

Network Process: How Topology Impacts the Dynamics of Epidemics and Cascading Failures

Submitted in partial fulfillment for the requirements for
the degree of
Doctor of Philosophy
in
Electrical and Computer Engineering

June Zhang

B.S., Electrical and Computer Engineering, Georgia Institute of Technology
M.S., Electrical Engineering, Stanford University

Carnegie Mellon University
Pittsburgh, PA

September 24, 2015

Strange is our situation here on Earth. Each of us comes for a short visit, not knowing why, yet sometimes seeming to divine a purpose. From the standpoint of daily life, however, there is one thing we do know: that man is here for the sake of other men.

—Albert Einstein

Abstract

Network processes model how information, virus, or failures interact and spread in a system or population. The network structure captures relationships or contacts between multiple, interactive agents. Network processes also account for the dynamical characteristics of these interactions. These models can give insights on how to respond to outbreaks as well as on how to design optimal distributed, network structures for engineering purposes. These processes are complex to study because of inherent dependencies between topology and dynamics. The inclusion of finite-size networks introduces combinatorial complexity, making analysis computationally intractable for large networks. Consequently, existing approaches rely on simulation or coarse approximations of the underlying network structures.

The thesis presents and studies the scaled SIS process, a model for which the exact network structure can be accounted for exactly. The scaled SIS process is an epidemics-like, binary-state, stochastic process on an arbitrary, undirected network whose nodes represent agents and edges represent contacts. Contagion of healthy agents by infected neighbors can potentially overcome the healing process leading to epidemics. Alternatively, the scaled SIS process can also be used to model cascading failures like blackouts in the power grid.

The scaled SIS process has a closed-form steady-state characterization (i.e., equilibrium distribution) of the Gibbs form, making it appropriate to study the effects of network topology and dynamics on the steady-state behavior of the process. We use the equilibrium distribution to formulate and study network vulnerability on 3 different scales: 1) individual agents, 2) substructures in the network, and 3) overall network.

With tools from discrete optimization and Monte Carlo sampling, we can solve or accurately approximate the solutions to these inference problems for network processes on arbitrary, real-world networks in polynomial-time. When infection and healing rates are extreme, the topology of the underlying network is unimportant. For the appropriate range of dynamics parameters, however, topology becomes an important factor in determining vulnerability; the macroscopic characterization of network vulnerability is not representative of the microscopic vulnerability of the individual agents. We show, using the scaled SIS

process, that, when the contagion rate is low, agents with more neighbors have higher probability of being infected at equilibrium. When the contagion rate is high, local characterizations such as the node degree is not sufficient; agent vulnerability depends on membership in dense subgraphs in the network.

Finally, we use the scaled SIS process to study other network process models such as the extended contact process as well as edge-centric network processes like the dynamic bond percolation process.

Keywords: complex networks, network processes, contact process, ϵ -SIS process, scaled SIS process, epidemics, cascading failures, contagion, SIS epidemics, continuous-time Markov process, network topology, graph density, dense subgraph, k -regular graphs, complete multipartite graphs, complete multipartite with k -regular islands, densest subgraph, maximum independent set, Perturb-and-MAP, loopy belief propagation, Gibbs distribution, equilibrium distribution, Markov network, undirected graphical model, Ising model, ground state, Max-Flow/Min-Cut, submodularity, maximum a posteriori, dynamic bond percolation process, maximum matching, motifs, marginal probability

Acknowledgements

The completion of this thesis was truly a group effort. This would not have been possible without the effort and guidance of my advisor, Prof. José M. F. Moura, whom I admire greatly for both the depth and breadth of his knowledge; I come away from every meeting with more. I will always remember the meeting we had where Prof. Moura held up my paper draft and said that ‘this is art.’ It was no masterpiece but I did slowly, painfully drag that paper from somewhere in my head and typed it into LaTeX.

I want to thank my committee members Prof. Sharad Goel, Prof. Pulkit Grover, Prof. Soumya Kar, and Prof. Jürgen Pfeffer for taking on the tasks of helping me with my research. I am also indebted to Prof. João P. Costeira and Prof. João M.F. Xavier of the Instituto Superior Técnico in Lisbon, Portugal, Prof. Cristopher Moore at the Santa Fe Institute, and Dr. Vinay Uday Prabhu for various technical helps and to Carolyn Patterson for assisting me with my travels and reimbursements.

And where would I be without my amazing friends and lab mates. Team Porter Basement B level and the lunch group! I learned so much from each and everyone: (Please remember me when you are all rich and famous!) Kyle Anderson, Nikos Arechiga, Kyri Baker, Rohan Chabukswar, Yuan Chen, Subhro Das, Joya Deri, Liangyan Gui, Joel Harley, Andrew Hsu, JY Joo, Stephen Kruzick, Jonathan Mei, Javad Mohammadi, Sergio Pequito, Nipun Popli, Anit Sahu, Augusto Santos, Aurora Schmidt, Evgeny Toropov, and so many more. And thanks to my friend Smitha Prasadh for giving me a place to stay these couple of months.

This thesis would also not have been possible without someone making it possible for me to eat. I am grateful for grants from Air Force Office of Scientific Research (FA95501010291) and National Science Foundation (CCF1011903 and CCF1018509) for sponsoring my research. I am also thankful for the Microsoft Azure Research Fellowship, which I used for many of the large scale computations presented in the thesis.

I am also grateful for my parents for planting the seeds for my love for learning ever since I was a child and for teaching me perseverance.

Contents

List of Figures	10
List of Tables	13
1 Introduction	15
1.1 Previous Work	17
1.2 Thesis Contributions	18
1.3 Thesis Outline	20
1.4 Mathematical Notations	22
2 Background	25
2.1 Epidemics	25
2.2 Graph Theory and Network Science	27
2.2.1 Neighborhood and Degree	28
2.2.2 Subgraphs	29
2.2.3 Graph Isomorphism	29
2.2.4 Graph and Network Models	29
2.3 Markov Process	37
2.3.1 Reversibility	38
2.3.2 Network Process as a Markov Process	39
2.4 Conclusion	39
3 Modeling Epidemics on Networks: Scaled SIS Process	41
3.1 Introduction	41
3.2 Scaled SIS Process	42
3.3 Continuous-Time Markov Process	45
3.4 Equilibrium Distribution	47
3.5 Parameter Regimes	50
3.6 Conclusion	52
4 Most-Probable Configuration	55

4.1	Introduction	55
4.2	Most-Probable Configuration Problem	56
4.3	Regime I) Healing Dominant & IV) Infection Dominant	57
4.4	Regime III) Exogenous Infection Dominant	57
4.5	Regime II) Endogenous Infection Dominant	61
4.6	Structured Network Topology	62
4.6.1	k -Regular Graph	62
4.6.2	Complete Bipartite Network	64
4.6.3	Complete Multipartite	65
4.6.4	Complete Multipartite with k -regular islands	67
4.6.5	Unstructured Networks	67
4.7	Ordering of Probability of \mathbf{x}^*	69
4.8	Conclusion	71
4.9	Appendix	72
5	Most-Probable Configuration: Non-Degenerate Solutions	81
5.1	Introduction	81
5.2	Submodular Function	81
5.2.1	Most-Probable Configuration: A Submodular Problem	82
5.2.2	Relationship to Ground State	85
5.3	Examples Using Real-World Networks	85
5.3.1	Computation Time	86
5.4	Induced Subgraphs and Graph Density	90
5.4.1	Equilibrium Distribution and Induced Subgraphs	91
5.4.2	Most-Probable Configuration and Induced Subgraphs	92
5.4.3	Most-Probable Configuration and the Densest Subgraph	95
5.5	Uniqueness of the Most-Probable Configuration	97
5.6	Regime II) Endogenous Infection Dominant	98
5.6.1	Subgraph Uniqueness	98
5.6.2	Uniqueness	99
5.7	Conclusion	100
5.8	Appendix	101
6	Characterizing Individual Susceptibility in Polynomial-Time	113
6.1	Introduction	113
6.2	Marginal Probability of Infection	114
6.3	Approximating Marginal Probability of Infection Using Sampling	114
6.3.1	Order-1 Perturb-and-MAP Accuracy	116
6.3.2	Comparing Perturb-and-MAP with Loopy Belief Propagation	121
6.4	Dependence of $P(x_i = 1)$ on Parameters and Topology	123
6.4.1	Simplest Scenario: $\gamma = 1$	123
6.4.2	Complex Scenario: $\gamma > 1$	124
6.4.3	Relating Marginal Probabilities to Centrality Measures	131
6.5	Conclusion	139

7	Expected Fraction of Infected Agents	141
7.1	Introduction	141
7.2	Relating Expected Fraction of Infected Agents to Marginal Probabilities	142
7.3	Dependence of $E[Y]$ on Parameters and Topology	143
7.4	Conclusion	151
8	Scaled SIS Process and the Contact Process (ϵ-SIS process)	153
8.1	Introduction	153
8.2	Contact Process	154
	8.2.1 Extended Contact Process	155
	8.2.2 Scaled SIS Process vs. Extended Contact Process	156
8.3	Equilibrium Behavior of the Extended Contact Process	157
8.4	Experimental Simulations	159
	8.4.1 Setup	159
	8.4.2 Results: $\pi_e(\mathbf{x})$ and $\pi_{\text{approx}}(\mathbf{x})$	160
	8.4.3 Results: TVD vs. Δ and $\frac{\Delta}{\mu}$	163
8.5	Most-Probable Configuration	166
8.6	Real-World Network	169
8.7	Conclusion	173
8.8	Appendix	173
9	Dynamic Bond Percolation Process	177
9.1	Introduction	177
9.2	Dynamic Bond Percolation Process	178
9.3	Equilibrium Distribution	181
	9.3.1 Reversibility and Equilibrium Distribution	183
	9.3.2 (SUD) Sum-Dependent Dynamic Bond Percolation Process	185
	9.3.3 (POD) Product-Dependent Dynamic Bond Percolation Process	185
9.4	At-Risk Edges and the Most-Probable Network Problem	187
9.5	Regime I) Recovery Dominant and Regime IV) Removal Dominant	188
9.6	Regime II) of the Most-Probable Network Problem	189
	9.6.1 SUD vs. POD Most-Probable Network	192
9.7	Regime III) of the Most-Probable Network Problem	194
	9.7.1 SUD Model and Maximum Matching	195
	9.7.2 POD Model and Maximum Star Matching	197
9.8	Conclusion	197
9.9	Appendix	198
10	Thesis Summary and Future Work	201
10.1	Future Works	204
	Bibliography	207

