


Finding patient zero in susceptible-infectious-susceptible epidemic processes

Robin Persoons * and Piet Van Mieghem 

Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology,
P.O. Box 5031, 2600 GA Delft, The Netherlands



(Received 12 July 2024; accepted 18 September 2024; published 18 October 2024)

Finding the source of an epidemic is important, because correct source identification can help to stop a budding epidemic or prevent new ones. We investigate the backward equations of the N -intertwined mean-field approximation susceptible-infectious-susceptible (SIS) process. The backward equations allow us to trace the epidemic back to its source on networks of sizes up to at least $N = 1500$. Additionally, we show that the source of the “more realistic” Markovian SIS model cannot feasibly be found, even in a “best-case scenario,” where the infinitesimal generator Q , which completely describes the epidemic process and the underlying contact network, is known. The Markovian initial condition $s(0)$, which reveals the epidemic source, can be found analytically when the viral state vector $s(t)$ is known at some time t as $s(0) = s(t)e^{-Qt}$. However, $s(0)$ can hardly be computed, except for small times t . The numerical errors are largely due to the matrix exponential e^{-Qt} , which is severely ill-behaved.

DOI: [10.1103/PhysRevE.110.044308](https://doi.org/10.1103/PhysRevE.110.044308)

I. INTRODUCTION

Finding the source, or patient zero, or the epicenter of an epidemic is a problem of great interest. The source of the COVID-19 epidemic, for example, has been investigated extensively [1–3]. Finding the source of an expanding epidemic could yield critical insights to stop the epidemic before a large portion of the population is affected. Related problems of interest are, for example, finding the source of a rumour or fake news or a computer virus.

Previous papers focus mainly on practical, but heuristic methods to *estimate* the source of an epidemic. Brockmann and Helbing [4] estimate the source of susceptible-infectious-recovered (SIR) processes by comparing the shape of the infection wavefront on the shortest path trees of different nodes. The epidemic spreads approximately uniformly from the source like a ripple in water. The nodes on the infection wavefront should therefore be approximately equidistant from the source on a shortest path tree. Prakash *et al.* [5,6] use Minimum Description Length to find the source nodes of discrete-time susceptible-infectious (SI) processes. Shah and Zaman [7] introduce the rumor centrality $R(v, G_N)$, which is the number of ways that the nodes of the subgraph G_N can be sequentially infected, starting with node v . They then explain several methods to estimate the epidemic source based on the rumor centrality. More recently, Shah *et al.* [8] used graph neural networks to find patient zero in mean-field SIR.

The problem of finding the source of an epidemic or rumour has received strangely little attention. In particular, the basic susceptible-infectious-susceptible (SIS) model, which is a cornerstone of network epidemiology [9] as most epidemics contain reinfections, is not considered, possibly because

finding the source of a SIS epidemic is significantly harder than finding the source of an SI or SIR epidemic. In this work, we reverse the governing equation(s) of the epidemic in time to find the backward equations that describe the epidemic backwards in time, to find patient zero. Like some of the papers mentioned above, we restrict ourselves to the best-case scenario: we assume complete knowledge of the underlying Markovian or N -intertwined mean-field approximation (NIMFA) [10] process, the contact network, the *exact* Markovian viral state vector $s(t)$, and the *exact* NIMFA viral state vector $V(t)$ at some known time t .

In Sec. II we introduce the Markovian and NIMFA epidemic models. We discuss both the forward and backward equations. In Sec. III we estimate the epidemic source of Markovian epidemics and in Sec. IV we trace back NIMFA SIS epidemics. We conclude in Sec. V.

II. EPIDEMIC PROCESSES

In this section we introduce the SI, SIS, and SIR Markovian processes and the NIMFA SIS process. We also discuss relevant similarities and differences between the Markovian and NIMFA SIS processes.

A. Markovian epidemics

We describe epidemics using Markovian compartmental models [11] on a static contact network, described by a graph $G(N, L)$ with N nodes and L links. The graph is fully defined by its adjacency matrix A , with elements $a_{ij} = 1$ if i and j are neighbors and $a_{ij} = 0$ otherwise. Each node represents an individual and the links represent contacts between the individuals. In a compartmental model, each node i is in a single compartment at each time t . In this work, only three compartments are considered: a node can be healthy but

*Contact author: r.d.l.persoons@tudelft.nl

susceptible (S), infectious (I), or removed (R) when the node is recovered and immune to the disease. The simplest model is the SI model, where susceptible nodes become infected by their infectious neighbours and never recover. In the SIS model, infectious nodes can cure and move back to the susceptible S compartment. Similarly, in the SIR model, nodes can cure, but they cannot be reinfected after recovery and are therefore placed in the removed R compartment instead of the susceptible S compartment. The spreading process (S→I) is a Poisson process on the links with infection rate $\tilde{\beta}_{ij}$ from node i to node j and the curing process (I→S or I→R) is a Poisson process on the nodes with curing rate δ_i for node i . All Poisson processes are independent. We define the infection rate matrix B with elements $\beta_{ij} = \tilde{\beta}_{ij}a_{ij}$ and the curing rate matrix $S = \text{diag}(\delta_1, \dots, \delta_N)$.

Each of the processes is fully defined [12] by its infinitesimal generator Q . The SI and SIS processes have 2^N possible epidemic configurations and the SIR process has 3^N configurations [13]. Therefore, the state vectors $s(t)$ of SI, SIS, and SIR are $2^N \times 1$, $2^N \times 1$ and $3^N \times 1$ vectors, respectively. The infinitesimal generators Q_{SI} , Q_{SIS} , and Q_{SIR} are $2^N \times 2^N$, $2^N \times 2^N$, and $3^N \times 3^N$ matrices, respectively. We define the states of the Markov processes and the infinitesimal generators Q_{SI} , Q_{SIS} , and Q_{SIR} as in Refs. [12–14]. For the SIS process, the explicit definition of the statespace and infinitesimal generator Q_{SIS} can also be found in Appendix A. The viral state vector $s(t)$ for each of the Markov processes is exactly given by the Chapman–Kolmogorov equation [12]:

$$s(t) = e^{Qt} s(0), \quad (1)$$

where $s(0)$ is the initial state vector. Each element $[s(t)]_i$ of the state vector $s(t)$ gives the probability that the process is in state i at time t . Since $s(t)$ is a probability distribution, the process is surely in one of the possible states, hence, $\sum_{i=1}^{|s(t)|} [s(t)]_i = 1$ for all times t , where $|s(t)|$ is the cardinality of the state vector $s(t)$. When the state vector $s(t)$ is zero everywhere except at index i , where it is 1, we write $s(t) = e_i$. The standard vector e_k has elements $(e_k)_i = \delta_{ki}$, thus $(e_k)_i = 0$, except when $k = i$, then $(e_k)_i = 1$. When $\sum_{i=1}^{|s(t)|} [s(t)]_i \neq 1$ or $[s(t)]_i \notin [0, 1]$ for some index i , we call the state vector $s(t)$ *nonphysical*.

