

Analysis of the survival time of the SIRS process via expansion*

Tobias Friedrich[†]  Andreas Göbel[†]  Nicolas Klodt[†] 
Martin S. Krejca[‡]  Marcus Pappik[†] 

Abstract

We study the SIRS process—a continuous-time Markov chain modeling the spread of infections on graphs. In this model, vertices are either susceptible, infected, or recovered. Each infected vertex becomes recovered at rate 1 and infects each of its susceptible neighbors independently at rate λ , and each recovered vertex becomes susceptible at a rate ρ , which we assume to be independent of the graph size. A central quantity of the SIRS process is the time until no vertex is infected, known as the survival time. Surprisingly though, to the best of our knowledge, all known rigorous theoretical results that exist so far immediately carry over from the related SIS model and do not completely explain the behavior of the SIRS process. We address this imbalance by conducting theoretical analyses of the SIRS process via the expansion properties of the underlying graph.

Our first result shows that the expected survival time of the SIRS process on stars is at most polynomial in the graph size for any value of λ . This behavior is fundamentally different from the SIS process, where the expected survival time is exponential already for small infection rates. This raises the question of which graph properties result in an exponential survival time. Our main result is an exponential lower bound of the expected survival time of the SIRS process on expander graphs. Specifically, we show that on expander graphs G with n vertices, degree close to d , and sufficiently small spectral expansion, the SIRS process has expected survival time at least exponential in n when $\lambda \geq c/d$ for a constant $c > 1$. Previous results on the SIS process show that this bound is almost tight. Additionally, our result holds even if G is a subgraph. Notably, our result implies an almost-tight threshold for Erdős–Rényi graphs and a regime of exponential survival time for complex network models. The proof of our main result draws inspiration from Lyapunov functions used in mean-field theory to devise a two-dimensional potential function and from applying a negative-drift theorem to show that the expected survival time is exponential.

*A preliminary version of this paper appeared as [15].

[†]Hasso Plattner Institute, University of Potsdam, Potsdam, Germany. E-mail: tobias.friedrich@hpi.de, andreas.goebel@hpi.de, nicolas.klodt@hpi.de, marcus.pappik@hpi.de

[‡]LIX, CNRS, Ecole Polytechnique, Institut Polytechnique de Paris, Palaiseau, France. E-mail: martin.krejca@polytechnique.edu

Keywords: epidemic models; contact process; expansion; survival time.

MSC2020 subject classifications: 60J60.

Submitted to EJP on January 30, 2024, final version accepted on May 5, 2024.

Supersedes arXiv:2205.02653.

1 Introduction

In the domain of modeling infectious diseases, a vast body of literature studying various stochastic processes on graphs exists (see, for example, the extensive survey by Pastor-Satorras, Castellano, Mieghem,²⁶ and Vespignani [31]). In this article, we focus on the SIRS process—a continuous-time Markov chain where each vertex is either susceptible, infected, or recovered. Each infected vertex becomes recovered at a normalized rate of 1 and infects each of its susceptible neighbors independently at an *infection rate* λ , while each recovered vertex becomes susceptible at a *deimmunization rate* ρ .

A question central to understanding the SIRS process is *how long* it takes until no vertex in the graph is infected, known as the *survival time*¹ of the process. Due to the relevance of the SIRS process, its survival time has been studied extensively. This includes empirical results [36, 24, 14], mean-field approaches [2, 33], and results that consider deterministic variants of the process [34] or generalized models [33]. However, surprisingly, to the best of our knowledge, all known rigorous theoretical results that exist so far immediately carry over from the related SIS model and do not completely explain the behavior of the SIRS process.

This lack of non-trivial theoretical results for the SIRS process stands in stark contrast to the plethora of theoretical results for a similar but slightly simpler process, known as the SIS process or contact process. In the SIS process, each vertex is either susceptible or infected. Each infected vertex becomes susceptible at rate 1 and infects each of its neighbors independently at an *infection rate* λ . Thus, with a grain of salt, the SIS process can be viewed as a special case of the SIRS process in which recovered vertices turn immediately susceptible (that is, the deimmunization rate ρ is infinite). The survival time of the SIS process is well-understood on a variety of graphs. Early results on the SIS process consider its survival time on \mathbb{Z}^d [19] and on infinite d -regular trees [32, 25, 35], while recent breakthroughs characterize the survival time on Galton–Watson trees [20, 5, 28]. On finite structures, the results of Nam, Nguyen,⁴⁶ and Sly [28] consider Erdős–Rényi graphs, while the SIS process has also been studied on scale-free graphs² [4, 7]. These results rely on the survival time on simple subgraphs, such as stars. Further, Ganesh, Massoulié, and Towsley [18] connect the survival time to the spectral radius and the isoperimetric constant of the host graph, which immediately translates to a variety of simple graphs.

The vast amount of rigorous results for the SIS process allows to carry over some, albeit limited, insights to the SIRS process. Most prominently, for the same graph, the survival time—which is a random variable—of a SIS process is an upper bound for the survival time of a SIRS process when starting with identical configurations, as the two processes can be coupled such that an infected vertex in the latter is also always infected in the former. However, our knowledge about the SIRS process remains in a very unsatisfactory state for multiple reasons. First, we only have upper bounds on the survival time for the SIRS process, which begs the question how tight they are. And second, more importantly, the survival time in the SIS process for a graph G is a lower bound for any graph H containing G as a subgraph, as adding more vertices does not

¹Sometimes also referred to as the *extinction time*.

²Generated by the preferential-attachment model [3].

reduce the number of infected vertices at any point in time. In contrast, it is not known whether the SIRS process also has this property. Adding more vertices to a graph in the SIRS process can lead to some vertices being earlier infected and thus potentially earlier recovered, which in turn can block an infection that would have occurred otherwise. Thus, it is not straightforward to generalize results for the SIRS process to supergraphs.

Our contribution. We conduct the first rigorous, theoretical study of the expected survival time of the SIRS process on a large variety of graph classes, most prominently expanders. In all of our results, we assume that the deimmunization rate is independent of the graph size and that the process starts with at least one infected vertex and no recovered vertices. Our results showcase the similarities and the differences between the SIS and the SIRS process, highlighting the impact of the state *recovered*. Furthermore, for our lower bounds, we prove that our results carry over to supergraphs of the graphs we analyze. This makes our results applicable to a great number of different graph classes.

More specifically, in Section 3, we show that the expected survival time of the SIRS process on stars is polynomial (in the number of vertices), regardless of the infection rate (Theorem 1.1). This strongly contrasts the SIS process, where the survival time is superpolynomial for already very small infection rates [18, Theorem 5.2]. This shows that recovered vertices can have a huge impact on the survival time. The reason for this drastic difference in the expected survival time between both processes is that the star is only connected through a single, central vertex. Thus, if the center is recovered, the infection only survives if not all leaves become recovered during this time interval. For the SIRS process, the latter event does not have sufficiently high probability of occurring for the infection to survive superpolynomially long.

In Section 4, we complement these findings by proving that the expected survival time of the SIRS process on expanders is at least exponential if the infection rate is greater than the inverse of the expander's average degree (Theorem 1.2). This result holds very similarly for the SIS process [18]. In contrast to stars, expanders have many edges between arbitrary subsets of vertices. Thus, if the number of infected vertices is sufficiently high, there exist enough edges between all susceptible and all infected vertices, regardless of the number of (remaining) recovered vertices. These edges give the process a high probability not to decrease the number of infected vertices, which leads to the overall long expected survival time.

Since we prove our result for expanders to carry over to supergraphs, this result implies respective expected survival times for other well-known graph classes, such as Erdős–Rényi graphs (Corollary 1.5) and complex networks exhibiting real-world properties (see, e.g., Corollary 1.6 for hyperbolic random graphs), which we discuss in Section 5. Combined, our results emphasize that while the SIRS and SIS process behave very differently on some of their subgraphs (namely stars), they have similar behavior if the graph is sufficiently connected. In the following, we discuss our results in more detail.

1.1 Expected survival time on stars

For stars, we prove the following upper bound on the expected survival time of the SIRS process.

Theorem 1.1. *Let G be a star with $n \in \mathbb{N}_{>0}$ leaves, and let C be a SIRS process on G with infection rate λ and with deimmunization rate ϱ . Let T be the survival time of C . Then for sufficiently large n , it holds that $\mathbb{E}[T] \leq (\ln(n) + 2)(4n^\varrho + 1) \in O(n^\varrho \ln(n))$.*

Note that this bound is independent of λ and that it results in a polynomial expected

survival time as long as ϱ is at most constant with respect to n . Although we only prove an upper bound, our bound matches, up to a logarithmic factor, empirical investigations of the star [14], suggesting that our bound is almost tight. Note that these experimental results consider the infection rate λ to be constant in terms of n , while our results apply for any λ . Our results also show a behavior similar to the deterministic variant of the process considered by Saif [34].

The analysis mainly relies on the method of investigating independent phases in which the center is not infected, bounding the probability of the infection process dying out during that time, as is common [7, 4]. A phase lasts at most until all leaves triggered their recovery at least once, which occurs in expectation after a time of about $\ln(n)$. Thus, if the center just recovered, it needs to become susceptible more quickly than that bound, as otherwise all leaves are recovered. Since deimmunization triggers at rate ϱ , the probability that the center does not become susceptible in this time interval is about $e^{-\varrho \ln n}$, resulting in a probability of about $n^{-\varrho}$ that the infection dies out. Since these phases are independent, the infection process survives, in expectation, about n^ϱ of these trials, each lasting about $\ln(n)$ time in expectation. By Markov's inequality, this bound on the survival time also holds with high probability.

Note that the deimmunization rate and the state *recovered* are important for this argument to hold. Without this additional state, that is, in the SIS process, it is quite likely that the center becomes quickly reinfected before all leaves are not infected, which leads to an exponential expected survival time once $\lambda \geq n^{-1/2+\varepsilon}$ in this setting [18], for all positive constants ε .

1.2 Expected survival time on expanders

Before we state our main result, we formally introduce the notion of expansion we use for our results. To this end, let $G = (V, E)$ be a graph with n vertices $\{v_i\}_{i=1}^n$, and let L be its normalized Laplacian, which is for all $i, j \in [n]$ defined as

$$L_{i,j} = \begin{cases} 1 & \text{if } i = j, \\ -\frac{1}{\sqrt{\deg(v_i)\deg(v_j)}} & \text{if there is an edge between } v_i \text{ and } v_j, \\ 0 & \text{otherwise.} \end{cases} \quad (1.1)$$

Let L have eigenvalues $\lambda_1 \leq \dots \leq \lambda_n$. The *spectral expansion* of L is defined as $\delta = \max_{i \geq 2} |1 - \lambda_i|$. We call G an $(n, (1 \pm \varepsilon_d)d, \delta)$ -expander if and only if it has n vertices, a spectral expansion of δ , and only vertices with degree between $(1 - \varepsilon_d)d$ and $(1 + \varepsilon_d)d$.

As noted above, in contrast to stars, expanders feature many edges between arbitrary subsets of vertices. The key property we require for our results from $(n, (1 \pm \varepsilon_d)d, \delta)$ -expanders is that the number of edges between any two sets X and Y of vertices is close to $\frac{d}{n}|X||Y|$ (see Theorems 2.6 and 2.7).

Our results hold for any expander G' that is a subgraph of a graph G that hosts a SIRS process C . In order to derive such a result, we define the *projection* C' of C onto G' to be the process on G' such that, at each point in time, each vertex of G' in C' is in the same state as it is in C . The survival time of a projected process is the first point in time that the projected process has no infected vertices. Given these definitions, our main result follows.

Theorem 1.2. *Let G be a graph, and let G' be a subgraph of G that is an $(n, (1 \pm \varepsilon_d)d, \delta)$ -expander. Let $d \rightarrow \infty$ and $\delta, \varepsilon_d \rightarrow 0$ as $n \rightarrow \infty$. Let C be the SIRS process on G with infection rate λ and with constant deimmunization rate ϱ . Further, let C start with at least one infected vertex in G' and no recovered vertices in G' . Last, let C' be the projection of C onto G' , and let T be the survival time of C' . If $\lambda \geq \frac{c}{d}$ for a constant $c \in \mathbb{R}_{>1}$, then for sufficiently large n , it holds that $\mathbb{E}[T] = 2^{\Omega(n)}$.*

We note that Theorem 1.2 is almost tight with respect to the range of λ . Ganesh, Massoulié, and Towsley [18, Theorem 3.1] show that the survival time of the SIS process is at most logarithmic in n when the spectral radius of a graph is less than $1/\lambda$. Note that the spectral radius of a graph is bounded from above by the maximum degree of the graph. This results in a logarithmic expected survival time of the process on $(n, (1 \pm \varepsilon_d)d, \delta)$ -expanders when $\lambda \leq \frac{1-\varepsilon}{d}$, for some constant ε . Recall our discussion earlier in the introduction that the expected survival time of the SIS process is an upper bound of the expected survival time of the SIRS process. Hence, the expected survival time of the SIRS process for $\lambda \leq \frac{1-\varepsilon}{d}$ is at most logarithmic in n on $(n, (1 \pm \varepsilon_d)d, \delta)$ -expanders.

The proof of Theorem 1.2 consists of two main parts. First, we prove that a linear number of vertices in G' becomes infected. Then, we show that the number of infected vertices stays linear for an expected exponential amount of time. For both parts, we make use of potential functions, which map the configuration of the process to a single real number that allows us to quantify how likely the process is to die out. In order to get the result on the projection of the process, we use that the influence of $G \setminus G'$ only increases the rate at which vertices in G' get infected. In the considered configurations, this rate increase only helps the process get into the desired region of the potential.

First part: reaching a linear number of infected vertices

For the first part, our key lemma shows that the process reaches a configuration with at least εn infected vertices with probability at least $\frac{1}{n+2}$. To this end, let I_{τ_t} be the number of infected vertices after the t -th change of the configuration of the process.

Lemma 1.3. Let G be a graph, and let G' be a subgraph of G that is an $(n, (1 \pm \varepsilon_d)d, \delta)$ -expander. Let $d \rightarrow \infty$ and $\delta, \varepsilon_d \rightarrow 0$ as $n \rightarrow \infty$. Let C be a SIRS process on G with infection rate λ and with constant deimmunization rate ρ . Further, let C start with at least one infected vertex in G' and no recovered vertices in G' . Consider the projection C' of C onto G' . If $\lambda \geq \frac{c}{d}$ for a constant $c \in \mathbb{R}_{>1}$, then there exists a constant $\varepsilon \in \mathbb{R}_{>0}$ such that for sufficiently large n , the probability that there exists a time step $t \in \mathbb{N}$ with $I_{\tau_t} \geq \varepsilon n$ is at least $\frac{1}{n+2}$.

Note that if this event does not occur, then the infection might die out fast. As the probability of the infection dying out in the first step is roughly $1/2$, the event of Lemma 1.3 does not have a high enough probability to give us super-polynomial survival time with high probability. To obtain the probabilistic lower bound of Lemma 1.3, we use a fairly simple potential H_t expressing the number of infected vertices minus ε times the recovered vertices. We show that H_t is a submartingale and then apply the optional-stopping theorem to H_t to conclude the proof of Lemma 1.3.