List of Figures

2.1	<i>k</i> -Regular Examples	30
2.2	Instantiations of Erdős-Rényi Graphs	31
2.3	Instantiations of Watts-Strogatz Graphs	31
2.4	Social Network of Drug Users ($N = 193, E = 273$)	33
2.5	US Western Power Grid ($N = 4941, E = 6595$)	34
2.6	Largest Component of Protein Network ($N = 1458, E = 1948$)	35
2.7	Sample of Facebook Network ($N = 4039, E = 88234$)	36
3.1	Examples of Healing Transitions (Blue = Infected, Red = Healthy)	43
3.2	Examples of Infection Transitions (Blue = Infected, Red = Healthy)	45
4.1	Example Network Topologies	68
4.2	Solution Space of the Most-Probable Configuration Problem	68
4.3	Non-Degenerate Most-Probable Configurations (Blue = Infected, Red = Healthy)	69
5.1	Most-Probable Configuration \mathbf{x}^* under Different $\left(\frac{\lambda}{\mu}, \gamma\right)$ Parameters (Blue = Infected, Red = Healthy)	88
5.2	Most-Probable Configuration \mathbf{x}^* under Different $\left(\frac{\lambda}{\mu}, \gamma\right)$ Parameters (Blue = Infected, Red = Healthy)	89
5.3	(a) Configuration $\mathbf{x}_1 = [0, 1, 1, 1, 0, 0, 0]^T$, (b) Induced Subgraph $F(\mathbf{x}_1) = F_1$	90
5.4	(a) Configuration $\mathbf{x}_2 = [0, 0, 0, 0, 1, 1, 1]^T$, (b) Induced Subgraph $F(\mathbf{x}_2) = F_2$	91
5.5	Most-Probable Configuration \mathbf{x}^* under Different $\left(\frac{\lambda}{\mu}, \gamma\right)$ Parameters (Blue = Infected, Red = Healthy)	95
5.6	Most-Probable Configuration \mathbf{x}^* under Different $\left(\frac{\lambda}{\mu}, \gamma\right)$ Parameters (Blue = Infected, Red = Healthy)	96
6.1	Small Networks	117
6.2	Network A: Error vs. Number of Samples	117
6.3	Network B: Error vs. Number of Samples	118
6.4	Network C: Error vs. Number of Samples	119

6.5	Network A: Error vs. Dynamics Parameters (Y-axis = $\frac{\lambda}{\mu}$, X-axis = γ)	119
6.6	Network B: Error vs. Dynamics Parameters (Y-axis = $\frac{\lambda}{\mu}$, X-axis = γ)	120
6.7	Network C: Error vs. Dynamics Parameters (Y-axis = $\frac{\lambda}{\mu}$, X-axis = γ)	120
6.8	Network A: $\frac{\lambda}{\mu} = 0.1, \gamma = 6, \text{error} = 0.06$	121
6.9	Network B: $\frac{\lambda}{\mu} = 0.1, \gamma = 6, \text{error} = 0.05$	122
6.10	Network C: $\frac{\lambda}{\mu} = 0.15, \gamma = 2, \text{error} = 0.07$	122
6.11	$P(x_i = 1), \forall i = 1, \dots, N$ vs $0 < \frac{\lambda}{\mu} \leq 1$ when $\gamma = 1$	125
6.12	$\hat{P}(x_i = 1)$ of the US Western Power Grid	126
6.13	$P(x_i = 1)$ of Network A	127
6.14	Network A: $P(x_i = 1)$ as a Function of $\gamma, \left(\frac{\lambda}{\mu} = 0.6895\right)$	128
6.15	Network A: $P(x_i = 1)$ as a Function of $\gamma, \left(\frac{\lambda}{\mu} = 0.1421\right)$	128
6.16	Network A: $P(x_i = 1)$ as a Function of $\gamma, \left(\frac{\lambda}{\mu} = 0.00025\right)$	129
6.17	Network A: $P(x_i = 1)$ as a Function of $\frac{\lambda}{\mu} (\gamma = 2.5053)$	129
6.18	Network A: $P(x_i = 1)$ as a Function of $\frac{\lambda}{\mu}, (\gamma = 4.8474)$	130
6.19	Network A: $P(x_i = 1)$ as a Function of $\frac{\lambda}{\mu}, (\gamma = 7.6579)$	130
6.20	Network A: $P(x_6 = 1)$ and $P(x_{10} = 1)$	130
6.21	Network A: Centrality and Correlation	134
6.22	Network A: Centrality and Correlation	135
6.23	US Western Power Grid: Centrality and Correlation	137
6.24	US Western Power Grid $P(x_i = 1)$ and Most-Probable Configuration when $\left(\frac{\lambda}{\mu} = 0.1, \gamma = 6\right)$	138
7.1	$E[Y]$ for Different $\frac{\lambda}{\mu}$ and γ for Four Networks	145
7.2	Projection of $E[Y]$ Onto $\frac{\lambda}{\mu}$ -Axis	147
7.3	Projection of $E[Y]$ Onto $\frac{\lambda}{\mu}$ -Axis	148
7.4	Projection of $E[Y]$ Onto γ -Axis	149
7.5	Projection of $E[Y]$ Onto γ -Axis	150
8.1	Agent V_3 Becomes Infected	157
8.2	Different Network Topologies with Different Maximum Degree	161
8.3	$\pi_e(\mathbf{x})$ and $\pi_{\text{approx}}(\mathbf{x})$ when $\frac{\lambda}{\mu} = 0.7, \Delta = 0.0023$	162
8.4	$\pi_e(\mathbf{x})$ and $\pi_{\text{approx}}(\mathbf{x})$ when $\frac{\lambda}{\mu} = 0.7, \Delta = 1.0496$	164
8.5	Dependence of $\text{TVD}(\pi_e, \pi_{\text{approx}})$ on Δ	165
8.6	Most-Probable Configuration when $\frac{\lambda}{\mu} = 0.9744$ and $\Delta = 0.02$ (Blue = Infected, Red = Healthy)	168
8.7	Extended Contact Process $\frac{\lambda}{\mu} = 0.9$ and $\gamma_e = 0.0387$	171
8.8	Extended Contact Process $\frac{\lambda}{\mu} = 0.9$ and $\gamma_e = 0.0585$	172
9.1	Maximal network represented by \mathbf{A}_{max} . Dashed edges (i.e., bonds) are the only possible edges in the network.	178
9.2	Invalid network state. Solid edges are <i>closed</i> . Dashed edges are <i>open</i> .	179
9.3	Example evolution of $\mathbf{A}(t)$. Solid edges are <i>closed</i> . Dashed edges are <i>open</i> .	179

9.4	Different edge removal scenarios. Solid edges are <i>closed</i> . Dashed edges are <i>open</i>	181
9.5	Example configurations \mathbf{A} and $T_{i,j}^- \mathbf{A}$. Solid edges are <i>closed</i> . Dashed edges are <i>open</i> . .	181
9.6	Example configurations \mathbf{A} and $T_{i,j}^+ \mathbf{A}$. Solid edges are <i>closed</i> . Dashed edges are <i>open</i> . .	181
9.7	P_3 graph	182
9.8	C_3 graph	182
9.9	P_4 graph	182
9.10	S_6 star graph	183
9.11	Matching and Star Matching	184
9.12	Most-Probable Network SUD (solid edges = closed, dashed edges = open)	192
9.13	Most-Probable Network POD (solid edges = closed, dashed edges = open)	193
9.14	Walks of Length 3 that are not Paths of Length 3: $i, e_1, k, e_2, i, e_3, j$	199
9.15	Walks of Length 3 that are not Paths of Length 3: $i, e_1, j, e_2, k, e_3, j$	199
9.16	Walks of Length 3 that are not Paths of Length 3: $i, e_1, j, e_3, i, e_2, j$	199

List of Tables

3.1	Parameter Regimes	51
5.1	Computation Time of the Most-Probable Configuration Problem	87
8.1	Most-Probable Configuration when $\frac{\lambda}{\mu} = 0.9744$ and $\Delta = 0.02$	167
8.2	Most-Probable Configuration when $\frac{\lambda}{\mu} = 0.7$ and $\Delta = 0.4333$	169

Introduction

Interaction characterizes many real-world phenomena. Infection spreads from infectives hosts to susceptible as a result of contacts between agents leading to epidemics. Similarly, rumors or information spread in a social network due to the interaction between a communicator and an adopter such as the ‘viralization’ of Internet memes. It is also often observed in networked systems that local failures of a few components can lead to cascading failures like blackouts in the power grid.

Models are useful for studying these phenomena because it is difficult to conduct experiments or make measurements on these types of complex, interactive systems due to their size and scale. While these processes can be simulated on a computer, simulations lack explanatory ability [1]. It is not feasible to simulate every possible scenario: What if a group of agents form a community amongst themselves? What if the population is larger than expected? What if healthy agents interact with infectives more often (i.e., the infection rate is high)? What if failed components recover faster? How will these changes affect the behavior of the process?

To answer these questions, a new type of models needs to be considered. Traditional models, such as the SIS (susceptible-infected-susceptible) compartmental model in epidemiology, assume homogenous-mixing, which means an agent may contact all other agents in the population [2]. In reality, homogenous-mixing is often not possible since agents may have a limited number of contacts; some infectives have more contacts with susceptible agents and are therefore super-spreaders [1, 3]. Networks can account for both the heterogeneity and finiteness of contacts [1]. In this respect, network process models, which assume that the structure of interactions amongst

multiple agents is described by a graph, are more realistic than traditional models.

For network processes, it is of interest to understand both the impact of dynamical parameters such as contagion rate and healing rate as well as the impact of the network structure on the dynamics of the process. The challenge is in developing feasible models that incorporate network structure as part of the model description; networks, being discrete structures, introduce combinatorial complexity in the analysis. This thesis presents and studies the **scaled SIS process** to model the stochastic, microscopic interactions between individual agents in a population, assuming healthy agents can become infected due to contagion from neighbors and that infected agents can heal. Unlike many existing models, the scaled SIS process accounts for the exact underlying network topology without resorting to approximations or assuming simpler network structures [3,4].