Since e^{-Qt} is the inverse of the matrix e^{Qt} , the governing Eq. (1) immediately shows that

$$s(0) = e^{-Qt} s(t). \quad (2)$$

We call Eq. (2) the Markovian backward equation. Analogously to the matrix e^{Qt} describing steps of length t forwards in time in Eq. (1), the matrix e^{-Qt} describes steps of length t backwards in time in Eq. (2). However, unlike Eq. (1), the right-hand side of Eq. (2) is numerically ill-conditioned, because e^{-Qt} is inherently unstable for large t . Due to the difference in sign the norm of the exponential e^{Qt} decreases in t , while the norm of e^{-Qt} grows exponentially. The matrix e^{Qt} can be written as $e^{Qt} = \sum_{k=1}^N e^{\lambda_k t} x_k y_k^T$, where λ_k is the k th, generally complex, eigenvalue of the infinitesimal generator Q and x_k and y_k are the right- and left-eigenvectors, respectively [15]. The eigenvalues λ_k have a nonpositive real part, as follows from Gershgorin’s circle theorem [16]. Gershgorin’s circle theorem states that any eigenvalue of a matrix A must lay within at least one of the closed Gershgorin

discs $D(a_{ii}, R_i)$ in the complex plane, with center at a_{ii} and radius $R_i = \sum_{j \neq i} |a_{ij}|$. Since $q_{ii} = -\sum_{j \neq i} |q_{ij}| \in \mathbb{R}$, all Gershgorin discs $D(q_{ii}, R_i)$ intersect the origin but otherwise stay in the negative half-space of the complex plane. Therefore, all eigenvalues of Q have a nonpositive real part. Analogously, in the matrix $e^{-Qt} = \sum_{k=1}^N e^{-\lambda_k t} x_k y_k^T$, the eigenvalue $-\lambda_k$ has nonnegative real part, such that $e^{-\lambda_k t}$ grows exponentially with t . The matrix exponential e^{At} is *notoriously hard* to compute numerically in a general way [17]. As a consequence, the equation $e^{-Qt} e^{Qt} s(0) = s(0)$ cannot be reproduced numerically in general. Specifically, for certain matrices Q , the error $\|e^{-Qt} e^{Qt} s(0) - s(0)\|_2$ grows exponentially for large enough times t when small rounding errors propagate into large errors. Therefore, we will employ different matrix exponentiation methods and compare their results. MATLAB provides the functions `expm` [18,19] and `expmv` [20] natively and `expv` through the EXPOKIT package [21]. All of these methods have increasing rounding error propagation issues in the calculation of e^{-Qt} for large t and, therefore, cannot accurately retrieve the exact initial state vector $s(0)$ from the state vector $s(t)$ when $t \gg 0$. We discuss why SIS is the more difficult case compared to SI and SIR in Appendix B and restrict ourselves to SIS from here on.

B. N -intertwined mean-field approximation

The N -intertwined mean-field approximation (NIMFA) of the Markovian SIS model [14] reduces the Markovian process to a system of N differential equations. The mean-field infection probability of node i is written as $v_i(t)$. The $N \times 1$ vector of the infection probabilities is $V(t) = [v_1(t), v_2(t), \dots, v_N(t)]^T$. The heterogeneous NIMFA governing equations are given by

$$\frac{d v_i(t)}{dt} = -\delta_i v_i(t) + [1 - v_i(t)] \sum_{j=1}^N \beta_{ij} v_j(t). \quad (3)$$

The NIMFA process has a phase transition determined by the basic reproduction number $R_0 = \rho(S^{-1}B)$, where $\rho(\cdot)$ indicates the spectral radius or largest eigenvalue [22,23]. Above the epidemic threshold ($R_0 > 1$) the infection probability vector $V(t) \neq 0$ converges for large time t to a nonzero steady-state vector V_∞ . Below and at the threshold ($R_0 \leq 1$) the steady-state infection probability vector V_∞ is the all-zero vector. The NIMFA SIS process converges slowest to the steady-state V_∞ around the epidemic threshold [24]. Because NIMFA has not been solved analytically in closed form, it is often solved approximately by repeatedly applying the forward Euler [25] method:

$$v_i(t+h) = v_i(t) + h \left(-\delta_i v_i(t) + [1 - v_i(t)] \sum_{j=1}^N \beta_{ij} v_j(t) \right). \quad (4)$$

We derive the NIMFA backward equations from the forward Euler equation $v_i(t+h) = v_i(t) + h \frac{d v_i(t)}{dt}$ by rewriting

$$v_i(t) = v_i(t+h) - h \frac{d v_i(t)}{dt} \text{ and shifting the time } t \text{ by } -h:$$

$$v_i(t-h) = v_i(t) - h \left(-\delta_i v_i(t-h) + [1 - v_i(t-h)] \sum_{j=1}^N \beta_{ij} v_j(t-h) \right). \quad (5)$$

The NIMFA backward equations (5) can be interpreted as a backward Euler method [25] going backwards in time. The backward Euler method has the form $v_i(t+h) = v_i(t) + h \frac{d v_i(t+h)}{dt}$ forwards in time. Backwards in time the method takes the same form as Eq. (5), namely, $v_i(t-h) = v_i(t) - h \frac{d v_i(t-h)}{dt}$. Because the backward Euler method is *implicit* [25], a nonlinear system of N equations with N unknowns has to be solved at each iteration step. Indeed, the unknown $v_i(t-h)$ is on both sides of Eq. (5).

It may seem straightforward to apply the forward Euler equations (but backwards in time) instead of Eq. (5). Specifically, applying a method of the form $v_i(t-h) = v_i(t) - h \frac{d v_i(t)}{dt}$, thus avoiding the need of solving a system of equations at each iteration step. However, like the Markovian case, the NIMFA system is ill-behaved backwards in time. The forward Euler method is, in fact, very unstable backwards in time. The backward Euler method (5), fortunately, is accurate when the number of steps is small enough. We found that the range of Eq. (5) can be extended beyond what is reasonably required by working with 128 bit variables (also known as quadruple accuracy).

C. Comparison between Markovian and NIMFA SIS

In this section we will discuss some similarities and differences of Markovian and NIMFA SIS processes relevant to the study.

