Simulating the contact process in heterogeneous environments

S. V. Fallert*
Department of Chemistry, University of Cambridge, Cambridge CB2 1EW, United Kingdom

J. J. Ludlam
Department of Chemistry, University of Cambridge, Cambridge CB2 1EW, United Kingdom

S. N. Taraskin
St. Catharine's College and Department of Chemistry, University of Cambridge, Cambridge CB2 1EW, United Kingdom
(Received 10 September 2007; revised manuscript received 20 December 2007; published 23 May 2008)

The one-dimensional contact process (CP) in a heterogeneous environment—a binary chain consisting of two types of site with different recovery rates—is investigated. It is argued that the commonly used random-sequential Monte Carlo simulation method which employs a discrete notion of time is not faithful to the rates of the contact process in a heterogeneous environment. Therefore, a modification of this algorithm along with two alternative continuous-time implementations are analyzed. The latter two are an adapted version of the n-fold way used in Ising model simulations and a method based on a modified priority queue. It is demonstrated that the commonly used (but incorrect as we believe) discrete-time method yields a different critical threshold from all other algorithms considered. Finite-size scaling of the lowest gap in the spectrum of the Liouville time-evolution operator for the CP gives an estimate of the critical rate which supports these findings. Further, a performance test indicates an advantage in using the continuous-time methods in systems with heterogeneous rates. This result promises to help in the analysis of the CP in disordered systems with heterogeneous rates in which simulation is a challenging task due to very long relaxation times.

DOI: 10.1103/PhysRevE.77.051125 PACS number(s): 05.70.Ln, 64.60.Ht, 02.50.Ey, 87.10.Tf

I. INTRODUCTION

The contact process (CP) [1] is a prototype process for the spread of epidemics in biological systems. It describes epidemics in a network where each node can be in one of two states: infected (I) or susceptible (S) (so-called SIS models). Infected nodes spread their infection to susceptible neighboring nodes at rate λ while recovery is spontaneous at rate ϵ. The CP is a continuous-time Markov process and exhibits a nonequilibrium absorbing state phase transition between an active and a nonactive regime of the disease, behaving at its critical point according to the directed percolation (DP) universality class. While the most part of the literature has considered the CP on ordered lattices with uniform infection and recovery rates, the influence of heterogeneity or even disorder on its behavior has recently attracted considerable interest [2–4]. The known behavior of the homogeneous CP has been established by a range of analytical and numerical techniques [2,3,5,6] such as renormalization group analysis [7], series expansions [8], Monte Carlo (MC) simulations [9,10] and diagonalization of the Liouville operator [11].

In particular, MC simulation has been an essential tool in the investigation of critical properties of the CP. In their seminal paper on the topic, Grassberger and de la Torre [9] introduced a discrete-time formulation of the continuous-time CP which is suited to computer simulation and has since become the standard method of such investigations [2]. At each time step of the simulation, an infected site is chosen at random followed by a choice of whether to spread the infection with probability $p=\lambda/(\lambda+\epsilon)$ or to recover with probability $1-p=\epsilon/(\lambda+\epsilon)$. The event of infection proceeds at a randomly selected neighboring site and is successful only if that site is in the susceptible state.

This investigation is concerned with the determination of the exact critical threshold rate, the precise knowledge of which is of prime importance in many cases, for instance, control of epidemics [12,13]. Also, the knowledge of the critical rate is vital when comparing to other methods like perturbative series expansions [14]. Even though the standard discrete-time method has been used in connection with heterogeneous systems [15,16], it must not be used in such systems if one is interested in the accurate critical rate. This is because it is not faithful to the rates in the definition of the corresponding CP which can be seen from the following qualitative arguments. In the simulation, once an infected site i has been chosen, infection or recovery are attempted with probabilities $1/(1+\epsilon_i)$ and $\epsilon_i/(1+\epsilon_i)$, respectively. This introduces an unwanted additional heterogeneity in the effective infection rates due to the normalization involving the recovery rate ϵ_i the value of which depends on the type of node i.

Therefore, a modification of the commonly used algorithm should be employed which remedies the problem and is faithful to the rates even in the presence of heterogeneity. Furthermore, it is possible to implement a continuous-time simulation of the CP in which at each step, possible events are randomly selected according to their rates along with appropriately distributed waiting times between subsequent events. Such an approach would naturally ensure a rate of unity for all possible infection events irrespective of the chosen site. This latter strategy has been employed for instance in connection with Ising model simulations [17].