The scaled SIS process assumes that the infection rate is exponentially dependent on the number of infected neighbors (i.e., contacts) whereas previous network process models such as the contact process [5,6] or the extended contact process (also known as the ϵ -SIS process) [7] consider the infection rate to be linearly dependent on the number of infected neighbors. It is possible for the infection rate to be nonlinear for some biological mechanisms [8]. The infection rate may be superlinear due to secondary effects such as nutritional stress. The assumption of exponential dependence that we make in this thesis is also one of convenience. We prove in the thesis that the steady-state behavior of the scaled SIS process can be characterized with a closed-form equilibrium distribution whereas the extended contact process can not.

Chapter 8 discusses how the scaled SIS process can also be used to study models with linear dependence. This follows from the binomial expansion of exponential functions; under small perturbations, the exponential rate of the scaled SIS process well-approximates the linear rate of the extended contact process. As a result, the equilibrium behavior of the extended contact process can be well approximated by the equilibrium behavior of a scaled SIS process.

With the scaled SIS process, we can address, for different network structures and different dynamics parameters, questions such as which individuals or communities in a network are the most vulnerable, what is the expected fraction of infected agents or failed agents. These questions can be answered quantitatively by performing inference on the equilibrium distribution. Although many inference problems are in general computationally NP-hard, for the scaled SIS process, we can solve them in polynomial-time for a range of parameter values of practical interest. This

thesis provides the first model for which the microscope, mesoscopic and macroscopic behaviors of networks processes over exact, arbitrary networks can be analyzed in polynomial-time; the results, which depend on subgraphs in the network, on the preference of individual agents, and on the strength of the coupling effects, show how complex network processes can be.

1.1 Previous Work

Networks appear in many domains. The potential of applying networks to describe relations among multiple elements in a system makes the study of networks popular in many different disciplines from social networks [9, 10], to power grids [11, 12], and to contact networks in epidemics [13–15]. Datasets of real-world networks include the KONECT database [16] and the SNAP database [17].

The study of networks is often related to the study of graphs, and understanding graph theory is fundamental to networks. Network science differs from traditional graph theory in that networks are often derived either from measurement or by design from real-world systems. These networks generally do not conform to standard graph classes. There exists a large corpus of literature on the statistical characterization of large-scale networks, particularly, the study of how the topology of the network reflects certain characteristics of the system [18, 19]. Much of this work is algorithmic in nature and focus in finding meaningful substructures such as communities or core structures [20, 21].

Network processes extend network science by considering dynamical processes on networks; the network is only one of the parameters of network processes. Network processes model the evolution of the states of agents in networks rather than the evolution of the population as a whole. The simplest model of network processes assumes that the state of each agent is binary and the transition from one state to another is dependent on the states of the agent’s neighbors in the network. There are many realizations of such models [3, 22–24], with various assumptions regarding for example topological dependencies. The common challenge these models face is the analytical difficulty in accounting for the network structure. The topology couples each agent’s evolution to the evolution of their neighbors, and the neighbors’ are coupled to their neighbors’ evolutions. As the networks of most interest in application are large (thousands or millions of nodes), keeping track of the evolution dynamics of all the individual agents becomes an intractable problem.

One approach, favored in statistical mechanics, to overcome this challenge is to assume that the number of agents, N , in the network grows to infinity, $N \rightarrow \infty$. This leads, under appropriate assumptions, to the **mean-field approximation** of the network process dynamics; by assuming the mean-field approximation, the dynamics of all the agents is summarized by a simpler set of ODEs (ordinary-differential equations) [3, 15, 25, 26]. Alternatively, many researchers assume that the agent states evolve stochastically and model the network process as Markov [6, 7, 27–29]. This approach is, in general, intractable for large networks since the number of states in the Markov process is exponential in the size, N , of the network. Assuming that each agent can only be in one of two states, a network with 30 agents results in a network process with 2^{30} , or roughly, 1 billion possible configurations. As a result, researchers focus on specific, often simple, network topologies such as the complete graph, which yield a reduced number of states through symmetry.

1.2 Thesis Contributions

This thesis considers the problem of modeling and analyzing network processes over arbitrary, static, undirected network topologies. We developed the scaled SIS process, which models the dynamics of the network process as a continuous-time Markov process. The advantage of the scaled SIS process over other similar models is that it is also a reversible Markov process, and we can derive its closed-form equilibrium distribution while accounting for arbitrary network structures. The equilibrium distribution characterizes the probabilistic behavior of a network process at equilibrium (i.e., the time-asymptotic behavior of the process). The expression of the equilibrium distribution shows that the sufficient statistics of the scaled SIS process at equilibrium are: the number of infected nodes and the number of edges where both end nodes are infected. Furthermore, the equilibrium distribution is decomposable into two terms: a topology-independent process, which accounts for the preference of individual agents, and a topology-dependent process, which determines if the network facilitates or impedes diffusion.

Vulnerability of the population to the epidemics can be characterized at different scales: at the **microscopic** scale, where we are interested in the vulnerability of individual agents; at the **mesoscopic** scale, where we are interested in finding the most vulnerable communities in the network; and at the **macroscopic** scale, where we want to characterize the susceptibility of the entire population. We show that only network processes models, like the scaled SIS process that

we study, can give analytical insights of the process at the microscopic and mesoscopic scale since they incorporate structure.

We study the mesoscopic vulnerability of the process by solving for the configuration with maximum equilibrium probability. For a given set of parameter values that we refer to as regime III) **Exogenous Infection Dominant**, we prove that structural vulnerability is related to the Maximum Independent Set Problem of the underlying graph [30]. In another parameter regime, one that is of more practical interest, regime II) **Endogenous Infection Dominant**, the more vulnerable communities are related to dense subgraphs in the network. The size of these vulnerable subgraphs shows nonlinear dependence on the dynamics parameters. Furthermore, although this is a combinatorial problem, we show that it can be solved exactly using polynomial-time algorithm due to the property of submodularity.

We characterize the microscopic vulnerability by solving for the marginal probability of infection for individual agents. The problem is NP-hard because the sample space of the equilibrium distribution grows exponentially with the size of the network. We use Monte Carlo sampling and show that it can effectively approximate the marginal probability of infection. Agents that have high probability of infection for one set of dynamical parameters may have low probability of infection for another set of parameter values. Depending on the dynamical parameters, it is *not* always the agents with large number of contacts that are more vulnerable to infection. The scaled SIS process is the *first* network process model for which we can present analytical results at this level of detail for arbitrary network topology.

Solving for the expected fraction of infected agents describes the vulnerability of the entire network to the epidemics. We prove that this macroscopic statistics is the average of the marginal probability of infection. Consequently, the expected fraction of infected agents can also be approximated effectively for large networks using Monte Carlo sampling methods.

The scaled SIS can be used to gain insights into other network process models. We consider two such models—the extended contact process and the dynamic bond percolation process. For the extended contact process [7,31], we prove the conditions under which, the equilibrium distribution of extended contact process can be described by an equivalent scaled SIS process. The dynamic bond percolation process considers the scenario where instead of infection or failures spreading from node to node, infections or failures spread from edge to edge. As a result, structural vul-

nerability relating to edges instead of nodes can be addressed using the dynamic bond percolation process. Under certain assumptions, we can show that the dynamic bond percolation process is also a reversible Markov process and can derive its closed-form equilibrium distribution. The model is more complex than the scaled SIS process; its sufficient statistics are related the number of small subgraphs in the network (i.e., motifs [32]).

1.3 Thesis Outline

The chapters are organized with their own introduction and conclusion. Some proofs are relegated to an appendix at the end of a chapter. Some results in the thesis have been published in [31, 33–39].

Chapter 2 introduces general background information on the theoretical foundations of the thesis: epidemics models, graph theory and network science, network processes and continuous-time Markov processes. The scaled SIS process is formally presented in chapter 3, which details the underlying assumptions. We then prove that the process is a reversible Markov process and derive in section 3.4 the closed-form equilibrium distribution of the scaled SIS process. In section 3.5 we introduce the 4 parameter regimes, corresponding to four different network process behaviors. Chapters 4 through 7 are related to solving various inference problems of the equilibrium distribution.

Chapters 4 and 5 discuss the most-probable configuration (i.e., the configuration with the maximum equilibrium probability). The Most-Probable Configuration Problem is introduced and defined in chapter 4. We then analyze the solution space of the Most-Probable Configuration Problem for all 4 parameter regimes to find the vulnerability of the network to infection or failure at the mesoscopic scale. In section 4.4, we prove, that for regime III) **Exogenous Infection Dominant**, the most-probable configuration is related to the Maximum Independent set. This means that solving for the most-probable configuration is NP-hard for general networks. In section 4.5, we prove that, for structured network topologies such as k -regular, complete multipartite, or complete-multipartite with k -regular islands, for regime II) **Endogenous Infection Dominant**, the solution of the Most-Probable Configuration Problem is either the configuration where all the agents are healthy or the configuration where all the agents are infected. For nonstructured topologies, the solution of the Most-Probable Configuration Problem in regime II) may be configurations where only

subsets of agents are infected. These non-degenerate most-probable configurations are discussed in detail in chapter 5.

Chapter 5 first shows that the Most-Probable Configuration Problem can be solved in polynomial-time in regime II). This is because within this range of parameter values, the combinatorial optimization function is submodular. Regardless of the underlying network structure or the degeneracy of the solution, the most-probable configuration can be found efficiently. We show that the most-probable configuration of a 4941 node network can be found in less than 1 sec on a standard desktop. The chapter proves that the non-degenerate most-probable configurations in unstructured networks are due to the existence of subgraphs denser than the overall network. There may be many non-degenerate solutions depending on these denser subgraph structures of the network. Finally, the chapter discusses the uniqueness of the most-probable configuration. We show that non-degenerate configurations are subgraph unique in that there are no other subgraphs isomorphic to subgraphs induced by non-degenerate most-probable configurations.

Chapters 6 and 7 discuss the vulnerability of the network to infection or failure at the microscopic and macroscopic scale, respectively. We study these vulnerabilities by discussing inference problems that are NP-hard but whose solution can be effectively approximated for regime II) **Endogenous Infection Dominant**. Chapter 6 studies vulnerability at the microscale through the marginal probability of infection of individual agents. It shows that this probability can be found using the Monte Carlo sampling technique known as Perturb-and-MAP [40]. The accuracy and convergence of the approximation is tested with small network examples. We show that Perturb-and-MAP sampling outperforms Loopy Belief Propagation [41]. Section 6.4.3 analyzes when the structural characteristics of each agent, such as its degree, is a reliable indicator of its vulnerability to infection by computing the correlation coefficients between the marginal probability with various nodal centrality measures.