Second part: retaining a linear number of infected vertices for exponential time

For showing that the infection survives exponentially long once at least εn vertices have been infected, we define a more involved potential function F_t than before, which increases when the number of infected vertices reduces. Our definition of F_t is based on a Lyapunov function f used by Korobeinikov and Wake [22], which they utilize in order to derive results on the global stability of the SIRS process via mean-field theory. We briefly overview this approach before we explain how we adjust it to our setting. To this end, let S_{τ_t} and I_{τ_t} denote the number of susceptible and of infected vertices, respectively, of the t -th change of the configuration of the process.

Korobeinikov and Wake [22] assume a fully mixed graph, which roughly corresponds to a clique for our process. In order to show global stability, the authors show a negative drift towards an equilibrium configuration with I^* infected and S^* susceptible vertices.

To this end, they use an auxiliary function $f: \mathbb{R}_{>0}^2 \rightarrow \mathbb{R}$ that satisfies for all $x, x^* \in \mathbb{R}_{>0}$ that $f(x^*, x) = x^* \left(\frac{x}{x^*} - \ln \frac{x}{x^*} - 1 \right)$. For a fixed x^* , the function has a global minimum at x^* and a derivative of $1 - \frac{x^*}{x}$, which is important for calculating the drift. They then define a Lyapunov function $F'(P_{\tau_t}, I_{\tau_t}) = f(P^*, P_{\tau_t}) + f(I^*, I_{\tau_t})$, where $P_{\tau_t} = S_{\tau_t} + \frac{\rho}{\lambda}$ and $P^* = S^* + \frac{\rho}{\lambda}$. Note that they use P_{τ_t} instead of S_{τ_t} in order for the drift not to be too large when S_{τ_t} is small. This function results in non-positive drift everywhere, which is enough for the setting of Korobeinikov and Wake [22].

The potential function of Korobeinikov and Wake [22] is not sufficient for our purposes, as its resulting drift is 0 for some configurations, whereas we require a constant negative drift in order to derive a rigorous lower bound for the expected survival time. Hence, we adjust the potential function of Korobeinikov and Wake [22] such that it creates a region in the potential that has a sufficiently large negative drift. We note that we do not need negative drift everywhere but only in configurations with less than εn infected vertices. We achieve this by changing the target of susceptible vertices from the equilibrium point to n . Further, we use a slightly different shift in our setting to adjust for the base graphs being expanders instead of cliques. Letting $n' = n + P^* - S^*$, we define the potential

$$F_t = F(P_{\tau_t}, I_{\tau_t}) = f(n', P_{\tau_t}) + f(I^*, I_{\tau_t}).$$

By the definition of f , the potential F_t has a global minimum for $n' = P_{\tau_t}$ and $I^* = I_{\tau_t}$, which roughly models the idealized (and impossible to reach) configuration of all vertices being susceptible while the number of infected vertices is as in the equilibrium.

For this new potential, we show that there is a region in which the process is a strict supermartingale with a constant negative drift. This is formally stated in the following lemma.

Lemma 1.4. Let G be a graph, and let G' be a subgraph of G that is an $(n, (1 \pm \varepsilon_d)d, \delta)$ -expander. Let $d \rightarrow \infty$ and $\delta, \varepsilon_d \rightarrow 0$ as $n \rightarrow \infty$. Let C be a SIRS process on G with infection rate λ and with constant deimmunization rate ρ . Consider the projection C' of C onto G' . Let $t \in \mathbb{N}$ and $\varepsilon_0, \varepsilon \in (0, 1)$ be sufficiently small constants. Assume that $\varepsilon_0 n \geq I_{\tau_t} \geq \varepsilon n$. If $\lambda \geq \frac{c}{d}$ for a constant $c \in \mathbb{R}_{>1}$, then there exists a constant $a \in \mathbb{R}_{>0}$ such that $D_t \leq -a$ for sufficiently large n .

We use the expansion properties of the base graph that guarantee that the infected vertices always have enough susceptible neighbors such that new vertices get infected and the potential decreases in expectation. This allows us to apply a concentration bound by Oliveto and Witt [29] (Theorem 2.2) for strict supermartingales, known as *negative-drift theorem*. The negative-drift theorem results in the lower exponential bound of the expected survival time.

1.3 Applicability of our main result

The generality of Theorem 1.2 makes it applicable to a broad range of graph classes, as the only requirement is for the base graph to contain a large expander as a subgraph. We illustrate this generality on the classical model of Erdős–Rényi graphs as well as on popular complex network models such as hyperbolic random graphs.

Erdős–Rényi graphs

The first random-graph model we are interested in is $G_{n,p}$ —the classical random-graph model of Erdős and Rényi [12]. The expansion properties of this model have been previously studied in literature. As Coja-Oghlan [10, Theorem 1.2] shows, Erdős–Rényi graphs have a very small spectral expansion. Furthermore, due to Chernoff bounds, the vertex degrees in Erdős–Rényi graphs are tightly distributed around their average

degree d . Therefore, Erdős–Rényi graphs fulfill, with high probability, our definition of an $(n, (1 \pm \varepsilon_d)d, \delta)$ -expander. Together with Theorem 1.2 and the upper bound that carries over from the SIS model (see Ganesh, Massoulié, and Towsley [18, Theorem 5.5]), this leads to the following corollary.

Corollary 1.5. *Let $G \sim G_{n,p}$ be an Erdős–Rényi graph with $(n-1)p \in \omega(\ln n)$. Consider the SIRS process C on G with constant deimmunization rate ϱ , and let T be the survival time of C when the process starts with at least one infected vertex. If $\lambda \geq \frac{c}{d}$ for a constant $c \in \mathbb{R}_{>1}$, then $\mathbb{E}[T] = 2^{\Omega(n)}$ asymptotically almost surely with respect to G . If $\lambda \leq \frac{c}{d}$ for a constant $c \in (0, 1)$, then $\mathbb{E}[T] \in O(\log n)$ asymptotically almost surely with respect to G .*

Comparing Corollary 1.5 with the respective result for the SIS process [18, Theorem 5.5] shows that the two processes, SIS and SIRS, behave similarly on Erdős–Rényi graphs.

Complex networks

A variety of random-graph models that exhibit properties found in real-world networks has appeared in network science [6]. Such network models provide a highly relevant structure for studying the survival time of the SIRS process. We focus our attention on three such complex network models that exhibit key properties required for applying Theorem 1.2. These models are Chung–Lu graphs [1], hyperbolic random graphs [23], and geometric inhomogeneous random graphs [21].

A common characteristic shared by these three network models is that the degrees of the vertices follow a power-law distribution. The exponent of the power-law is controlled by a parameter γ . The interesting parameter range is $\gamma \in (2, 3)$, as beyond this range, these models lose key properties present in real-world networks. When $\gamma \in (2, 3)$, two key properties hold on all three models, namely, a polynomially-sized clique as a subgraph and a polylogarithmic upper bound on the network’s diameter. See Chung and Lu [9] for the proofs for Chung–Lu graphs, Friedrich and Krohmer [17, 16] for hyperbolic random graphs, and Keusch [21] for geometric inhomogeneous random graphs.

As cliques fulfill the definition of expander required by Theorem 1.2, the aforementioned results ensure that the expander subgraph is large enough for the SIRS process to survive superpolynomially long on such graph models. Additionally, the polylogarithmic bound on the diameter of the graphs generated by these models suffices to show that the infection reaches a vertex of the expander subgraph with sufficiently high probability for the expected survival time to be superpolynomial. Following this line of argumentation on hyperbolic random graphs for example, we arrive at the following result.

Corollary 1.6. *Let G be a hyperbolic random graph with n vertices that follows a power-law degree distribution with exponent $\gamma \in (2, 3)$, and let C be the SIRS process on G with infection rate λ and with constant deimmunization rate ϱ . Further, let C start with at least one infected vertex in the giant component and no recovered vertices, and let T be the survival time of C . Then there exists a constant $c \in \mathbb{R}_{>0}$ such that if $\lambda \geq cn^{(\gamma-3)/2}$, then $\mathbb{E}[T] = 2^{\Omega(n^{(3-\gamma)/2})}$.*

Similar statements can be shown for the other two models, i.e., Chung–Lu graphs and geometric inhomogeneous random graphs.

1.4 Outlook

Although our results cover already a great range of interesting graph classes, this article is just the first step to understanding the SIRS process more thoroughly. Our

analyses pose exciting new challenges for different scenarios, which we briefly delineate in the following.

Our upper bound of the expected survival time on stars (Theorem 1.1) is off from empirical results [14] by a logarithmic factor. This shows that there is potential for improvement in the analysis. Ideally, proving a matching lower bound would answer the question for the exact expected survival time.

Combined, our results for stars (Theorem 1.1) and expanders (Theorem 1.2) show that adding edges to a graph leads, eventually, from a polynomial expected survival time to an exponential one. However, it is not clear so far when this transition happens. An interesting next step is to look into connected stars instead of single stars. Connected stars appear as subgraphs in important real-world network models, most prominently, the preferential-attachment model [3], but also in Chung–Lu graphs [9], for which our initial results could be improved, motivating this research question.

With respect to expanders with vertex degrees concentrated around d , our result (Theorem 1.2) implies that $1/d$ is the threshold for the infection rate λ at which the expected survival time transitions from logarithmic to exponential. However, our bounds require λ to be bounded away from $1/d$ by a constant. It is not clear, given a sequence of values $\varepsilon_n \in o(1)$, what happens if $\lambda = \frac{1 \pm \varepsilon_n}{d}$. A more detailed analysis could provide insights into how rapidly the transition at the threshold occurs.

A different extension of our results is to consider deimmunization rates that are dependent on the graph size. Comparing the behavior of the SIS and the SIRS process on stars suggests that an increased deimmunization rate leads to far longer expected survival times. Thus, an interesting question is whether the survival time exhibits a threshold behavior with respect to the deimmunization rate.

Multi-dimensional potentials, as the one we use for the SIRS process on expanders, are rare in the analysis of stopping times of stochastic processes. Our approach draws inspiration from Lyapunov stability to devise a potential function for the stochastic process under study and then applies drift theory to convert this into a rigorous proof. Lyapunov functions are used in mean-field theory to show stable points of dynamical systems [26], and epidemic processes constitute only a glimpse of their applicability. We believe that our approach might inspire further rigorous results of determining stopping times of other stochastic processes, not limited to epidemic models.

2 Preliminaries

We study the SIRS process, which is a continuous-time Markov chain on graphs in which the vertices change between different states, following events triggered by Poisson processes. We analyze how this process behaves asymptotically in the number of vertices n of the graph. Especially, when we use big-O notation or refer to variables as constants, this is with respect to n . When we use big-O notation inside of a term in a relation, this means that there exists a function from the big-O expression such that the relation holds, for example, the equation $a = 2^{\Omega(n)}$ means that there exists a function $f \in \Omega(n)$ such that $a = 2^{f(n)}$ holds. If not stated otherwise, all variables we consider may depend on n . Whenever we talk about Poisson processes, we refer to one-dimensional Poisson point processes that output a random subset of the non-negative real numbers.

We first define our infection models and some related terms that we use throughout the paper. We then state the probabilistic tools we use in the proofs.

2.1 Infection processes

Let $G = (V, E)$ be a finite graph with vertex set V and edge set E . Further, let $\lambda, \varrho \in \mathbb{R}_{>0}$. In the SIRS process, for each edge $e \in E$, we define a Poisson process M_e

with parameter λ , and for each vertex $v \in V$, we define the two Poisson processes N_v with parameter 1 and O_v with parameter ρ . We refer to these processes as *clocks*, and when an event occurs in one of them, we say that the relevant clock *triggers*. We use Z to denote the set of all of these clocks, that is, $Z = (\bigcup_{e \in E} \{M_e\}) \cup (\bigcup_{v \in V} \{N_v, O_v\})$. Let P denote the stochastic process in which all of the clocks in Z evolve simultaneously and independently, starting at time 0. Note that almost surely there is no time point at which two clocks trigger at once. There are almost surely a countably infinite number of trigger times in P , which we index by the increasing sequence $\{\gamma_i\}_{i \in \mathbb{N}_{>0}}$, where $\gamma_0 = 0$.

A SIRS process $C = (C_t)_{t \in \mathbb{R}_{\geq 0}}$ has an underlying graph $G = (V, E)$, an infection rate λ , a deimmunization rate ρ , and an initial partition of V into susceptible, infected, and recovered vertices with the respective sets S'_0, I'_0 , and R'_0 . Note that we do not need to specify a healing rate as we normalized that to 1. At every time $t \in \mathbb{R}_{\geq 0}$, the configuration C_t is a partition of V into S'_t, I'_t , and R'_t . The configuration only changes at times in P . Let $i \in \mathbb{N}_{>0}$. We consider the following configuration transitions in γ_i :

- If for some $e = \{u, v\} \in E$ we have $\gamma_i \in M_e$, $u \in I'_{\gamma_{i-1}}$, and $v \in S'_{\gamma_{i-1}}$, then $S'_{\gamma_i} = S'_{\gamma_{i-1}} \setminus \{v\}$, $I'_{\gamma_i} = I'_{\gamma_{i-1}} \cup \{v\}$, and $R'_{\gamma_i} = R'_{\gamma_{i-1}}$. We say that v *gets infected* at time point γ_i by u .
- If for some $v \in V$ we have $\gamma_i \in N_v$ and $v \in I'_{\gamma_{i-1}}$ then $S'_{\gamma_i} = S'_{\gamma_{i-1}}$, $I'_{\gamma_i} = I'_{\gamma_{i-1}} \setminus \{v\}$ and $R'_{\gamma_i} = R'_{\gamma_{i-1}} \cup \{v\}$. We say that v *recovers* at time point γ_i .
- If for some $v \in V$ we have $\gamma_i \in O_v$ and $v \in R'_{\gamma_{i-1}}$, then $S'_{\gamma_i} = S'_{\gamma_{i-1}} \cup \{v\}$, $I'_{\gamma_i} = I'_{\gamma_{i-1}}$ and $R'_{\gamma_i} = R'_{\gamma_{i-1}} \setminus \{v\}$. We say that v *gets susceptible* at time point γ_i .

If none of the above three cases occurs, the configuration of C at γ_i is the same as the configuration of C at γ_{i-1} . Note that at all times between γ_{i-1} and γ_i , C retains the same configuration as in γ_{i-1} .

In our proofs, we only consider the time points in P at which the configuration changes. To this end, let $P' = \{\gamma_0\} \cup \{\gamma_i \mid i \in \mathbb{N}_{>0} \wedge C_{\gamma_i} \neq C_{\gamma_{i-1}}\}$. We index the times in P' by the increasing sequence $\{\tau_i\}_{i \in \mathbb{N}}$. For all $i \in \mathbb{N}$, we call τ_i the i -th *step* of the process.

