Published in final edited form as:

Comput Phys Commun. 2019 July ; 240: 30–37. doi:10.1016/j.cpc.2019.02.008.

Efficient sampling of spreading processes on complex networks using a composition and rejection algorithm

Guillaume St-Onge^{a,b}, Jean-Gabriel Young^{a,b}, Laurent Hébert-Dufresne^{a,c}, Louis J. Dubé^{a,b}

^aDépartement de Physique, de Génie Physique, et d'Optique, Université Laval, Québec (Québec), Canada, G1V 0A6

^bCentre interdisciplinaire de modélisation mathématique de l'Université Laval, Québec (Québec), Canada, G1V 0A6

^cDepartment of Computer Science and Vermont Complex Systems Center, University of Vermont, Burlington, VT 05401, USA

Abstract

Efficient stochastic simulation algorithms are of paramount importance to the study of spreading phenomena on complex networks. Using insights and analytical results from network science, we discuss how the structure of contacts affects the efficiency of current algorithms. We show that algorithms believed to require $\mathcal{O}(\log N)$ or even $\mathcal{O}(1)$ operations per update—where N is the number of nodes—display instead a polynomial scaling for networks that are either dense or sparse and heterogeneous. This significantly affects the required computation time for simulations on large networks. To circumvent the issue, we propose a node-based method combined with a composition and rejection algorithm, a sampling scheme that has an average-case complexity of $\mathcal{O}[\log(\log N)]$ per update for general networks. This systematic approach is first set-up for Markovian dynamics, but can also be adapted to a number of non-Markovian processes and can enhance considerably the study of a wide range of dynamics on networks.

Keywords

Spreading process; Complex network; Stochastic simulation algorithm

1. Introduction

Stochastic processes in a discrete state space are useful models for many natural and human-related complex systems. Considering that complex and heterogeneous connectivity patterns form the backbone of these systems, the study of dynamical processes on networks has grown in popularity in the last decades. Spreading processes are a prime example [1–3]. They are used to model a wide range of phenomena related to propagation among a

guillaume.st-onge.4@ulaval.ca (Guillaume St-Onge)

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

population. While disease epidemics [4] is the most commonly studied process, other important examples include social contagions [5, 6] and even beneficial epidemics [7].

Many analytical approaches have been developed to study the outcome of spreading processes on complex networks, using mean field [8–11], moment closure [12–15], percolation mapping [16–18] and message passing techniques [19, 20] to name a few (see Refs. [1–3] for recent overviews). However, most results hold only for random networks and tree-like structures, or stand as approximations for general networks. This undeniably contributes to our understanding of real systems, but any conclusions drawn from these approaches need to be supported by exact, robust, numerical simulations.

However, an important limitation of numerical simulations is certainly their computational cost, which generally increases with the size of the system. Spreading processes are a class of nonequilibrium models that exhibit a phase transition in the thermodynamic limit (infinite size systems). This particular feature has motivated the need for simulation algorithms that can handle network with very large number of nodes N . This is especially true for the study of anomalous phenomena, such as the localization of the epidemic [21–24], Griffith phases and smeared phase transitions [15, 25–29], typically observed in large-size networks. Therefore, the development of stochastic simulation algorithms that can efficiently tackle large networks of varied structures is unescapable, although it offers a considerable, albeit fascinating, computational challenge. It is also critical to the foundations of the field of spreading processes, one of the cornerstones of network science.

One standard formulation of spreading on networks is in terms of a continuous-time stochastic process. A widely used numerical approach involves a temporal discretization using a finite time step. Despite its simplicity, this is both computationally inefficient—requiring $\mathcal{O}(N)$ operations per time step—and prone to discretization errors [30]. Instead, the state-of-the-art methods are based on the Doob-Gillespie’s algorithm [31], which produces statistically exact sequences of states. Many studies have been dedicated to the improvement of its efficiency [32–35], and from these ideas have emerged some implementations for spreading processes [3, 36–39]. However, these approaches can still be inefficient for certain structures, for instance dense or sparse and heterogeneous networks.

In this work, we propose an efficient method for the simulation of Markovian spreading processes that builds on the Doob-Gillespie’s algorithm. It hinges on a composition and rejection scheme [33] that requires only $\mathcal{O}[\log(\log N)]$ operations per update for general networks. We demonstrate that it provides a formidable improvement compared to other algorithms whose computation time can scale polynomially with the network size. An implementation of the algorithm is available online [40].

2. Markovian spreading processes on networks

For the sake of simplicity, we consider simple graphs of N nodes and M edges, although the methods discussed can all be applied to directed, weighted, and multi-layer networks. We consider spreading processes defined as continuous Markov processes with discrete states $\mathbf{x}$. We can focus on canonical epidemiological models to understand the methods—other

compartmental models can also be accommodated, for instance by adding accessible states for the nodes, but numerical approaches are similar. Let us therefore denote the state of each node i as $x_i \in \{S, I, R\}$, susceptible, infected or recovered respectively. If infected node i and susceptible node j are connected, node i transmits the disease to node j at rate λ . Node i spontaneously recovers at rate μ , and becomes recovered if it develops immunity against the spreading disease, or susceptible otherwise. From this framework, we distinguish two models with different properties in the limit $t \rightarrow \infty$.

SIS model.

Infected nodes become directly susceptible after recovery (they do not develop immunity). For finite size networks, the absorbing state with all nodes susceptible, where the system remains trapped, is eventually visited. However, for infinite size networks¹, there exists some threshold λ_c for fixed μ such that the system reaches an endemic state, where a stationary density of the nodes remains infected on average, for all $\lambda > \lambda_c$ [see Fig. 1(a)]. This stationary density of infected nodes is hereafter referred as the *prevalence*.

SIR model.