*sf287@cam.ac.uk

In this paper, we first demonstrate the existence and extent of the unwanted and uncontrolled heterogeneity in effective rates. To this end, the CP is analyzed as defined on the simplest heterogeneous one-dimensional system—a binary chain consisting of two types of site A and B with identical infection rates $\lambda=1$ but different recovery rates ϵ_A and ϵ_B, respectively. A comparison between results produced by the commonly used simulation method and the modified version which rectifies the problem is carried out to quantify the deviation.

Further, we present two implementations of event-driven simulations in continuous time. First, an adopted version of the n-fold way algorithm [18] used in Ising model simulations. Second, a simulation method based on a modified priority queue. This method takes advantage of the heap data structure to pick and perform the event that happens next, and uses the locality of the CP to ensure that the influence of the chosen event only impacts minimally the heap of next events [19].

An efficiency comparison between the different simulation algorithms indicates a performance gain when using the continuous-time methods. A Liouville operator scaling approach to the problem confirms the necessity of an algorithm faithful to the rates for a quantitative determination of the critical point.

In Sec. II, the different MC simulation methods are explained in more detail. Section III defines the Liouville operator description of the CP and introduces finite-size scaling. Following on, we present our results and their discussion in Sec. IV. The findings are summarized in Sec. V.

II. MONTE CARLO SIMULATION METHODS

In the following, we present a description of three different types of simulation method for the CP in heterogeneous environments.

A. The discrete-time algorithm

As mentioned above, in the commonly used method of discrete-time simulation [2,9], the definition of time as a continuous variable is abandoned. In this method, at each step of the simulation, an infected site is chosen randomly from a list of infected sites and the next event chosen with the appropriate probabilities defined above. For infection, a random nearest neighbor is selected and the proliferation to this neighbor proceeds only if it is in the susceptible state. The time increment between such update steps is chosen as $\Delta t = 1/N_{\text{infected}}$ such that a series of N_{infected} updates constitute an integer time step. Thereby, on average each infected site has a chance to either proliferate its infection or recover during each time step. The algorithm is easy to implement and efficient as it keeps a list of infected sites only. Unfortunately, as argued in the previous section, it is unsuitable for the precise determination of the critical point in the presence of heterogeneous rates.

Therefore, the following scheme, which remedies the inappropriate normalization ought to be used. One defines $Q = \max(1+\epsilon_A,\ 1+\epsilon_B)$ and, as in the previous method, selects an active site at each simulation step. Subsequently, one of the events "infection" (I), "recovery" (R), or "wait" (W) are executed with probabilities

$$p_I = \frac{1}{Q}, \quad p_R = \frac{\epsilon_i}{Q}, \quad p_W = 1 - \frac{1+\epsilon_i}{Q} \tag{1}$$

given that active site i has been selected. This scheme reproduces the correct probabilities for the transitions between states of the associated embedded Markov jump chain. The time increment is now chosen as $\Delta t = 1/QN$. As in the previous approach, the algorithm is easy to implement and keeps a list of infected sites only. However, faithful simulation comes at the price of a waiting step, the probability of which increases with the difference between recovery rates, and thus wastes computing time.

Both schemes represent an approximation of the true CP, as the continuous-time Markov jump process is replaced by a Markov jump chain with constant holding times. Grassberger remarks [9] that by discretizing the CP, strictly speaking one loses the ability to compare the resulting critical rates to those in the original definition of the CP. However, the fixation of holding times is a good approximation at long times and the discrete-and continuous-time processes share the same stationary properties and long-time dynamics thus in particular leading to identical critical behavior [2,20].

B. The n-fold way algorithm

In order to simulate the CP in continuous time, we propose to adapt a variant of the n-fold way algorithm [17,18] which was introduced as an efficient tool for Ising model simulations. The basic idea of the algorithm is that it is event driven: at each step the simulation chooses one of the possible events, executes it and increments the clock by an appropriate waiting time. In the n-fold way, one identifies the M possible different situations a site could be in and the associated possible events, the number of which is usually small compared to the number of sites in the system. Thereby, depending on the state they are in, these sites can be grouped into M different event classes where class i has a rate r_i corresponding to the associated event and contains n_i elements; in the CP a class would for instance be an infected class with the associated event being recovery. At each simulation step, one first determines the total probability that an event will occur, $Q_M = \sum_{j=1}^{M} n_j\ r_j$. Subsequently, the class i which the next event will come from is chosen by generating a random number between 0 and Q_M and determining which partial sum $Q_i = \sum_{j=1}^{i \le M} n_j\ r_j$ it belongs to. Finally, an element from that class is randomly selected and the event executed.

