, and provided some insights of epidemiological relevance. Section 5 indicates directions for further development.

2 EXISTING MODELING APPROACHES

Two common types of models for the spread of infection are compartmental models and stochastic simulations. Each has benefits and drawbacks.

2.1 Compartmental Models

Compartmental models (for example, see Jacquez 1996) are a widely used method for studying the steady state and transient dynamics of infection. A desirable attribute of compartmental models is that analytical results can often be obtained for a given model, resulting in valuable insights into the infection process. In this context, a compartmental model is a set of differential equations that describe how the number of infected and susceptible individuals changes through time. For example, for an infection where individuals follow the pattern susceptible-infected-susceptible (SIS):

$$\frac{dS}{dt} = -c\beta\frac{SI}{S+I} + \frac{1}{d}I - mS + R$$
$$\frac{dI}{dt} = c\beta\frac{SI}{S+I} - \left(m + \frac{1}{d}\right)I$$

where S and I denote the number of susceptible and infected individuals at time t, d denotes the mean duration of infection, c denotes the number of contacts per unit time, β denotes the transmission probability per contact, m denotes an underlying mortality rate, and R denotes the recruitment rate for new susceptibles. See Kurtz (1981) and Jacquez (1996) for a discussion of how deterministic compartmental models relate to stochastic population models for arbitrarily large population sizes, and when the duration of infection is exponentially distributed.

A disadvantage of compartmental models is that the identity of individuals is not maintained. This implies that compartmental models cannot analyze surveillance and intervention strategies that require knowledge of characteristics of individuals. This includes the strategy of contact tracing, the identification of individuals that are known to be infected.

Furthermore, it may be difficult for a compartmental model to account for several factors that greatly influence the dynamics of infection. This motivates microsimulation of populations, as described in Section 2.2.

2.2 Discrete-Individual Simulations

Microsimulations have been studied as means of overcoming some of the drawbacks of compartmental models. Microsimulations are discrete-event stochastic simulations of individuals in a host population. Opportunities for the transmission of infections are explicitly modeled as part of the population dynamics.

For example, Morris and Kretzschmar (1997) indicate that partnership concurrency in a sexually active population has a tremendous effect on the rate of epidemic rise for HIV infection. The work of Welch, Chick, and Koopman (1998) extends this result by indicating that the rate of sexual contacts as a function of the number of partners is also a significant factor for infection dynamics for Gonorrhea (GC). Ghani, Swinton and Garnett (1997) also study the effects of partnership networks on the epidemiology of GC. Both Longini et al. (1998) and Adams et al. (1998) describe simulations that quantify the value of information about the partners of infected individuals, where the objective is to evaluate HIV vaccine trial designs. Also see van der Ploeg et al. (1998) and Habbema et al. (1996).

While microsimulation has been used to gain many insights for individual phenomena, such as those listed above, there are some difficulties. First, many studies focus on a relatively limited set of factors that affect the infection dynamics. Second, the structure of the mathematical models underlying the simulation is typically so complex that it is difficult to determine analytical solutions for special cases, and it is difficult to identify the real-world meaning of some input parameters.

3 GERMS MODEL FOR POPULATION AND INFECTION DYNAMICS

GERMS is a microsimulation that is innovative in several ways. First, many previous microsimulation studies focus on a few input factors that describe infection and population dynamics. In contrast, GERMS incorporates a wide range of aspects, thereby allowing for the study of a broad set of interactions among input factors. GERMS also permits complex mixing of individuals so that many of the complexities known to influence the transmission of infection are modeled. Furthermore, the histories of contacts and infections for each simulated individual in the population are maintained, allowing surveillance programs such as contact tracing to be modeled. In addition, the model is designed so that for special cases closed form solutions for a number of relevant quantities can be determined, including quasi-equilibrium prevalence levels (Jacquez and Simon 1993). This allows for verification of correct computer implementation of the mathematical models for population dynamics and infection transmission. GERMS is designed to provide insight by simulating a broad class of infection transmission systems. Our initial analysis is therefore not validated with respect to a specific population. The general structure of the model, however, is designed to have face validity for modelers of infectious diseases.