Recovered nodes develop immunity. After a certain time, all nodes are either susceptible or recovered. Similarly to the SIS model, there exists a threshold value λ_c above which even an infinitesimal initial density of infected nodes will ultimately lead to a macroscopic number of recovered nodes that scales with network size [see Fig. 1(b)]. The final density of recovered nodes is hereafter referred as the *final size*.

Threshold estimates are provided in Appendix A for uncorrelated random networks with an arbitrary degree distribution.

3. Exact simulation algorithms

Spreading processes can be decomposed into a set $\Omega(\mathbf{x}) = \{u_a\}$ of independent processes for each state $\mathbf{x}$ of the system. A process u_a can be the transmission of a disease from node i to j or the recovery of an infected node. We can distinguish the set of transmission processes $\Omega_T(\mathbf{x})$ and the set of recovery processes $\Omega_R(\mathbf{x})$ such that $\Omega_T \cap \Omega_R = \emptyset$ and $\Omega_T \cup \Omega_R = \Omega$. To establish a coherent framework for the stochastic simulation algorithms, we consider that the execution of a process is an event that *may* change the state $\mathbf{x}$ of the system, but it is not required to do so. For instance, the transmission of a disease from node i to j is a process that implies $x_j \mapsto I$ if $x_j = S$. However, if $x_j \in \{I, R\}$ prior to the transmission, then the state of the system stays the same. Considering or not the transmission of a disease between two infected nodes as an ongoing process is only an algorithmic decision; from our definition, an infected node transmits the disease to all of its neighbors irrespectively of their states.

Since every possible transmission and recovery for a state $\mathbf{x}$ is an independent process u_a with rate $w_a \in \{\lambda, \mu\}$, the total rate of events is defined as

¹In practice, this is observed for large (finite size) networks.

$$W(\mathbf{x}) = \sum_{u_\alpha \in \Omega(\mathbf{x})} w_\alpha. \quad (1)$$

The inter-event time $\tau(\mathbf{x})$ before the execution of any process is exponentially distributed

$$P(\tau|\mathbf{x}) = \exp(-W(\mathbf{x})\tau)d\tau, \quad (2)$$

and the executed process is chosen proportionally to its *propensity*, i.e. its rate w_α

$$P(u_\alpha|\mathbf{x}) = \frac{w_\alpha}{W(\mathbf{x})}. \quad (3)$$

In this section, we introduce and discuss different numerical methods to sample sequences of states based on the previous equations. To compare the methods and give insights on their behavior for large system, we will provide expressions for the theoretical number of operations required on average to perform a state transition, i.e. $\mathbf{x} \mapsto \mathbf{x}'$ with $\mathbf{x}' \neq \mathbf{x}$. The expressions are *upper bounds* to the number of operations required on average and are expressed using the big $\mathcal{O}$ notation.

3.1. Doob-Gillespie's algorithm (direct method)

The direct method consists of the determination of the next process u_α and the time at which it will be executed. It can be summarized by these steps:

1. Determine the inter-event time $\tau(\mathbf{x})$ using Eq. (2).
2. Determine the executed process u_α using Eq. (3).
3. Update the state $\mathbf{x} \mapsto \mathbf{x}'$ according to the chosen process u_α (if $\mathbf{x}' \neq \mathbf{x}$), and update the time $t \mapsto t + \tau(\mathbf{x})$.
4. Update, create, or remove processes resulting from the execution of process u_α .
5. Update $W(\mathbf{x})$ and $P(u_\alpha|\mathbf{x})$, then return to step 1.

There have been various implementations of this direct procedure that have been discussed for spreading processes on static networks [3, 30, 37, 38].

For general networks, the number of ongoing processes is $\mathcal{O}(M + N)$. Efficient implementations typically use a binary-tree data structure to store the processes and keep updated $P(u_\alpha|\mathbf{x})$ [32, 38]. This allows the insertion and the retrieval of processes in $\mathcal{O}(\log N)$ operations. The costly processes involve the infection of a degree k_I node. To perform this update, the required number of operations is

$$\mathcal{O}(k_I \log N) \quad (4)$$

and is associated with the storage of future transmission processes.

A fact often overlooked is that the degree k_I of a newly infected node can be large on average. For dense networks, it scales as $k_I \sim N$; for sparse and heterogeneous networks near

the phase transition, $k_I \sim N^\xi$ on average with $0 < \xi < 1$ (see Appendix A). This is problematic, since phase transitions are central to many studies.

3.2. Node-based method

Another type of implementation has been suggested [37, 39]. To sample among the infection processes, one selects a node i among the infected nodes, proportionally to its degree k_i , then selects one of its neighbors $j \in \mathcal{N}_i$ randomly and infects it if $x_j = S$. The sampling of recovery processes is simply done by selecting an infected node uniformly. Using this scheme, one does not have to look through all the neighbors of an infected node to check for new processes u_α , associated with the $\mathcal{O}(k_i)$ operations in the direct method.

We introduce what we call the *node-based* method in the spirit of this scheme. The idea is to regroup the propensity of all processes associated with a node into a single propensity in order to prevent their enumeration. The probability of selecting a process u_α performed by node i is then factorized as $P(u_\alpha, i | \mathbf{x}) = P(u_\alpha | i, \mathbf{x}) P(i | \mathbf{x})$. The probability of selecting a node is written

$$P(i | \mathbf{x}) = \frac{\omega_i}{W(\mathbf{x})}, \quad (5)$$

where

$$\omega_i(\mathbf{x}) = \delta(\mathbf{x}_i, I)(\lambda k_i + \mu) \quad (6)$$

is the total propensity for node i to be involved in a transmission or recovery event, $\delta(a, b)$ is the Kronecker delta and k_i is the degree of node i . The probability of selecting a process u_α performed by node i is formally written

$$P(u_\alpha | i, \mathbf{x}) = \frac{w_\alpha}{\omega_i(\mathbf{x})}. \quad (7)$$

In practice, we use the probability that the chosen process is a transmission

$$P(u_\alpha \in \Omega_T | i, \mathbf{x}) = \frac{\lambda k_i}{\lambda k_i + \mu}, \quad (8)$$

to select the type of process. If it is a transmission, we select a random neighbor and infect it; if it is a recovery, the node becomes susceptible or recovered depending on the spreading model used. Note that the total rate of events [Eq. (1)] has an equivalent definition in terms of the propensities of the nodes, namely

$$W(\mathbf{x}) = \sum_{i=1}^n \omega_i(\mathbf{x}). \quad (9)$$

An important property of this scheme in the context of spreading processes is that Eq. (6) is only dependent on the state of the node i . Hence, when the state of a node i changes due to

the execution of a process, we do not need to update the propensity ω_j of neighboring nodes $j \in \mathcal{I}$, but only ω_i . We discuss two implementations for this method.