Chapter 7 studies macroscopic vulnerability by computing the expected fraction of infected agents or failed components at equilibrium. The chapter shows that the expected fraction of infected agents is the average of the marginal probabilities of infection. Using approximations computed with Perturb-and-MAP sampling, we plot the nonlinear dependence of the expected fraction of infected agents with respect to the dynamics parameters.

Chapters 8 and 9 analyze two other network processes related to the scaled SIS process. Chap-

ter 8 considers the extended contact process, which also models SIS epidemics as a continuous-time Markov process, but with linear instead of exponential dependencies [6, 7]. However, the equilibrium distribution can only be solved numerically, therefore making the model intractable for large networks. We prove in section 8.3 that for a range of parameter values, the equilibrium distribution of the contact process with exogenous infection can be well approximated by that of an equivalent scaled SIS process. This allows us to gain insights on the long-term behavior of the extended contact process.

Chapter 9 presents and studies a new model, the dynamic bond percolation process, for which the states of edges, instead of nodes, change in time. It is similar to the scaled SIS process in that it exhibits exponential dependencies and is also a reversible Markov process with closed-form equilibrium distribution. This closed-form equilibrium distribution is more complex than the equilibrium distribution of the scaled SIS process. This is due to complication of defining neighborhood for edges; we show two possible methods of accounting for the total number of neighboring edges: SUD (sum-dependent) in section 9.3.2 and POD (product-dependent) in section 9.3.3. Section 9.4 discusses the computational complexities of finding the most-probable configuration for the dynamic bond percolation process.

Chapter 10 concludes the thesis with a summary of the thesis contributions and discussion of future directions.

1.4 Mathematical Notations

- $G(V, E)$, a simple, static, unweighted, undirected graph representing relationship between multiple agents
- V , set of nodes in a graph
- E , set of edges in a graph
- N , total number of nodes in a graph
- $|E|$, total number of edges in a graph
- $d(G)$, density of graph $G(V, E)$
- μ , healing rate
- λ , exogenous (i.e., spontaneous) infection rate
- γ , endogenous (i.e., contagion) infection rate

- \mathbf{A} , adjacency matrix of a graph
- $\mathbf{1}^T = [1, 1 \dots 1]^T$
- x_i , state of the i th agent in a network
- d_i , total number of neighbors of agent i
- m_i , total number of infected neighbors of agent i
- $\mathbf{x} = [x_1, x_2, \dots, x_N]^T$, the configuration of the state of all the agents in the network
- $\mathbf{x}^N = [1, 1, \dots, 1]^T$, configuration where all the agents are infected
- $\mathbf{x}^0 = [0, 0, \dots, 0]^T$, configuration where all the agents are healthy
- \mathbf{Q} , transition rate matrix of a continuous-time Markov process
- $q(i, j) = Q_{ij}$, transition rate of going from Markov state i to Markov state j
- \mathcal{X} , the set of all possible configurations/states of a continuous-time Markov process
- $\pi(\mathbf{x})$, equilibrium distribution of a continuous-time Markov process
- $H(\mathbf{x})$, the Hamiltonian of a probability distribution
- Z , partition function of a probability distribution
- \mathbf{x}^* , configuration with the maximum equilibrium probability (i.e., most-probable configuration)
- $F(\mathbf{x})$, subgraph induced by the configuration \mathbf{x}
- $d(F(\mathbf{x}))$, density of the subgraph induced by the configuration \mathbf{x}
- $F(\mathbf{x}^*)$, subgraph induced by the most-probable configuration
- \overline{F} , subgraph with the highest density
- $E[Y]$, expected fraction of infected agents at equilibrium
- $P(x_i = 1)$, probability that agent i is infected at equilibrium
- $\widehat{E}[Y]$, approximate expected fraction of infected agents at equilibrium
- $\widehat{P}(x_i = 1)$, approximate probability that agent i is infected at equilibrium
- K , total number of Monte Carlo samples
- \mathbf{x}_k , the k th sample configuration,
- $(\mathbf{x}_k)_i$, the state of the i th agent in the k th sample configuration

Background

This chapter reviews basic concepts from epidemics modeling, graph theory, and continuous-time Markov processes. Additional background information is in subsequent individual chapters and presented when needed. Due to the interdisciplinary nature of networks, different fields may refer to the same concept differently. For example, the forward equations of a Markov process are called master equations in physics. We will note this when appropriate. In section 2.1, we review the common concepts and terminologies used in epidemics modeling. Section 2.2 reviews the fundamentals of graph theory and network science such as degree, subgraphs, and graph models like Erdős-Rényi graphs. Markov processes and network processes are discussed in section 2.3.

2.1 Epidemics

Epidemics models study the progression of disease in a population. Two fundamental frameworks in mathematical epidemiology are the **susceptible-infected-susceptible (SIS)** model and the **susceptible-infected-recovered (SIR)** model.

The SIS model assumes that agents in the population can be in one of two states 1) **healthy and susceptible** or 2) **infected and contagious**. A healthy agent is infected due to contagion from an infection source; infected agents may heal and be reinfected at a later time. The SIR model on the other hand, assumes that once an infected agent has recovered, it can no longer be reinfected. An agent in the SIR model can be in one of three states 1) healthy but susceptible, 2) infected and contagious, or 3) recovered and no longer susceptible.

The generality of these basic frameworks make them suitable as first step approximations for

other diffusion phenomena in networks such as the **spread of information/memes in social networks** or **cascading failures in technological networks** such as the power grid or computer networks. In the former case, the infection process is akin to the spread of new ideas or opinions whereas the healing process is reverting back to a former opinion. Susceptible agents are would-be adopters and infected agents are spreaders. In the latter case, infection is like a component failure process whereas healing is the recovery of a failed component.

Traditional epidemics models often assume **full-mixing**, also called **random-mixing** or **homogenous mixing**, behavior [3, 18]. This means that an agent has an equal chance of interacting with all the other agents in the network. Without considering the behavior of individual agents, these models analyze the evolution of **macroscopic** quantities such as the fraction of infected agents. A more realistic assumption is **heterogeneous-mixing**, in that agents interact with only a subset of the population and/or have heterogeneous rates of interactions. The result of these interactions over time is a graph structure called the **contact network**. Network-based epidemics avoid the full-mixing assumption by accounting for the topology of the contact network as a parameter in the model. Further, network-based models often take into account the **microscopic** interactions between individual agents.

Epidemics models may be deterministic, in which case, the progression of the disease is modeled by differential equations of the proportion of healthy or infected agents in the population over time [3, 18]. The epidemics model may also be stochastic, in which case, progression is modeled as a stochastic process. In this case, the variations between infection time and healing time of different individuals is also accounted for. The model we present and study in this thesis further divides the infection into two different types depending on the source of infection. An infection is **exogenous** if the source of infection is external to the network; this may be due to the network being only an approximation of the population or if an agent spontaneously develops an infection or failure. An infection is considered **endogenous** if the infection originates from other agents in the network.

Studying network-based epidemics requires knowledge of both graph theory and dynamical processes, which we review briefly in the following sections.

2.2 Graph Theory and Network Science

Information on graphs can be found in [18,19,42,43]. A graph, $G(V, E)$, is defined by the set of **nodes** (also called vertices) V , and the set of **edges**, E . Graphs are well-studied in the mathematical field of graph theory. In the context of networks, the nodes often represent different agents or components under consideration whereas the set of edges represents relationships, contacts, or dependencies between agents. For a network with N agents, the cardinality of the set V is N . If an edge exists between node i and node j , we write $(i, j) \in E$.

When these dependencies exist only in one direction, then $G(V, E)$ is a **directed** graph (also called a **digraph**). If all the edges are bidirectional, then $G(V, E)$ is an **undirected** graph. If we associate different numerical values (weights) to the edges, the graph $G(V, E)$ is called a **weighted** graph. Graphs where all the edges are associated with the same or no numerical value are **unweighted** graphs. An unweighted, undirected graph is **simple** if there is only one edge between two different nodes i and j and no edge between a node and itself. If V and E remain unchanged, the graph is **static**. If V or E changes over time, the graph is **time-varying**.

A **path** is a sequence of edges in a graph that goes from some node i to a node j ; with the exception of the starting and ending node, all the other nodes in the path are distinct. A path where the starting and ending nodes are the same is a **cycle**. A **walk** is also a sequence of edges in a graph that goes from node i to node j . Unlike a path, a walk may have repeated nodes. A graph is **connected** if there is a path between any pair of nodes in V .

Graphs can also be represented as matrices. The **adjacency matrix**, $\mathbf{A} = [A_{ij}]$, of a simple graph $G(V, E)$ is a $N \times N$ matrix where

$$A_{ij} = \begin{cases} 1, & \text{if } (i, j) \in E \\ 0, & \text{otherwise.} \end{cases}$$

When the graph is undirected, then $A_{ij} = A_{ji}$ and the adjacency matrix, \mathbf{A} , is symmetric. The total number of edges can be found using the adjacency matrix:

$$|E| = \frac{1}{2} \mathbf{1}^T \mathbf{A} \mathbf{1},$$

where $\mathbf{1} = [1, 1, \dots, 1]^T$ is the N -length vector of ones.

2.2.1 Neighborhood and Degree

The neighborhood of node i is the set of nodes that are adjacent (i.e., directly connected) to node i :

$$\mathcal{N}_i = \{j \in V \mid (i, j) \in E\}.$$

These nodes are called the **neighbors** of node i . The total number of neighbors of node i is the **degree** of node i . The degree of node i in a graph is given by

$$d_i = \sum_{j=1}^N A_{ij},$$

where $\mathbf{A} = [A_{ij}]$ is the adjacency matrix that describes the graph. If the graph is directed, then the degree of a node can be further classified as the in-degree or the out-degree depending on if the edges are going in or exiting the node. The **degree distribution** is the collection of the degrees of all the nodes in the graph. It is the most commonly used statistical characterization of the structure of large networks.