GERMS explicitly models:

- The identity of individuals, including infection history and personal attributes
- Contact patterns for partnership formation
- Transmission of infection

3.1 Formal Mathematical Specification

The model can abstractly be thought of as a network-valued stochastic process. Each node represents an individual in the population. Arcs may be added or removed through time as relationships between individuals are formed and dissolved.

3.1.1 Identity of Individuals

Non-homogenous populations can be specified by assigning different parameter values to each simulated individual. Some of these parameters include:

Gender: Male or Female
Sexual Preference: Male or Female or Both
Rate, λ, of seeking new partnerships when unpartnered
Damping factor, θ, for seek rate, per concurrent partner
Partnership Profile: Monogamous or Polygamous
Geographical home (Cartesian coordinates)
Social group: (one or more of 32 groups that represent membership in a socio-economic category that may be important to distinguish for an epidemiologic study)

Most parameters are readily understood, except perhaps the damping factor θ. When $\theta = 1$, individuals constantly seek new partners, regardless of their current number of partners. Individuals with $\theta=0$ are monogamous, and when θ is in between, the rate of seeking new partners declines as a function of the current number of partners (see equation 1 below for a formal definition).

Additional parameters include probabilities of seeking treatment, given that symptoms are noticed; the probability of reporting a given partner to medical authorities when reporting is requested; and others.

During a simulation run, state variables are maintained to record information for each individual, such as (1) the identity of each partner, if any, (2) the infection status, (3) the identity of the most recent few partners, as well as timing information, and (4) the most recent infection and recovery times, if applicable.

3.1.2 Contact Patterns

Contact patterns determine the manner in which partnerships are formed and dissolved through time. Previous work has indicated that many aspects of the complex partnering patterns observed in the field may have a strong influence on the dynamics of the infection in a population (as discussed briefly in Section 2).

Contact patterns are determined in GERMS by letting individuals "seek" partnerships in "activity settings", known in GERMS as *bins*. Each bin has both geographic (Cartesian coordinates) and social coordinates (e.g., only individuals from certain social groups can form a relationship in a given bin). The bin concept was motivated by empirical models of Jacquez et al. (1989) to describe the complex mixing processes observed in the field. Bins allow for a rich combination of assortive ("birds of a feather") and disassortive ("opposites attract") mixing patterns.

Each individual splits his/her seeking of a partnership amongst one or more bins, so that a fraction, f_{ij}, of individual j's time is spent seeking a partner in bin i. The f_{ij} will be positive for bins that allow an individual to enter, and may be greater for geographically "close" bins (there is flexibility for assigning the f_{ij}). The seeking rate for individual j depends on their base seek rate λ_j, their damping factor θ_j, and their current number of partners n_j. The seek rate of individual j in bin i is then:

$$\xi_{ij} \equiv f_{ij} \lambda_j \theta_j^{n_j} \qquad (1)$$

(we use $\theta_j = 0$ for monogamy, and allow no seeking for partnered monogamous individuals).

The rate of partnership formation between two individuals is assumed to be a function of the seek rates of each individual. The rate, r_{ijk}, of partnership formation between individuals j and k in bin i is defined as follows:

$$r_{ijk} \equiv \frac{\xi_{ij} + \xi_{ik}}{2}$$

assuming that relationship formation is permissible, considering monogamy contraints, and $r_{ijk} = 0$ otherwise. The decision to define r_{ijk} in terms of the arithmetic mean of individual seek rates was motivated by the fact that this relationship satisfies the Fredrickson/McFarland properties (Castillo-Chavez 1989) for partnership mixing and allows for relatively fast simulations, as described in Section 3.2 below.

In GERMS, partnerships form between explicitly modeled individuals and persist over a finite time span. This time span is divided into two phases; the courtship phase and the relationship phase. The relationship phase is defined by the onset of sexual activity and is the period during which infection transmission can occur between two partners. The length of each phase has a Gamma distribution.

3.1.3 Transmission of Infection

Transmission can occur only within the context of a relationship between two individuals. Our model allows for the probability of transmission to depend on the genders of the infector and infectee, as well as the rate of contacts, which is a parameter of the bin.