3.2.1. Rejection sampling—A simple approach to select a node i according to $P(\mathcal{I}|\mathbf{x})$ is to use rejection sampling. First, one needs to determine the maximal propensity possible $\omega_{\max}$ for a node, in this case being

$$\omega_{\max} = \lambda k_{\max} + \mu, \quad (10)$$

where $k_{\max}$ is the maximal degree of a node in the network. Second, one needs to keep updating an array of pairs of elements $\mathcal{A}$, the pairs being node labels and their propensity (i, ω_i)—we keep only pairs with non-zero propensities ω_i . Finally, one draws a pair (i, ω_i) at random from $\mathcal{A}$ and accepts it with probability

$$P[(i, \omega_i) | \omega_{\max}] = \frac{\omega_i}{\omega_{\max}}. \quad (11)$$

The complete procedure for the execution of a process using the node-based method with rejection sampling is presented in Algorithm 1. Although the SIS model was considered, one only needs to change step 19 by $S \mapsto R$ for it to be compliant with the SIR model.

It is straightforward to verify that most steps are $\mathcal{O}(1)$ in Algorithm 1, except possibly the rejection sampling portion (steps 3 to 9). The average acceptance probability for a state $\mathbf{x}$ is $\langle \omega_i | \mathcal{I}; \mathbf{x} \rangle / \omega_{\max}$, where

$$\langle \omega_i | \mathcal{I}; \mathbf{x} \rangle = \lambda \langle k | \mathcal{I}; \mathbf{x} \rangle + \mu, \quad (12)$$

is the average propensity for infected nodes in state $\mathbf{x}$ and $\langle k | \mathcal{I}; \mathbf{x} \rangle$ stands for the average degree of infected nodes in state $\mathbf{x}$. Since the number of required operations follows a geometric distribution, the average number of operations for an update from a state $\mathbf{x}$ is

$$\mathcal{O}(\omega_{\max} / \langle \omega_i | \mathcal{I}; \mathbf{x} \rangle). \quad (13)$$

For networks with a homogeneous degree distribution, we have $\langle \omega_i | \mathcal{I}; \mathbf{x} \rangle \sim \omega_{\max}$ for all states, leading to a small number of operations. However, this rejection sampling scheme is vulnerable to cases where the propensity ranges over multiple scales—more specifically, when $\langle \omega_i | \mathcal{I}; \mathbf{x} \rangle \ll \omega_{\max}$ for typical states $\mathbf{x}$. This happens when $k_{\max}$ is large, which is common for heterogeneous degree distribution, such as $P(k) \sim k^{-\gamma}$ (see Appendix A). To circumvent this, we must update and sample efficiently from a heterogeneous distribution of propensities $P(\mathcal{I}|\mathbf{x})$ [Eq. (5)].

3.2.2. Composition and rejection for multiscale propensity—The high rejection probability of the rejection sampling for heterogeneous propensities is illustrated by the gray portion in Fig. 2(a). The problem is the uniform proposal distribution on $\mathcal{A}$ (step 3 of Algorithm 1), which is a poor choice for a heterogeneous distribution $P(\mathcal{I}|\mathbf{x})$. To improve upon

rejection sampling, we need an algorithm that *systematically* constructs a proposal distribution that upper bounds the rejection probability, as illustrated in Fig. 2(b).

Algorithm 1 Node-based method with rejection sampling for an update of the SIS model

Input: $x, \mathcal{A}, W$

Output: τ and updated input

```

1: choose  $r_1 \in [0, 1)$  uniformly at random
2:  $\tau \leftarrow -\frac{\ln r_1}{W}$ 
3: select  $(i, \omega_i) \in \mathcal{A}$  uniformly at random
4: choose  $r_2 \in [0, 1)$  uniformly at random
5: if  $r_2 < \frac{\omega_i}{\omega_{\max}}$  then
6:   choose node  $i$ 
7: else
8:   go back to step 3
9: end if
10: choose  $r_3 \in [0, 1)$  uniformly at random
11: if  $r_3 < \frac{\lambda k_i}{\lambda k_i + \mu}$  then ▷ transmission process
12:   choose  $j \in \partial i$  uniformly at random
13:   if  $x_j = S$  then
14:      $x_j \leftarrow I$ 
15:      $\mathcal{A} \leftarrow \mathcal{A} \cup \{(j, \omega_j)\}$ 
16:      $W \leftarrow W + \omega_j$ 
17:   end if
18: else ▷ recovery process
19:    $x_i \leftarrow S$ 
20:    $\mathcal{A} \leftarrow \mathcal{A} \setminus \{(i, \omega_i)\}$ 
21:    $W \leftarrow W - \omega_i$ 
22: end if

```

We propose to use a method of *composition and rejection*, similarly to Ref. [33] for biochemical reaction networks, inspired by Ref. [42]. It is also a direct improvement over the 2-group method proposed in Ref. [37]. The idea is to create a partition of nodes $\mathcal{P} = \{g_\beta\}$ with similar propensities $\omega > 0$, with $\beta \in \{1, 2, \dots, q\}$. Once a node i gets infected (or more generally gets a propensity $\omega_i > 0$), it is assigned to a group g_β —in our implementation, the pair (i, ω_i) is stored in an array $\mathcal{A}_\beta$. The probability to select a node i is then factorized as $P(i|\mathbf{x}) = P(i|g_\beta, \mathbf{x})P(g_\beta|\mathbf{x})$.