Considering the forward equation of the stochastic process [17] one can show that the appropriate distribution of waiting times Δt is exponential, $P(\Delta t) \sim \exp(-Q_M \Delta t)$. Therefore, after having executed the chosen event, the clock is updated by a sample from this distribution, $\Delta t = -\ln(u)/Q_M$, where u is a uniform random number between 0 and 1.

For the one-dimensional CP, we define an infected class and two corresponding susceptible classes. The infected class contains the infected sites which could potentially recover (at rate ϵ) while the two susceptible classes contain susceptible

sites with one or two infected neighbors which could potentially become infected (at rate λ and 2λ, respectively). For a system with heterogeneous rates, additional classes can be defined and are easily incorporated in an analogous fashion. Every time an infection occurs, the chosen site along with its nearest neighbors have to change class appropriately. This book-keeping renders the *n*-fold way more involved than the straightforward discrete-time version but allows one to systematically incorporate heterogeneous rates into the CP. In comparison to the modified random-sequential algorithm, one expects a performance gain as now every update attempt will succeed.

C. The priority queue algorithm

Similarly to the *n*-fold way method, the priority queue algorithm is based on the idea of picking the next event to happen. In this case however, the times for every single possible individual event are picked from the relevant distributions, and of these the earliest calculated time is used to decide the next event to perform. Since the CP is a local process, the times already picked for events that take place far from the chosen event will be unaffected by the change, and these therefore do not need to be recalculated. The heap data structure, alongside a reverse map from lattice position to position in the heap, help minimize the effort required to perform the events efficiently.

A heap is a binary tree-based data structure, where the elements of the tree satisfy the heap condition, whereby the key value (in this case the time of the event) at each point in the tree is less than or equal to the key values in the node's left and right children [19]. Of the standard operations on a heap, we use four: Removal of the top node of the tree, insertion of a new node, alteration of an existing node, and the removal of an existing node. We also maintain a map from lattice position to the position of the relevant event in the heap. This ensures that for a given site, we can immediately locate its event in the heap without searching. In the nodes of the tree, we store the time of the event, the transition that is to happen (S to I or I to S for the CP), and the lattice site on which the event is to be performed.

The simulation proceeds as follows. After populating the heap with the initial possible events, the node at the top of the tree contains the event with the minimal time; the next event to perform. This event is removed from the heap, which is reordered to maintain the heap condition. Performing the event will cause several changes. First, a new event might now be possible at the site. If so, a time should be picked from the distribution and this new event inserted into the heap. Secondly, for each neighboring site the possible events are likely to change, and so consequently the relevant events must be either inserted into the heap, removed from the heap, or the heap must be reordered to maintain the heap condition. This is when the map from lattice position to heap-node is used in order to find quickly the relevant event in the tree. Once these updates have been performed, the event at the top of the tree is once again the next event to perform. The whole procedure is then repeated until a time cutoff, or there are no more events to perform (an absorbing state has been reached).

One advantage of this approach is that it becomes possible to use different distributions for the waiting-times for events. The exponential waiting-time distribution of the CP is particularly well suited to this as it is "memoryless"—the distribution of the time to the next event at a site is independent of its history. Other distributions would need more care to implement as they might require knowledge of the history at each site. A notable exception is the process that occurs at a fixed time, for example, the fixed infection period of Grassberger's *SIR* model [21]. One further advantage to using this method is that heterogeneous disordered rates can be incorporated easily into the algorithm. In particular, there is no need for a separate class for each rate as in the *n*-fold way. The main disadvantage of this method is in the increased complexity of the book-keeping required, since although libraries exist that implement generic heap-based structures, the additional map that must be kept in sync with the heap complicates the implementation.

III. SPECTRAL ANALYSIS OF THE LIOUVILLE OPERATOR

A. Master equation description of the CP

The following outlines the master equation description of the CP which, when written in matrix form, gives rise to the Liouville operator. For the CP on a linear chain of size N with sites $i=1,\ldots,N$ one denotes the two possible states of site i as $s_i=1$ (infected) or $s_i=0$ (susceptible). A microstate of the system can be defined as a vector $\mathbf{S}=(s_1,\ldots,s_N)^T$ and the probability of finding the system in a specific microstate at time t is denoted as $P(\mathbf{S},t)$. Assuming the transition rates between microstates $\mathbf{S}$ and $\mathbf{S}'$ to be $w_{\mathbf{S}\to\mathbf{S}'}$, the time evolution of this probability is expressed by the master equation [6],

$$\partial_t P(\mathbf{S},t) = \sum_{\mathbf{S}'} w_{\mathbf{S}'\to\mathbf{S}} P(\mathbf{S}',t) - \sum_{\mathbf{S}'} w_{\mathbf{S}\to\mathbf{S}'} P(\mathbf{S},t). \tag{2}$$