3.2 Discrete-Event Simulation Design and Implementation

GERMS models a population of individuals who mingle in various activity settings, form sexual partnerships that persist over time, and who may infect or become infected by a partner. These processes of partnership formation and infection, along with some of their algorithmic complexities, are described below.

3.2.1 Partnership Life-Cycle

Partnership formation events are sampled by first determining a time of next partnership formation, then by determining the bin and specific individuals involved in the partnership. Since each partnership occurs at a given rate, and the rate may change through time as other partnerships are formed or terminated in the population, partnership formation times are randomly sampled in accordance with a non-homogenous Poisson process. The implementation essentially inverts the cumulative of a non-homogenous Poisson process by distributing calculations through time, performing one calculation per partnership formation or breakup event.

A next potential partnership formation time is sampled at the start time, t_a, of every partnership. At that time each bin is queried for the rate at which partnerships are forming in it and these rates are summed to yield an overall partnership formation rate, R_1. The time, Δt, until the next partnership begins is then sampled from an exponential distribution with mean $1/R_1$ and the next partnership is scheduled to begin at time $t_s = t_a + \Delta t$. A problem with t_s arises if another partnership ends at time $t_b < t_s$. When a partnership ends, the partnership formation rate increases in each bin in which the partners circulate, yielding a new overall rate, R_2. Since Δt was generated for a random process having rate R_1, it would be incorrect to schedule an event at t_s when the process now has rate R_2. Therefore t_s is re-scaled by the ratio of the old and new rates as follows.

$$t_s := t_b + (t_s - t_b)\frac{R_1}{R_2}$$

This type of re-scaling occurs for each partnership that terminates prior to t_s, until the next partnership formation event actually occurs.

3.2.2 Randomly Sampling Partners

The probability of partners being sampled from a particular bin is proportional to that bin's contribution to the overall partnership formation rate, R. Once the bin is determined, the partners must be sampled. In the discussion that follows, only a two-sided bin (which allows for partnerships with two distinct roles, e.g. heterosexual partnerships) is considered. However, an analogous process can be constructed for one-sided bins for homosexual partnerships.

The probability of two people being sampled from a given bin is proportional to the rate at which they are currently forming partnerships. Partners are sampled by first sampling the first person, then sampling the second person conditional on the identify of the first person. In principle, this process requires checking each pair of individuals circulating in the bin to determine the rate at which they seek each other, assuming they are still compatible. Two people are *compatible* (a) if they are not already in a partnership with one another and (b) if neither person is monogamous and already in a partnership. Once the set of potential partners is determined for the first sampled person, the probability distribution over the set must be calculated and the second partner sampled.

An exact implementation of the process described above has a computational time complexity (Knuth 1973) of $O(N^2)$ per partnership formed, where N denotes the population size. Because there are $O(N)$ partnerships per unit simulated time, the overall time would be $O(N^3)$. A more computationally efficient approach is needed. The probability, p_{ij}, of selecting person j from bin i, for individuals j that can form partnerships in bin i, can be approximated as:

$$p_{ij} \propto \left[\frac{\xi_{ij}}{2} + \frac{1}{2}\sum_{k|side2}\frac{\xi_{ik}}{w_{i1}}\right]$$

where w_{i1} is the number of people on side one of bin i, excluding partnered monogamous people. This approximation is exact for entirely monogamous populations. However, it becomes less accurate as the fraction of polygamous individuals increases because it counts as compatible those people on side two with whom person j is already involved. An overall simulation time complexity of $O(N \log N)$ is then achieved by combining this approximation with (1) a balanced binary containing a node for each person, and (2) a rejection sampling method to reject partnerships that are incompatible, but were inappropriately sampled because of the approximation. The number of rejections until a compatible partner is found is geometrically distributed and for large populations and/or low partnership formation rates, a partner is quickly determined.