The probability of selecting a group g_β is

$$P(g_\beta|\mathbf{x}) = \frac{W_\beta(\mathbf{x})}{W(\mathbf{x})}, \quad (14)$$

where $W_\beta(\mathbf{x}) \equiv \sum_{i \in g_\beta} \omega_i(\mathbf{x})$ is the propensity associated to the group of nodes. In practice, to select an array $\mathcal{A}_\beta$ proportionally to W_β , we implement a binary decision tree $\mathcal{T}$ as illustrated in Fig. 3, where each leaf points to an array $\mathcal{A}_\beta$. Starting from the root, it takes $\mathcal{O}(\log q)$ operations to choose one of the leaves. Once a process u_α is chosen and executed, we update the array $\mathcal{A}_\beta$, the propensity W_β and recursively the parent values in the tree—one notes that the root value is in fact $W(\mathbf{x})$. Again, $\mathcal{O}(\log q)$ operations are needed for this task.

The probability of selecting a node within a group is

$$P(i|g_\beta, \mathbf{x}) = \frac{\omega_i(\mathbf{x})}{W_\beta(\mathbf{x})}. \quad (15)$$

If the partition $\mathcal{P}$ is wisely chosen, nodes are selected efficiently using rejection sampling, replacing $\omega_{\max}$ in Eq. (11) by a group specific maximal propensity $\omega_{\beta\max}$.

A systematic approach to construct the partition $\mathcal{P}$ is to impose a minimum acceptance probability, say one half. This leads to an average number of operations upper-bounded by 2. Let us define the minimal propensity as

$$\omega_{\min} = \lambda k_{\min} + \mu, \quad (16)$$

where $k_{\min}$ is the minimal degree in the network. We impose that the β -th group allows nodes with propensity $\omega \in [2^{\beta-1}\omega_{\min}, 2^\beta\omega_{\min})$, except for the last group allowing $\omega \in [2^{q-1}\omega_{\min}, \omega_{\max})$. The group specific maximal propensity is thus

$$\omega_{\beta\max} = \min(2^\beta\omega_{\min}, \omega_{\max}), \quad (17)$$

and the number of group required is

$$q = \max\left(1, \left\lceil \log_2\left(\frac{\omega_{\max}}{\omega_{\min}}\right) \right\rceil\right). \quad (18)$$

The complete procedure for the execution of a process using the node-based method with composition and rejection sampling is presented in Algorithm 2.

For general networks with $\lambda/\mu < \infty$, the ratio of extreme propensities $\omega_{\max}/\omega_{\min} = \mathcal{O}(N)$.

Therefore, the average-case complexity per update is

$$\mathcal{O}[\log(\log N)]. \quad (19)$$

For networks with a maximal degree $k_{\max}$ independent of N , the complexity is $\mathcal{O}(1)$.

It is worth mentioning that our choice to impose a minimum acceptance probability of one half is not unique : one could impose an acceptance probability of a^{-1} with $a > 1$ and obtain the same average-case complexity. For $a \rightarrow 1$, it increases the acceptance probability, but it also increases the number of groups required q . Therefore, there is a trade-off to consider when one tries to minimize the required number of operations. We made several trials with $a = 2$, but it has never resulted in noticeable improvements for the computation time.

Algorithm 2 Node-based method with composition and rejection sampling for an update of the SIS model

Input: $x, \mathcal{T}$

Output: τ and updated input

```

1: get  $W$  from the root of  $\mathcal{T}$ 
2: choose  $r_1 \in [0, 1)$  uniformly at random
3:  $\tau \leftarrow -\frac{\ln r_1}{W}$ 
4: choose  $r_2 \in [0, 1)$  uniformly at random
5: select  $\mathcal{A}_\beta$  from  $\mathcal{T}$  using  $r_2$ 
6: select  $(i, \omega_i) \in \mathcal{A}_\beta$  uniformly at random
7: choose  $r_3 \in [0, 1)$  uniformly at random
8: if  $r_3 < \frac{\omega_i}{\omega_{\beta\max}}$  then
9:   choose node  $i$ 
10: else
11:   go back to step 6
12: end if
13: choose  $r_4 \in [0, 1)$  uniformly at random
14: if  $r_4 < \frac{\lambda k_i}{\lambda k_i + \mu}$  then ▷ transmission process
15:   choose  $j \in \partial i$  uniformly at random
16:   if  $x_j = S$  then
17:      $x_j \leftarrow I$ 
18:      $\mathcal{A}_{\beta'} \leftarrow \mathcal{A}_\beta \cup \{(j, \omega_j)\}$  ▷  $j \in g_{\beta'}$ 
19:     update  $\mathcal{T}$ 
20:   end if
21: else ▷ recovery process
22:    $x_i \leftarrow S$ 
23:    $\mathcal{A}_\beta \leftarrow \mathcal{A}_\beta \setminus \{(i, \omega_i)\}$ 
24:   update  $\mathcal{T}$ 
25: end if

```

3.3. Event-driven method

Another type of approach for the simulation of spreading processes has been considered lately [3, 39]. The philosophy is similar to the next reaction method [32], originally proposed for the simulation of chemical reaction networks to improve upon the original Doob-Gillespie's algorithm. The principal concept of this scheme is to draw an interevent time τ_a for each of the specific process u_a , and execute the latter at time $t_a + \tau_a$, where t_a is the absolute time when the process u_a was created. Therefore, one focus on the execution time of each process independently, instead of inferring the global inter-event time $\tau(x)$ and the first process to be executed among $\Omega(x)$, as in standard Gillespie method.