We can introduce a time-dependent probability state vector for the system by

$$|P(t)\rangle = \sum_{\sigma} P(|\sigma\rangle,t)|\sigma\rangle \tag{3}$$

such that $\langle\sigma|P(t)\rangle = P(|\sigma\rangle,t)$ with $|\sigma\rangle$ the orthonormal basis diagonal in the occupation number representation. For normalization, the elements of the probability state vector have to sum to unity.

The master equation Eq. (2) can then be recast in compact form

$$\partial_t|P(t)\rangle = -\hat{\mathcal{L}}|P(t)\rangle, \tag{4}$$

where $\hat{\mathcal{L}}=\Sigma_{\sigma,\sigma'}\hat{\mathcal{L}}_{\sigma'\sigma}|\sigma'\rangle\langle\sigma|$ is the non-Hermitian Liouville operator, the matrix elements of which, $\hat{\mathcal{L}}_{\sigma'\sigma}\equiv\langle\sigma'|\hat{\mathcal{L}}|\sigma\rangle$, coincide with the transition rates from state σ to state $\sigma'\neq\sigma$ and $\hat{\mathcal{L}}_{\sigma\sigma}=-\Sigma_{\sigma'\neq\sigma}\hat{\mathcal{L}}_{\sigma'\sigma}$. If the number of occupied nodes in state σ' is greater by 1 than in state σ then these states are coupled by the transmission rate between the affected pairs of sites i and j, $\hat{\mathcal{L}}_{\sigma'\sigma}\propto-\lambda$, while in the case of a recovery of

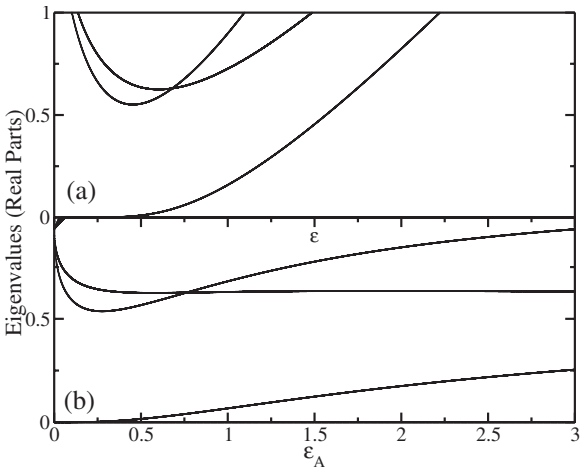

FIG. 1. The low-lying spectrum (modulus of real parts of eigenvalues) of the Liouville operator for the CP on a linear chain of length $L=8$ for the homogeneous case [(a), upper panel] and the case of heterogeneous recovery rates ϵ_A and ϵ_B for sites of type A and B, respectively, where ϵ_B has been fixed to the homogeneous critical rate ϵ_c [(b), lower panel].

an infected node i in the state σ, $\hat{\mathcal{L}}_{\sigma'\sigma} \propto -\epsilon_i$. Eq. (4) can be solved, in principle, by performing direct diagonalization of the $2^N \times 2^N$ real sparse (for lattice topologies) nonsymmetric Liouville matrix.

As the Liouville operator is non-Hermitian, it is not necessarily diagonalizable and its eigenvalues are not guaranteed to be real. We have verified diagonalizability of $\hat{\mathcal{L}}$ for the systems considered and thus assume the existence of a set of eigenvectors $|e_0\rangle, |e_1\rangle, \ldots, |e_{2^N-1}\rangle$ with eigenvalues $\mu_0, \mu_1, \ldots, \mu_{2^N-1}$ that are complete, $1=\Sigma_i |e_i\rangle\langle e_i|$. Therefore, the formal solution of Eq. (4) can be expressed as

$$|P(t)\rangle = \sum_i e^{-\mu_i t} \langle e_i|P(0)\rangle |e_i\rangle. \tag{5}$$