3.2.3 Infection Transmission and Natural History

GERMS currently implements what is termed a Susceptible-Infected-Susceptible (SIS) model of infection transmission. In such a model, individuals are susceptible to infection, become infected and subsequently return to the susceptible state. The modeled transmission mechanism is sexual contact. An example of a sexually transmitted SIS infection is Gonorrhea.

When a partnership enters its relationship phase (defined by the onset of sexual activity) the possibility exists for infection transmission. If exactly one of the partners is infected, a *potential* time to transmission is sampled from an exponential distribution with mean depending on the rate of contacts and the probability of infection per contact. An infection event is scheduled if the resulting transmission time occurs during the partnership's relationship phase and before the infected partner is scheduled to recover.

When a person's infection begins, each of the person's partners is examined. If a partner is not already infected, then a potential time to transmission is sampled and conditionally scheduled as described above. Finally, the potential exists for re-infection from an infected partner when an infection clears. At that time each of the newly cured person's partners is examined and if a partner is infected, then a potential time to transmission is sampled and conditionally scheduled.

In the near future GERMS will be expanded to include additional notions of infection surveillance and intervention. If an infected individual is cured by some external intervention, such as the administration of antibiotics, then all scheduled transmissions for which that individual is responsible are cancelled.

3.2.4 Simulation Engine and Operational Issues

Previous experience (Adams et al. 1998, Welch et al. 1998) suggested that an object-oriented design is well-suited to the type of infectious-disease modeling being explored here. Other requirements for a simulation package included a complete programming language, the ability to link in C/C++ objects, a mature code base and the availability of technical support. After reviewing both commercial and freely available packages, MODSIM III from CACI Products Company running on Microsoft Windows NT 4.0 was finally selected.

The experience with MODSIM III has thus far been positive. Technical support has been adequate and no problems have been discovered in the MODSIM III code base. Another desirable feature of MODSIM III is that one specifies processes rather than individual events. Partnership and infection life-cycles are very naturally expressed with this abstraction. A minor deficit in using MODSIM III is that time and space optimization efforts have been hampered by the lack of suitable profiling tools.

GERMS is run from the command line, reads simulation parameters from several input files and writes several output files. The input files specify the characteristics of the infection, people and bins. Other input parameters allow GERMS to be placed in various special modes that facilitate testing and the construction of regression test suites for verifying the simulation code. The output files optionally contain a record of all simulation events and a report containing measures of interest collected at regular intervals during the simulation run. In addition, the complete state of the simulation is saved at the end and can be read in for a later run. This allows one to run GERMS for a burn-in period, save the state, and start future simulation experiments from this state.

3.3 Population Editor Interface

Previous experience and early requirements analysis for the current work indicated that some sort of graphical input file generator would be extremely useful. Such a tool would allow for easier creation and visualization of a population (and the associated input files) on which to base a set of simulation runs. The Population Editor Interface was created for this purpose.

The Population Editor allows for the easy creation of an arbitrary number of bins and sub-groups of individuals. Using a point-and-click interface, bins can be precisely placed on a geographical map and their attributes (i.e. seek rate for individuals circulating in the bin) modified. Similarly, groups of individuals can be created with similar attributes (i.e. gender, sexual preference, etc.) and randomly distributed over the map. Clicking on a specific individual displayed on the map pops up a dialog box containing the values of the individual's attributes. This dialog box allows one to fine-tune attributes of individuals.

Several features make it easy to identify individuals or groups of individuals of interest. A "zoom" feature makes it easy to zero-in on individuals in a given geographic area. For identifying groups of individuals with similar attributes, the Population Editor allows the construction of simple queries against which the population is searched. Individuals matching the query can be color-coded for easy identification.

In addition to allowing for the specification of bins and populations, the Population Editor also allows one to specify parameters associated with the infection process; i.e. infection duration, probability of transmission per sex act, and recovery rate. Finally, the Population Editor allows one to specify the number of reporting intervals and the duration of each interval.

The Population Editor Interface was developed using Microsoft Visual Basic 5.0 on the Microsoft Windows NT 4.0 platform. The ability to rapidly prototype graphical interfaces was a main consideration in the choice of Visual Basic.