The total number of edges in an undirected graph can also be found from the degree:

$$|E| = \frac{1}{2} \sum_{i=1}^N d_i.$$

The **density** [44] of a graph, $d(G)$, is

$$d(G) = \frac{|E|}{|V|}.$$

The density is equal to half the average degree since

$$d(G) = \frac{1}{2} \frac{\sum_{i=1}^N d_i}{N}.$$

An alternate definition for density (also called conductance) [45] of the graph is the number of edges divided by the maximum possible number of edges,

$$\frac{2|E|}{N(N-1)}.$$

A graph is considered a **dense graph** if the number of edges, $|E| \approx N^2$. A graph is considered a **sparse graph** if the number of edges, $|E| \approx N^\alpha$, $\alpha < 2$ [3].

2.2.2 Subgraphs

A graph $G'(V', E')$ is a **subgraph** of $G(V, E)$ if and only if all the nodes in V' are also in V and all the edges in E' are in E . The subgraph is an **induced subgraph** when it satisfies the additional condition that, for all $i, j \in V'$, the edge $(i, j) \in E'$ if and only if $(i, j) \in E$.

2.2.3 Graph Isomorphism

Two graphs, G and F , are equivalent if they are isomorphic: there is a bijection between the vertex sets of G and F , $f : V(G) \rightarrow V(F)$, such that any two vertices u and v of G are adjacent in G if and only if $f(u)$ and $f(v)$ are adjacent in F . This means that edge (u, v) is in G if and only if $(f(u), f(v))$ is in F . We denote the equivalence between two different graphs as

$$F \simeq G$$

2.2.4 Graph and Network Models

Graphs are divided into two major classes: tree and non-tree graphs. A graph is a **tree** if it does not contain any cycle. In addition, common classes of graphs that we will reference in the thesis are:

- **Complete graph** is a graph with each node connected to every other node in the network. The complete graph with N nodes is denoted by K_N . The degree of each node in a complete graph is $N - 1$. The total number of edges in a simple complete graph is then $\frac{N(N-1)}{2}$. The complete graph is the densest possible connected, simple graph.
- **Cycle graph** is a graph where each node is connected to two other nodes. The cycle graph with N nodes is denoted by C_N . The degree of each node in a cycle graph is 2. The total number of edges in a cycle graph is N .
- A graph is **k -regular graph** if and only if each node has degree k . The complete graph and the cycle graph are both examples of k -regular graphs, with $k = N - 1$ and $k = 2$, respectively.
- A graph is a **complete multipartite graph** if and only if the nodes can be partitioned into m sets, $V = \{V_1, \dots, V_m\}$, such that no nodes in the same set are connected; a node in set V_i

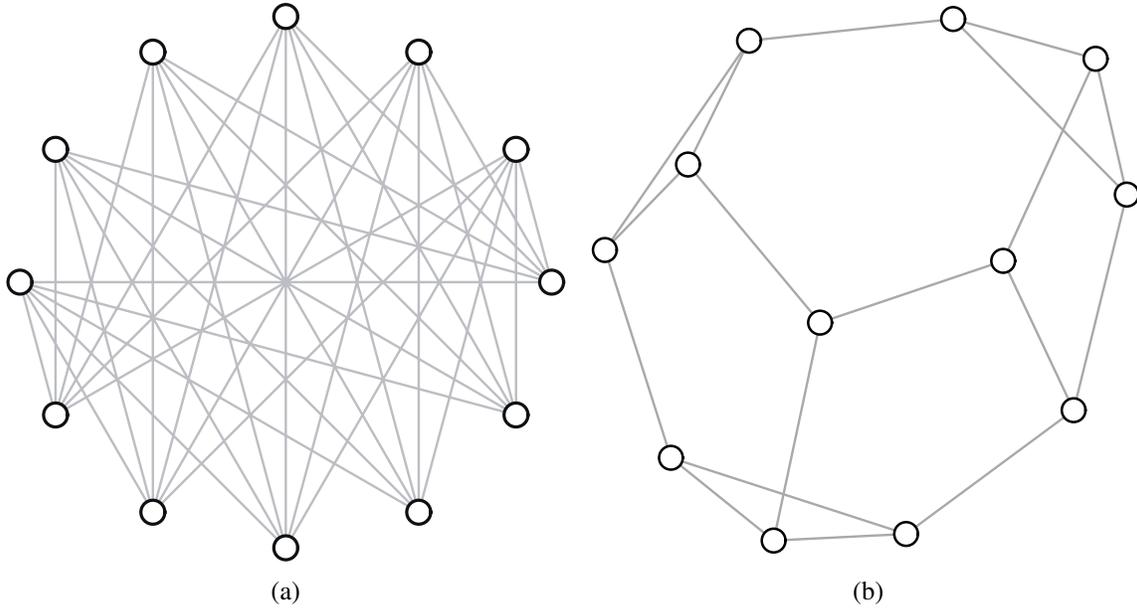


Figure 2.1: k -Regular Examples

is connected to all the other nodes in set V_j , $j \neq i$. The sets V_i are islands or partitions. A **complete bipartite graph** is partitioned into 2 islands.

- A graph is a **complete multipartite graph with k -regular islands** if and only if it is a complete multipartite graph and, within the i th island, each node is connected to d_i^{intra} other nodes in the same set. We will refer to d_i^{intra} as the *intra-degree* of a node. The *inter-degree* of a node is the total number of connections to nodes belonging to different sets. The total degree of a node in $G(V, E)$ is its intra-degree plus its inter-degree.

Graphs can also be specified probabilistically. Different instantiations of a **random graph** may have different structures but the statistical characteristics remain the same.

- The **Erdős-Rényi (ER)** [18] model is the most common random graph model. For a pair of nodes i, j out of the N nodes of the graph, the edge (i, j) forms with probability p , independently of all the other edges. It is referred as $G(N, p)$. An alternative formulation of the model is to choose uniformly at random a graph from the set of all possible graphs with N nodes and M edges. Figure 2.2 shows examples of Erdős-Rényi graphs.

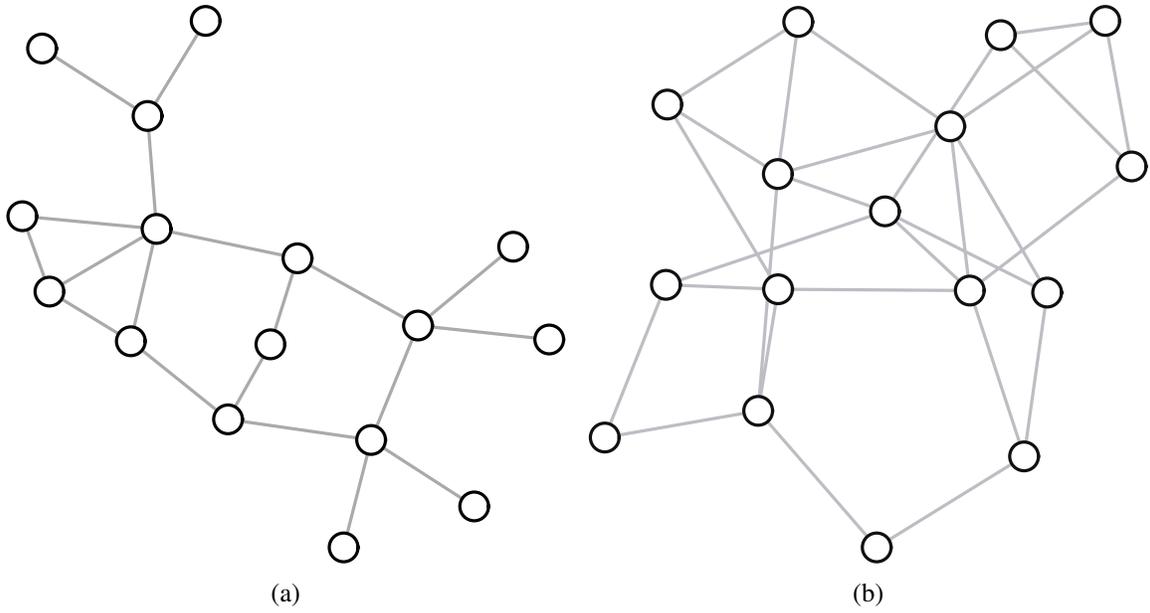


Figure 2.2: Instantiations of Erdős-Rényi Graphs

- The **Watts-Strogatz (WS)** [18] model is a random graph that demonstrates small average distance between nodes and high clustering. This is known as the small-world phenomena in networks. Figure 2.3 shows examples of Watts-Strogatz networks.

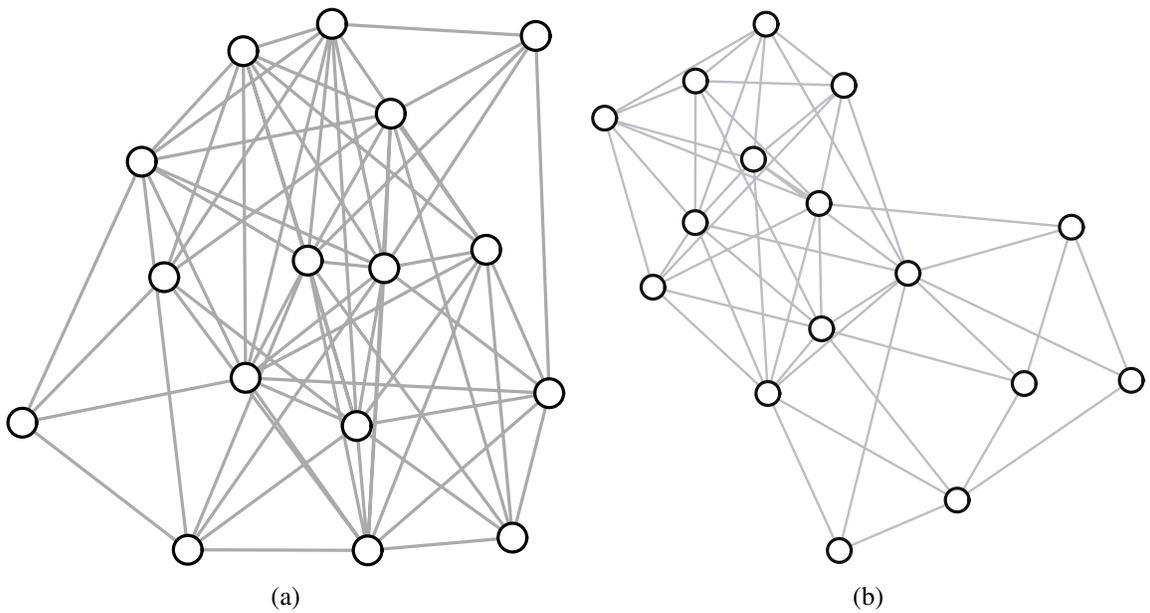


Figure 2.3: Instantiations of Watts-Strogatz Graphs

Real-world networks are derived from real-world systems. They can arise from measurements (e.g., biological networks or contact networks), or they describe engineered systems (e.g., computer networks, the power grid). Real-world networks do not in general conform to having nice deterministic or statistical characteristics. It has been pointed out, though subject to debate, that many real-world networks demonstrate **scale-free** degree distribution [3]. Databases of real-world networks are in [16,17]. These are some of the real-world networks that we are interested in studying in this thesis:

- Figure 2.4 shows a real-world social network of drug users. The 193 nodes represent different users. The network was determined through interviews of drug users in Hartford, CT. Reference [46] looked for influential agents in the network by considering it as a graph connectivity problem. However, they did not consider a dynamical model of influence. Assuming that we can model drug habits as an epidemic (i.e., there is a social contagion aspect to the behavior), we applied the scaled SIS process to this network and solved for the most-probable configuration under different parameters to find influential network structures.
- Figure 2.5 shows the 4941-node, 6595-edge network abstraction of the US Western power grid used by Watts and Strogatz in [12]. They showed through simulation of the SIR (susceptible-infected-removed) epidemics model on the western power grid that small-world networks like the western power grid are more conducive to spreading infection/failures than lattice networks. This is useful for explaining why failures propagate so quickly in a blackout. However, they can not identify *which* components in the power grid are more vulnerable to the epidemics with their approach. Here, we model the blackout as a SIS epidemic by assuming that failures and recoveries of grid components (e.g., power stations, substations, generators, switches, lines) are intermittent; a failed component may return to power, possibly failing again, as often happens in practice. Using the scaled SIS process, we can identify the most vulnerable substructures in the network.
- Figure 2.6 shows the largest connected component of a network of protein interactions in the yeast *Saccharomyces cerevisiae* [16]. The largest connected component contains 1458 nodes and 1948 edges. The nodes represent individual proteins; the edges are direct physical interaction between protein pairs as measured in experiments. The network was used in [47]

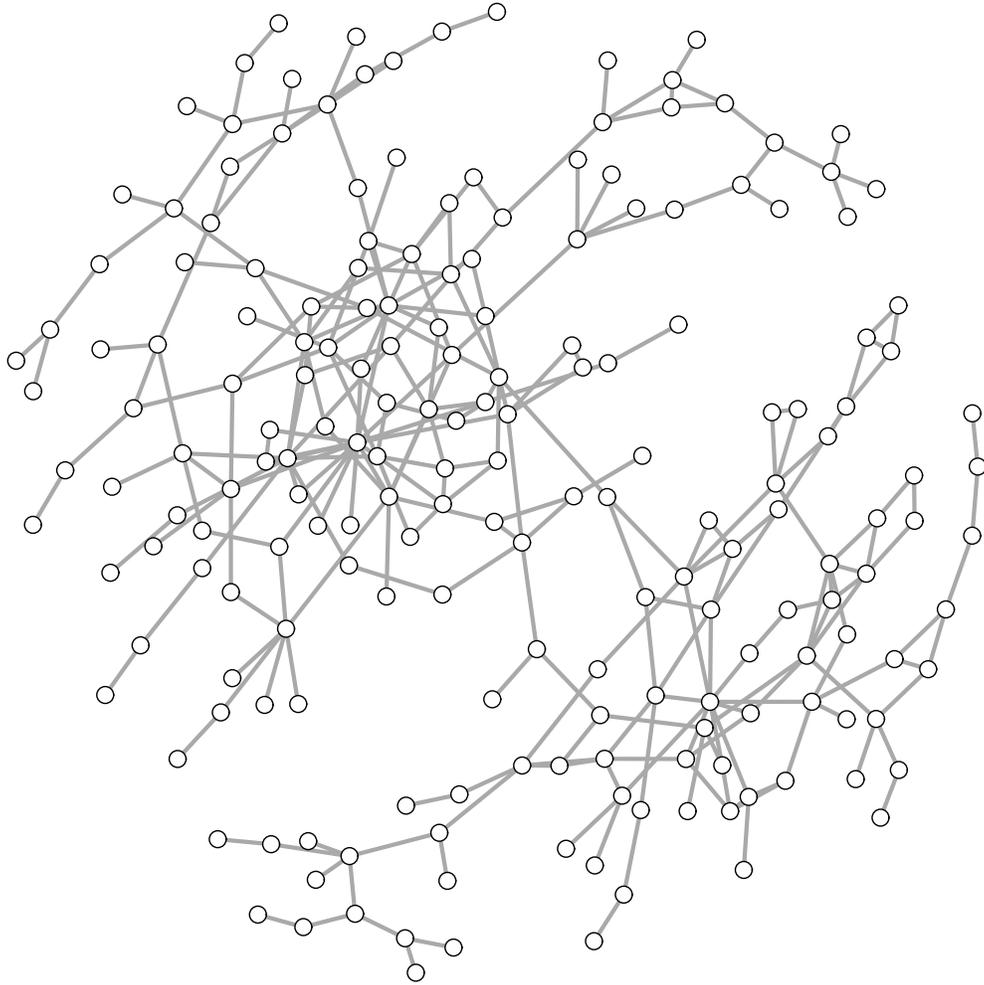


Figure 2.4: Social Network of Drug Users ($N = 193$, $|E| = 273$)

to demonstrate the connection between nodal centrality and protein lethality. The paper argued that nodes with high degree are more crucial to the function of the yeast (i.e., their removal is lethal to the yeast). The analysis has been called into question due to the measurement method with which the network was obtained. Reference [48] argued that studies and measurements artificially bias degree with lethality.

- Figure 2.7, obtained from [17], shows a small sample of the Facebook network. The nodes are Facebook users and edges denote relationships between users. Due to the large number of edges, the visualization of the networks itself becomes a complicated issue.

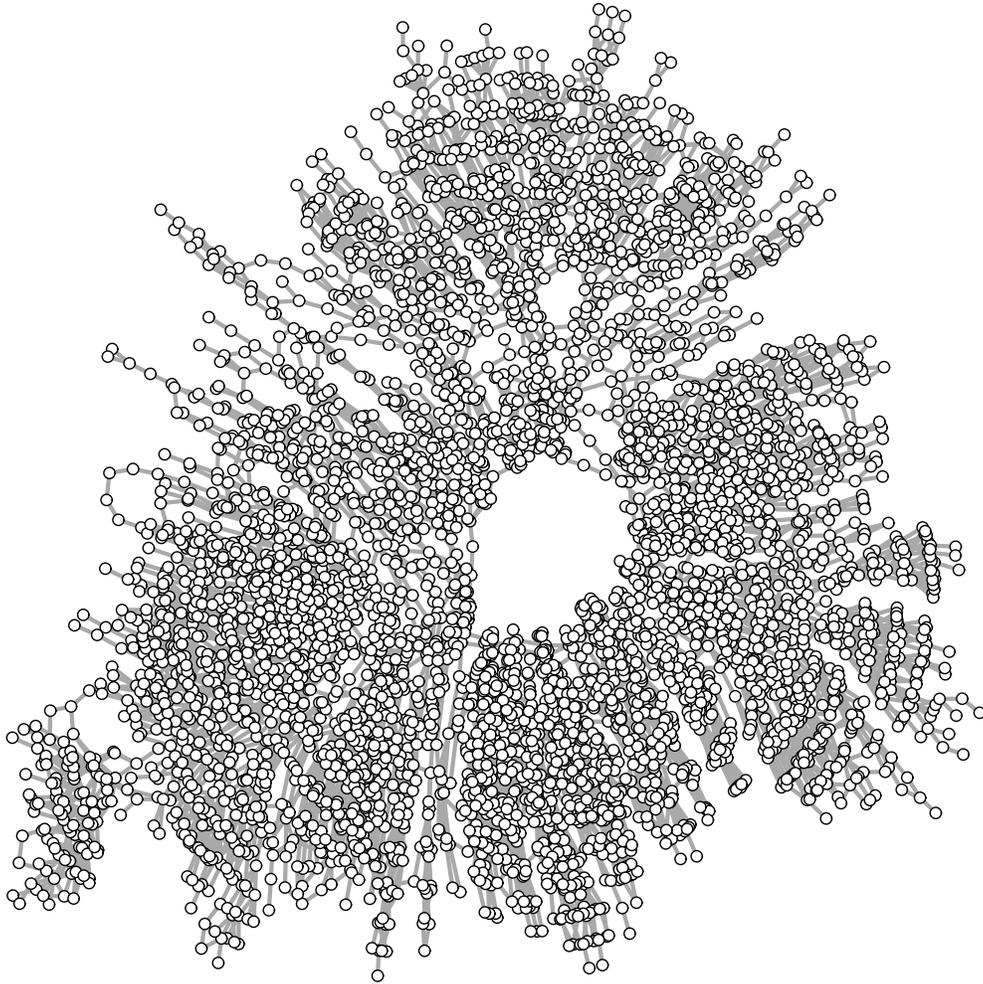


Figure 2.5: US Western Power Grid ($N = 4941$, $|E| = 6595$)