The trivial solution $|e_0\rangle$ of the eigenproblem for the Liouville operator with $\mu_0=0$ corresponds to the absorbing state of the system. All other eigenvectors $|e_i\rangle$ in finite systems have eigenvalues with positive real parts and thus decay exponentially with time. In the thermodynamic limit ($N \to \infty$), there is one eigenstate $|e_1\rangle$ with corresponding eigenvalue μ_1, which is zero in the active and nonzero in the nonactive phase. In a finite system, the value of μ_1 in the active (nonactive) regime approaches a zero (nonzero) value with increasing N, thereby signaling the phase transition. This behavior is illustrated in Fig. 1(a) for the CP on a linear chain of length 8. The eigenvalue that governs the decay of the active state μ_1 effectively governs the temporal correlations which is why one makes the identification $\xi_\parallel^{-1}=\mathrm{Re}(\mu_1)$ where $\xi_\parallel$ denotes the temporal correlation length.

B. Finite-size scaling

As in a finite system the spectrum does not by itself yield accurate predictions for the critical rates of the CP, finite-size scaling (FSS) has to be employed to extrapolate these rates from finite-lattice data to the thermodynamic limit. In the following we describe a procedure following Refs. [11,22]

In the contact process various quantities are known to follow scaling relations in the vicinity of the critical point [6]. The correlation lengths in the time and space directions $\xi_\parallel$ and $\xi_\perp$ diverge according to

$$\xi_\parallel \sim \xi_\perp^\theta \sim |\Delta|^{-\nu_\parallel} \sim |\Delta|^{-\nu_\perp \theta}, \tag{6}$$

where $\Delta=\epsilon-\epsilon_c$ (with ϵ_c the critical recovery rate in the homogeneous system) is the deviation from criticality, $\nu_\parallel$ and $\nu_\perp$ are universal critical exponents and $\theta=\nu_\parallel/\nu_\perp$ is the dynamical critical exponent. These scaling relations are only strictly valid in the limit of an infinite system $N \to \infty$.

However, one can consider the spatial correlation length $\xi_{\perp,L}$ in a finite system of linear size L. In the spirit of the finite-size renormalization group technique by Nightingale [23], the change of $\xi_{\perp,L}$ under real space rescaling $L \to L'$ defines the renormalized rate ϵ'

$$\xi_{\perp,L'}(\epsilon') = \left(\frac{L'}{L}\right) \xi_{\perp,L}(\epsilon). \tag{7}$$

The fixed point of this transformation ϵ_L^* defined by condition

$$\xi_{\perp,L'}(\epsilon_L^*) = \left(\frac{L'}{L}\right) \xi_{\perp,L}(\epsilon_L^*), \tag{8}$$

can be shown to converge to the critical rate in the infinite system $\epsilon_L^* \to \epsilon_c$ as $L \to \infty$ [24].

As the spatial and temporal correlation lengths are related via the dynamical exponent θ defined above, the finite-size version of these $\xi_{\parallel,L}$ and $\xi_{\perp,L}$ are related via θ_L, the dynamical exponent in a system of size L. Thus, using the scaling behavior of the correlation lengths (6) and the finite-size renormalization transformation at its fixed point (8), one obtains

$$\frac{\xi_{\parallel,L'}(\epsilon_L^*)}{\xi_{\parallel,L}(\epsilon_L^*)} = \left(\frac{\xi_{\perp,L'}(\epsilon_L^*)}{\xi_{\perp,L}(\epsilon_L^*)}\right)^{\theta_L} = \left(\frac{L'}{L}\right)^{\theta_L}. \tag{9}$$

Therefore, by comparing three system sizes L, L', L'' one obtains finite-size estimates for the critical rate ϵ_c in the infinite system and the corresponding dynamical exponent θ. By exploiting the fact that $\hat{\mathcal{L}}$ is sparse which allows the use of specialized techniques such as the Arnoldi method, we can determine the lowest nontrivial eigenvalue (previously μ_1) for one-dimensional systems of up to size $L=22$ [27]. From a sequence of eigenvalues for systems of increasing size and employing the BST extrapolation scheme [25,26] we are able to make predictions about the behavior in the infinite system.