1. State vector

The Markovian $2^N \times 1$ state vector $s(t)$ describes the probability of each of the 2^N possible network configurations of the N nodes being in either the I or S compartment. The NIMFA $N \times 1$ state vector $V(t)$ describes the infection probability of each of the N nodes. The NIMFA state vector can also be interpreted as the fraction of infectious individuals in a group when a node represents a building, city or country. While the infection probability of an individual and the probability of each of the 2^N possible Markovian states are (almost) impossible to estimate accurately during a real-world epidemic, the fraction of infected individuals in cities or countries is not. We argue that assuming a good estimation of $V(t)$ is realistic for real-world epidemics when nodes represent groups of individuals.

2. Governing equations

Both forward equations (1) and (4) are numerically stable. Both backward equations (2) and (5) are not. In the following we will see that the NIMFA backward equations are accurate for much larger times t than the Markovian ones. The other main difference between the Markovian backward equation (2) and the NIMFA backward equations (5) is that the Markovian backward equation maps the state vector $s(t)$




FIG. 1. Left: Accuracy of the estimation of patient zero z in Markovian SIS when t is known exactly on the left y axis and the average probability of being in the steady state $\xi(t)$ on the right y axis versus time. The accuracy and steady-state percentage are averaged over 1000 epidemic processes. The contact graph is the complete graph on $N = 8$ nodes K_8 , the infection rates β_{ij} are random uniform on the interval $[0,1]$ and the curing rate $\delta = 1$. All three exponentiation methods are shown: expm in blue (X marker), expmv in red (circle marker), and expv in black (square marker). The steady-state fraction $\xi(t)$ is shown in dotted green. Right: Norm of the error $s(0) - e^{-Qt} s(t)$ versus time for the three exponentiation methods.

to the initial state $s(0)$ in one step, while the NIMFA backward equations solve the entire trajectory of the epidemic.

3. Steady-state information loss

The Markovian SIS steady state is the all-healthy state, which corresponds to the basic vector e_1 . The first row of the infinitesimal generator Q is all-zero (see Ref. [12] or Appendix A), which implies that the first row of e^{Qt} is all-zero for every $t \in \mathbb{R}$ and thus $e_1 = e^{Qt} e_1$ for all times t (including negative t). Similarly, in the NIMFA case, the all-healthy state $V(t) = 0$ gives $V(t-h) = 0$ after solving Eq. (5) and the nonzero steady-state $V(t) = V_\infty$ gives $V(t-h) = V_\infty$ after solving Eq. (5). This means that after converging to the steady state the history of neither process can be retrieved.

III. TIME REVERSAL IN MARKOVIAN SIS

A. Estimating patient zero when the elapsed time is known

Since we are interested in estimating patient zero z , we assume that $s(0) = e_i$, where $i = 2^{z-1}$, which is the state corresponding to every node being susceptible, but only node z is infectious (see Appendix A). We try to recover patient zero z from the state vector $s(t)$ by determining $s(0) = e^{-Qt} s(t)$. Patient zero is then estimated to be the node l corresponding to the state $i = 2^{l-1}$ which has the largest component in the approximate initial state $s(0)$ out of the states $i = 2^{n-1}$ with exactly one infected node $n = 1, \dots, N$.

Figure 1 shows the accuracy of estimating patient zero averaged over 1000 SIS processes. Here, the contact graph is the complete graph on $N = 8$ nodes K_8 , the infection rates β_{ij} are random uniform on the interval $[0,1]$ and the curing rate $\delta = 1$. Appendix C discusses the simulation parameters in more detail. The left y axis shows the accuracy, which is the fraction of trials where patient zero was estimated correctly. The accuracy is shown for all three exponentiation methods: expm in blue (X marker), expmv in red (circle marker) and expv in black (square marker). On the right y axis, in dotted green, the steady-state fraction $\xi(t)$ is shown, which is the average probability of being in the steady state at time t . More

specifically, the steady-state fraction $\xi(t)$ is the average of $[s(t)]_i = \Pr[\text{all-healthy state}]$ over the 1000 trials. The right panel shows the norm of the error $s(0) - e^{-Qt}s(t)$.

The left panel of Fig. 1 indicates that patient zero can sometimes still be found by `expv` at large times $t > 30$. The accuracy of `expm` and `expmv` falls off, but also stays above zero significantly longer than the error in the right panel stays small. Interestingly, for times $t > 3$, the accuracy of all three methods is slightly above $\frac{1}{N} = 0.125$, which would be the accuracy of a random node choice. The fact that the steady-state fraction $\xi(t)$ has not reached 1 when the accuracy decreases around $t = 3$ suggests that information loss through the steady state is not the main cause of the accuracy decrease. We were able to determine the cause of the second drop in accuracy of the `expm` and `expmv` methods between $t = 30$ and $t = 35$. At these times the entries of the vector $e^{-Qt}s(t)$ grow so large that MATLAB interprets them as infinite. Then, each of the N nodes shares the same entry $[s(0)]_i$ (where $i = 2^{k-1}$) and we consider these cases a failure to prevent an unrealistically high accuracy when multiple entries tend to zero or infinity. This reason also explains why the right panel of Fig. 1 stops at $t = 30$. Indeed, for higher values of t the norm of the error is infinite according to MATLAB. The right panel of Fig. 1 also shows that the estimated initial state vector $s(0)$ is nonphysical from at least $t \approx 3$.

We conclude from Fig. 1 that, although the errors in the matrix exponential and in the initial state vector $s(0)$ increase exponentially, estimation can still work at times $t > 10$. To show more explicitly that the low accuracy is mainly due to numerical errors in the matrix exponentiation we considered quadruple accuracy, that approximately doubled the time for which the numerical methods were accurate to $t = 6$.

B. Finding the source node when t is unknown

The assumption that the elapsed time t is known is rather unrealistic. In this section, we investigate the case when the time t is unknown. If we denote by $t^* > 0$ the estimated elapsed time, then we can estimate the initial viral state $s(0)$ as $s(t - t^*)$, when $t^* \approx t$:

$$s(0) \approx s(t - t^*) = e^{-Qt^*} s(t). \quad (6)$$

Equation (6) follows from the semigroup property of the exponential $s(t - t^*) = e^{Q(t-t^*)} s(0) = e^{-Qt^*} e^{Qt} s(0)$ and from Eq. (1) leading to $s(t - t^*) = e^{-Qt^*} s(t)$.

Since e^{-Qt} has nonnegative eigenvalues, the norm of $s(t - t^*)$ will explode for large t^* . Analytically, the vector $s(t - t^*)$ must be a probability vector for $t - t^* \geq 0$ and therefore this explosion cannot happen before $t = t^*$. Under our assumption that $s(0) = e_i$, where e_i corresponds to an initial state with only patient zero z infected, simulations suggest the explosion starts exactly at $t^* = t$:

Conjecture 1. Denote with $s(0) = e_i$ the initial state of an Markovian epidemic process, with finite infection rates β_{ij} and finite curing rates δ_i . If e_i is not an absorbing state, then it holds for $t > 0$ that the norm $\|e^{-Qt}s(0)\|_2 = \|e^{-Qt}e_i\|_2 > 1$.