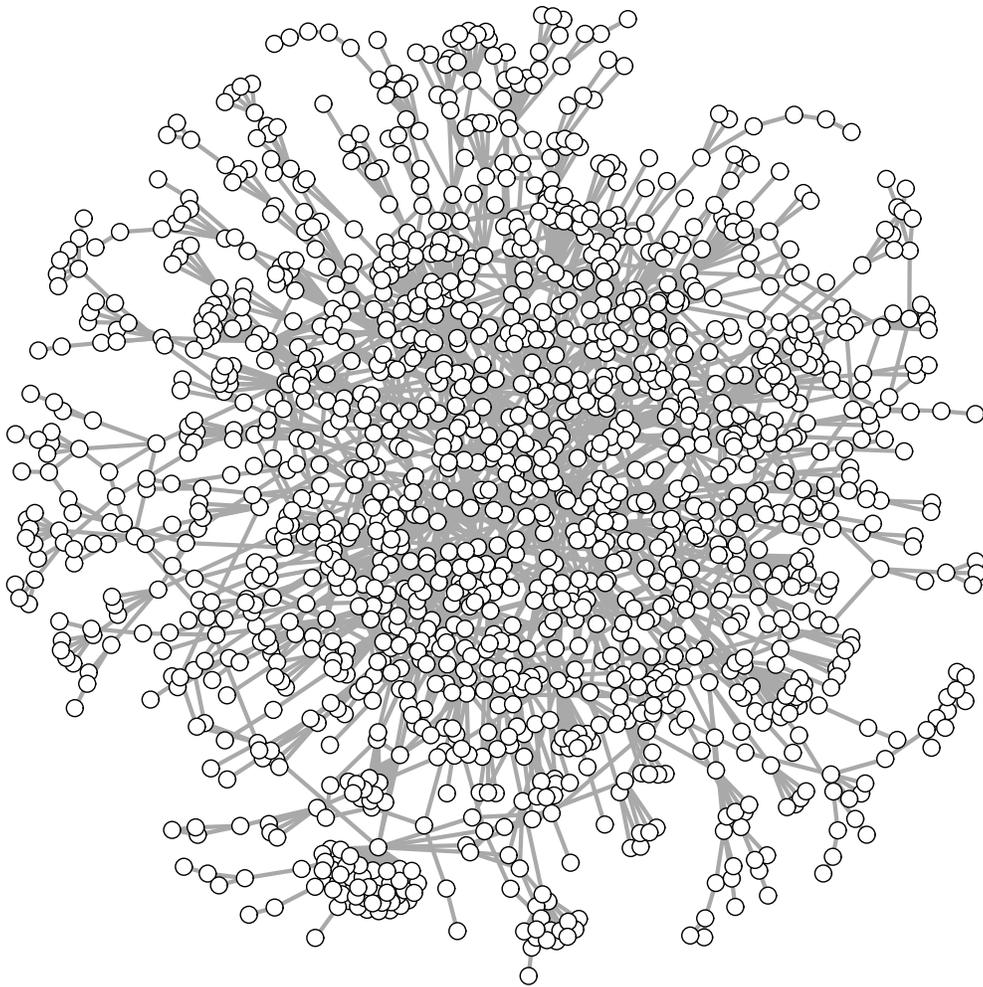


Figure 2.6: Largest Component of Protein Network ($N = 1458$, $|E| = 1948$)

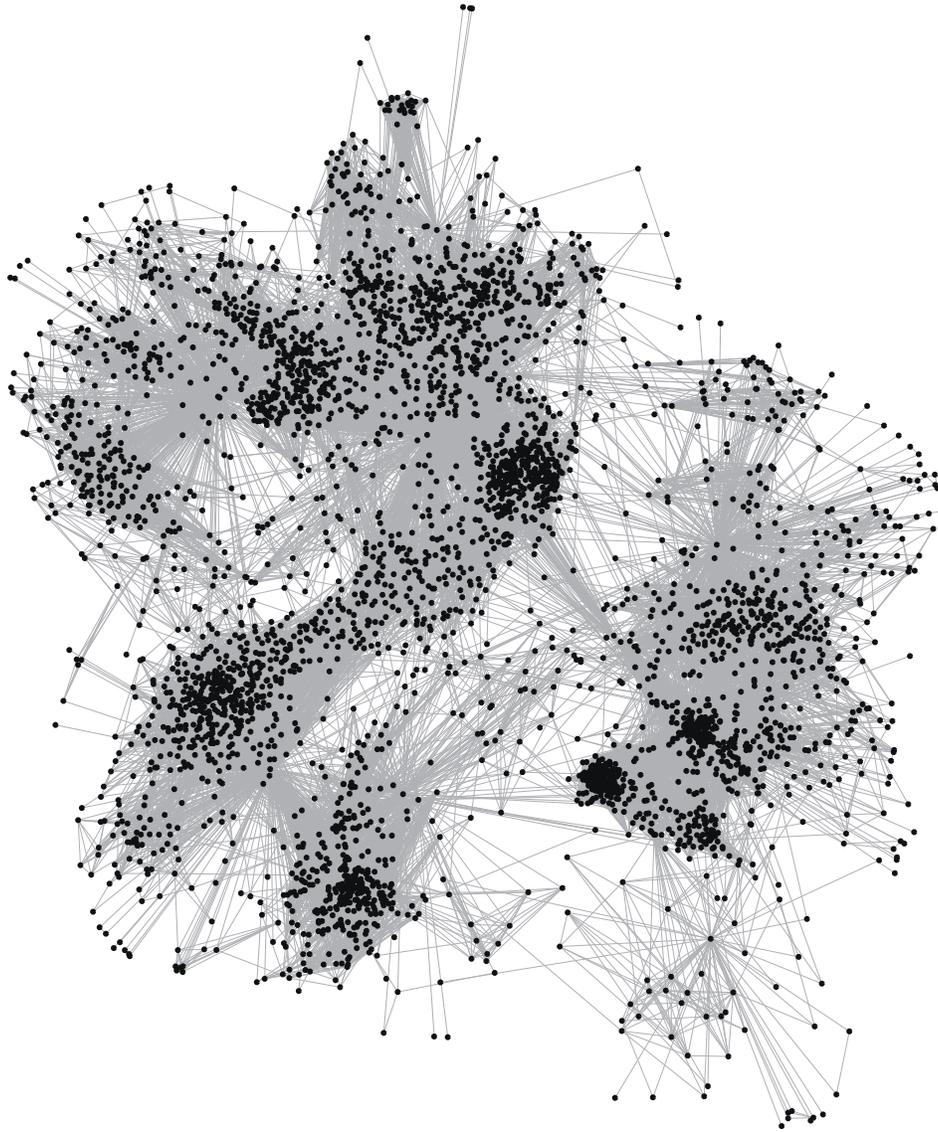


Figure 2.7: Sample of Facebook Network ($N = 4039, |E| = 88234$)

2.3 Markov Process

A stochastic process (also called random process) describes the random evolution of a system in time. A **Markov process**, $\{X(t), t \geq 0\}$, is a stochastic process that takes value over a countable set of states, \mathcal{X} , and for which

$$P(X(t_{n+1}) = i_{n+1} \mid X(t_n) = i_n, \dots, X(0) = i_0) = P(X(t_{n+1}) = i_{n+1} \mid X(t_n) = i_n).$$

There are many references regarding Markov processes [49,50]. We summarize the main concepts that we will be using. Markov processes are often used in queueing theory.

A Markov process can be **discrete-time** or **continuous-time**. A discrete-time Markov process, often called a Markov chain, changes state at set times $t = \{0, 1, 2, \dots\}$, whereas a continuous-time Markov process changes states at random times $t = [0, \infty)$. A Markov process is **stationary** if $\forall t_1, \dots, t_n, \tau$, and n ,

$$P(X(t_1 + \tau), X(t_2 + \tau), \dots, X(t_n + \tau)) = P(X(t_1), X(t_2), \dots, X(t_n)).$$

A Markov process is **homogenous** if $\forall t, \tau$,

$$P(X(t + \tau) = j \mid X(\tau) = i) = P(X(t) = j \mid X(0) = i).$$

It is **irreducible** if every state in \mathcal{X} can be reached from every other state. The process is **ergodic** if the proportion of the time the process spends in state k during the period $[0, t]$ converges to $\pi(k)$ as $t \rightarrow \infty$.

The **transition rate** from state $i \in \mathcal{X}$ to state $j \in \mathcal{X}$ for a homogenous Markov process is

$$q(i, j) = \lim_{\tau \rightarrow 0} \frac{P(X(t + \tau) = j \mid X(t) = i)}{\tau}, \quad i \neq j.$$

The evolution of the homogenous, continuous-time Markov process is completely specified by the **transition rate matrix**, $\mathbf{Q} = [Q_{ij}]$. It is also called the infinitesimal generator of the process. The transition rate matrix is a $|\mathcal{X}| \times |\mathcal{X}|$ matrix and has the following properties

- $Q_{ij} = q(i, j)$
- $Q_{ij} \geq 0$ for all $i \neq j$
- $\sum_{j \in \mathcal{X}} Q_{ij} = 0$ for all i .

- $Q_{ii} = -\sum_{j \neq i} Q_{ij} < \infty$ for all i .

Let $P_{ij}(t) = P(X(t + \tau) = j \mid X(\tau) = i)$, then **Kolmogorov's forward equation** of the Markov process, also called the **master equation (ME)** in physics, is

$$\frac{d}{dt}P_{ij}(t) = \sum_{k \neq j} Q_{kj}P_{ik}(t) - P_{ij}(t) \sum_{k \neq j} Q_{kj},$$

which, with $\mathbf{P}(t) = [P_{ij}(t)]$, can be written as

$$\frac{d}{dt}\mathbf{P}(t) = \mathbf{P}(t)\mathbf{Q}.$$

The equilibrium distribution, π , is the probability distribution over \mathcal{X} for which

$$\frac{d}{dt}\mathbf{P}(t) = 0.$$

Finite-size Markov processes (i.e., $|\mathcal{X}| < \infty$) always have an equilibrium distribution. The equilibrium distribution does not depend on the initial condition of the process $X(0)$ and is also the limiting distribution

$$\lim_{t \rightarrow \infty} P(X(t) = k \mid X(0) = j) = \pi(k).$$

The equilibrium distribution of the continuous-time Markov process satisfies the full balance condition

$$\pi(j) \sum_{k \in \mathcal{X}} q(i, k) = \sum_{k \in \mathcal{X}} \pi(k)q(k, j), \quad j \in \mathcal{X}$$

and is the left eigenvector corresponding to the 0 eigenvalue of the transition rate matrix:

$$0 = \pi\mathbf{Q}$$

with

$$\sum_{i \in \mathcal{X}} \pi(i) = 1.$$

2.3.1 Reversibility

A stochastic process is **reversible** if the behavior of the process is statistically the same forward in time as it is backward in time. This means that $(X(t_1), X(t_2), \dots, X(t_n))$ has the same probability distribution as $(X(\tau - t_1), X(\tau - t_2), \dots, X(\tau - t_n))$ for $t, \tau = [0, \infty)$. A reversible process is by definition also stationary. The equilibrium distribution of a reversible, continuous-time Markov process satisfies the **detailed balance conditions**

$$\pi(j)q(j, k) = \pi(k)q(k, j), \quad j, k \in \mathcal{X}.$$

2.3.2 Network Process as a Markov Process

Network processes model the evolution of individual agent states. For a network with N agents, the state of the network process at some time t , which we call the **network configuration** is

$$\mathbf{x} = [x_1, x_2, \dots, x_N]^T,$$

where $x_i, i = 1, \dots, N$ is the state of the individual agents. When the size of the network is finite, the set of all possible network configurations, \mathcal{X} , is also finite. The number of configurations,

$$|\mathcal{X}| = 2^N,$$

is exponential in N .