IV. RESULTS AND DISCUSSION

In this investigation, we have considered the CP on a periodic binary chain of alternating sites of A and B type $\ldots ABAB \ldots$ with rates as defined above. MC simulations starting from a single infection seed were carried out for a range of values of ϵ_A and ϵ_B (3×10^6 realizations to a maximum of 3×10^3 time steps, system size much larger than size of active cluster). The critical point was located following [9]

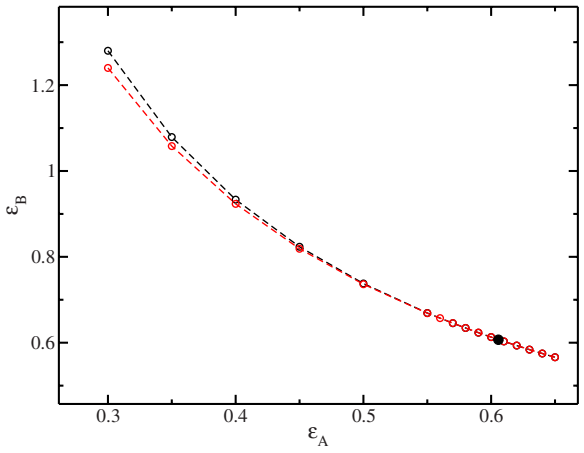

FIG. 2. (Color online) The phase diagram for the CP on a binary chain of two species A and B with recovery rates ϵ_A and ϵ_B, respectively, obtained using the commonly used discrete-time simulation algorithm (lower curve) and the modified discrete-time algorithm (upper curve) as well as the two continuous-time methods described in the text (upper curve, all data points coincide at this scale). For $\epsilon_A > \epsilon_c$, the critical line is symmetric about the bisector (not shown). The homogeneous critical point at $\epsilon_A = \epsilon_B = 0.606456$ is marked with a solid dot.

by comparison to asymptotic dynamical scaling forms. The resulting line of critical points is shown in Fig. 2 for all methods.

At the homogeneous point $\epsilon_A = \epsilon_B = \epsilon_c = 0.606456(4)$ [2] (marked with a solid dot in Fig. 2), all methods are found to agree within measurement uncertainty. When ϵ_A is moved away either way from the homogeneous critical point, the commonly used discrete-time implementation, as expected, increasingly deviates from the other methods with increasing distance from ϵ_c giving a wrong critical point. The deviations δ as a function of distance from the homogeneous critical point (for $\epsilon_A < \epsilon_c$) are found numerically to follow to leading order $\delta \approx (\epsilon_c - \epsilon_A)^\alpha$, with $\alpha \simeq 2.4$. The bottom rows of Table I list the critical values of ϵ_B obtained for the four specific points $\epsilon_A = 0.4, 0.45, 0.5, 0.55$.

In order to compare the modified discrete-time method and the continuous-time implementations, a performance test was carried out. To this end, the computing time required for each algorithm to reach a certain value of the mean density of infected sites was analyzed. Specifically, we considered a fixed number of runs at the critical point in systems of increasing heterogeneity in rates. This performance measure was chosen due to the fact that MC time in the different methods cannot be simply related to each other. For such a test to be meaningful, we require that (a) fluctuations in the simulation processes are comparable and (b) that the MC time needed to reach the chosen average density of infection is not very sensitive to the difference between rates. Indeed, fluctuations are found to be comparable between simulation methods with a similar number of infection and recovery events required. Also, the second requirement was found to be met with the MC time changing by not more than 3% for the differences in rates under consideration. In practice, the test was carried out by initially determining the MC time required to reach a fixed value for the average density of infected sites for all methods. Subsequently, the corresponding real computing time was recorded. The resulting times for the process (in arbitrary units) are shown in Fig. 3 for a range of critical points along the phase separation line with increasing difference in rates $\epsilon_B - \epsilon_A$.

As can be seen (see Fig. 3), the continuous-time methods are found to perform marginally better than the discrete-time version even at the homogeneous critical point. Moreover, with increasing difference between the rates, the computation time required by the discrete-time simulation rises while it remains constant for both continuous-time methods (whose times are found to coincide). The difference in times at the homogeneous critical point can be explained by the fact that, unlike in the continuous-time methods, in the discrete-time simulation procedure infection can fail if the randomly cho-

TABLE I. Critical rate ϵ_B for various ϵ_A obtained through the extrapolation of finite lattice data and simulation results for both the commonly used discrete time (wrong MC) as well as all other methods (correct MC) methods where ω is the free parameter of the BST extrapolation scheme.