Given that the initial state vector $s(0) = e_i$ and that the time $0 < t < \infty$, the state vector $s(t)$ is a probability vector with norm $\|s(t)\|_2 < 1$. The norm of the initial state vector $\|s(0)\|_2 = \|e_i\|_2 = 1$ and thus Conjecture 1 states that for




FIG. 2. Left: Accuracy of the estimation of patient zero z in Markovian SIS when t is not known on the left y axis and the average probability of being in the steady state $\xi(t)$ on the right y axis versus time. The accuracy and steady-state percentage are averaged over 100 epidemic processes. The contact graph is the complete graph on $N = 8$ nodes K_8 , the infection rates β_{ij} are random uniform on the interval $[0,1]$ and the curing rate $\delta = 1$. All three exponentiation methods are shown: `expm` in blue (X marker), `expmv` in red (circle marker), and `expv` in black (square marker). The steady-state fraction $\xi(t)$ is shown in dotted green. Right: Norm of the error $s(0) - e^{-Qt^*} s(t)$ versus time for the three exponentiation methods.

negative t the norm $\|s(t)\|_2 > 1$. Therefore, we can use a binary search algorithm [26] to find t using the following criteria on $s(t - t^*)$:

$$\begin{aligned} \|s(t - t^*)\|_2 &< 1 && \text{if } t^* < t, \\ \|s(t - t^*)\|_2 &= 1 && \text{if } t^* = t, \\ \|s(t - t^*)\|_2 &> 1 && \text{if } t^* > t. \end{aligned} \quad (7)$$

Here, we consider the 2-norm $\|\cdot\|_2$ instead of the 1-norm $\|\cdot\|_1$, which is usually more natural in probability theory, because $\|s(t - t^*)\|_1 = 1$ when $t^* < t$, which leads to a step function in Eq. (7) that hinders the binary search. The binary search algorithm is explained in more detail in Appendix D.

Figure 2 shows the estimation of patient zero when the elapsed time t is unknown. The estimation is shown for all three exponentiation methods: `expm` in blue (X marker), `expmv` in red (circle marker), and `expv` in black (square marker). The averages are taken over 100 SIS simulations. Comparison of the left panels in Figs. 1 and 2 indicates that *not knowing* the elapsed time t has little impact on the accuracy because the left panels of Figs. 1 and 2 are quite similar.

The binary search algorithm can fail when numerical errors make the relations in Eq. (7) unreliable. While accurate up to $t = 2$, all methods estimated $t^* \approx 2$ for all $t \geq 2$ and `expv` additionally estimated $t^* \approx 50$ for all $t \geq 15$. It is unclear why the accuracy in the left panel of Fig. 2 is not impacted by these significant errors. It is also unclear why `expm` and `expmv` have such different errors in the right panel of Fig. 2, while they have the same estimated elapsed time $t^* \approx 2$ for times $t \geq 2$. We hypothesized that a remnant of the initial state $s(0) = e_i$ is present in $s(t)$ such that $[s(t - t^*)]_i$ is the largest of the entries of $s(t - t^*)$ representing states with a single infected node for a large range of times t^* . Then, accurate estimation of the elapsed time t is relatively unimportant when estimating patient zero. However, this is not true as simulations (not shown here) indicate that the index of the maximum entry from the set $\{[s(t - t^*)]_i | i = 2^{k-1}\}$, fluctuates significantly with t^* in a single epidemic. The right panel of Fig. 2 suggests that the estimated initial state $s(0)$ is nonphysical for `expmv` and `expv`.



FIG. 3. Schematic visualization of the behavior of the backward NIMFA equations. Top: Regular double (64-bit) precision. Bottom: Quadruple (128-bit) precision. Proportions of the different regimes are loosely based on simulations of homogeneous NIMFA SIS slightly above the epidemic threshold ($R_0 = 1.1$) on Barabási-Albert random graphs of size $N = 50$. The curing rate was chosen $\delta = 1$ and the infection rate β was chosen such that the basic reproduction number $R_0 = 1.1$.

IV. TIME REVERSAL IN NIMFA SIS

In NIMFA SIS, the initial state $V(0) = e_z$ corresponds to node z being infected with probability 1 while all other nodes are not infected, i.e., node z is patient zero. The NIMFA backward equations (5) are computationally nicer than the Markovian backward equation (2). First, solving the system of equations (5) is computationally easier than calculating the matrix exponent in Eq. (2). In addition, the system has size N instead of 2^N . Therefore, significantly larger graphs can be investigated. Second, the assumption that $s(t)$ is known in the Markovian case is not very realistic. In a real-world epidemic, it is implausible to estimate the probability of each of the 2^N states accurately. Conversely, assuming that $V(t)$ is known is reasonable, because the infection probabilities $v_i(t)$ can be estimated as the fraction of infected people in a node when nodes represent buildings, cities or countries instead of individuals. Last, the backward Euler method (5) determines the infection probability vector V at all times in the interval $[0, t]$. This means information about the trajectory of the epidemic is also retrieved at no additional cost.

A. Observations

The NIMFA backward equations show three regimes of accuracy, that are visualized in Fig. 3. For small times, the method is accurate and the backward epidemic corresponds with the forward epidemic. An example is shown in the left panel of Fig. 4, where each line represents the infection probability $v_i(t)$ of one of the N nodes. There is one thick line for each of the $N = 10$ nodes that shows the value of $v_i(t)$ in the forward epidemic. The backward epidemic is shown with + symbols. The line and + symbols of a single node have the same color. In the left panel the + symbols overlap with

the lines, showing the epidemic is traced accurately. When the elapsed time increases errors start propagating and after a time $t = \mathcal{T}$, the backward epidemic starts diverging from the forward epidemic. This divergence is shown in the middle panel of Fig. 4, where around $t = 10$ all nodal states diverge. The backward process starts at $t = 25$ and is thus accurate for approximately 15 time units until $t = 10$. The divergence time \mathcal{T} seems largely determined by R_0 and the precision of the viral state $V(t)$. We explore this finding in more detail below. The last regime in Fig. 4 is when the process is in the steady state. The right panel of Fig. 4 shows that the backward process is stuck in the steady state. The convergence time of NIMFA is mainly [24] determined by the basic reproduction number R_0 . It is also influenced by the numerical precision, as illustrated in the bottom panel of Fig. 3. A graph with $N = 10$ nodes is shown in Fig. 4 such that the figure is still clear with all $v_i(t)$ shown, but the figure is representative of larger graphs as well. For example, we found that the divergence time \mathcal{T} , for $R_0 = 1.1$ (like in Fig. 4), is around 12 for graphs of various sizes including as high as $N = 1500$.