If at any point in time no vertex is infected, then from that point onward, no vertex is infected. We say that the infection *dies out* or *goes extinct* at the first (random) time T with $I'_T = \emptyset$. We call T the *survival time* of the SIRS process.

We only keep track of the number of vertices in each of the sets. To this end, we define for all $t \in \mathbb{R}_{\geq 0}$ the random variables $S_t = |S'_t|$, $I_t = |I'_t|$, and $R_t = |R'_t|$. These random variables change depending on the clocks in P . We say that an event *happens* at a rate of $r \in \mathbb{R}_{>0}$ if and only if the set of clocks that cause this event when they trigger has a sum of rates equal to r .

We define the *projection* C' of C onto G' as the process on G' such that, at each point in time, each vertex of G' in C' is in the same state as it is in C . When considering such a projection, we use S_t, I_t , and R_t to only count the vertices of C' in the corresponding state. Also $\{\tau_i\}_{i \in \mathbb{N}}$ only contains times at which the state of a vertex in C' changes. The survival time of a projected process is the first point in time that the projected process has no infected vertices. Note that the survival time T' of C' is a lower bound for the survival time T of C , as all infected vertices of C' are also infected in C .

We use stochastic domination to transfer results from one random variable to another. We say that a random variable $(X_t)_{t \in \mathbb{R}}$ *dominates* another random variable $(Y_t)_{t \in \mathbb{R}}$ if and only if there exists a coupling $(X'_t, Y'_t)_{t \in \mathbb{R}}$ in a way such that for all $t \in \mathbb{R}_{\geq 0}$ we have $X'_t \geq Y'_t$.

2.2 Probabilistic tools

We use general concepts from probability theory (see for example [13, 27]). In addition, we use the following theorems.

We use the optional-stopping theorem for submartingales to bound the probability of reaching a specific configuration. For an event E , the symbol $\mathbb{1}_E$ denotes the indicator random variable that is 1 if E is true and 0 otherwise.

Theorem 2.1 (Optional stopping [27, Theorem 13.2]). *Let $(X_t)_{t \in \mathbb{N}}$ be a submartingale and T a stopping time, both with respect to a filtration $(\mathcal{F}_t)_{t \in \mathbb{N}}$. Assume that the following two conditions hold:*

1. $\mathbb{E}[T] < \infty$.
2. *There is a $c \in \mathbb{R}$ such that for all $t \in \mathbb{N}$ we have $\mathbb{E}[|X_{t+1} - X_t| \mid \mathcal{F}_t] \cdot \mathbb{1}_{t < T} \leq c \cdot \mathbb{1}_{t < T}$.*

Then $\mathbb{E}[X_T] \geq \mathbb{E}[X_0]$.

We use the following theorem in order to show an exponential expected survival time for the SIRS process. We state it in a fashion that better suits our purposes.

Theorem 2.2 (Negative drift [29, Theorem 4] [30]). *Let $(X_t)_{t \in \mathbb{N}}$ be a random process over \mathbb{R} , adapted to a filtration $(\mathcal{F}_t)_{t \in \mathbb{N}}$. Let there be an interval $[a, b] \subseteq \mathbb{R}$, two constants $\delta, \varepsilon \in \mathbb{R}_{>0}$ and, possibly depending on $l := b - a$, a function $r(l)$ satisfying $1 \leq r(l) \in o(l/\log(l))$. Let $T = \inf\{t \in \mathbb{N} \mid X_t \geq b\}$. Suppose that for all $t \in \mathbb{N}$ the following two conditions hold:*

1. $\mathbb{E}[X_{t+1} - X_t \mid \mathcal{F}_t] \cdot \mathbb{1}_{a < X_t < b} \leq -\varepsilon \cdot \mathbb{1}_{a < X_t < b}$.
2. *For all $j \in \mathbb{R}_{\geq 0}$ we have $\Pr[|X_{t+1} - X_t| \geq j \mid \mathcal{F}_t] \cdot \mathbb{1}_{t < T} \leq \frac{r(l)}{(1+\delta)^j} \cdot \mathbb{1}_{t < T}$.*

Then there exists a constant $c \in \mathbb{R}_{>0}$ such that

$$\Pr[T \leq 2^{cl/r(l)} \mid \mathcal{F}_0] \cdot \mathbb{1}_{X_0 \leq a} = 2^{-\Omega(l/r(l))} \cdot \mathbb{1}_{X_0 \leq a}.$$

The following theorem bounds the expected value of the maximum of n exponentially distributed random variables.

Theorem 2.3 ([27, Lemma 2.10]). *Let $n \in \mathbb{N}_{>0}$, and let $\{X_i\}_{i \in [n]}$ be independent random variables that are each exponentially distributed with parameter $\lambda \in \mathbb{R}_{>0}$. Let $Y = \max_{i \in [n]} X_i$, and let H_n be the n -th harmonic number. Then*

$$\mathbb{E}[Y] = \frac{H_n}{\lambda} < \frac{1 + \ln(n)}{\lambda}.$$

We use the following version of Wald's equation, which does not require the addends to be independent.

Theorem 2.4 (Generalized Wald's equation [11, Theorem 5]). *Let $c, c' \in \mathbb{R}$, and let $(X_t)_{t \in \mathbb{N}}$ be a random process over $\mathbb{R}_{\geq c}$ such that $\sum_{i \in [S]} X_i$ has a finite expectation. Furthermore, let $(\mathcal{F}_t)_{t \in \mathbb{N}}$ be a filtration, and let S be a stopping time with respect to $(\mathcal{F}_t)_{t \in \mathbb{N}}$. If for all $i \in \mathbb{N}$, it holds that $\mathbb{E}[X_{i+1} \mid \mathcal{F}_i] \leq c'$, then*

$$\mathbb{E}\left[\sum_{i \in [S]} X_i \mid \mathcal{F}_0\right] = \mathbb{E}\left[\sum_{i \in [S]} \mathbb{E}[X_i \mid \mathcal{F}_{i-1}] \mid \mathcal{F}_0\right].$$

2.3 Expander graphs

There are many notions of how to define expander graphs. We use algebraic expanders in which all but one of the eigenvalues of the normalized Laplacian of the graph are very close to 1. These graphs have some nice properties that let us bound the number of edges between infected and susceptible vertices. Formally, let $G = (V, E)$

be a graph with n vertices $\{v_i\}_{i=1}^n$, and let L be its normalized Laplacian as defined in Equation (1.1). Let L have eigenvalues $\lambda_1 \leq \dots \leq \lambda_n$. The *spectral expansion* of L is defined as $\delta = \max_{i \geq 2} |1 - \lambda_i|$. We call G an $(n, (1 \pm \varepsilon_d)d, \delta)$ -expander if and only if it has n vertices, a spectral expansion of δ and only vertices with degree between $(1 - \varepsilon_d)d$ and $(1 + \varepsilon_d)d$.

For two vertex sets $X, Y \subseteq V$, let $E(X, Y)$ denote the number of edges between X and Y . For a vertex set X , let $\text{vol}(X)$ denote the sum of the vertex degrees of all vertices in X . Using this notation, we have the following theorem

Theorem 2.5 ([8, Theorem 5.2]). *Let $G = (V, E)$ be a graph with spectral expansion δ and let $X, Y \subseteq V$. Then*

$$\left| |E(X, Y)| - \frac{\text{vol}(X) \cdot \text{vol}(Y)}{\text{vol}(V)} \right| \leq \delta \cdot \frac{\sqrt{\text{vol}(X)\text{vol}(\bar{X})\text{vol}(Y)\text{vol}(\bar{Y})}}{\text{vol}(V)}.$$

Applying Theorem 2.5 to expanders, we get the following two corollaries..

Corollary 2.6. *Let $G = (V, E)$ be a $(n, (1 \pm \varepsilon_d)d, \delta)$ -expander, and let $X \subseteq V$. Then*

$$|E(X, \bar{X})| \geq (1 - \delta)(1 - 3\varepsilon_d)d \frac{|X| \cdot |\bar{X}|}{n}.$$

Proof. Because the vertex degrees of all vertices in G are bounded, we know that for each $S \subseteq V$ holds

$$(1 - \varepsilon_d)d|S| \leq \text{vol}(S) \leq (1 + \varepsilon_d)d|S|. \tag{2.1}$$

Plugging that into the result of Theorem 2.5 gives us

$$\begin{aligned} |E(X, \bar{X})| &\geq \frac{\text{vol}(X) \cdot \text{vol}(\bar{X})}{\text{vol}(V)} - \delta \cdot \frac{\sqrt{\text{vol}(X)\text{vol}(\bar{X})\text{vol}(\bar{X})\text{vol}(\bar{X})}}{\text{vol}(V)} \\ &= (1 - \delta) \frac{\text{vol}(X) \cdot \text{vol}(\bar{X})}{\text{vol}(V)} \\ &\geq (1 - \delta) \frac{(1 - \varepsilon_d)d|X| \cdot (1 - \varepsilon_d)d|\bar{X}|}{(1 + \varepsilon_d)dn} \\ &\geq (1 - \delta)(1 - 3\varepsilon_d)d \frac{|X| \cdot |\bar{X}|}{n}. \quad \square \end{aligned}$$

Corollary 2.7. *Let $G = (V, E)$ be a $(n, (1 \pm \varepsilon_d)d, \delta)$ -expander, and let $X, Y \subseteq V$. If $\varepsilon_d \leq 1/5$, then*

$$\left| |E(X, Y)| - d \frac{|X| \cdot |Y|}{n} \right| \leq 4\varepsilon_d d \frac{|X| \cdot |Y|}{n} + 2\delta d \sqrt{|X| \cdot |Y|}.$$

Proof. Because the vertex degrees of all vertices in G are bounded, we know that for each $S \subseteq V$, Equation (2.1) holds. Theorem 2.5 gives us both an upper and a lower bound for $|E(X, Y)| - \frac{\text{vol}(X) \cdot \text{vol}(Y)}{\text{vol}(V)}$. We solve them for $|E(X, Y)|$ and bound them separately using that $\varepsilon_d \leq 1/5$.

$$\begin{aligned} |E(X, Y)| &\geq \frac{\text{vol}(X) \cdot \text{vol}(Y)}{\text{vol}(V)} - \delta \cdot \frac{\sqrt{\text{vol}(X)\text{vol}(\bar{X})\text{vol}(Y)\text{vol}(\bar{Y})}}{\text{vol}(V)} \\ &\geq \frac{\text{vol}(X) \cdot \text{vol}(Y)}{\text{vol}(V)} - \delta \cdot \sqrt{\text{vol}(X)\text{vol}(Y)} \end{aligned}$$

$$\begin{aligned}
 &\geq \frac{(1 - \varepsilon_d)d|X| \cdot (1 - \varepsilon_d)d|Y|}{(1 + \varepsilon_d)dn} - \delta(1 + \varepsilon_d)d\sqrt{|X| \cdot |Y|} \\
 &\geq (1 - 4\varepsilon_d)d\frac{|X| \cdot |Y|}{n} - 2\delta d\sqrt{|X| \cdot |Y|} \text{ as well as} \\
 |E(X, Y)| &\leq \frac{\text{vol}(X) \cdot \text{vol}(Y)}{\text{vol}(V)} + \delta \cdot \frac{\sqrt{\text{vol}(X)\text{vol}(\bar{X})\text{vol}(Y)\text{vol}(\bar{Y})}}{\text{vol}(V)} \\
 &\leq \frac{\text{vol}(X) \cdot \text{vol}(Y)}{\text{vol}(V)} + \delta \cdot \sqrt{\text{vol}(X)\text{vol}(Y)} \\
 &\leq \frac{(1 + \varepsilon_d)d|X| \cdot (1 + \varepsilon_d)d|Y|}{(1 - \varepsilon_d)dn} + \delta(1 + \varepsilon_d)d\sqrt{|X| \cdot |Y|} \\
 &= \left(1 + \frac{3\varepsilon_d + \varepsilon_d^2}{1 - \varepsilon_d}\right)d\frac{|X| \cdot |Y|}{n} + \delta(1 + \varepsilon_d)d\sqrt{|X| \cdot |Y|} \\
 &\leq (1 + 4\varepsilon_d)d\frac{|X| \cdot |Y|}{n} + 2\delta d\sqrt{|X| \cdot |Y|}.
 \end{aligned}$$

Subtracting $d\frac{|X| \cdot |Y|}{n}$ from both inequalities and combining them proves the corollary. \square

3 SIRS on stars

We show that the expected survival time of the SIRS process on stars is bounded from above by a polynomial in the number of vertices that is independent of the infection rate (Theorem 1.1). To this end, we bound the number of times that the center gets infected and the time between two infections of the center. We use that while the center is not infected, no leaf gets infected. Hence, if all of the leaves recover before the center gets susceptible after it recovered, the infection dies out.

We first bound the expected time that it takes for all of the leaves to recover. We refer to each clock at a vertex whose rate is the recovery (of 1) rate as *recovery clock*.

Lemma 3.1. Let G be a star with $n \in \mathbb{N}_{>0}$ leaves, and let C be a SIRS process on G with infection rate λ and with deimmunization rate ϱ . Let T be the time that it takes for all recovery clocks of the leaves to trigger at least once. Then $\mathbb{E}[T] \leq \ln(n) + 1$.

Proof. The star has n leaves, which all have a clock that recovers them at a rate of 1. For each of the clocks, the time until the first trigger happens is exponentially distributed with parameter 1. Hence, T is calculated as the maximum of the n exponential distributions of the independent clocks. By Theorem 2.3, $\mathbb{E}[T] \leq \ln(n) + 1$. \square

We now use Lemma 3.1 to bound the time it takes from one infection of the center until it gets infected again or until the infection dies out.

Lemma 3.2. Let G be a star with $n \in \mathbb{N}_{>0}$ leaves, and let C be a SIRS process on G with infection rate λ and with deimmunization rate ϱ . Let $t_0 \in \mathbb{R}_{\geq 0}$ be a time at which the infection has not died out yet. Further, let $T \in \mathbb{R}_{\geq 0}$ be the first time after t_0 at which either the center gets infected after being susceptible or the infection dies out. Then $\mathbb{E}[T - t_0] \leq \ln(n) + 2$.

Proof. If the center starts infected, in order for either the center to get infected again after being susceptible or the infection to die out, the center has to recover first. Let $T' \in \mathbb{R}$ be the first time after t_0 at which the center recovers. As all vertices recover at a rate of 1, the random variable $T' - t_0$ is exponentially distributed with a parameter of 1.

Between T' and T , no leaf gets infected, as the center is not infected and all edges are incident to the center. Hence, when all recovery clocks of the leaves trigger in this time

interval at least once, the infection dies out. Therefore, the first point in time after T' at which all of these recovery clocks triggered at least once is an upper bound for T . By Lemma 3.1, the expected time for this last trigger to happen is at most $\ln(n) + 1$. That gives us

$$\begin{aligned} \mathbb{E}[T - t_0] &= \mathbb{E}[T - T' + T' - t_0] \\ &= \mathbb{E}[T - T'] + \mathbb{E}[T' - t_0] \\ &\leq \ln(n) + 2. \end{aligned} \quad \square$$

Next, we bound the probability from below that when starting with an infected center, the infection dies out before the center gets infected again. We use this later to get an upper bound on the number of times that the center gets infected in total.