For Markovian dynamics, the inter-event time before the execution of a process u_a is exponentially distributed

$$P(\tau_a) = \exp(-w_a \tau_a) d\tau_a. \quad (20)$$

However, it is important to stress that this approach can also be applied to non-Markovian spreading processes [3, 39].

To store and retrieve the processes efficiently, one can use a priority queue, where the highest priority element corresponds to the process u_a with the lowest absolute execution time $t_a + \tau_a$. Recovery processes are eventually executed and can be stored directly in the priority queue. Transmission processes are stored if the inter-event time for the transmission

is smaller than the inter-event time for the recovery of the infected node. Depending on the implementation, one can also verify *a priori* if a neighbor node will already be infected or recovered, and prevent the storage of the transmission process.

To compare this approach with our node-based schemes, we used the implementation provided by Ref. [3, 43], called the *event-driven method*, for which a detailed pseudocode is available for both the SIS and the SIR models. The interface is modular and general enough to be used for any networks, a requirement for our benchmark calculations. Let us mention that there are alternative implementations available, for instance the efficient implementation of Ref. [44] for the SIR model (see Supplementary Material [45] for a comparison with our algorithm).

As for the standard Gillespie algorithm, a costly process for the event-driven method involves the infection of a degree k_I node, for which

$$\mathcal{O}(k_I \log N), \quad (21)$$

operations are required. While this quantity is probably not the most representative for the average computation time associated with the algorithm [see Fig. 4], we can clearly identify the different impact of the term k_I for dense or sparse and heterogeneous networks.

4. Efficiency of the stochastic simulation algorithms

We have described different schemes for the simulation of spreading processes, with different associated complexity for the update of the state $\mathbf{x} \mapsto \mathbf{x}'$. In this section, we show and discuss the impact on the computation time using synthetic networks. In Fig. 4 we compile the results of our benchmark calculations, using the SIS model in the stationary state and complete sequences of the SIR model, starting with a small infected node density. We fix the recovery rate $\mu = 1$ without loss of generality. Furthermore, we scale the transmission rate parameter $\lambda = \eta \lambda_c$ according to the underlying threshold λ_c of the dynamics. The intent is to always have a similar prevalence (SIS) or final size (SIR) in our simulations as we tune the network structure: as we discuss in Appendix B, different values for the order parameter can affect the expected number of required operations.

Implementations for the two node-based methods—hereafter referred as to *rejection* and *composition and rejection*—are in C++, while the implementation of the event-driven method is in Python [3, 43]. This discrepancy of programming languages causes a multiplicative factor overhead of roughly 100 for the computation times of the event-driven method. To provide a fairer comparison, and since we are mostly interested in the scaling of the algorithms with the size of the network, we have rescaled the computation times obtained with the event-driven method to match the first marker of the composition and rejection method in each panel of Fig. 4. For instance, if $T_{ED}(N)$ [$T_{CR}(N)$] represents the computation time of the event-driven (composition and rejection) method for different sizes N , we perform the transformation $T_{ED}(N) \mapsto T_{CR}(N_{\min}) T_{ED}(N) / T_{ED}(N_{\min})$, with $N_{\min} = 10^3$ in Fig. 4. A comparison without rescaling is also available in the Supplementary Material [45].

4.1. Computation time for homogeneous networks

As a model of homogeneous networks, we used the $G(N, M)$ random graphs ensemble [41]. In the limit $N \rightarrow \infty$ to, the degree distribution is a binomial with $\langle k \rangle = 2M/N$ and the degree of nodes are well represented by the mean value $\langle k \rangle$.

We observe in Fig. 4(a) that the average computation time for the event-driven method roughly scales linearly with the average degree, in agreement with Eq. (21). For the SIR model in Fig. 4(d), the dependence is less important, except for very large average degree.

As a comparison, both node-based methods are completely independent of the average degree [Fig. 4(a) and 4(d)]. This is in line with the two steps selection procedure for infection processes. A node-based scheme should therefore be privileged whenever one wants to sample networks with large average degree, such as dense networks where $\langle k \rangle = \mathcal{O}(N)$.

One also observes that the rejection method is slightly more efficient [Fig. 4(a)] than the composition and rejection method : this is due to the simpler implementation of the rejection sampling, without the need for a composition step. As discussed in Sec. 3.2.2, the average number of operations $\mathcal{O}(\omega_{\max}/\langle \omega_i I; \mathbf{x} \rangle)$ is only problematic when propensities span multiple scales; for homogeneous networks, the ratio $\omega_{\max}/\langle \omega_i I; \mathbf{x} \rangle$ is expected to be $\mathcal{O}(1)$ for all states.

4.2. Computation time for heterogeneous networks

As a model of heterogeneous networks, we used random graphs with an expected degree sequence $\{\kappa_i\}$ [46, 47], also called Chung-Lu graphs, where κ_i is the expected degree of node i . We used sequences drawn from a power-law distribution $P(\kappa) \sim \kappa^{-\gamma}$ to generate heterogeneous networks, with $\gamma = 2.25$ in Fig. 4(b) and Fig. 4(e), and $\gamma = 3$ in Fig. 4(c) and Fig. 4(f).

We observe in Fig. 4(b) that the computation time for the event-driven method scales polynomially with the number of nodes. In Fig. 4(c), the computation time slightly increases, but with a much smaller exponent. This is explained by Eq. (21), with

$$k_I \sim N^{(3-\gamma)/2} \quad (22)$$

near the phase transition (see Appendix A). For the SIR model in Fig. 4(e) and 4(f), the computation times are less influenced by the size of the network. In this case, the number of required operations predicted in Eq. (21) by the most costly processes overestimates the average computation time (as also noted for homogeneous networks).