B. Divergence time \mathcal{T}

The divergence time \mathcal{T} is the largest time before the backward process diverges. Moreover, simulations suggest that the divergence time is independent of the starting time of the backward process (assuming the steady state is not reached). Hence, the divergence time is also approximately the largest starting time from which the nodal probability vector $V(t)$ can be retraced back to $V(0)$ accurately.

As mentioned above, the divergence time \mathcal{T} is determined by the basic reproduction number R_0 and the numerical precision of the forward and backward processes. Since the basic




FIG. 4. Nodal state v_i for each of the $N = 10$ nodes over time, forward epidemic (thick line) and backward epidemic (+ symbol) values have the same color for each node. Different times of starting the reversal corresponding to the different regimes. The forward epidemic starts at $t = 0$, while the backward epidemic starts at $t = 10$ (left), $t = 25$ (middle), and $t = 40$ (right). The three figures correspond with the three regimes: accurate (left), divergence (middle), and stuck in the steady state (right). The contact graph has $N = 10$ nodes, the curing rate $\delta = 1$ and the infection rate β is chosen such that the basic reproduction number $R_0 = 1.1$. The x axis is in logscale for clarity.

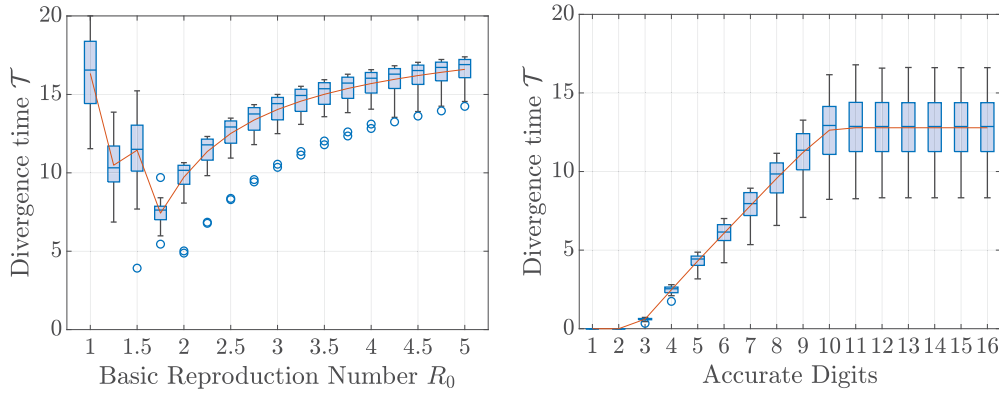


FIG. 5. Left: Divergence time \mathcal{T} versus basic reproduction number R_0 . The infection rates $\tilde{\beta}_{ij} = \beta$ are chosen homogeneous and such that the basic reproduction number R_0 is what we choose. Right: Divergence time \mathcal{T} versus the number of accurate digits of the state vector $V(t = 20)$ (the initial state of the backward equations) with basic reproduction number $R_0 = 1.1$, guaranteeing that the process has not converged at $t = 20$. The state vector $V(t)$ is rounded to have the right amount of accurate digits. For both panels the sample is taken over 50 random graphs of size $N = 50$, including Barabási-Albert and Erdős-Rényi random graphs, with curing rate $\delta = 1$.

reproduction number R_0 influences the convergence time, it could be possible that for large values of R_0 the diverging regime in Fig. 3 is replaced by the steady-state regime. However, the basic reproduction number R_0 influences the divergence time \mathcal{T} in such a way that this does not happen. Indeed, even if the process converges very quickly all three regimes as sketched in Fig. 3 exist. Figure 5 shows the divergence time \mathcal{T} for different homogeneous infection rates $\tilde{\beta}_{ij} = \beta$ corresponding to specific R_0 values (left) and for different numbers of accurate digits in $V(t)$ (right). The number of accurate digits is varied by rounding the viral state $V(t)$ to the nearest decimal number with that amount of digits after the period. For example, if one of the elements of $V(t)$ is $v_i(t) = 0.123456789$, then with four accurate digits it would be $v_i(t) = 0.1235$ and with seven accurate digits $v_i(t) = 0.1234568$. The divergence time \mathcal{T} is estimated as the time of the first step in the backward epidemic where the difference between the forward and backward value of any of the N infection probabilities $v_i(t)$ is more than 10^{-3} . For both panels in Fig. 5 the epidemic is reversed from $t = 20$. The time $t = 20$ is larger than \mathcal{T} , but also small enough to avoid convergence at $R_0 = 1.1$. By fixing the starting time, instead of searching for the smallest time where $V(0)$ is retrieved without an error of more than 10^{-3} in any of the N infection probabilities $v_i(t)$, we save computation time. The left panel shows that from $R_0 = 2$ on, the divergence time (and thus the accuracy) increases. For $R_0 > 2$, the backward process is stuck in the steady state on the entire interval. The accuracy increases because, for higher values of R_0 , the forward process reaches the steady state faster and is in the steady state for a larger part of the interval $[0, 20]$. Therefore, the backward process, that is stuck in the steady state during the entire interval, deviates from the forward process after a longer time when R_0 increases. As eluded to before, for $R_0 < 2$, the divergence time decreases with R_0 . Therefore, even though for $R_0 = 1.5$ the process converges faster than for $R_0 = 1.1$, there still is a diverging regime. The right panel of Fig. 5 shows that the backward equations stay accurate longer when the state vector $V(t)$ is known more precisely. We are

uncertain why the improvement halts at 10 accurate digits, while MATLAB should have 16 decimal digits of precision as default. Presumably, the forward Euler method has already introduced errors that make $V(t)$ accurate in only 10 decimal digits. The requirement of a large number of accurate digits emphasizes the need for accurate and high-resolution measurements. In quadruple precision computations, we found that, even near the epidemic threshold, an epidemic reaching times large enough to show convergence at double precision could still be retraced accurately. The accuracy regimes at quadruple precision are illustrated in the bottom part of Fig. 3. Quadruple precision may not be practical however, because 32 digits of accuracy is not realistic in a real-world setting.