Lemma 3.3. Let G be a star with $n \in \mathbb{N}_{>0}$ leaves, and let C be a SIRS process on G with infection rate λ and with deimmunization rate ϱ . Let $t_0 \in \mathbb{R}_{\geq 0}$ be a time at which the center is infected. Further, let E_0 be the event that the infection dies out after t_0 before the center gets infected again (after being recovered in between). Then for sufficiently large n , it holds that $\Pr[E_0] \geq \frac{1}{4}n^{-\varrho}$.

Proof. In order for either the center to get infected again after being susceptible or for the infection to die out, the center has to recover first. Let $t_1 \in \mathbb{R}$ be the first time after t_0 at which the center recovers. As long as the center is in the recovered state, no vertex gets infected, as all edges of the graph are incident to the center. If all leaves recover before the center gets susceptible, the infection dies out. In order to bound the probability of this event, we consider the first time $T \in \mathbb{R}$ after t_1 at which the center gets susceptible, and we also consider the first time $T' \in \mathbb{R}$ after t_1 at which all of the recovery clocks of the leaves trigger at least once in the interval $(t_1, T']$. In particular, we use that all leaves recover before the center gets susceptible if $T' - t_1 < \ln(n)$ and $T - t_1 \geq \ln(n)$.

Each vertex recovers after a time that is exponentially distributed with parameter 1. By definition of T' , it is the maximum of n exponentially distributed random variables. In order for $T' - t_1 < \ln(n)$, all of those random variables have to be smaller than $\ln(n)$. As all of them are independent, we get that, for sufficiently large n ,

$$\begin{aligned} \Pr[T' - t_1 < \ln(n)] &= \Pr[\text{Exp}(1) < \ln(n)]^n \\ &= \left(1 - e^{-1 \ln(n)}\right)^n \\ &= \left(1 - \frac{1}{n}\right)^n \\ &\geq \frac{1}{4}. \end{aligned}$$

All vertices lose their immunity at a rate of ϱ . Hence, $T - t_1$ is exponentially distributed with parameter ϱ . Using the exponential probability distribution, we get

$$\begin{aligned} \Pr[T - t_1 \geq \ln(n)] &= e^{-\varrho \ln(n)} \\ &= n^{-\varrho}. \end{aligned}$$

Now using the fact that the infection dies out when all leaves recover before the center gets susceptible and that $T - t_1$ and $T' - t_1$ are independent, we get

$$\begin{aligned} \Pr[E_0] &\geq \Pr[T' - t_1 < T - t_1] \\ &\geq \Pr[T' - t_1 < \ln(n) \wedge T - t_1 \geq \ln(n)] \end{aligned}$$

$$\begin{aligned}
 &= \Pr[T' - t_1 < \ln(n)] \cdot \Pr[T - t_1 \geq \ln(n)] \\
 &\geq \frac{1}{4}n^{-e}. \quad \square
 \end{aligned}$$

Using the previous bounds, we now derive an upper bound on the expected survival time of a SIRS process on a star.

Theorem 1.1. *Let G be a star with $n \in \mathbb{N}_{>0}$ leaves, and let C be a SIRS process on G with infection rate λ and with deimmunization rate ϱ . Let T be the survival time of C . Then for sufficiently large n , it holds that $\mathbb{E}[T] \leq (\ln(n) + 2)(4n^e + 1) \in O(n^e \ln(n))$.*

Proof. Let S be the random variable that counts the number of times that the center gets infected before the infection dies out. For all $i \in \mathbb{N}_{\leq S+1}$, let X_i be the i -th time at which either the center gets infected or the infection dies out (we define $X_0 = 0$). It then holds that $T = X_{S+1} = \sum_{i=0}^S X_{i+1} - X_i$. We aim to bound the expectation of this value using the generalized Wald's equation (Theorem 2.4).

Let $(\mathcal{F}_t)_{t \in \mathbb{R}_{\geq 0}}$ be the natural filtration of C . By Lemma 3.2, it holds for all $i \in \mathbb{N}_{\leq S}$ that $0 \leq \mathbb{E}[X_{i+1} - X_i \mid \mathcal{F}_{X_i}] \leq \ln(n) + 2$. Hence, the expectations of all of the summed random variables are bounded. By Lemma 3.3, for all $i \in \mathbb{N}_{\geq 1}$, the i -th infection of the center has a probability of at least $\frac{1}{4}n^{-e}$ to be the last one if there is an i -th infection of the center. Therefore, S is dominated by a geometrically distributed random variable $A \sim \text{Geom}(\frac{1}{4}n^{-e})$. Hence, $\sum_{i=0}^S X_{i+1} - X_i$ is integrable. By Theorem 2.4, we get

$$\begin{aligned}
 \mathbb{E}[T \mid \mathcal{F}_0] &= \mathbb{E}\left[\sum_{i=0}^S X_{i+1} - X_i \mid \mathcal{F}_0\right] \\
 &= \mathbb{E}\left[\sum_{i=0}^S \mathbb{E}[X_{i+1} - X_i \mid \mathcal{F}_{X_i}] \mid \mathcal{F}_0\right] \\
 &\leq \mathbb{E}\left[\sum_{i=0}^S \ln(n) + 2 \mid \mathcal{F}_0\right] \\
 &= (\ln(n) + 2)\mathbb{E}\left[\sum_{i=0}^S 1 \mid \mathcal{F}_0\right] \\
 &\leq (\ln(n) + 2)(4n^e + 1). \quad \square
 \end{aligned}$$

4 SIRS on expanders

We consider the SIRS process on graphs that have expanders as subgraphs. In particular, we show an exponential expected survival time for the projection of the SIRS process onto the expander when the deimmunization rate is constant and the infection rate is sufficiently high (Theorem 1.2). Note that the exponential expected survival time and the required infection rate depend only on the size and vertex degrees of the expander. In Section 4.1, we begin by analyzing basic properties of the process, such as the transition rates between all of the states.

In Section 4.2, we show that the expected survival time of the considered SIRS processes is exponential if $\lambda \geq \frac{c}{d}$ for a constant $c \in \mathbb{R}_{>1}$. We first prove that the process reaches a configuration with at least εn infected vertices with sufficiently high probability. We then provide a lower bound for the expected survival time starting at such a configuration. To this end, we define a potential over the configuration space that has in a specific region a constant negative drift away from the configuration with no infected vertices. We then translate this region into bounds for the potential, allowing us to apply the negative-drift theorem (Theorem 2.2) to get an exponential expected survival time.

4.1 The SIRS process

Let $G = (V, E)$ be a graph and let $G' = (V', E')$ be a subgraph of G that is an $(n, (1 \pm \varepsilon_d)d, \delta)$ -expander. Let C be a SIRS process with infection rate $\lambda \geq \frac{c}{d}$ for a constant $c \in \mathbb{R}_{>1}$ and deimmunization rate ϱ on G . Consider the projection C' of C onto G' . We define for all $t \in \mathbb{N}$ the random variable $P_{\tau_t} = S_{\tau_t} + \frac{\varrho}{c}n$. We use P_{τ_t} to define the potential later. Roughly, using P_{τ_t} instead of S_{τ_t} has the effect that changes of S_{τ_t} have a lower impact on the potential. Note that, at all times t , it holds that $S_{\tau_t} + I_{\tau_t} + R_{\tau_t} = n$, since every vertex of G is always in exactly one of these three sets. Additionally, $P_{\tau_t} + I_{\tau_t} + R_{\tau_t} = n + \frac{\varrho}{c}n = n'$.

For all $t \in \mathbb{N}, t < T$, one of the following four events occurs at step $t + 1$ (i.e., τ_{t+1}): either a susceptible vertex is infected through an edge outside of G' , which we call $E_{o,t}$; or a susceptible vertex is infected through an edge inside of G' , which we call $E_{si,t}$; or an infected vertex recovers in the event $E_{ir,t}$; or a recovered vertex loses its immunity, which we call $E_{rs,t}$.

For each time point τ_t , let $E_{\tau_t}(I, S)$ be the number of edges from the infected to the susceptible vertices in G' and let $E_{\tau_t}(I + R, S)$ be the number of edges from the infected and recovered to the susceptible vertices in G' . At the time point τ_t , vertices get infected by other vertices via edges inside G' at a rate of $r_{si,t} = \lambda E_{\tau_t}(I, S)$, because every infected vertex infects each susceptible vertices at a rate of λ . Vertices recover from an infection at a rate of $r_{ir,t} = I_{\tau_t}$ and get susceptible at a rate of $r_{rs,t} = \varrho R_{\tau_t}$. As we only consider the states of the vertices in G' , we cannot calculate the rate $r_{o,t}$ at which susceptible vertices get infected through edges outside of G' , we only know that it is non-negative. Now let $r_t = r_{o,t} + r_{si,t} + r_{ir,t} + r_{rs,t}$. We get

$$\begin{aligned} p_{o,t} &= \Pr[E_{o,t}] = \frac{r_{o,t}}{r_t} \geq 0, \\ p_{si,t} &= \Pr[E_{si,t}] = \frac{r_{si,t}}{r_t} = \frac{\lambda E_{\tau_t}(I, S)}{r_t}, \\ p_{ir,t} &= \Pr[E_{ir,t}] = \frac{r_{ir,t}}{r_t} = \frac{I_{\tau_t}}{r_t}, \text{ and} \\ p_{rs,t} &= \Pr[E_{rs,t}] = \frac{r_{rs,t}}{r_t} = \frac{\varrho R_{\tau_t}}{r_t}. \end{aligned}$$

Note that we only consider these probabilities in configurations in which at least one vertex is infected, hence $r_t \neq 0$ and the above probabilities are well-defined. We now define

$$I^* = \frac{\varrho(c-1)}{(1+\varrho)c}n.$$

This value is the number of infected vertices in an equilibrium configuration of a SIRS process on a clique with n vertices and an infection rate of $\frac{c}{n}$. The equilibrium configuration is obtained by making the probabilities of infecting, healing and deimmunizing a vertex equally large. The process tends to drift towards that configuration. A clique and the expanders we consider behave very similarly, thus, I^* is also a good estimate for the number of infected vertices that C tends to have on G' .

4.2 Exponential survival time

We now show that the infection becomes epidemic if $\lambda \geq \frac{c}{d}$ for a constant $c \in \mathbb{R}_{>1}$. We start by proving that, when starting with one infected vertex inside of the expander, the infection reaches a configuration with at least εn infected vertices with sufficiently large probability.

Lemma 1.3. Let G be a graph, and let G' be a subgraph of G that is an $(n, (1 \pm \varepsilon_d)d, \delta)$ -expander. Let $d \rightarrow \infty$ and $\delta, \varepsilon_d \rightarrow 0$ as $n \rightarrow \infty$. Let C be a SIRS process on G with infection rate λ and with constant deimmunization rate ϱ . Further, let C start with at least one infected vertex in G' and no recovered vertices in G' . Consider the projection C' of C onto G' . If $\lambda \geq \frac{c}{d}$ for a constant $c \in \mathbb{R}_{>1}$, then there exists a constant $\varepsilon \in \mathbb{R}_{>0}$ such that for sufficiently large n , the probability that there exists a time step $t \in \mathbb{N}$ with $I_{\tau_t} \geq \varepsilon n$ is at least $\frac{1}{n+2}$.

Proof. Let $c' = c - 1$. Note that c' is positive because $c > 1$. Let $\varepsilon_H, \varepsilon_S \in \mathbb{R}_{>0}$ be constants that we specify later. We define for all $t \in \mathbb{N}$ the potential $H_t = H(I_{\tau_t}, R_{\tau_t}) = I_{\tau_t} - \varepsilon_H R_{\tau_t}$. Additionally, we define the stopping time $T = \inf\{t \in \mathbb{N} \mid H_t \leq 0 \vee S_{\tau_t} < (1 - \varepsilon_S)n\}$ and the natural filtration $(\mathcal{F}_t)_{t \in \mathbb{R}_{\geq 0}}$ of C . We aim to show that $(H_t)_{t \in \mathbb{N}}$ is a sub-martingale until T . This allows us to apply the optional-stopping theorem (Theorem 2.1) to bound $\mathbb{E}[H_T]$ from below. The law of total expectation then yields a lower bound of $\frac{1}{n+2}$ for $\Pr[H_T > 0]$. We conclude the proof by showing that if $H_T > 0$, then $I_{\tau_T} \geq \varepsilon n$.

We first bound $r_{si,t}$ using Theorem 2.6 for all times $t < T$. We get

$$\begin{aligned} r_{si,t} &= \lambda E_{\tau_t}(I, S) \\ &\geq \lambda(E_{\tau_t}(I + R, S) - (1 + \varepsilon_d)dR_{\tau_t}) \\ &\geq \lambda\left((1 - \delta)(1 - 3\varepsilon_d)\frac{d(I_{\tau_t} + R_{\tau_t})S_{\tau_t}}{n} - (1 + \varepsilon_d)dR_{\tau_t}\right) \\ &\geq \lambda((1 - \delta)(1 - 3\varepsilon_d)(1 - \varepsilon_S)d(I_{\tau_t} + R_{\tau_t}) - (1 + \varepsilon_d)dR_{\tau_t}) \\ &\geq \frac{c}{d}((1 - \delta - 3\varepsilon_d - \varepsilon_S)d(I_{\tau_t} + R_{\tau_t}) - (1 + \varepsilon_d)dR_{\tau_t}) \\ &\geq cI_{\tau_t} - (\delta + 4\varepsilon_d + \varepsilon_S)c(I_{\tau_t} + R_{\tau_t}). \end{aligned}$$

We now bound for all $t \in \mathbb{N}$ the drift $\mathbb{E}[(H_{t+1} - H_t) \cdot \mathbb{1}_{t < T} \mid \mathcal{F}_{\tau_t}]$. To improve readability, we omit the multiplicative $\mathbb{1}_{t < T}$ in all of the terms.

$$\begin{aligned} \mathbb{E}[H_{t+1} - H_t \mid \mathcal{F}_{\tau_t}] &= (p_{si,t} + p_{o,t}) \cdot (H(I_{\tau_t} + 1, R_{\tau_t}) - H_t) \\ &\quad + p_{ir,t}(H(I_{\tau_t} - 1, R_{\tau_t} + 1) - H_t) + p_{rs,t}(H(I_{\tau_t}, R_{\tau_t} - 1) - H_t) \\ &= p_{si,t} + p_{o,t} - p_{ir,t}(1 + \varepsilon_H) + p_{rs,t}\varepsilon_H \\ &\geq p_{si,t} - p_{ir,t}(1 + \varepsilon_H) + p_{rs,t}\varepsilon_H \\ &\geq (cI_{\tau_t} - (\delta + 4\varepsilon_d + \varepsilon_S)c(I_{\tau_t} + R_{\tau_t}) - I_{\tau_t}(1 + \varepsilon_H) + \varrho R_{\tau_t}\varepsilon_H)/r_t \\ &= \frac{(c' - \varepsilon_H - (\delta + 4\varepsilon_d + \varepsilon_S)c)I_{\tau_t} + (\varrho\varepsilon_H - (\delta + 4\varepsilon_d + \varepsilon_S)c)R_{\tau_t}}{r_t} \\ &\geq 0. \end{aligned}$$

The last inequality holds by first choosing ε_H such that $\varepsilon_H < c'$ and $\varepsilon_H \leq 1$ and then choosing ε_S small enough. Then for sufficiently small δ and ε_d , both of the summands are positive. Note that we do not need $\varepsilon_H \leq 1$ here, but we use that later to bound the change of H in one step.