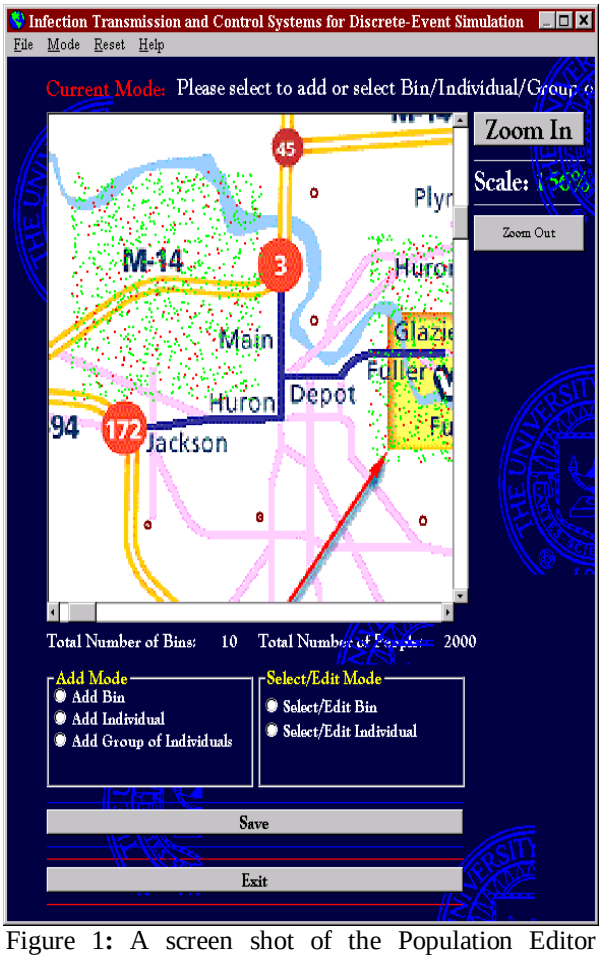

Figure 1: A screen shot of the Population Editor Interface after 2000 individuals (small dots) and 10 bins (small circles) are defined.

4 OBSERVATIONS FROM INITIAL ANALYSIS

One claimed advantage of GERMS is that closed-form solutions are available for simplified input parameter settings. During the verification phase of our simulation development, we ran simulation experiments on test populations created with the Population Editor to test whether or not the simulation response indeed corresponded to theoretical values. We describe here some results that confirm correct implementation of the population dynamics and SIS infection process.

For the simplified test populations, we assumed a closed population (no recruitment or departures) of N=1000 males and N=1000 females that seek monogamous relationships in a single bin. All individuals have the same rate ξ=1/14 days. Partnerships break up at rate σ=1/14 days, the rate of contact during a partnership is ϕ=3/7 (3 per week), with per-contact transmission probability h=0.3 when exactly one partner is infected. The infection duration is exponentially distributed with mean $1/\rho$=55 days, a value that is reasonable for gonorrhea.

For this scenario, the stationary distribution has 0 prevalence (infection dies out). However, infection levels hover around a pseudo-equilibrium endemic level of infection for extended periods of time (the stationary zero prevalence level may take extremely long time periods to be obtained). For these experiments, the pseudo-equilibrium level is the prevalence level such that a newly infected individual infects one additional person, in expectation. Chick et al. (1999) shows that

$$30N\,\frac{2\sigma\xi(\xi+\rho)}{(\xi+\sigma)\xi}\,\frac{\sigma(\sigma+2h\phi+2\rho)-\left(\sigma^2+\sigma h\phi+3\sigma\rho+2\rho^2\right)\left(\frac{\xi+\rho}{\xi}\right)}{\sigma h\phi}$$

people are infected at the time of partnership breakup, per month, assuming the units of rates are per day.

Table 1 (below) summarizes the first set of initial experiments with the above parameters, as well as with some variations on the parameters. Variations on the base case allowed for variations on one parameter individually. Simulation analysis had the population mixing for 1 year, infection was introduced into the population at time 1 year, and initial analysis indicated that pseudo-equilibrium levels were reached in approximately 3-5 years. Simulation statistics are based on batch mean analysis with 16 batches of 1 year, and give a 95% confidence level/credible set using a t-statistic approximation. The analysis was roughly the same for 32 batches of 6 months. Statistical tests for the independence of batches indicated that most correlations are not significant, but that there may be a potential positive correlation between adjacent batches for the Base Case, so the CI may be slightly overconfident for that case.