V. CONCLUSIONS

Finding the source of a continuous-time Markovian epidemic via $s(0) = e^{-Qt}s(t)$ in Eq. (2) is significantly hindered, because the matrix exponent e^{-Qt} is ill-conditioned and cannot be accurately calculated numerically unless the time t is small. We investigate how well the source node, or patient zero z , can be estimated when (1) the elapsed time t and the state vector $s(t)$ are known and (2) only the state vector $s(t)$ is known. Both scenarios show very similar results, because both methods are significantly impacted by the numerical errors. We introduce a binary search algorithm, based on conditions in Eq. (7), which can estimate the elapsed time t when the initial state $s(0) = e_i$. Even though numerical errors should prevent the algorithm from working, it still performs well and the estimation of patient zero is effectively unchanged when the elapsed time t is unknown.

While we claim that our time reversal method based on Eq. (2) is analytically and (ignoring numerical challenges) algorithmically a best-case scenario for Markovian SIS, heuristic algorithms can outperform its accuracy due to its exponential vulnerability to numerical errors. Our Markovian method is not suitable for finding patient zero in real-world situations even if the entire Q matrix and state vector $s(t)$ are known (a best-case scenario).

Given that time reversal of the Markovian process is unfeasible, the NIMFA mean-field model is one of the best alternative models. To find the source of a NIMFA SIS epidemic, we repeatedly apply the backward equations (5) to the state vector $V(t)$ until the initial state $V(0)$ is reached. We observe three regimes: for short elapsed times t the method is accurate, for intermediate elapsed times t the backward epidemic diverges and for large elapsed times t the backward epidemic is stuck in the steady state. The basic reproduction number R_0 and the accuracy of the initial state of the backward epidemic $V(t)$ influence the accuracy of the backward epidemic. Increasing R_0 decreases the divergence time \mathcal{T} and causes the diverging regime to exist for all R_0 , even if the epidemic converges very fast at large R_0 . The precision of $V(t)$ has a large positive effect on the accuracy of the backward epidemic. Using quadruple precision we can retrace the epidemics accurately from times t that are large enough such that the process is converged at that time t , if one had used double precision instead. Unlike the assumption of a known Markovian state vector $s(t)$, the assumption that the NIMFA state vector $V(t)$ is known is realistic. Additionally, because the backward NIMFA method is more stable and scales better with the network size N than the Markovian method, it is applicable for graphs with size $N = 1500$ or more. Measurements of the required precision will be incredibly rare in real-world scenarios. However, this is not necessarily a problem because, for small times, the backward equations are

very accurate and less dependent on the precision. Times like $t = 10$ may seem small, but the time t is defined in terms of the curing rate δ and a time of $t = 10\frac{1}{\delta}$ could reasonably correspond to multiple weeks for some epidemics.

ACKNOWLEDGMENTS

We are grateful to Kees Vuik for his useful insights and a reviewer for providing comments that have improved this paper. This research has been funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant Agreement No. 101019718).

APPENDIX A: MARKOVIAN SIS STATESPACE AND INFINITESIMAL GENERATOR

The state i in the Markovian SIS process is given by

$$i = \sum_{k=1}^N x_k(i)2^{k-1},$$

where $x_k(i) \in \{0, 1\}$ is the viral state of node k in state i and $x_k(i) = 0$ if node k is susceptible in state i and $x_k(i) = 1$ if node k is infectious in state i . The state i can equivalently be written as $x_N(i)x_{N-1}(i) \dots x_2(i)x_1(i)$, which corresponds to the binary representation of the number i . The $2^N \times 2^N$ infinitesimal generator Q_{SIS} is given by [12]

$$q_{ij} = \begin{cases} \sum_{k=1}^N \beta_{mk}x_k(i) & \text{if } j = i + 2^{m-1}; m = 1, 2, \dots, N \text{ and } x_m(i) = 0, \\ \delta_m & \text{if } j = i - 2^{m-1}; m = 1, 2, \dots, N \text{ and } x_m(i) = 1, \\ -\sum_{k=0; k \neq j}^{2^N-1} q_{kj} & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases}$$

$$Q_{\text{SIS}} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \delta_1 & -\delta_1 - \beta_{12} & 0 & \beta_{12} & 0 & 0 & 0 & 0 \\ \delta_2 & 0 & -\delta_2 - \beta_{21} - \beta_{23} & \beta_{21} & 0 & 0 & \beta_{23} & 0 \\ 0 & \delta_2 & \delta_1 & -\delta_1 - \delta_2 - \beta_{23} & 0 & 0 & 0 & \beta_{23} \\ \delta_3 & 0 & 0 & 0 & -\delta_3 - \beta_{32} & 0 & \beta_{32} & 0 \\ 0 & \delta_3 & 0 & 0 & \delta_1 & -\delta_1 - \delta_3 - \beta_{12} - \beta_{32} & 0 & \beta_{12} + \beta_{32} \\ 0 & 0 & \delta_3 & 0 & \delta_2 & 0 & -\delta_2 - \delta_3 - \beta_{21} & \beta_{21} \\ 0 & 0 & 0 & \delta_3 & 0 & \delta_2 & \delta_1 & -\delta_1 - \delta_2 - \delta_3 \end{bmatrix}.$$

We also gave an explicit example of the $2^3 \times 2^3$ infinitesimal generator for the path graph on $N = 3$ nodes. Nodes 1 and 3 are the endpoints of the path graph, node 2 is in the middle. Therefore, $\beta_{13} = \beta_{31} = 0$, since $a_{13} = a_{31} = 0$. Figure 6 visualizes the states and transitions of the SIS process on a graph with $N = 4$ nodes. The states 2^z , with integer z , on the second row are the states corresponding to a single infected node z . Exact computations with the $2^N \times 2^N$ infinitesimal generator Q_{SIS} are limited due to exponentially growing matrix size with N . On a normal personal computer, the infinitesimal generator can be computed up to $N \approx 25$, but linear calculations are feasible only up to $N \approx 18$. In this work, we require repeated matrix exponentiation, which is feasible up to $N \approx 10$.

APPENDIX B: COMPARISON OF SI, SIR, AND SIS

Figure 7 shows the average error $\|e^{Qt}e^{-Qt} - I\|_2$ for different times t , where $\|\cdot\|_2$ indicates the matrix 2-norm. The average error is calculated with the `expm` function, since the other methods do not calculate e^{Qt} or e^{-Qt} directly. Figure 7 shows the average error for $Q = Q_{\text{SI}}$, $Q = Q_{\text{SIS}}$, and $Q = Q_{\text{SIR}}$ averaged over 1000 different infinitesimal generators, with heterogeneous infection rates b_{ij} i.i.d. uniform between 0 and 1. The curing rate $\delta = 0$ for SI, while the curing rate $\delta = 1$ for SIS and SIR. The underlying contact graph was the complete graph K_6 for SIR and K_8 for SI and SIS, because of the exponential scaling 3^N versus 2^N .




FIG. 6. The state diagram of SIS on a graph with $N = 4$ nodes, from Ref. [12].