We observe that the rejection method scales polynomially with the number of nodes as well, but this time the scaling exponent is larger for moderately heterogeneous networks [Fig. 4(c) and 4(f)] than for very heterogeneous networks [Fig. 4(b) and 4(e)]². This is roughly explained by Eq. (13), with $\{\omega_i I; \mathbf{x}\} \sim \langle k^2 \rangle$ (see Appendix A) and $\omega_{\max} \sim k_{\max}$, leading to

²A larger dispersion for the degree distribution implies a more heterogeneous network.

$$\frac{\omega_{\max}}{\langle \omega_i | I; \mathbf{x} \rangle} \sim N^{(\gamma-2)/2}. \quad (23)$$

Finally, we see that the computation time for the composition and rejection method is, for all practical purposes, independent of the number of nodes.

5. Concluding remarks

We have introduced a stochastic simulation algorithm for the simulation of spreading processes on networks, combining a node-based perspective with the efficiency of composition and rejection sampling [40]. This algorithm requires $\mathcal{O}[\log(\log N)]$ operations per update, providing a superior scaling (or at worst, equivalent) to the other state-of-the-art methods. It is particularly well suited for the sampling of large and heterogeneous networks, since its average computation time is, for all practical purposes, independent of the network size or the density of edges.

For the SIR model specifically, our benchmark calculations reveal that the event-driven method scales efficiently as well, except for very dense networks. If computation time is a crucial issue, we suggest to compare the speed of our implementation [40] with highly optimized implementations of the event-driven method, such as the one provided by Ref. [44], or the recent hybrid approach proposed by Ref. [48], and choose the one that performs better for the particular case study.

Note that there is a complete branch of literature whose concern is the efficient sampling of the quasistationary distribution of states for processes with an absorbing state, such as the SIS model (see Refs. [37, 49–51] for instance). Indeed, finite size analysis of the critical phenomenon requires the sampling of sequences that do not fall on the absorbing state. In this paper, we have focused on the stochastic simulation algorithms that provide statistically exact sequences of states, which is also fundamental to sample the quasistationary distribution. Therefore, our work contributes indirectly to this line of study.

Despite the fact that we have considered explicitly Markovian spreading processes, our composition and rejection scheme can be directly applied to certain classes of non-Markovian processes with completely monotone survival function, using the Laplace transform, in the spirit of Ref. [38]. It can also directly be used with a variety of spreading processes on time-varying networks, where the structure evolves independently from the dynamics [15, 52], or co-evolves with it [53, 54]. Extension of this method for complex contagion [6] would further be an interesting avenue.

Finally, from a more general perspective, we can argue that the idea behind composition and rejection is a systematic and efficient regrouping of independent processes, especially suited for multiscale propensities, as discussed in Sec. 3.2. It would be simple to exploit this scheme for many of the stochastic processes studied in network science, such as multistate dynamical processes [55, 56] or network generation models [57, 58].

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank Petter Holme for useful comments. We acknowledge Calcul Quebec for computing facilities. This research was undertaken thanks to the financial support from the Natural Sciences and Engineering Research Council of Canada (NSERC), the Fonds de recherche du Québec—Nature et technologies (FRQNT) and the Sentinel North program, financed by the Canada First Research Excellence Fund.

Appendix A.: Some results for uncorrelated random networks

Over the last two decades, many analytical results have been obtained for spreading processes on uncorrelated random networks with an arbitrary degree distribution $P(k)$, i.e. the probability of having an edge with endpoint nodes of degree k, k' is $P(k, k') \propto k k' P(k) P(k')$. To generate such uncorrelated networks, one must respect the structural cut-off [59], i.e. $k_{\max} < \langle k \rangle N^{1/2}$.

In the infinite size limit, threshold estimates λ_c are obtained for spreading processes. For the SIS model, a good approximation is [15]

$$\lambda_c^{\text{SIS}} = \min \left(\sqrt{\frac{2}{k_{\max}}}, \frac{\langle k \rangle}{\langle k^2 \rangle - \langle k \rangle} \right), \quad (\text{A.1})$$

while for the SIR model, the exact result is [47, 60]

$$\lambda_c^{\text{SIR}} = \frac{\langle k \rangle}{\langle k^2 \rangle - 2\langle k \rangle}. \quad (\text{A.2})$$

Equations (A.1) and (A.2) were used in the simulations of Fig. 4.

Using the heterogeneous mean-field theory [8, 9], we can also approximate the degree distribution of infected nodes near the phase transition. On the one hand, for the SIS model in the stationary state, we have

$$P(k|I)^* \sim \frac{\lambda \theta^* k}{\mu + \lambda \theta^* k} P(k), \quad (\text{A.3})$$

with degree distribution $P(k)$, and where θ^* is the average fraction of infected neighbors for a susceptible node. Near the phase transition, $\theta^* \rightarrow 0$, which means that the average degree of infected nodes is

$$\langle k|I \rangle^* \sim \langle k^2 \rangle. \quad (\text{A.4})$$

On the other hand, if we consider a complete sequence of the SIR model, the average degree of recovered nodes is a good proxy, which is also $\sim \langle k^2 \rangle$ near the phase transition. For power-law degree distribution $P(k) \sim k^{-\gamma}$, maximal degree $k_{\max} < \langle k \rangle N^{1/2}$ and $2 - \gamma < 3$, the second moment of the degree distribution scales with the number of nodes as

$$\left\langle k^2 \right\rangle \sim \int_{k_{\min}}^{k_{\max}} k^{2-\gamma} dk \sim N^{(3-\gamma)/2}. \quad (\text{A.5})$$

Appendix B.: Overhead factor due to phantom events

An element that we have discarded in our complexity analysis is the average fraction of *phantom processes* [37] ϕ , i.e. processes u_a that do not change the state $\mathbf{x}$. One example is the transmission to an already infected node. It is assumed in our analysis of Sec. 3 that ϕ is upper-bounded by a constant $C < 1$, independent of the number of nodes in the system—this is well supported by the results of the composition and rejection scheme in Fig. 4. It is worth pointing that the results of Fig. 4 are unbiased, since we counted the number of real transitions $\mathbf{x} \mapsto \mathbf{x}'$ with $\mathbf{x}' = \mathbf{x}$ to evaluate the empirical average computation time.