Note that $\mathbb{E}[T] < \infty$ because in each step $t \in \mathbb{N}_{<T}$, there is a non-zero probability (independent of t) to recover a vertex, hence there is always a non-zero probability to recover all vertices within the next n steps, which stops the process. Therefore, by applying the optional-stopping theorem (Theorem 2.1), we get $\mathbb{E}[H_T] \geq \mathbb{E}[H_0]$.

By the law of total expectation, we get that

$$\begin{aligned} \mathbb{E}[H_T] &= \mathbb{E}[H_T \mid H_T \leq 0] \cdot \Pr[H_T \leq 0] + \mathbb{E}[H_T \mid H_T > 0] \cdot \Pr[H_T > 0] \\ &= \mathbb{E}[H_T \mid H_T \leq 0] \cdot (1 - \Pr[H_T > 0]) + \mathbb{E}[H_T \mid H_T > 0] \cdot \Pr[H_T > 0]. \end{aligned}$$

Because of the definition of T and the fact that H changes by at most $1 + \varepsilon_H \leq 2$ in one step, we get that $H_T \geq -2$. We also know that $H_T \leq n$ as $I_{\tau_T} \leq n$. By definition of C , it holds that $H_0 \geq 1$. By substituting $\mathbb{E}[H_T]$ in $\mathbb{E}[H_T] \geq \mathbb{E}[H_0]$ and solving for $\Pr[H_T > 0]$, we get

$$\begin{aligned} \Pr[H_T > 0] &\geq \frac{1 - \mathbb{E}[H_T \mid H_T \leq 0]}{\mathbb{E}[H_T \mid H_T > 0] - \mathbb{E}[H_T \mid H_T \leq 0]} \\ &\geq \frac{1}{n + 2}. \end{aligned}$$

Now assume $H_T > 0$. By the definition of T , it then holds that $S_{\tau_T} < (1 - \varepsilon_S)n$. Therefore,

$$I_{\tau_T} + R_{\tau_T} = n - S_{\tau_T} > \varepsilon_S n.$$

With $H_T > 0$, we then get $I_{\tau_T} > \varepsilon_H R_{\tau_T}$, which implies

$$(1 + \varepsilon_H^{-1})I_{\tau_T} > \varepsilon_S n.$$

Choosing ε accordingly concludes the proof. □

Note that the proof of the previous lemma does not directly use that $d \rightarrow \infty$. However, that is a necessary condition for $\delta \rightarrow 0$, so we included it into the lemma statement to make sure that the lemma does not work for constant d . We do the same in following lemmas.

To show that the infection survives long from that point onward, we define a potential function that assigns a real number to each configuration of the process, and we analyze its drift. The potential function is an adjusted version of the Lyapunov function of Korobeinikov and Wake [22]. We first define an auxiliary function f .

Definition 4.1. Let $f: (\mathbb{R}_{>0})^2 \rightarrow \mathbb{R}$ be such that, for all $x, x^* \in \mathbb{R}_{>0}$, we have

$$f(x^*, x) = x^* \left(\frac{x}{x^*} - \ln \frac{x}{x^*} - 1 \right).$$

Note that the derivative $\frac{df(x^*, x)}{dx} = 1 - \frac{x^*}{x}$. Hence, for a given $x^* \in \mathbb{R}_{>0}$, the value $x = x^*$ is the only local optimum of $f(x^*, x)$, and it is a global minimum. The x^* therefore acts like a target value and $f(x^*, x)$ gets larger the further x is away from that target. We now define the potential function that we use in the following lemmas.

Definition 4.2. Let G be a graph and let G' be a subgraph of G that is an $(n, (1 \pm \varepsilon_d)d, \delta)$ -expander. Let C be a SIRS process on G with infection rate $\lambda \geq \frac{c}{d}$ for a constant $c \in \mathbb{R}_{>1}$ and with deimmunization rate ϱ . Consider the projection C' of C onto G' . Let $n' = (1 + \frac{\varrho}{c})n$. For all $t \in \mathbb{N}$, we define F_t as

$$F_t = F(P_{\tau_t}, I_{\tau_t}) = f(n', P_{\tau_t}) + f(I^*, I_{\tau_t}).$$

Further, let $(\mathcal{F}_t)_{t \in \mathbb{R}_{\geq 0}}$ be the natural filtration of C . We define for all $t \in \mathbb{N}$ the drift D_t as

$$D_t = \mathbb{E}[F_{t+1} - F_t \mid \mathcal{F}_{\tau_t}].$$

The potential F becomes very large when the infection is close to dying out. We aim to show that the process tends to drift away from that high-potential region when we ignore the impact of the vertices outside of the considered subgraph and that there is a region in which the extra vertices only enlarge that drift. To calculate the differences of the F values in the drift, we first have a look at f .

Lemma 4.3. Let $x^* \in \mathbb{R}_{>0}$ and $x \in \mathbb{R}_{>2}$. Then

$$f(x^*, x + 1) - f(x^*, x) \leq 1 - \frac{x^*}{x} + \frac{x^*}{x(x + 1)} \text{ and}$$

$$f(x^*, x - 1) - f(x^*, x) \leq -\left(1 - \frac{x^*}{x} - \frac{x^*}{x(x - 1)}\right).$$

Proof. We use that for all $y \in \mathbb{R}_{>1}$, it holds that

$$\frac{1}{y + 1} < \ln(y + 1) - \ln(y) < \frac{1}{y}.$$

Together with the definition of f , we have

$$\begin{aligned} f(x^*, x + 1) - f(x^*, x) &= x^* \left(\frac{x + 1}{x^*} - \ln \frac{x + 1}{x^*} - 1 \right) - x^* \left(\frac{x}{x^*} - \ln \frac{x}{x^*} - 1 \right) \\ &= 1 - x^* (\ln(x + 1) - \ln x) \\ &\leq 1 - \frac{x^*}{x + 1}. \end{aligned}$$

For the second part, we get

$$\begin{aligned} f(x^*, x - 1) - f(x^*, x) &= x^* \left(\frac{x - 1}{x^*} - \ln \frac{x - 1}{x^*} - 1 \right) - x^* \left(\frac{x}{x^*} - \ln \frac{x}{x^*} - 1 \right) \\ &= -1 + x^* (\ln x - \ln(x - 1)) \\ &\leq -\left(1 - \frac{x^*}{x - 1}\right). \end{aligned}$$

Noting that $\frac{x^*}{x+1} = \frac{x^*}{x} - \frac{x^*}{x(x+1)}$ and $\frac{x^*}{x-1} = \frac{x^*}{x} + \frac{x^*}{x(x-1)}$ concludes the proof. \square

To bound the drift, we first show that there is an $\varepsilon \in \mathbb{R}_{>0}$ such that if there are less than εn infected vertices, the drift is upper bounded by a term that is independent of $r_{o,t}$. That is needed as we do not assume any information about $r_{o,t}$ other than it being non-negative.

Lemma 4.4. Let G be a graph, and let G' be a subgraph of G that is an $(n, (1 \pm \varepsilon_d)d, \delta)$ -expander. Let C be a SIRS process on G with infection rate λ and with constant deimmunization rate ϱ . Consider the projection C' of C onto G' . Let $E(I, S)$ be the amount of edges between the infected and the susceptible vertices at time t , and let $r'_t = \frac{c}{d}E(I, S) + r_{ir,t} + r_{rs,t}$. If $\lambda \geq \frac{c}{d}$ for a constant $c \in \mathbb{R}_{>1}$, then there exists a constant $\varepsilon \in \mathbb{R}_{>0}$ such that, for all $t \in \mathbb{N}$ and sufficiently large n , if $2 \leq I_{\tau_t} \leq \varepsilon n$, then

$$\begin{aligned} r'_t \cdot D_t &\leq \frac{c}{d}E(I, S) \cdot (F(P_{\tau_t} - 1, I_{\tau_t} + 1) - F(P_{\tau_t}, I_{\tau_t})) \\ &\quad + r_{ir,t} \cdot (F(P_{\tau_t}, I_{\tau_t} - 1) - F(P_{\tau_t}, I_{\tau_t})) \\ &\quad + r_{rs,t} \cdot (F(P_{\tau_t} + 1, I_{\tau_t}) - F(P_{\tau_t}, I_{\tau_t})). \end{aligned}$$

Proof. Let $t \in \mathbb{N}$. For easier notation, we first define

$$\begin{aligned} F_{si,t} &= F(P_{\tau_t} - 1, I_{\tau_t} + 1) - F(P_{\tau_t}, I_{\tau_t}), \\ F_{ir,t} &= F(P_{\tau_t}, I_{\tau_t} - 1) - F(P_{\tau_t}, I_{\tau_t}) \\ \text{and } F_{rs,t} &= F(P_{\tau_t} + 1, I_{\tau_t}) - F(P_{\tau_t}, I_{\tau_t}). \end{aligned}$$

We know that $r_{si,t} = \lambda E(I, S) = \frac{c}{d}E(I, S) + r_{c,t}$ for some $r_{c,t} \in \mathbb{R}_{\geq 0}$. By the definition of D_t and the fact that $r_t = r'_t + r_{o,t} + r_{c,t}$, we get that

$$r'_t \cdot D_t = r'_t \cdot \frac{r_{o,t}F_{si,t} + r_{si,t}F_{si,t} + r_{ir,t}F_{ir,t} + r_{rs,t}F_{rs,t}}{r_t}$$

Analysis of the survival time of the SIRS process via expansion

$$\begin{aligned}
 &= r'_t \cdot \frac{r_{o,t}F_{si,t} + r_{c,t}F_{si,t} + \frac{c}{d}E(I, S)F_{si,t} + r_{ir,t}F_{ir,t} + r_{rs,t}F_{rs,t}}{r_t} \\
 &= \frac{c}{d}E(I, S)F_{si,t} + r_{ir,t}F_{ir,t} + r_{rs,t}F_{rs,t} \\
 &\quad + \frac{r_{o,t} + r_{c,t}}{r_t} \left(r'_t F_{si,t} - \frac{c}{d}E(I, S)F_{si,t} - r_{ir,t}F_{ir,t} - r_{rs,t}F_{rs,t} \right)
 \end{aligned}$$

As $\frac{r_{o,t} + r_{c,t}}{r_t}$ is non-negative, to prove the lemma it is sufficient to show that

$$r'_t F_{si,t} - \frac{c}{d}E(I, S)F_{si,t} - r_{ir,t}F_{ir,t} - r_{rs,t}F_{rs,t} \leq 0.$$

By Lemma 4.3, we know that for all $x^* \in \mathbb{R}_{>0}$ and $x \in \mathbb{R}_{>2}$ holds

$$1 - \frac{x^*}{x} \leq f(x^*, x+1) - f(x^*, x) \leq 1 - \frac{x^*}{x+1}.$$

Using these bounds, we get that

$$\begin{aligned}
 -F_{ir,t} &= -(F(P_{\tau_t}, I_{\tau_t} - 1) - F(P_{\tau_t}, I_{\tau_t})) \\
 &= -(f(I^*, I_{\tau_t} - 1) - f(I^*, I_{\tau_t})) \\
 &\leq 1 - \frac{I^*}{I_{\tau_t}}, \text{ that} \\
 -F_{rs,t} &= -(F(P_{\tau_t} + 1, I_{\tau_t}) - F(P_{\tau_t}, I_{\tau_t})) \\
 &= -(f(n', P_{\tau_t} + 1) - f(n', P_{\tau_t})) \\
 &\leq \frac{n'}{P_{\tau_t}} - 1, \text{ and that} \\
 F_{si,t} &= F(P_{\tau_t} - 1, I_{\tau_t} + 1) - F(P_{\tau_t}, I_{\tau_t}) \\
 &= f(n', P_{\tau_t} - 1) - f(n', P_{\tau_t}) + f(I^*, I_{\tau_t} + 1) - f(I^*, I_{\tau_t}) \\
 &\leq 1 - \frac{I^*}{I_{\tau_t} + 1} - \left(1 - \frac{n'}{P_{\tau_t} - 1} \right) \\
 &= \frac{n'}{P_{\tau_t} - 1} - \frac{I^*}{I_{\tau_t} + 1}.
 \end{aligned}$$

Note that n' is in $\Theta(n)$ and P_{τ_t} is bounded from below by $\frac{\varrho}{c}n$, therefore $\frac{n'}{P_{\tau_t} - 1}$ is bounded from above by a constant a . Assume that $I_{\tau_t} + 1 \leq \varepsilon n$. Let $b = \frac{\varrho(1-c)}{(1+\varrho)c}$. Note that $b > 0$ is constant and $I^* = bn$. Using the bounds from above we get

$$\begin{aligned}
 &r'_t F_{si,t} - \frac{c}{d}E(I, S)F_{si,t} - r_{ir,t}F_{ir,t} - r_{rs,t}F_{rs,t} \\
 &= (r_{ir,t} + r_{rs,t})F_{si,t} - r_{ir,t}F_{ir,t} - r_{rs,t}F_{rs,t} \\
 &\leq (r_{ir,t} + r_{rs,t}) \left(a - \frac{bn}{\varepsilon n} \right) + r_{ir,t} \left(1 - \frac{bn}{\varepsilon n} \right) + r_{rs,t}(a - 1) \\
 &\leq (r_{ir,t} + r_{rs,t}) \left(2a - \frac{b}{\varepsilon} \right).
 \end{aligned}$$

We know that $(r_{ir,t} + r_{rs,t})$ is non-negative, therefore we can choose ε small enough such that the right-hand side of the previous equation is always at most 0, which concludes the proof. \square

We now show that the drift D_t is bounded from above by a negative constant for configurations in which the number of infected vertices is very small but still linear in n .

Lemma 1.4. Let G be a graph, and let G' be a subgraph of G that is an $(n, (1 \pm \varepsilon_d)d, \delta)$ -expander. Let $d \rightarrow \infty$ and $\delta, \varepsilon_d \rightarrow 0$ as $n \rightarrow \infty$. Let C be a SIRS process on G with infection rate λ and with constant deimmunization rate ϱ . Consider the projection C' of C onto G' . Let $t \in \mathbb{N}$ and $\varepsilon_0, \varepsilon \in (0, 1)$ be sufficiently small constants. Assume that $\varepsilon_0 n \geq I_{\tau_t} \geq \varepsilon n$. If $\lambda \geq \frac{c}{d}$ for a constant $c \in \mathbb{R}_{>1}$, then there exists a constant $a \in \mathbb{R}_{>0}$ such that $D_t \leq -a$ for sufficiently large n .