L', L, L''	$\epsilon_A = \epsilon_B$ ϵ_B (critical)	$\epsilon_A = 0.55$	$\epsilon_A = 0.5$	$\epsilon_A = 0.45$	$\epsilon_A = 0.4$
			ϵ_B (critical)		
6,8,10	0.6139938759	0.6861248737	0.7566459788	0.8446025135	0.9572746371
8,10,12	0.6116343230	0.6808438382	0.7508606811	0.8382581629	0.9503119002
10,12,14	0.6102010124	0.6776329261	0.7473059648	0.8342843197	0.9458127224
12,14,16	0.6092784404	0.6755659904	0.7450044211	0.8316843257	0.9428206243
14,16,18	0.6086534110	0.6741658321	0.7434397946	0.8299054657	0.9407534279
16,18,20	0.6082117797	0.6731766609	0.7423318390	0.8286404733	0.9392739992
18,20,22	0.6078888141	0.6724533770	0.7415203690	0.8277112543	0.9381824293
∞	0.6067(3)	0.6699(6)	0.7387(6)	0.8244(7)	0.9339(8)
ω	0.681	0.666	0.643	0.593	0.483
correct MC	0.606457(3)	0.6694(1)	0.7382(1)	0.8238(1)	0.9332(2)
wrong MC	0.606456(3)	0.6688(1)	0.7362(1)	0.8188(1)	0.9234(2)

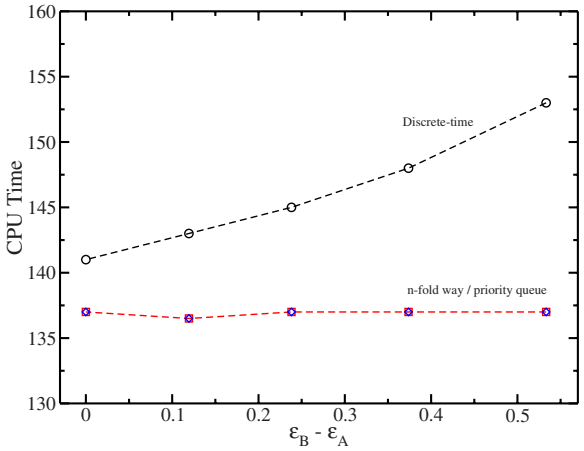

FIG. 3. (Color online) CPU Time (arbitrary units) required for the discrete-time (black circles), the *n*-fold way (red squares), and the priority queue algorithms (blue diamonds) to reach a fixed average number of infected sites in 10^5 simulation runs for a range of values of ϵ_A at criticality.

sen nearest neighbor is not in the susceptible state. The increase in computation time with larger heterogeneity for the discrete-time simulation, on the other hand, can be understood by the fact that the waiting step in Eq. (1) becomes more probable. This fact renders the continuous-time methods very useful for simulation of, e.g., truly disordered systems with heterogeneous rates which suffer from very long relaxation times [2].

As a second means of analysis, the Liouville operator for a binary AB chain of up to size $L=22$ sites with periodic boundary conditions was diagonalized. Due to symmetry considerations only even system sizes are useful for the analysis. The results of the FSS scheme outlined in Sec. III B are listed in Table I for the homogeneous critical point and the previously mentioned four other points that gradually move away from this point towards greater heterogeneity of rates. The Arnoldi method which was used to carry out the calculation of the lowest nontrivial eigenvalue of $\hat{\mathcal{L}}$ was found to produce eigenvalues numerically precise to 10^{-11} for the systems considered. This limits the precision to which the ϵ_L^* in Eq. (9) can be determined. Carefully considering the propagation of errors, we report the system-size dependent critical rate ϵ_L^* defined in Eq. (8), up to ten decimal places which we are confident are numerically precise. The value of ϵ_L^* in the limit $N\to\infty$ has been extrapolated via the BST extrapolation scheme where the parameter ω has, as usual, been chosen to minimize the difference between penultimate entries in the extrapolation tableaux [25,26].

For the homogeneous case, the extrapolation yields a value which agrees with both the established critical threshold rate as well as the MC result to within 3×10^{-4}. Concerning the heterogeneous systems, the method confirms that the commonly used discrete-time method predicts the wrong critical rate ("wrong MC" in Table I) as opposed to all other methods ("correct MC" in Table I). Despite the smaller number of available finite-lattice data points as compared to studies of homogeneous systems [11] the precision is found to be sufficient for our comparison.

As a by-product, the dynamical critical exponent θ was determined and found to be 1.56(2). This is in agreement with the known DP value [1.580759(6) [2]], as previously observed for other exponents in heterogeneous systems [14].