For different values of prevalence (SIS) or final size (SIR), phantom processes can lead to a certain overhead. We show in Fig. B.5 the expected multiplicative factor on the number of operations required for an update of the state, due to a certain fraction ϕ of phantom processes.

Near the phase transition, the overhead factor is negligible, but it can get important for prevalence or final size near 1 where infected nodes are mostly surrounded by infected or recovered nodes. It would be possible to reduce the fraction of phantom processes ϕ in the node-based methods for the SIR model : we could count the number $\tilde{k}_i \leq k_i$ of susceptible neighbors for a newly infected node, and modify the propensity to

$$\tilde{\omega}_i = \delta(x_i, I) (\lambda \tilde{k}_i + \mu). \quad (\text{B.1})$$

However, the update involving the infection of a degree k_I node would now require

$$\mathcal{O}[\log(\log N) + k_I], \quad (\text{B.2})$$

which could be worst in certain cases according to our study. Since $\mathcal{O}[\log(\log N)]$ —or $\mathcal{O}(1)$ for upper-bounded maximal degree—is a small number of operations, it is probably safer (and simpler) to keep the schemes introduced in Sec. 3.2.

Figure B.5:

Expected overhead factor for the update time of a state due to phantom events, using node-based methods. This factor is estimated by $1/(1 - \phi)$ and ϕ is measured for (a) the SIS model in the stationary state and (b) complete sequence of the SIR model. We sampled over 10 realizations of random graphs ensembles : the $G(N, M)$ ensemble with $N = 10^5$ and $M = 5 \times 10^5$; random graphs with an expected degree sequence, a power law expected degree distribution $P(k) \sim k^{-\gamma}$ with $\gamma = 2.5$, a natural cut-off $k_{\max} < N$ and $N = 105$.

References

- [1]. Pastor-Satorras R, Castellano C, Van Mieghem P, and Vespignani A, Rev. Mod. Phys 87, 925 (2015).
- [2]. Wang W, Tang M, Stanley HE, and Braunstein LA, Rep. Prog. Phys 80, 036603 (2017). [PubMed: 28176679]
- [3]. Kiss IZ, Miller JC, and Simon PL, Mathematics of Epidemics on Networks: From Exact to Approximate Models, Vol. 46 (Springer, 2017).
- [4]. Anderson RM and May RM, Infectious Diseases of Humans: Dynamics and Control (Oxford Science Publications, 1991).
- [5]. Morgan AC, Economou D, Way SF, and Clauset A, arXiv:1805.09966 (2017).
- [6]. Lehmann S and Ahn Y-Y, Complex Spreading Phenomena in Social Systems (Springer, 2018).
- [7]. Berdahl A, Brelsford C, De Bacco C, Dumas M, Ferdinand V, Grochow JA, Hebert-Dufresne L, Kallus Y, Kempes CP, Kolchinsky A, et al., arXiv:1604.02096 (2016).
- [8]. Boguñá M and Pastor-Satorras R, Phys. Rev. E 66, 047104 (2002).
- [9]. Moreno Y, Pastor-Satorras R, and Vespignani A, The European Physical Journal B - Condensed Matter and Complex Systems 26, 521 (2002).
- [10]. Van Mieghem P, Omic J, and Kooij R, IEEE/ACM Trans. Netw 17, 1 (2009).
- [11]. Barrat A, Barthélemy M, and Vespignani A, Dynamical Processes on Complex Networks (Cambridge University Press, 2008).
- [12]. Eames KTD and Keeling MJ, Proc. Natl. Acad. Sci. USA 99, 13330 (2002). [PubMed: 12271127]
- [13]. Mata AS, Ferreira RS, and Ferreira SC, New J. Phys 16, 053006 (2014).
- [14]. Sharkey KJ, Kiss IZ, Wilkinson RR, and Simon PL, Bulletin of Mathematical Biology 77, 614 (2015). [PubMed: 24347252]