Table 1: Theoretical and Simulation Estimates for Number Infected at End of Partnership, Per Month

Experiment	**Theoretical**	**95% Sim Est.**
Base Case	710.6	(686, 737)
Base. but ξ=1/7	1692.4	(1664, 1705)
Base, but σ=1/21	572.7	(542, 587)
Base, but h=0.5	1228.0	(1203, 1235)
Base, but ρ=1/50	457.3	(454, 504)

All the theoretical values fell within the interval estimates from simulation, and therefore supported the assertion that the coding and the analysis are correct.

Additional experiments (results not shown) further verified the simulation code, and gave an initial feel for the

simulated infection dynamics when more realistic infection and contact parameters were used. Results indicate that the degree of concurrency with sexual partners and both parameters of the gamma distribution for infection duration strongly influence infection dynamics.

5 FUTURE DIRECTIONS

This article documents work in progress. The work already allows for the modeling of a number of behavioral changes that may be a result of intervention programs. For instance, increased condom usage can be modeled by a reduction in the probability of transmission per contact. We wish to implement a number of further extensions, however, before comparing surveillance and control programs in detail. Some extensions have already been specified and implementation is under way. Additional extensions are presently being specified. The extensions include:

1) Passive and active surveillance and control activities, including (a) self reporting of infectiousness when symptoms are noted, (b) self-reporting of recent partners to public health programs for treatment, (c) the attempt of public health workers to identify infecteds by seeking in geographically important areas or by contact tracing.
2) More complicated infection processes than SI and SIS models. First in line are syphilis and HIV infection.
3) Multiple strains of infectious agents.
4) The development of adaptive immune responses.

These changes are necessary to adequately model challenges faced by the CDC for resource allocation decisions.

6 SUMMARY

The development of GERMS is driven by decision-support needs of public health agencies that must evaluate surveillance and control programs for infectious diseases. We have developed a model that (1) for simple parameter settings, allows for closed-form solutions of figures of interest and (2) is sufficiently complex to allow for a broader set of population and infection dynamics than do most microsimulations to date. Initial simulation analysis verified the correct operation of the simulation code. The project is a work in progress, and implementation of additional complexities is underway. We hope that this work will result in a valuable resource for decision makers for controlling infectious diseases.

ACKNOWLEDGMENTS

The generous support of the Center for Disease Control, Atlanta GA, is gratefully acknowledged. The help of Szu Hui Ng is also appreciated.

REFERENCES

Adams, A.L., D.C. Barth-Jones, S.E. Chick and J.S. Koopman. 1998. Simulations to Evaluate HIV Vaccine Trial Designs. *Simulation*, 71(4):228-241.

Castillo-Chavez, C. Busenberg, S., Gerow, K. 1991. Pair formation in structured populations. In Differential equations with applications in biology, physics, and engineering, Goldstein, J.A., Kappel, F. Schappacher, W. (eds.), Lecture notes in pure and applied mathematics, 133. New York, Basel, Hong Kong, Marcel Dekker, Inc.

Chick S.E., A.L. Adams, M. Becker, M. Wilson, J.S. Koopman. 1999. The Effect of Reinfection and Partnership Formation Dynamics on the Spread of Infections With Recovery. Working document.

Kurtz, T.G. 1981. *Approximation of population processes.* Society for Industrial and Applied Mathematics, Philadelphia, Pa.

Ghani, A. C., S. Swinton, and G.P. Garnett. 1997. "The Role of Sexual Partnership Networks in the Epidemiology of Gonorrhea," *Sexually Transmitted Diseases* 24(1):45-56.

Habbema, J.D.F., S.J. de Vlas, A.P. Plaisier, G.J. van Oortmarssen. 1996. The microsimulation approach to epidemiologic modeling of helmenthic infections, with special reference so schistosomiasis. American *Journal of Tropical Medicine and Hygeine* 55(5, Supp):165-169.