We hypothesize that the reason why the average error in Fig. 7 grows faster in time for SIS than for SI and SIR is due to the structure of the infinitesimal generator Q_{SIS} . The infinitesimal generators Q_{SI} and Q_{SIR} are upper triangular matrices, which are numerically more suited for exponentiation than Q_{SIS} , which is not triangular. Triangular matrices have the property that their powers are also triangular. The matrix exponent e^A can be written as $e^A = \sum_{k=0}^{\infty} \frac{1}{k!} A^k$ and therefore the matrix exponent e^T of a triangular matrix T is also triangular. This means that $\frac{N(N-1)}{2}$ elements of e^T are 0 and will neither contribute to nor get influenced by propagating numerical errors, reducing the average error.

The triangular infinitesimal generators of SI and SIR are due to the property that in both processes a state i can never be reached again after the process has left state i (i.e., no states communicate).

We explain with an example why finding patient zero in SIS is more difficult compared to SI and SIR. Suppose node z is the only infected node. Then, after infecting its neighbor w , both z and w are infected. Now in the SIS process, node z can cure leaving only node w infected. Due to the memoryless property of the Markov process, the remainder of this hypothetical SIS process will be identical to one that started with

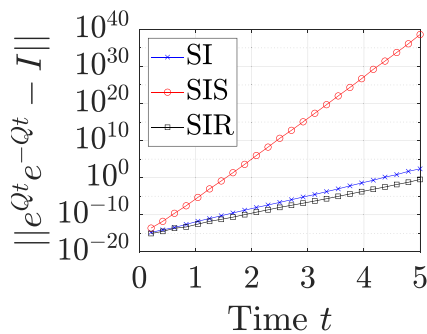


FIG. 7. Development of the error $\|e^{Qt}e^{-Qt} - I\|_2$ over time. Here, $\|\cdot\|_2$ represents the matrix 2-norm. The error is shown for SI in blue (X marker), SIS in red (circle marker), and SIR in black (square marker). We emphasize that SIR is on the graph K_6 instead of K_8 like SI and SIS, which could make the comparison skewed.

patient zero w , although patient zero was in fact node z . In SI and SIR a similar situation cannot occur: in SI node z can cure and in SIR node z would be in compartment R instead of compartment S. Both SI and SIR cannot reach the state where only node w is infected from the state where only node z is infected, which makes it easier to find patient zero for these processes. We emphasize that the problem from the example is an obstacle in general: both heuristic models as well as exact ones are obstructed.

APPENDIX C: SIMULATION PARAMETERS

1. Curing rates

Since we investigate the impact of the elapsed time t when estimating the initial state $s(0)$ or $V(0)$ from the state vector $s(t)$ or $V(t)$, we require the time to be dimensionless. In a homogeneous Markovian process, where each node i has the same curing rate $\delta_i = \delta$, the average curing time $\frac{1}{\delta}$ is the characteristic time. The system can be made dimensionless by considering the infinitesimal generator $\tilde{Q} = \frac{1}{\delta}Q$ and dimensionless time $\tilde{t} = \delta t$, where the time \tilde{t} is specified in units of $\frac{1}{\delta}$. The dimensionless system allows us to compare the impact of the elapsed time between processes and removes one parameter, because $\delta = 1$ can be fixed without loss of generality in the dimensionless system. The dimensionless Markovian process directly translates to a dimensionless NIMFA process when the curing rate matrix $S = I$. In this work, we consider the dimensionless system, with heterogeneous infection rates β_{ij} , but homogeneous curing rates $\delta_i = \delta = 1$.

2. Infection rates

The accuracy of the Markovian process for large times t depends on the infinitesimal generator Q and thus on the infection rate matrix B and the curing rate matrix S . In particular, some methods perform poorly when there are small differences in the eigenvalues of the infinitesimal generator Q (see Refs. [12,27] about eigenvalues of multiplicity larger than 1). Therefore, we will consider heterogeneous infection rates on the complete graph, to reduce degeneracy of the process due to high multiplicity of the eigenvalues of the infinitesimal generator Q . Specifically, we took $\tilde{\beta}_{ij}$ uniform between 0 and 1 on the complete graph K_8 for the Markovian simulations. This choice guarantees that the epidemic does not die-out rapidly. For the NIMFA simulations the multiplicity of the eigenvalues of the infinitesimal generator Q do not matter and therefore we investigated different cases as well. In particular we chose a homogeneous infection rate β such that the basic reproduction number can be controlled since $R_0 = \beta\rho(A)$.

3. Step-size h

The size of the step-size h used to numerically solve the NIMFA backward equations is a parameter of large importance. A “sweet spot” has to be found, because both large and small h reduce the accuracy of the solver. If the step-size is too large, then the numerical solution is inaccurate. However, the process will also escape the steady state easier if the step-size h is increased, because $h \frac{dV(t)}{dt}$ increases in magnitude and is thus further away from zero. If the step-size h is too small,

then the process gets stuck in the steady state easier. In addition, errors accumulate faster (because there are more steps) and the increase in accuracy of the numerical method does not necessarily compensate this. We obtained the best results around $h = 0.01$, which is the value used in all simulations in this paper.

4. Solver accuracy

We investigated whether increasing the accuracy of the solver that solves the system of equations in Eq. (5) every iteration step improved the solution of the backward Euler method. This had no impact on the accuracy of the method as the improvement was very small and the divergence is thus likely caused by the inherent instability of the backward Euler equation (5).

APPENDIX D: BINARY SEARCH ALGORITHM

The binary-search Algorithm 1 explained below finds t and $s(0)$ as follows. Assume Q and $s(t)$ are known and the time t lies in an interval $[l, h]$. Then, the binary-search loop (lines 5 through 17) will, at each iteration, set the estimated time t^* to $\frac{l+h}{2}$ (line 7) and then calculate $s(0) = e^{-Qt^*} s(t)$ (line 8). The interval $[l, h]$ is narrowed to $[t^*, h]$ or to $[l, t^*]$ using criteria (7) (lines 9 through 16). [As a minor optimization, $\sum_{i=1}^{|s(t-t^*)|} [s(t-t^*)]_i^2$ is calculated instead of $\|s(t-t^*)\|_2$, because taking the square root on both sides does not change the criteria in Eq. (7) because $\sqrt{1} = 1$.] The binary-search loop is then repeated until $|\sum_{i=1}^{|s(t-t^*)|} [s(t-t^*)]_i^2 - 1| < \epsilon$. Numerical errors due to the matrix exponential can prevent convergence of the binary search at large times t , because the

ALGORITHM 1. Finding $s(0)$ and t with perfect knowledge of $s(t)$.