Proof. Let $E(I, S)$ be the amount of edges between the infected and the susceptible vertices at time t , and let $r'_t = \frac{c}{d}E(I, S) + r_{ir,t} + r_{rs,t}$. For this proof, we first use the law of total expectation and Lemma 4.3 to get a large formula as an upper bound for $r'_t D_t$. We split this bound into multiple parts and bound each part separately. We show that, with the given conditions, one of the parts is bounded from above by $-mn$ for a constant $m \in \mathbb{R}_{>0}$, and the other part is in $o(n)$, so it is asymptotically much smaller in absolute values than the other part. We conclude the proof by bounding r'_t and dividing the obtained bound for $r'_t D_t$ by it.

Using Lemma 4.4 and Lemma 4.3, we get

$$\begin{aligned} r'_t \cdot D_t &\leq \frac{c}{d}E(I, S) \cdot (F(P_{\tau_t} - 1, I_{\tau_t} + 1) - F(P_{\tau_t}, I_{\tau_t})) \\ &\quad + r_{ir,t} \cdot (F(P_{\tau_t}, I_{\tau_t} - 1) - F(P_{\tau_t}, I_{\tau_t})) \\ &\quad + r_{rs,t} \cdot (F(P_{\tau_t} + 1, I_{\tau_t}) - F(P_{\tau_t}, I_{\tau_t})) \\ &\leq \frac{c}{d}E(I, S) \cdot \left(\left(1 - \frac{I^*}{I_{\tau_t}} + \frac{I^*}{I_{\tau_t}(I_{\tau_t} + 1)} \right) - \left(1 - \frac{n'}{P_{\tau_t}} - \frac{n'}{P_{\tau_t}(P_{\tau_t} - 1)} \right) \right) \\ &\quad + r_{ir,t} \cdot \left(- \left(1 - \frac{I^*}{I_{\tau_t}} - \frac{I^*}{I_{\tau_t}(I_{\tau_t} - 1)} \right) \right) \\ &\quad + r_{rs,t} \cdot \left(\alpha \left(1 - \frac{n'}{P_{\tau_t}} + \frac{n'}{P_{\tau_t}(P_{\tau_t} + 1)} \right) \right) \\ &= \left(1 - \frac{I^*}{I_{\tau_t}} \right) \left(\frac{c}{d}E(I, S) - r_{ir,t} \right) + \left(1 - \frac{n'}{P_{\tau_t}} \right) \left(r_{rs,t} - \frac{c}{d}E(I, S) \right) \\ &\quad + \frac{\frac{c}{d}E(I, S)I^*}{I_{\tau_t}(I_{\tau_t} + 1)} + \frac{\frac{c}{d}E(I, S)n'}{P_{\tau_t}(P_{\tau_t} - 1)} + \frac{r_{ir,t}I^*}{I_{\tau_t}(I_{\tau_t} - 1)} + \frac{r_{rs,t}n'}{P_{\tau_t}(P_{\tau_t} + 1)}. \end{aligned}$$

Note that with the given conditions, all values of P_{τ_t} , I_{τ_t} , n' , and I^* are in $\Theta(n)$. All of $\frac{c}{d}E(I, S)$, $r_{ir,t}$, and $r_{rs,t}$ are in $O(n)$. Therefore, the terms in the second row of the last sum are in $O(1)$, thus we only need to bound the first part.

We know the exact values of $r_{rs,t}$ and $r_{ir,t}$. However, the value of $\frac{c}{d}E(I, S)$ depends on which vertices are infected. We use the expander properties of G' and Theorem 2.7 to bound this number. Note that both $\left(1 - \frac{I^*}{I_{\tau_t}} \right)$ and $\left(1 - \frac{n'}{P_{\tau_t}} \right)$ are negative and lower bounded by some constant. We get for sufficiently large n that

$$\begin{aligned} &\left(1 - \frac{I^*}{I_{\tau_t}} \right) \left(\frac{c}{d}E(I, S) - r_{ir,t} \right) + \left(1 - \frac{n'}{P_{\tau_t}} \right) \left(r_{rs,t} - \frac{c}{d}E(I, S) \right) \\ &\leq \left(1 - \frac{I^*}{I_{\tau_t}} \right) \left(\frac{c}{n}I_{\tau_t}S_{\tau_t} - r_{ir,t} - 4\varepsilon_d \frac{c}{n}I_{\tau_t}S_{\tau_t} - 2c\delta\sqrt{I_{\tau_t}S_{\tau_t}} \right) \\ &\quad + \left(1 - \frac{n'}{P_{\tau_t}} \right) \left(r_{rs,t} - \frac{c}{n}I_{\tau_t}S_{\tau_t} - 4\varepsilon_d \frac{c}{n}I_{\tau_t}S_{\tau_t} - 2c\delta\sqrt{I_{\tau_t}S_{\tau_t}} \right) \\ &\leq \left(1 - \frac{I^*}{I_{\tau_t}} \right) \left(\frac{c}{n}I_{\tau_t}S_{\tau_t} - r_{ir,t} \right) + \left(1 - \frac{n'}{P_{\tau_t}} \right) \left(r_{rs,t} - \frac{c}{n}I_{\tau_t}S_{\tau_t} \right) \\ &\quad + \left(\frac{I^*}{I_{\tau_t}} + \frac{n'}{P_{\tau_t}} \right) \left(2c\delta\sqrt{I_{\tau_t}S_{\tau_t}} + 4\varepsilon_d \frac{c}{n}I_{\tau_t}S_{\tau_t} \right). \end{aligned}$$

Note that $\frac{I^*}{I_{\tau_t}} + \frac{n'}{P_{\tau_t}}$ is in $\Theta(1)$, hence the last part of the last sum is in $O((\delta + \varepsilon_d)n)$. As $\delta + \varepsilon_d$ goes towards 0, this is asymptotically smaller than the rest of the drift, which we show now.

$$\begin{aligned} & \left(1 - \frac{I^*}{I_{\tau_t}}\right) \left(\frac{c}{n} I_{\tau_t} S_{\tau_t} - r_{ir,t}\right) + \left(1 - \frac{n'}{P_{\tau_t}}\right) \left(r_{rs,t} - \frac{c}{n} I_{\tau_t} S_{\tau_t}\right) \\ &= \left(1 - \frac{I^*}{I_{\tau_t}}\right) \left(\frac{c}{n} I_{\tau_t} P_{\tau_t} - \varrho I_{\tau_t} - I_{\tau_t}\right) + \left(1 - \frac{n'}{P_{\tau_t}}\right) \left(\varrho R_{\tau_t} - \frac{c}{n} I_{\tau_t} P_{\tau_t} + \varrho I_{\tau_t}\right) \\ &= \left(1 - \frac{I^*}{I_{\tau_t}}\right) \left(\frac{c}{n} I_{\tau_t} P_{\tau_t} - (1 + \varrho) I_{\tau_t}\right) + \left(1 - \frac{n'}{P_{\tau_t}}\right) \left(\varrho n' - \varrho P_{\tau_t} - \frac{c}{n} I_{\tau_t} P_{\tau_t}\right) \\ &= \frac{c}{n} I_{\tau_t} P_{\tau_t} - (1 + \varrho) I_{\tau_t} - \frac{c}{n} I^* P_{\tau_t} + (1 + \varrho) I^* + 2\varrho n' - \varrho P_{\tau_t} - \frac{c}{n} I_{\tau_t} P_{\tau_t} - \varrho \frac{n'^2}{P_{\tau_t}} + \frac{c}{n} I_{\tau_t} n' \\ &= (c - 1) I_{\tau_t} - \frac{\varrho(c - 1)}{1 + \varrho} P_{\tau_t} + \frac{\varrho(c - 1)}{c} n + 2\varrho n' - \varrho P_{\tau_t} - \varrho \frac{n'^2}{P_{\tau_t}} \\ &= \varrho \left(\frac{(c - 1) I_{\tau_t}}{\varrho} + \frac{c - 1}{c + \varrho} n' + 2n' - \frac{c + \varrho}{1 + \varrho} P_{\tau_t} - \frac{n'^2}{P_{\tau_t}} \right). \end{aligned}$$

We aim to bound this term from above. To this end, we bound $-\frac{(c+\varrho)}{1+\varrho} P_{\tau_t} - \frac{n'^2}{P_{\tau_t}}$ from above. This term has exactly one maximum for positive P_{τ_t} which is at $P_{\tau_t} = n' \sqrt{\frac{1+\varrho}{c+\varrho}}$. We also bound $I_{\tau_t} \leq \frac{\varepsilon_0 c}{c+\varrho} n'$ from above. We get

$$\begin{aligned} & \varrho \left(\frac{(c - 1) I_{\tau_t}}{\varrho} + \frac{c - 1}{c + \varrho} n' + 2n' - \frac{c + \varrho}{1 + \varrho} P_{\tau_t} - \frac{n'^2}{P_{\tau_t}} \right) \\ & \leq \varrho n' \left(\frac{(c - 1)c}{\varrho(c + \varrho)} \varepsilon_0 + \frac{c - 1}{c + \varrho} + 2 - 2\sqrt{\frac{c + \varrho}{1 + \varrho}} \right). \end{aligned}$$

The expression in the brackets is a constant, and we aim to show that it is negative for sufficiently small ε_0 . We achieve this by showing that $\frac{c-1}{c+\varrho} + 2 - 2\sqrt{\frac{c+\varrho}{1+\varrho}}$ is negative and then choosing ε_0 small enough. As both ϱ and $c - 1$ are positive, we get

$$\begin{aligned} & \frac{c - 1}{c + \varrho} + 2 - 2\sqrt{\frac{c + \varrho}{1 + \varrho}} < 0 \\ \Leftrightarrow & \frac{c - 1}{c + \varrho} + 2 < 2\sqrt{\frac{c + \varrho}{1 + \varrho}} \\ \Leftrightarrow & \frac{(c - 1)^2}{(c + \varrho)^2} + 4\frac{c - 1}{c + \varrho} + 4 < 4\frac{c + \varrho}{1 + \varrho} \\ \Leftrightarrow & \frac{(c - 1)^2}{(c + \varrho)^2} + 4\frac{c - 1}{c + \varrho} < 4\frac{c - 1}{1 + \varrho} \\ \Leftrightarrow & \frac{(c - 1)^2}{(c + \varrho)^2} < 4\frac{(c - 1)^2}{(1 + \varrho)(c + \varrho)} \\ \Leftrightarrow & 1 + \varrho < 4(c + \varrho). \end{aligned}$$

The last line holds because $c > 1$. Taking everything together, we get that $r'_t \cdot D_t$ is bounded from above by the sum of a constant, a term that is in $\Theta((\delta + \varepsilon_d)n)$, and $-b\varrho n'$, where b is a positive constant for sufficiently small ε_0 .

We know that $r'_t = \frac{c}{d} E(I, S) + I_{\tau_t} + \varrho R_{\tau_t} \leq cn(1 + \varepsilon_d) + n + \varrho n = (c(1 + \varepsilon_d) + 1 + \varrho)n$. As also $r'_t \geq I_{\tau_t} \geq \varepsilon n > 0$, by dividing the inequality for $r'_t \cdot D_t$ by r'_t , we get that there exists a constant $a \in \mathbb{R}_{>0}$ such that $D_t \leq -a$, concluding the proof. \square

We aim to apply the negative-drift theorem (Theorem 2.2) to bound the expected survival time of the infection. In Lemma 1.4, we showed a constant negative drift of the potential in a region of the configuration space. To apply the drift theorem, we first transform the configuration space restrictions into restrictions on the value of the potential. The first lemma shows that if there is at least a constant amount of infected vertices, the potential does not get too large.

Lemma 4.5. Let G be a graph, and let G' be a subgraph of G that is an $(n, (1 \pm \varepsilon_d)d, \delta)$ -expander. Let $d \rightarrow \infty$ and $\delta, \varepsilon_d \rightarrow 0$ as $n \rightarrow \infty$. Let C be a SIRS process on G with infection rate λ and with constant deimmunization rate ϱ . Consider the projection C' of C onto G' . Let $t \in \mathbb{N}$ and $\varepsilon \in (0, 1)$ be constants such that $I_{\tau_t} \geq \varepsilon n$. If $\lambda \geq \frac{c}{d}$ for a constant $c \in \mathbb{R}_{>1}$, then $F_t \in O(n)$.

Proof. We aim to bound F_t from above by writing it as a sum and bounding the individual summands. To this end, we first bound the terms that appear in the summands. By the definition of our random variables and the fact that there are only n vertices, we get

$$\begin{aligned} \max(P_{\tau_t}, I_{\tau_t}, I^*) &\leq n', \\ \min(P_{\tau_t}, I_{\tau_t}) &\geq \min(\varepsilon, \varrho/c)n. \end{aligned}$$

Applying these bounds to F_t results in

$$\begin{aligned} F_t &= f(n', P_{\tau_t}) + f(I^*, I_{\tau_t}) \\ &= n' \left(\frac{P_{\tau_t}}{n'} - \ln \frac{P_{\tau_t}}{n'} - 1 \right) + I^* \left(\frac{I_{\tau_t}}{I^*} - \ln \frac{I_{\tau_t}}{I^*} - 1 \right) \\ &\leq P_{\tau_t} + n' \ln \frac{n'}{P_{\tau_t}} + I_{\tau_t} + I^* \ln \frac{I^*}{I_{\tau_t}} \\ &\leq 2 \cdot \left(n' + n' \ln \frac{n'}{\min(\varepsilon, \varrho/c)n} \right). \end{aligned}$$

As $n' = (1 + \varrho/c)n$, the calculated upper bound for F_t is linear in n . Thus, $F_t \in O(n)$. \square

The next lemma shows that when the number of vertices becomes small, the potential gets rather large. Together with the previous lemma, this shows that having few infected vertices and having a high drift is more or less the same.