- [15]. St-Onge G, Young J-G, Laurence E, Murphy C, and Dubé LJ, Phys. Rev. E 97, 022305 (2018). [PubMed: 29548152]
- [16]. Newman MEJ, Phys. Rev. E 66, 016128 (2002).
- [17]. Kenah E and Robins JM, Phys. Rev. E 76, 036113 (2007).
- [18]. Parshani R, Carmi S, and Havlin S, Phys. Rev. Lett 104, 258701 (2010). [PubMed: 20867419]
- [19]. Karrer B and Newman MEJ, Phys. Rev. E 82, 016101 (2010).
- [20]. Shrestha M, Scarpino SV, and Moore C, Phys. Rev. E 92, 022821 (2015).
- [21]. Goltsev AV, Dorogovtsev SN, Oliveira JG, and Mendes JFF, Phys. Rev. Lett 109, 128702 (2012). [PubMed: 23006000]
- [22]. Castellano C and Pastor-Satorras R, Sci. Rep 2, 371 (2012). [PubMed: 22523634]
- [23]. Pastor-Satorras R and Castellano C, J. Stat. Phys 173, 1110 (2018).
- [24]. Liu Q and Van Mieghem P, IEEE Trans. Netw. Sci. Eng, 1 (2019).
- [25]. Munoz MA, Juhász R, Castellano C, and Ódor G, Phys. Rev. Lett 105, 128701 (2010). [PubMed: 20867681]
- [26]. Ódor G, Phys. Rev. E 90, 032110 (2014).
- [27]. Mata AS and Ferreira SC, Phys. Rev. E 91, 012816 (2015).
- [28]. Cota W, Ferreira SC, and Ódor G, Phys. Rev. E 93, 032322 (2016). [PubMed: 27078381]
- [29]. Cota W, Ódor G, and Ferreira SC, arXiv:1801.06406 (2018).
- [30]. Fennell PG, Melnik S, and Gleeson JP, Phys. Rev. E 94, 052125 (2016). [PubMed: 27967171]
- [31]. Gillespie DT, J. Comput. Phys 22, 403 (1976).
- [32]. Gibson MA and Bruck J, J. Phys. Chem. A 104, 1876 (2000).
- [33]. Slepoy A, Thompson AP, and Plimpton SJ, J. Chem. Phys 128, 05B618 (2008).
- [34]. Goutsias J and Jenkinson G, Phys. Rep 529, 199 (2013).
- [35]. Yates CA and Klingbeil G, J. Chem. Phys 138, 094103 (2013). [PubMed: 23485273]
- [36]. Vestergaard CL and Genois M, PLOS Comput. Biol 11, 1 (2015).
- [37]. Cota W and Ferreira SC, Comput. Phys. Commun 219, 303 (2017).
- [38]. Masuda N and Rocha LE, SIAM Review 60, 95 (2018).
- [39]. de Arruda GF, Rodrigues FA, and Moreno Y, arXiv:1804.08777 (2018).
- [40]. St-Onge G, “spreading_cr,” https://github.com/gstonge/spreading_CR (2018).
- [41]. Erdős P and Rényi A, Publicationes Mathematicae (Debrecen) 6, 290 (1959).
- [42]. Devroye L, Non-Uniform Random Variate Generation (Springer, 1986).
- [43]. Miller JC, “Mathematics-of-epidemics-on-networks,” <https://github.com/springer-math/Mathematics-of-Epidemics-on-Networks> (2018).
- [44]. Holme P, “sir,” <https://github.com/pholme/sir> (2018).
- [45]. See Supplementary Material at [URL will be inserted by publisher] for additional comparisons.
- [46]. Chung F and Lu L, Annals of Combinatorics 6, 125 (2002).
- [47]. Miller JC and Hagberg A, in Algorithms and Models for the Web Graph, edited by Frieze A, Horn P, and Pralat P (Springer Berlin Heidelberg, Berlin, Heidelberg, 2011) pp. 115–126.
- [48]. Großmann G and Wolf V, arXiv:1812.10845 (2018).
- [49]. de Oliveira MM and Dickman R, Phys. Rev. E 71, 016129 (2005).
- [50]. Sander RS, Costa GS, and Ferreira SC, Phys. Rev. E 94, 042308 (2016). [PubMed: 27841497]
- [51]. Macedo-Filho A, Alves GA, Costa Filho RN, and Alves TFA, J. Stat. Mech. Theory Exp 2018, 043208 (2018).
- [52]. Taylor M, Taylor TJ, and Kiss IZ, Phys. Rev. E 85, 016103 (2012).
- [53]. Gross T, D’Lima CJD, and Blasius B, Phys. Rev. Lett 96, 208701 (2006). [PubMed: 16803215]
- [54]. Scarpino SV, Allard A, and Hebert-Dufresne L, Nat. Phys 12, 1042 (2016).
- [55]. Gleeson JP, Phys. Rev. X 3, 021004 (2013).
- [56]. Fennell PG and Gleeson JP, arXiv:1709.09969 (2017).
- [57]. Krapivsky PL, Redner S, and Leyvraz F, Phys. Rev. Lett 85, 4629 (2000). [PubMed: 11082613]

- [58]. Hébert-Dufresne L, Allard A, Marceau V, Noel P-A, and Dubé LJ, Phys. Rev. Lett 107, 158702 (2011). [PubMed: 22107326]
- [59]. Catanzaro M, Boguñá M, and Pastor-Satorras R, Phys. Rev. E 71, 027103 (2005).
- [60]. Volz E, J. Math. Biol 56, 293 (2008). [PubMed: 17668212]
- [61]. Tange O, ;login: The USENIX Magazine 36, 42 (2011).

Figure 1:

Typical phase transition of the order parameters for the two canonical models of spreading. The recovery rate is fixed to $\mu=1$ while the transmission rate λ is varied. Simulations were performed using the composition and rejection algorithm (see Sec. 3.2.2) on a single realisation of the $G(N, M)$ random graph ensemble [41], with $N = 105$ and $M = 5 \times 105$. **(a)** Average prevalence in the endemic state for the SIS model. **(b)** Average final size for an outbreak of the SIR model, starting with an initial infected node density of 10^{-3} . In both cases, the error bars associated to the standard deviation are hidden by the markers.

Figure 2: Rejection area (gray) compared to the acceptance area (blue), for the selection of a node with the node-based method. **(a)** Rejection sampling. **(b)** Composition and rejection sampling.

Figure 3:
Decision tree $\mathcal{T}$ used to select the arrays $\{\mathcal{A}_\beta\}$.

Figure 4:

Average computation time for a single update of the state for spreading processes using different algorithms. Each marker is averaged over 10 simulations on 10 realizations of a random graph ensemble. The computation time of the event-driven method was rescaled to match the first marker of the composition and rejection method. **(Upper row)** Average over sequences of 106 state transitions for the SIS model in the stationary state. The systems have been thermalized using 106 transitions beforehand. **(Lower row)** Average over 106/N complete sequences of the SIR model, starting with an initial infected node density of 10–2. **(a) and (d)** $G(N, M)$ random graphs with fixed number of nodes $N = 10^4$ and different average degree $\langle k \rangle$. The recovery and transmission rates used are $\mu = 1$ and $\lambda = 1.05\lambda_c$. **(b) and (e)** Random graphs with an expected degree sequence [46, 47]. We used a power law expected degree distribution $P(k) \sim k^{-\gamma}$ with $\gamma = 2.25$, $k_{\min} = 3$ and $k_{\max} < \langle k \rangle N^{1/2}$. The recovery and transmission rates used are $\mu = 1$ and $\lambda = 3\lambda_c$. **(c) and (f)** Same as (b) and (e), but with $\gamma = 3$.