Input: State vector $s(t)$, infinitesimal generator Q , bounds on the time l and h , tolerance ϵ and maximal number of iterations K

Output: Estimated elapsed time t^* and estimated initial state vector $s(0)$

```

1  $t_l \leftarrow l$ ;
2  $t_h \leftarrow h$ ;
3  $E \leftarrow 2$ ;
4  $i \leftarrow 0$ ;
5 While  $|E - 1| > \epsilon$  &&  $i < K$  do
6    $i \leftarrow i + 1$ ;
7    $t^* \leftarrow \frac{t_h + t_l}{2}$ ;
8    $s(0) \leftarrow e^{-Qt^*} s(t)$ ;
9    $E \leftarrow \sum_{i=1}^N [s(0)]_i^2$ ;
10  if  $E > 1$  then
11     $t_h = t^*$ ;
12  else
13    if  $E < 1$  then
14       $t_l = t^*$ ;
15    end
16  end
17 end

```

estimated state vectors $e^{-Qt} s(t)$ can be nonphysical. Therefore, we also output t^* and $s(0)$ after some number of iterations K to guarantee termination. Algorithm 1 only finds the elapsed time t if $s(0) = e_i$ for some i because the basic vectors e_i are the only probability vectors with 2-norm equal to 1.

-
- [1] M. Worobey, J. I. Levy, L. Malpica Serrano, A. Crits-Christoph, J. E. Pekar, S. A. Goldstein, A. L. Rasmussen, M. U. G. Kraemer, C. Newman, M. P. G. Koopmans *et al.*, The Huanan seafood wholesale market in Wuhan was the early epicenter of the COVID-19 pandemic, *Science* **377**, 951 (2022).
 - [2] J. E. Pekar, A. Magee, E. Parker, N. Moshiri, K. Izhikevich, J. L. Havens, K. Gangavarapu, L. M. M. Serrano, A. Crits-Christoph, N. L. Matteson *et al.*, The molecular epidemiology of multiple zoonotic origins of SARS-COV-2, *Science* **377**, 960 (2022).
 - [3] Senate Homeland Security Committee hearing—Origins of COVID-19: An examination of available evidence, <https://www.hsgac.senate.gov/hearings/origins-of-covid-19-an-examination-of-available-evidence/> (2024), accessed July 5, 2024.
 - [4] D. Brockmann and D. Helbing, The hidden geometry of complex, network-driven contagion phenomena, *Science* **342**, 1337 (2013).
 - [5] B. A. Prakash, J. Vreeken, and C. Faloutsos, Spotting culprits in epidemics: How many and which ones? in *Proceedings of the IEEE 12th International Conference on Data Mining (IEEE, Piscataway, NJ, 2012)*, pp. 11–20.
 - [6] B. A. Prakash, J. Vreeken, and C. Faloutsos, Efficiently spotting the starting points of an epidemic in a large graph, *Knowl. Inf. Syst.* **38**, 35 (2014).
 - [7] D. Shah and T. Zaman, Rumors in a network: Who’s the culprit? *IEEE Trans. Inf. Theory* **57**, 5163 (2011).
 - [8] C. Shah, N. Dehmamy, N. Perra, M. Chinazzi, A.-L. Barabási, A. Vespignani, and R. Yu, Finding patient zero: Learning contagion source with graph neural networks, [arXiv:2006.11913](https://arxiv.org/abs/2006.11913).
 - [9] R. Pastor-Satorras, C. Castellano, P. Van Mieghem, and A. Vespignani, Epidemic processes in complex networks, *Rev. Mod. Phys.* **87**, 925 (2015).
 - [10] P. Van Mieghem, The N-intertwined SIS epidemic network model, *Computing* **93**, 147 (2011).
 - [11] F. D. Sahneh, C. Scoglio, and P. Van Mieghem, Generalized epidemic mean-field model for spreading processes over multilayer complex networks, *IEEE/ACM Trans. Netw.* **21**, 1609 (2013).
 - [12] P. Van Mieghem, *Performance Analysis of Complex Networks and Systems* (Cambridge University Press, Cambridge, UK, 2014).
 - [13] M. A. Achterberg and P. Van Mieghem, Analytic solution of markovian epidemics without re-infections on heterogeneous networks, [arXiv:2311.16721](https://arxiv.org/abs/2311.16721).
 - [14] P. Van Mieghem, J. Omic, and R. Kooij, Virus spread in networks, *IEEE/ACM Trans. Networking* **17**, 1 (2008).
 - [15] P. Van Mieghem, *Graph Spectra for Complex Networks*, 2nd ed. (Cambridge University Press, Cambridge, UK, 2023).

- [16] S. A. Gershgorin, Uber die abgrenzung der eigenwerte einer matrix, *Bull. Acad. Sci. USSR Cl. Sci. Math. Nat.* **6**, 749 (1931).
- [17] C. Moler and C. Van Loan, Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later, *SIAM Rev.* **45**, 3 (2003).
- [18] N. J. Higham, The scaling and squaring method for the matrix exponential revisited, *SIAM J. Matrix Anal. Appl.* **26**, 1179 (2005).
- [19] A. H. Al-Mohy and N. J. Higham, A new scaling and squaring algorithm for the matrix exponential, *SIAM J. Matrix Anal. Appl.* **31**, 970 (2010).
- [20] A. H. Al-Mohy and N. J. Higham, Computing the action of the matrix exponential, with an application to exponential integrators, *SIAM J. Sci. Comput.* **33**, 488 (2011).
- [21] R. B. Sidje, Expokit: A software package for computing matrix exponentials, *ACM Trans. Math. Software* **24**, 130 (1998).
- [22] P. Van den Driessche and J. Watmough, Reproduction numbers and sub-threshold endemic equilibria for compartmental models of disease transmission, *Math. Biosci.* **180**, 29 (2002).
- [23] B. Prasse and P. Van Mieghem, Time-dependent solution of the NIMFA equations around the epidemic threshold, *J. Math. Biol.* **81**, 1299 (2020).
- [24] R. Persoons, M. Sensi, B. Prasse, and P. Van Mieghem, Transition from time-variant to static networks: Timescale separation in N -intertwined mean-field approximation of susceptible-infectious-susceptible epidemics, *Phys. Rev. E* **109**, 034308 (2024).
- [25] C. Vuik, F. J. Vermolen, M. B. van Gijzen, and M. J. Vuik, *Numerical Methods for Ordinary Differential Equations*, 2nd ed. (Delft Academic Press, Delft, The Netherlands, 2015).
- [26] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms* (MIT Press, Cambridge, MA, 2022).
- [27] P. Van Mieghem, Properties of a triangular matrix and its linear differential equation, Delft University of Technology, Technical Report No. 20230703 (2023).