Lemma 4.6. Let G be a graph, and let G' be a subgraph of G that is an $(n, (1 \pm \varepsilon_d)d, \delta)$ -expander. Let $d \rightarrow \infty$ and $\delta, \varepsilon_d \rightarrow 0$ as $n \rightarrow \infty$. Let C be a SIRS process on G with infection rate λ and with constant deimmunization rate ϱ . Consider the projection C' of C onto G' . Let $t \in \mathbb{N}$ and $\varepsilon \in (0, I^*/n)$ be constants such that $1 \leq I_{\tau_t} \leq \varepsilon n$. If $\lambda \geq \frac{c}{d}$ for a constant $c \in \mathbb{R}_{>1}$, then

$$F_t \geq I^* \left(\ln \frac{1}{\varepsilon} + \ln \frac{I^*}{n} - 1 \right).$$

Proof. We aim to bound F_t from below by bounding the f values in its definition. Recall that for a given $x^* \in \mathbb{R}_{>0}$, the function $f(x^*, x)$ is minimized for $x = x^*$, which is the only local extreme value for $x \in \mathbb{R}_{>0}$. Therefore, we get for all $x, x^* \in \mathbb{R}_{>0}$

$$f(x^*, x) \geq f(x^*, x^*) = 0.$$

Using $1 \leq I_{\tau_t} \leq \varepsilon n$ and that for all $x^* \in \mathbb{R}_{>0}$, the function $f(x^*, x)$ decreases monotonically in x while $x < x^*$, we conclude

$$F_t = f(n', P_{\tau_t}) + f(I^*, I_{\tau_t})$$

$$\begin{aligned} &\geq 0 + f(I^*, \varepsilon n) \\ &\geq I^* \left(\frac{\varepsilon n}{I^*} - \ln \frac{\varepsilon n}{I^*} - 1 \right) \\ &\geq I^* \left(\ln \frac{1}{\varepsilon} + \ln \frac{I^*}{n} - 1 \right). \end{aligned} \quad \square$$

The next lemma shows that while the process has at least a constant fraction of vertices in the infected state, each potential next step only changes the potential by at most a constant.

Lemma 4.7. Let G be a graph, and let G' be a subgraph of G that is an $(n, (1 \pm \varepsilon_d)d, \delta)$ -expander. Let $d \rightarrow \infty$ and $\delta, \varepsilon_d \rightarrow 0$ as $n \rightarrow \infty$. Let C be a SIRS process on G with infection rate λ and with constant deimmunization rate ϱ . Consider the projection C' of C onto G' . Let $t \in \mathbb{N}$ and $\varepsilon \in (0, \varrho/c)$ be constants. Assume that $I_{\tau_t} \geq \varepsilon n$. Further, let $\Delta P, \Delta I \in \{-1, 0, 1\}$. If $\lambda \geq \frac{c}{d}$ for a constant $c \in \mathbb{R}_{>1}$, then for sufficiently large n , it holds that

$$|F(P_{\tau_t} + \Delta P, I_{\tau_t} + \Delta I) - F(P_{\tau_t}, I_{\tau_t})| \leq 2(1 + 2(1 + \varrho/c)\varepsilon^{-1}).$$

Proof. We aim to use the triangle inequality to bound the absolute change in the F -values from above by the sum of the absolute changes in the f -values. We use that for all $x \in \mathbb{R}_{>1}$ holds that

$$\frac{1}{x+1} < \ln\left(\frac{x+1}{x}\right) < \frac{1}{x}.$$

Further, for all $x, x^* \in \mathbb{R}_{>2}$ and $\Delta x \in \{-1, 0, 1\}$ holds that

$$\begin{aligned} |f(x^*, x + \Delta x) - f(x^*, x)| &= \left| x^* \left(\frac{x + \Delta x}{x^*} - \ln \frac{x + \Delta x}{x^*} - 1 \right) - x^* \left(\frac{x}{x^*} - \ln \frac{x}{x^*} - 1 \right) \right| \\ &= \left| \Delta x - x^* \ln\left(\frac{x + \Delta x}{x}\right) \right| \\ &\leq |\Delta x| + \left| x^* \ln\left(\frac{x + \Delta x}{x}\right) \right| \\ &\leq 1 + \frac{x^*}{x-1}. \end{aligned}$$

We apply this inequality to bound the absolute change in potential from above. Note that by the choice of ε it holds $\min(P_{\tau_t}, I_{\tau_t}) \geq \varepsilon n$. Hence, for sufficiently large n , it holds that $\min(P_{\tau_t} - 1, I_{\tau_t} - 1) \geq \varepsilon n/2$. We conclude

$$\begin{aligned} &|F(P_{\tau_t} + \Delta P, I_{\tau_t} + \Delta I) - F(P_{\tau_t}, I_{\tau_t})| \\ &= |f(n', P_{\tau_t} + \Delta P) + f(I^*, I_{\tau_t} + \Delta I) - f(n', P_{\tau_t}) - f(I^*, I_{\tau_t})| \\ &\leq |f(n', P_{\tau_t} + \Delta P) - f(n', P_{\tau_t})| + |f(I^*, I_{\tau_t} + \Delta I) - f(I^*, I_{\tau_t})| \\ &\leq \left(1 + \frac{n'}{P_{\tau_t} - 1} \right) + \left(1 + \frac{I^*}{I_{\tau_t} - 1} \right) \\ &\leq 2\left(1 + \frac{n'}{\varepsilon n/2} \right) \\ &\leq 2(1 + 2(1 + \varrho/c)\varepsilon^{-1}). \end{aligned} \quad \square$$

We now have the tools to apply the negative-drift theorem (Theorem 2.2) to bound the survival time of the infection.

Lemma 4.8. Let G be a graph, and let G' be a subgraph of G that is an $(n, (1 \pm \varepsilon_d)d, \delta)$ -expander. Let $d \rightarrow \infty$ and $\delta, \varepsilon_d \rightarrow 0$ as $n \rightarrow \infty$. Let C be a SIRS process on G with infection rate λ and with constant deimmunization rate ϱ . Consider the projection C' of C onto G' . Let $\varepsilon_0 \in (0, 1)$ be a constant and let E_{ε_0} be the event that there exists a time step $t \in \mathbb{N}$ such that $I_{\tau_t} \geq \varepsilon_0 n$. Let T be the first time after τ_t with $I_{\tau_t} = 0$. If $\lambda \geq \frac{c}{d}$ for a constant $c \in \mathbb{R}_{>1}$, then $\mathbb{E}[T \mid E_{\varepsilon_0}] = 2^{\Omega(n)}$.

Proof. We assume that E_{ε_0} occurs. Let $(\mathcal{F}_t)_{t \in \mathbb{R}_{\geq 0}}$ be the natural filtration of C , and let $t \in \mathbb{N}$ be such that $I_{\tau_t} \geq \varepsilon_0 n$. We aim to apply the negative-drift theorem (Theorem 2.2) to get the desired bound. To this end, we define a stopping time that is dominated by the number of steps until T , and we use the previous lemmas to show that all of the conditions for the drift theorem are satisfied. Note that we shift the time to start at t instead of 0. We then translate the bound on the number of steps into a bound on the survival time.

As $I_{\tau_t} \geq \varepsilon_0 n$, by Lemma 4.5, there exists a constant $a_0 \in \mathbb{R}_{>0}$ such that $F_t \leq a_0 n$. Let ε_c be the minimum of the ε 's from Lemma 4.4 and Lemma 1.4. By the contraposition of Lemma 4.5, there exists a constant $a_1 \in \mathbb{R}_{>0}$ such that $F_t \geq a_1 n$ implies that $I_{\tau_t} \leq \varepsilon_c n$. We define $a = \max(a_0, a_1)$ and $T_1 = \inf\{i \in \mathbb{N}_{\geq t} \mid F_i > 2an\}$.

We first show that for all $i \in \mathbb{N}$ with $t \leq i < T_1$ holds that I_{τ_i} is large enough such that Lemma 4.7 is applicable. Let $\varepsilon_1 \in (0, I^*/n)$ be a constant low enough such that $\frac{I^*}{n} \left(\ln \frac{1}{\varepsilon_1} + \ln \frac{I^*}{n} - 1 \right) > 2a$. Such an ε_1 exists, as $\frac{I^*}{n}$ and a are positive constants. Then by the contraposition of Lemma 4.6, for all $i \in \mathbb{N}$, it follows that $F_i \leq 2an$ implies that $I_{\tau_i} \geq \varepsilon_1 n$.

To show that condition 2 of Theorem 2.2 is satisfied, let $s = 2(1 + 2(1 + \varrho/c)\varepsilon_1^{-1})$. For all $i \in \mathbb{N}$ with $t \leq i < T_1$ holds $F_i \leq 2an$ and therefore $I_{\tau_i} \geq \varepsilon_1 n$. Hence, by Lemma 4.7, for all $i \in \mathbb{N}_{\geq t}$ holds that $|F_{i+1} - F_i| \cdot \mathbb{1}_{i < T_1} \leq s \cdot \mathbb{1}_{i < T_1}$. Thus, for all $i \in \mathbb{N}_{\geq t}$ and $j \in \mathbb{R}_{>0}$ holds that $\Pr[|F_{i+1} - F_i| \geq j \mid \mathcal{F}_{\tau_i}] \cdot \mathbb{1}_{i < T_1} \leq \frac{2^s}{2^j} \cdot \mathbb{1}_{i < T_1}$. This is true as for $j > s$ the probability is zero, so the inequality holds and for $j \leq s$ the term $\frac{2^s}{2^j}$ is at least one, which is a trivial upper bound for a probability. Note that 2^s is a constant.

We now show that condition 1 is satisfied as well. We already showed that for all $i \in \mathbb{N}$ with $an < F_i < 2an$ holds $\varepsilon_1 n \leq I_{\tau_i} \leq \varepsilon_c n$. Hence, the conditions for Lemma 1.4 are satisfied, and we get that there exists a constant $r \in \mathbb{R}_{>0}$ such that for all $i \in \mathbb{N}$ holds that $\mathbb{E}[F_{i+1} - F_i \mid \mathcal{F}_{\tau_i}] \cdot \mathbb{1}_{an < F_i < 2an} \leq -r \cdot \mathbb{1}_{an < F_i < 2an}$.

Now all of the conditions of Theorem 2.2 are satisfied, and we get that there exists a constant $c^* \in \mathbb{R}_{>0}$ such that

$$\Pr\left[T_1 - t \leq 2^{c^* an/2^s} \mid \mathcal{F}_{\tau_t}\right] \cdot \mathbb{1}_{F_t \leq an} = 2^{-\Omega(an/2^s)} \cdot \mathbb{1}_{F_t \leq an}.$$

Note that this probability goes towards 0 as n goes towards infinity. Hence, we get $\mathbb{E}[T_1 \mid \mathcal{F}_{\tau_t}] \cdot \mathbb{1}_{F_t \leq an} = 2^{\Omega(n)} \cdot \mathbb{1}_{F_t \leq an}$. Remember that $I_{\tau_t} \geq \varepsilon_0 n$ implies $F_t \leq an$. We therefore get $\mathbb{E}[T_1 \mid \mathcal{F}_{\tau_t}] \cdot \mathbb{1}_{I_{\tau_t} \geq \varepsilon_0 n} = 2^{\Omega(n)} \cdot \mathbb{1}_{I_{\tau_t} \geq \varepsilon_0 n}$.

We showed that for all $i \in \mathbb{N}$ with $t \leq i < T_1$ holds that $I_{\tau_i} \geq \varepsilon_1 n > 0$. Thus, T dominates τ_{T_1} . Note that clocks in C trigger at an arbitrarily high rate, as we do not have an upper bound on r_t . However, the amounts of recovery triggers, infection triggers, and deimmunization triggers that occur until τ_{T_1} differ by at most n , pairwise by type, so each of them also has an exponentially large expectation. As we only consider n recovery clocks, they trigger at a rate of at most n , and the expected time between each trigger is at least $\frac{1}{n}$. By Wald's equation (Theorem 2.4), we get that

$$\begin{aligned} \mathbb{E}[T \mid \mathcal{F}_0] \cdot \mathbb{1}_{E_{\varepsilon_0}} &\geq \mathbb{E}[\tau_{T_1} \mid \mathcal{F}_0] \cdot \mathbb{1}_{E_{\varepsilon_0}} \\ &\geq \frac{1}{n} 2^{\Omega(n)} \cdot \mathbb{1}_{E_{\varepsilon_0}}. \end{aligned} \quad \square$$

We now prove our main result.

Theorem 1.2. *Let G be a graph, and let G' be a subgraph of G that is an $(n, (1 \pm \varepsilon_d)d, \delta)$ -expander. Let $d \rightarrow \infty$ and $\delta, \varepsilon_d \rightarrow 0$ as $n \rightarrow \infty$. Let C be the SIRS process on G with infection rate λ and with constant deimmunization rate ϱ . Further, let C start with at least one infected vertex in G' and no recovered vertices in G' . Last, let C' be the projection of C onto G' , and let T be the survival time of C' . If $\lambda \geq \frac{c}{d}$ for a constant $c \in \mathbb{R}_{>1}$, then for sufficiently large n , it holds that $\mathbb{E}[T] = 2^{\Omega(n)}$.*

Proof. For all constants $\varepsilon \in (0, 1)$, let E_ε be the event that there exists a time step $t \in \mathbb{N}$ such that $I_{\tau_t} \geq \varepsilon n$. By Lemma 1.3, there exists an $\varepsilon \in \mathbb{R}_{>0}$ such that for sufficiently large n holds that $\Pr[E_\varepsilon] \geq \frac{1}{n+2}$. By Lemma 4.8, it holds that $\mathbb{E}[T \mid E_\varepsilon] = 2^{\Omega(n)}$. Using the law of total expectation, we get

$$\begin{aligned} \mathbb{E}[T] &= \Pr[E_\varepsilon]\mathbb{E}[T \mid E_\varepsilon] + \Pr[\overline{E_\varepsilon}]\mathbb{E}[T \mid \overline{E_\varepsilon}] \\ &\geq \Pr[E_\varepsilon]\mathbb{E}[T \mid E_\varepsilon] \\ &\geq \frac{1}{n+2}2^{\Omega(n)} \\ &= 2^{\Omega(n)}. \end{aligned} \quad \square$$

5 Applications

We illustrate the applicability of our theorem by considering the SIRS process on Erdős–Rényi graphs as well as models of real-world networks.

5.1 Erdős–Rényi graphs

In order to apply Theorem 1.2 to Erdős–Rényi graphs, we make use of the following result.

Theorem 5.1 ([10, Theorem 1.2]). *Let $G \sim G_{n,p}$ be an Erdős–Rényi graph with $(n-1)p \geq c_1 \ln(n)$ for a sufficiently large constant $c_1 \in \mathbb{R}_{>0}$. Then asymptotically almost surely, for the spectral expansion δ of the Laplacian of G holds $\delta \in O((p(n-1))^{-1/2})$.*

By Chernoff bounds, it holds that the vertex degrees in Erdős–Rényi graphs are tightly distributed around the average degree d if $d \in \omega(\ln n)$. Therefore, Erdős–Rényi graphs satisfy with high probability our definition of an $(n, (1 \pm \varepsilon_d)d, \delta)$ -expander. Combining this with Theorem 1.2 and the upper bound that carries over from the SIS model (see Ganesh, Massoulié, and Towsley [18, Theorem 5.5]), we obtain the following corollary.