Institute of Medicine (U.S.). 1997. Committee on Prevention and Control of Sexually Transmitted Diseases. The hidden epidemic: confronting sexually transmitted Diseases. Institute of Medicine. Division of Health Promotion and Disease Prevention.

Jacquez, J.A. 1996. *Compartmental Analysis in Biology and Medicine*. 3rd ed., BioMedware, Ann Arbor, Michigan.

Jacquez J.A., J.S. Koopman and C.P. Simon. 1989. Structured Mixing: Heterogeneous Mixing by the Definition of Activity Groups. In: Castillo-Chavez, C. (ed.). Mathematical and statistical approaches to AIDS epidemiology. Lecture Notes in Biomathematics, 83, Springer-Verlag, Berlin, Heidelberg, New York, London, Paris, Tokyo, Hong Kong.

Jacquez J.A. and C.P. Simon. 1993. The Stochastic SI Model with Recruitment and Deaths I. Comparison with the Closed SIS Model. *Mathematical Biosciences* 117:77-125.

Knuth, D. 1973. *The Art of Computer Programming. v1: Fundamental Algorithms.* 2nd edition. Addison-Wesley. Reading, MA.

Longini, I.M., K. Sagatelian, W.N. Rida, M.E. Halloran. 1998. Optimal vaccine trial design when estimating vaccine efficacy for susceptibility and infectiousness from multiple populations. *Statistics in Medicine*. 17(10):1121-1136.

Morris M and M. Kretzschmar. 1997. "Concurrent partnerships and the spread of HIV," *AIDS* 11(5): 641-648.

van der Ploeg, C. B. P., C. van Vliet, S.J. de Vlas, J.O. Ndinya-Achola, L. Fransen, G.J. van Oortmarssen, J.D.F. Habbema. 1998. STDSIM: A Microsimulation Model for Decision Support in STD Control. *Interfaces* 28:84-100.

Welch, G., S. E. Chick and J. S. Koopman. 1998. "Effect of Concurrent Partnerships and Sex-Act Rate on Gonorrhea Prevalence," *Simulation*, 71(4):242-249.

AUTHOR BIOGRAPHIES

ANDREW L. ADAMS is a software engineer at the University of Michigan Department of Epidemiology and at Merit Network, Inc., housed at the University of Michigan. He was the lead architect/developer for GERMS. Mr. Adams holds a BSE in Computer Engineering and an MSE in Computer Science, both from the University of Michigan. His research interests include stochastic and agent-based modeling and simulation, and Internet routing and management technologies.

STEPHEN E. CHICK is an assistant professor of the Department of Industrial and Operations Engineering at the University of Michigan, Ann Arbor. In addition to stochastic simulation, his research interests include Bayesian statistics, reliability, and decision analysis. His research is motivated by projects in manufacturing and health care, particularly epidemiology. His work experience includes several years of material handling system design for the automotive industry using simulation analysis.

JAMES S. KOOPMAN is professor of epidemiology at the University of Michigan, where he received is medical degree. After pursuing a residency in pediatrics, he began his career in epidemiology as an epidemic intelligence service officer with the Centers for Disease Control. He then received a Master's degree in Public Health. He contributed to the effort to eradicate smallpox in India in 1973 and instituted an epidemiology program in 1984 that is still ongoing in Mexico. He now dedicates himself to developing a science of transmission system analysis. In this endeavor, he helps organize public health department efforts to control infectious diseases and he designs and analyzes transmission system models. He views the science of transmission system analysis as one that has pursued divergent paths within the disciplines of applied mathematics and field epidemiology. He now works with mathematicians, operations engineers, and computer scientists to help pull this science together.

PETER J. YU is a graduate student in the Department of Industrial and Operations Engineering at the University of Michigan, Ann Arbor, concentrating on Information Systems and Simulation Analysis. He holds MS in Industrial and Operations Engineering and BS in Computer Science also from the University of Michigan. He was largely responsible for the design and implementation of the Population Editor Interface.