Lastly, we consider the low-lying spectrum of the Liouville operator as shown in Fig. 1 for a homogeneous system (a) and a binary chain (b) both of length $L=8$. The lowest nontrivial eigenvalue is found to behave rather differently in these two situations: While for the homogeneous system it grows from (close to) zero with increasing ϵ in the heterogeneous system it resembles a saturation curve with an inflexion point with increasing ϵ_A and fixed value of $\epsilon_B=\epsilon_c$. Physically, this point of inflexion marks a point at which the effect that an increase in ϵ_A has on the decay of the active state starts to diminish (we are approaching a saturation regime). At large values of ϵ_A the eigenvalue seems to asymptotically approach a constant value which is solely set by the value of ϵ_B. This is intuitively clear because when the recovery rate of one type of site is very large, only the recovery rate of the respective other type of site controls the decay of the active state.

V. CONCLUSION

We have simulated the one-dimensional contact process on periodic binary chains with heterogeneous recovery rates. Four different simulation methods were presented; the commonly used discrete-time algorithm, a modified version of it, and two continuous-time implementations. The CP in an environment with heterogeneous recovery rates was simulated by each of these confirming that the commonly used discrete-time method predicts a wrong critical threshold. Therefore, the commonly used discrete-time algorithm must not be used in heterogeneous systems if one is interested in the precise critical thresholds. A performance comparison between the latter three methods showed an advantage in using the continuous-time methods, particularly in the presence of strong heterogeneity. This may help in the analysis of disordered systems with heterogeneous rates in which simulation is a challenging task due to very long relaxation times. Results from considering the Liouville operator for finite versions of the simulated systems and applying finite-size scaling to the lowest gap of its eigenspectrum supported the predictions of the continuous-time simulation methods.

ACKNOWLEDGMENTS

We would like to thank Malte Henkel for helpful remarks on the BST extrapolation method. Also, we thank a referee for bringing the modified random-sequential algorithm to our attention. MC simulations were performed on the Cambridge University Condor Grid. S.V.F. acknowledges financial support from the UK EPSRC and the Cambridge European Trust.

[1] T. Harris, Ann. Probab. **2**, 969 (1974).
[2] J. Marro and R. Dickman, *Nonequilibrium Phase Transitions in Lattice Models* (Cambridge University Press, Cambridge, 1999).
[3] G. Ódor, Rev. Mod. Phys. **76**, 663 (2004).
[4] T. Vojta, J. Phys. A **39**, R143 (2006).
[5] T. M. Liggett, *Interacting Particle Systems* (Springer-Verlag, New York, 1985).
[6] H. Hinrichsen, Adv. Phys. **49**, 815 (2000).
[7] J. Hooyberghs, F. Iglói, and C. Vanderzande, Phys. Rev. Lett. **90**, 100601 (2003).
[8] R. Dickman and I. Jensen, Phys. Rev. Lett. **67**, 2391 (1991).
[9] P. Grassberger and A. De La Torre, Ann. Phys. (N.Y.) **122**, 373 (1979).
[10] P. Grassberger, J. Phys. A **22**, 3673 (1989).
[11] J. R. G. de Mendonça, J. Phys. A **32**, L467 (1999).
[12] W. Otten, J. A. N. Filipe, and C. A. Gilligan, Ecology **86**, 1948 (2005).
[13] G. A. Forster and C. A. Gilligan, Proc. Natl. Acad. Sci. U.S.A. **104**, 4984 (2007).
[14] C. J. Neugebauer, S. V. Fallert, and S. N. Taraskin, Phys. Rev. E **74**, 040101(R) (2006).
[15] T. Vojta and M. Dickison, Phys. Rev. E **72**, 036126 (2005).
[16] M. Dickison and T. Vojta, J. Phys. A **38**, 1199 (2005).
[17] D. Landau and K. Binder, *A Guide to Monte Carlo Simulations in Statistical Physics* (Cambridge University Press, New York, 2000).
[18] A. Bortz, M. Kalos, and J. Lebowitz, J. Comput. Phys. **17**, 10 (1975).
[19] D. Knuth, *The Art of Computer Programming* (Addison Wesley, Redwood City, CA, 1998), Vol. 3.
[20] W. Feller, *An Introduction to Probability Theory and its Applications* (Wiley, New York, 1971).
[21] P. Grassberger, Math. Biosci. **63**, 157 (1982).
[22] W. Kinzel and J. M. Yeomans, J. Phys. A **14**, L163 (1981).
[23] M. P. Nightingale, Physica A **83**, 561 (1975).
[24] R. R. dos Santos and L. Sneddon, Phys. Rev. B **23**, 3541 (1981).
[25] R. Bulirsch and J. Stoer, Numer. Math. **6**, 413 (1964).
[26] M. Henkel and G. Schutz, J. Phys. A **21**, 2617 (1988).
[27] We used the ARPACK package available from http://www.caam.rice.edu/software/ARPACK/