Corollary 1.5. *Let $G \sim G_{n,p}$ be an Erdős–Rényi graph with $(n-1)p \in \omega(\ln n)$. Consider the SIRS process C on G with constant deimmunization rate ϱ , and let T be the survival time of C when the process starts with at least one infected vertex. If $\lambda \geq \frac{c}{d}$ for a constant $c \in \mathbb{R}_{>1}$, then $\mathbb{E}[T] = 2^{\Omega(n)}$ asymptotically almost surely with respect to G . If $\lambda \leq \frac{c}{d}$ for a constant $c \in (0, 1)$, then $\mathbb{E}[T] \in O(\log n)$ asymptotically almost surely with respect to G .*

5.2 Complex networks

A variety of random graph models that exhibit properties found in real-world networks has appeared in network science [6]. To apply Theorem 1.2 in such models, we require that the graph has a large expander as a subgraph and that the infection reaches a vertex of this subgraph with sufficiently high probability. To this end, two key properties have been shown to hold on popular complex network models, namely, a polynomially-sized clique as a subgraph and a polylogarithmic diameter, both in terms of the number of vertices of the graph. These two properties hold for example on Chung–Lu graphs [9],

on hyperbolic random graphs (Theorems 5.2 and 5.3) and on geometric inhomogeneous random graphs [21].

We formally prove our claim for hyperbolic random graphs. The following two theorems state the two key properties we require.

Theorem 5.2 ([17, Theorem 1]). *Let G be a hyperbolic random graph with n vertices that follows a power-law degree distribution with exponent $\gamma \in (2, 3)$. Then the diameter of the giant component of G is $O((\log n)^{2/(3-\gamma)})$ with probability $1 - O(n^{-3/2})$.*

Theorem 5.3 ([16]). *Let G be a hyperbolic random graph with n vertices that follows a power-law degree distribution with exponent $\gamma \in (2, 3)$. Then the size of the largest clique of G is in $\Theta(n^{(3-\gamma)/2})$ with high probability.*

We first use the poly-logarithmic diameter to show that the infection reaches the largest clique with a sufficient probability when the process starts with at least one infected vertex.

Lemma 5.4. *Let G be a hyperbolic random graph with n vertices that follows a power-law degree distribution with exponent $\gamma \in (2, 3)$, and let C be an SIRS process on G with infection rate λ and with constant deimmunization rate ϱ . Further, let C start with at least one infected vertex in the giant component and no recovered vertices in the giant component. If $\lambda \geq cn^{(\gamma-3)/2}$ for a constant $c \in \mathbb{R}_{>0}$, then the probability that the infection reaches a configuration in which a vertex in the largest clique is infected is at least $\exp(-(\ln n)^{3/(3-\gamma)})$ for sufficiently large n .*

Proof. Let v be a vertex that starts infected, and let d be the shortest distance from v to any vertex of the largest clique. Note that d is bounded from above by the diameter of the giant component. Therefore, by Theorem 5.2, there exists a constant $a \in \mathbb{R}_{>0}$ such that for sufficiently large n with a probability of at least $\frac{1}{2}$, it holds that $d \leq a(\ln n)^{2/(3-\gamma)}$.

For all $i \in \mathbb{N}$, let E_i be the event that C reaches a configuration with an infected vertex that has a distance of i to the largest clique. Consider for all $i \in \mathbb{N}_{<d}$ the probability $\Pr[E_i \mid E_{i+1}]$. Each vertex with a distance of $i + 1$ to the largest clique has a neighbor that has a distance of i to the clique. With a probability of $\frac{\lambda}{1+\lambda}$, an infected vertex infects a specific neighbor before recovering. Therefore, $\Pr[E_i \mid E_{i+1}] \geq \frac{\lambda}{1+\lambda} \geq \frac{c}{2}n^{(\gamma-3)/(2)}$ for sufficiently large n .

With a probability of at least $\frac{1}{2}$, it holds that $d \leq a(\ln n)^{2/(3-\gamma)}$. This yields for sufficiently large n that

$$\begin{aligned} \Pr[E_0] &= \prod_{i=0}^{d-1} \Pr[E_i \mid E_{i+1}] \\ &\geq \prod_{i=0}^{d-1} \frac{c}{2}n^{\frac{\gamma-3}{2}} \\ &= \left(\frac{c}{2}n^{\frac{\gamma-3}{2}}\right)^d \\ &\geq \left(\frac{c}{2}\right)^d \left(n^{\frac{\gamma-3}{2}}\right)^{a(\ln n)^{\frac{2}{3-\gamma}}} \\ &= e^{\frac{\gamma-3}{2}a(\ln n)^{\frac{5-\gamma}{3-\gamma}} + d \ln(c/2)} \\ &\geq e^{-(\ln n)^{\frac{3}{3-\gamma}}}. \quad \square \end{aligned}$$

When the infection reaches the largest clique of a hyperbolic random graph, Theorem 1.2 yields an exponential expected survival time for a sufficiently large infection rate.

Corollary 1.6. *Let G be a hyperbolic random graph with n vertices that follows a power-law degree distribution with exponent $\gamma \in (2, 3)$, and let C be the SIRS process on G with infection rate λ and with constant deimmunization rate ρ . Further, let C start with at least one infected vertex in the giant component and no recovered vertices, and let T be the survival time of C . Then there exists a constant $c \in \mathbb{R}_{>0}$ such that if $\lambda \geq cn^{(\gamma-3)/2}$, then $\mathbb{E}[T] = 2^{\Omega(n^{(3-\gamma)/2})}$.*

Proof. Let k be the size of the largest clique of G . By Theorem 5.3, there exists a constant $a \in \mathbb{R}_{>0}$ such that with high probability it holds that $k \geq an^{(3-\gamma)/2}$. Let $c = a^{-1} + 1$ such that with high probability it holds that $\lambda \geq \frac{1+a}{k}$. Let E be the event that there exists a configuration in which a vertex in the largest clique of G is infected. By Lemma 5.4, it holds that $\Pr[E] \geq \exp(-(\ln n)^{3/(3-\gamma)})$ for sufficiently large n . Note that a clique with k vertices is a $(k, (1 \pm k^{-1})k, (k-1)^{-1})$ -expander. Hence, by Theorem 1.2, it holds that $\mathbb{E}[T | E] = 2^{\Omega(k)}$, as the infection survives that long on the clique alone after its first vertex gets infected.

By the law of total expectation and that with high probability $k \geq an^{(3-\gamma)/2}$, we conclude

$$\begin{aligned} \mathbb{E}[T] &\geq \Pr[E] \cdot \mathbb{E}[T | E] \\ &\geq e^{-(\ln n)^{\frac{3}{3-\gamma}}} \cdot 2^{\Omega(n^{(3-\gamma)/2})} \\ &= 2^{\Omega(n^{(3-\gamma)/2})}. \quad \square \end{aligned}$$

Note that similar statements can be proven for Chung–Lu graphs and geometric inhomogeneous random graphs for an appropriate choice of the respective parameters of these models.

References

- [1] William Aiello, Fan Chung, and Linyuan Lu, *A random graph model for power law graphs*, *Experimental Mathematics* **10** (2001), no. 1, 53–66. MR1822851
- [2] Jean-Daniel Bancal and Romualdo Pastor-Satorras, *Steady-state dynamics of the forest fire model on complex networks*, *The European Physical Journal B* **76** (2010), no. 1, 109–121.
- [3] Albert-László Barabási and Réka Albert, *Emergence of scaling in random networks*, *Science* **286** (1999), no. 5439, 509–512. MR2091634
- [4] Noam Berger, Christian Borgs, Jennifer T. Chayes, and Amin Saberi, *On the spread of viruses on the internet*, *Symposium on Discrete Algorithms (SODA)* (2005), 301–310. MR2298278
- [5] Shankar Bhamidi, Dannz Nam, Oanh Nguyen, and Allan Sly, *Survival and extinction of epidemics on random graphs with general degree*, *The Annals of Probability* **49** (2021), no. 1, 244–286. MR4203338
- [6] Marian Boguna, Ivan Bonamassa, Manlio De Domenico, Shlomo Havlin, Dmitri Krioukov, and M Ángeles Serrano, *Network geometry*, *Nature Reviews Physics* **3** (2021), no. 2, 114–135.
- [7] Christian Borgs, Jennifer Chayes, Ayalvadi Ganesh, and Amin Saberi, *How to distribute antidote to control epidemics*, *Random Structures & Algorithms* **37** (2010), no. 2, 204–222. MR2676029
- [8] Fan R. K. Chung, *Spectral graph theory*, American Mathematical Society, 1997. MR1421568
- [9] Fan R. K. Chung and Linyuan Lu, *The average distance in a random graph with given expected degrees*, *Internet Mathematics* **1** (2003), no. 1, 91–113. MR2076728
- [10] Amin Coja-Oghlan, *On the laplacian eigenvalues of $g_{n,p}$* , *Combinatorics, Probability and Computing* **16** (2007), no. 6, 923–946. MR2351691
- [11] Carola Doerr and Martin S. Krejca, *Run time analysis for random local search on generalized majority functions*, *IEEE Transactions on Evolutionary Computation* (2022), In press.

- [12] Paul Erdős and Alfréd Rényi, *On random graphs i*, *Publicationes Mathematicae* **6** (1959), no. 1, 290–297. MR0120167
- [13] William Feller, *An introduction to probability theory and its applications*, 3 ed., vol. 1, John Wiley & Sons, 1968. MR0228020
- [14] Silvio C. Ferreira, Renan S. Sander, and Romualdo Pastor-Satorras, *Collective versus hub activation of epidemic phases on networks*, *Physical Review E* **93** (2016), no. 3, 032314. MR3657790
- [15] Tobias Friedrich, Andreas Göbel, Nicolas Klodt, Martin S. Krejca, and Marcus Pappik, *The irrelevance of influencers: Information diffusion with re-activation and immunity lasts exponentially long on social network models*, *AAAI Conference on Artificial Intelligence*, 2024, p. To appear.
- [16] Tobias Friedrich and Anton Krohmer, *Cliques in hyperbolic random graphs*, *Conference on Computer Communications (INFOCOM)*, 2015, pp. 1544–1552.
- [17] Tobias Friedrich and Anton Krohmer, *On the diameter of hyperbolic random graphs*, *SIAM Journal on Discrete Mathematics* **32** (2018), no. 2, 1314–1334. MR3813231
- [18] Ayalvadi Ganesh, Laurent Massoulié, and Don Towsley, *The effect of network topology on the spread of epidemics*, *International Conference on Computer Communications (INFOCOM)*, 2005, pp. 1455–1466.
- [19] T. E. Harris, *Contact interactions on a lattice*, *The Annals of Probability* **2** (1974), no. 6, 969–988. MR0356292
- [20] Xiangying Huang and Rick Durrett, *The contact process on random graphs and Galton-Watson trees*, *Latin American Journal of Probability and Mathematical Statistics* **17** (2020), 159–182. MR4057187
- [21] Ralph Keusch, *Geometric inhomogeneous random graphs and graph coloring games*, *Doctoral thesis*, ETH Zurich, Zurich, 2018.
- [22] Andrei Korobeinikov and Graeme C. Wake, *Lyapunov functions and global stability for sir, sirs, and sis epidemiological models*, *Applied Mathematics Letters* **15** (2002), no. 8, 955–960. MR1925920
- [23] Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguñá, *Hyperbolic geometry of complex networks*, *Physical Review E* **82** (2010), 036106–1–036106–18. MR2787998
- [24] Marcelo Kuperman and Guillermo Abramson, *Small world effect in an epidemiological model*, *Physical Review Letters* **86** (2001), 2909–2912.
- [25] Thomas M. Liggett, *Multiple transition points for the contact process on the binary tree*, *The Annals of Probability* **24** (1996), no. 4, 1675–1710. MR1415225
- [26] Aleksandr Mikhailovich Lyapunov, *The general problem of the stability of motion*, *International Journal of Control* **55** (1992), no. 3, 531–534. MR1154209
- [27] Michael Mitzenmacher and Eli Upfal, *Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis*, 2 ed., Cambridge University Press, 2017. MR3674428
- [28] Danny Nam, Oanh Nguyen, and Allan Sly, *Critical value asymptotics for the contact process on random graphs*, *Transactions of the American Mathematical Society* **375** (2022), no. 12. MR4419050
- [29] Pietro S. Oliveto and Carsten Witt, *Simplified drift analysis for proving lower bounds in evolutionary computation*, *Algorithmica* **59** (2011), no. 3, 369–386. MR2770464
- [30] Pietro S. Oliveto and Carsten Witt, *Erratum: Simplified drift analysis for proving lower bounds in evolutionary computation*, *CoRR* **abs/1211.7184** (2012). MR2770464
- [31] Romualdo Pastor-Satorras, Claudio Castellano, Piet Van Mieghem, and Alessandro Vespignani, *Epidemic processes in complex networks*, *Reviews of Modern Physics* **87** (2015), no. 3, 925–979. MR3406040
- [32] Robin Pemantle, *The contact process on trees*, *Annals of Probability* **20** (1992), 2089–2116. MR1188054

- [33] B. Aditya Prakash, Deepayan Chakrabarti, Nicholas C Valler, Michalis Faloutsos, and Christos Faloutsos, *Threshold conditions for arbitrary cascade models on arbitrary networks*, Knowledge and Information Systems **33** (2012), 549–575.
- [34] M. Ali Saif, *Epidemic threshold for the sirs model on the networks*, Physica A: Statistical Mechanics and its Applications **535** (2019), 122251.1–122251.7.
- [35] Alan M. Stacey, *The existence of an intermediate phase for the contact process on trees*, The Annals of Probability **24** (1996), no. 4, 1711 – 1726. MR1415226
- [36] Yi Wang, Jinde Cao, Ahmed Alsaedi, and Tasawar Hayat, *The spreading dynamics of sexually transmitted diseases with birth and death on heterogeneous networks*, Journal of Statistical Mechanics: Theory and Experiment **2017** (2017), no. 2, 1–26. MR3637253

Acknowledgments. We would like to thank Silvio Ferreira for bringing [14] to our attention and for insightful comments that helped us improve the expected survival time of Theorem 1.1.

Andreas Göbel was funded by the project PAGES (project No. 467516565) of the German Research Foundation (DFG). Marcus Pappik was funded by the HPI Research School on Data Science and Engineering.

Electronic Journal of Probability

Electronic Communications in Probability

Advantages of publishing in EJP-ECP

- Very high standards
- Free for authors, free for readers
- Quick publication (no backlog)
- Secure publication (LOCKSS¹)
- Easy interface (EJMS²)

Economical model of EJP-ECP

- Non profit, sponsored by IMS³, BS⁴, ProjectEuclid⁵
- Purely electronic

Help keep the journal free and vigorous

- Donate to the IMS open access fund⁶ (click here to donate!)
- Submit your best articles to EJP-ECP
- Choose EJP-ECP over for-profit journals

¹LOCKSS: Lots of Copies Keep Stuff Safe <http://www.lockss.org/>

²EJMS: Electronic Journal Management System: <https://vtex.lt/services/ejms-peer-review/>

³IMS: Institute of Mathematical Statistics <http://www.imstat.org/>

⁴BS: Bernoulli Society <http://www.bernoulli-society.org/>

⁵Project Euclid: <https://projecteuclid.org/>

⁶IMS Open Access Fund: <https://imstat.org/shop/donation/>