A network process that is Markov assumes that the network configuration at time t_{n+1} depends only on the network configuration at time t_n ,

$$P(X(t_{n+1}) = \mathbf{x}(t_{n+1}) \mid X(t_n) = \mathbf{x}(t_n)).$$

The dynamics of the continuous-time Markov process are specified by the transition rate of the network process going from a network configuration, \mathbf{x} , to another network configuration, \mathbf{x}' ,

$$q(\mathbf{x}, \mathbf{x}'), \mathbf{x}, \mathbf{x}' \in \mathcal{X}.$$

Typically, network processes assume that the transition rate between two configurations is nonzero if only *one* agent changes state from \mathbf{x} to \mathbf{x}' . The transition rate matrix, \mathbf{Q} , of a network process is a $|\mathcal{X}| \times |\mathcal{X}|$ with many zero entries.

2.4 Conclusion

This chapter reviews basic theoretic concepts used in this thesis from epidemiology, graph theory and network science, and Markov processes and networks processes.

Modeling Epidemics on Networks: Scaled SIS Process

3.1 Introduction

The chapter introduces a network process model that allows for tractable analysis of contagion effects in arbitrary heterogeneous networks. The details of the dynamics and assumptions of the scaled SIS (susceptible-infected-susceptible) process are presented in this chapter. We will present the scaled SIS process from different but equivalent perspectives: 1) as a microscopic model that accounts for the interactions among N individual agents; and 2) as a continuous-time Markov process where each Markov state corresponds to a particular configuration of healthy and infected agents in the network. Therefore, the number of Markov states grows exponentially with the number of agents N . This illustrates the challenge of using general Markov processes to study epidemics on large networks of arbitrary structure.

Rather than having to solve the intractable eigenvalue-eigenvector problem of the transition rate matrix, \mathbf{Q} , to find the equilibrium distribution $\pi(\mathbf{x})$ of the process, we develop the scaled SIS process. In section 3.4, we prove that this process is a reversible Markov process for which we can find the equilibrium distribution in closed-form regardless of network size and structure. This allows us to study how arbitrary network structure and dynamics affect the equilibrium or the time-asymptotic ($t \rightarrow \infty$) behavior of the network process. Section 3.5 discusses the four different parameter regimes of the process.

3.2 Scaled SIS Process

The relationships between N agents are described by a static, unweighted, undirected, simple graph $G(V, E)$, where V represents the agents, and E is the set of edges representing the dependencies, contacts, or relationships amongst the agents. The structure of the network is described by its adjacency matrix, \mathbf{A} . We assume that the graph G is undirected, unweighted, simple, connected, and static.

The state of the i th agent in the network is denoted by x_i . When $x_i = 0$, the agent is healthy but susceptible to infection. When $x_i = 1$, the agent is infected and capable of spreading the virus to others. The state of the population is the N -tuple collection of the states of all the agents in the network,

$$\mathbf{x} = [x_1, x_2, \dots, x_N]^T.$$

We refer to \mathbf{x} as either the **network state** or **configuration**. The set of all possible configurations is $\mathcal{X} = \{\mathbf{x}\}$. Since each agent can be in one of two states, the cardinality

$$|\mathcal{X}| = 2^N.$$

The **scaled SIS process**, $\{X(t), t \geq 0\}$, models the evolution of a stochastic process on the network $G(V, E)$. The dynamics of the process is based on the SIS epidemics framework. The scaled SIS process assumes that infected agents can heal and healthy agents will become infected in a random amount of time. As a result, the process will transition from one configuration to another configuration as individual agents heal or become infected. Since infected agents can pass virus to their healthy neighbors, how long the agent remains infected depends on the state of their neighbors. The scaled SIS process assumes the following microscopic dynamics:

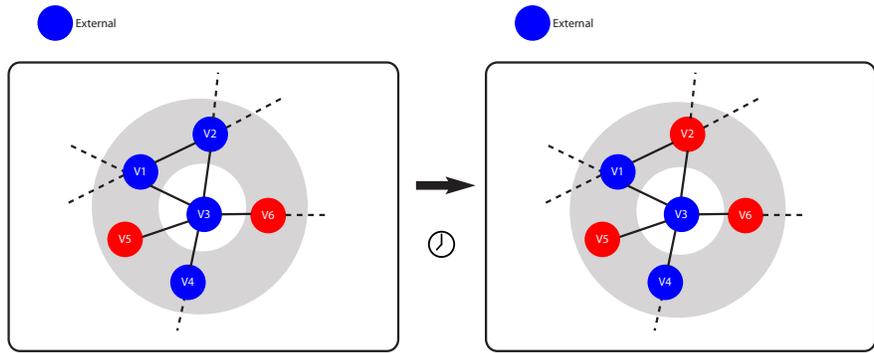
- Only one agent can change its state at any time.
- The amount of time it takes for the i th agent to heal (assuming that it was infected) is exponentially distributed with rate

$$\mu > 0.$$

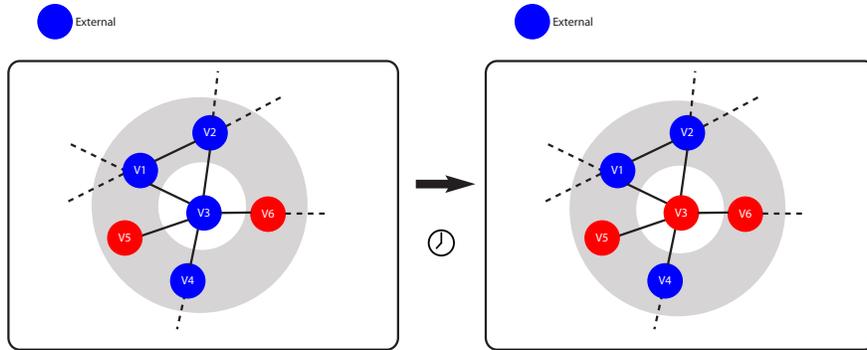
The parameter μ is the **healing rate**. By the properties of the exponential distribution. The expected time it takes for an agent to heal is

$$\bar{T}_h = \frac{1}{\mu}.$$

The scaled SIS process assumes that all the infected agents in the network have the same expected healing time \bar{T}_h . For example, the expected healing time for agent V_2 in Figure 3.1(a) and for agent V_3 in Figure 3.1(b) are the same.



(a) Agent V_2 Heals



(b) Agent V_3 Heals

Figure 3.1: Examples of Healing Transitions (Blue = Infected, Red = Healthy)

- The amount of time it takes for the i th agent to become infected (assuming that it was healthy) is exponentially distributed with rate

$$\lambda \gamma_i^m > 0,$$

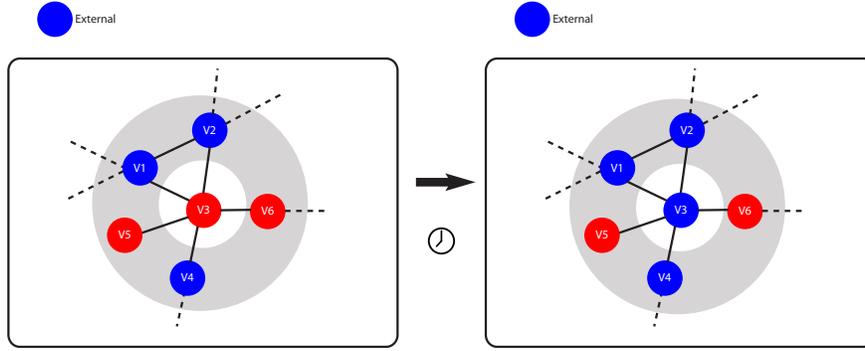
where m_i is the total number of infected neighbors of the i th agent.

- The rate $\gamma > 0$ is the **endogenous infection rate** or the contagion rate. It characterizes a contagion effect when an agent has infected neighbors that may potentially spread the infection to the agent. When $\gamma = 1$, the infection rate does not depend on the number of infected neighbors; the structure does not affect the dynamics. When $\gamma > 1$, the infection rate increases as the number of infected neighbors increases, which models a typical epidemics scenario.
- The rate $\lambda > 0$ is the **exogenous infection rate** or the spontaneous infection rate. It is the infection rate for which a healthy agent becomes infected from sources outside the network (i.e, **not** by its infected neighbors).

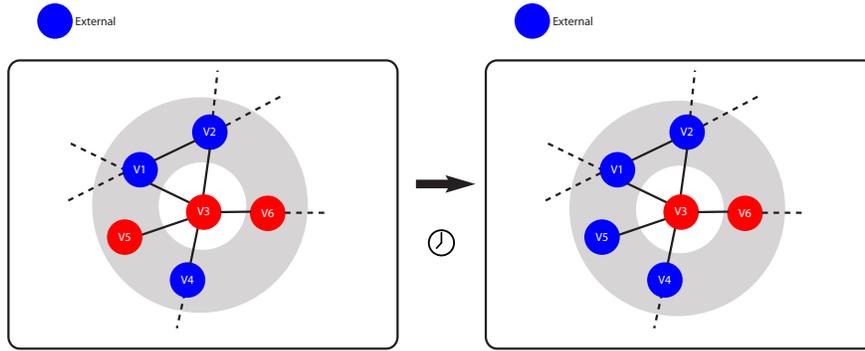
The expected time it takes for the i th agent to become infected is

$$\bar{T}_i = \frac{1}{\lambda\gamma^{m_i}}.$$

Only healthy agents with the same number of infected neighbors, m_i , have the same expected infection time \bar{T}_i . The expected infection time for agent V_3 in Figure 3.2(a) is not the same as the expected infection time for agent V_5 in Figure 3.2(b) since V_3 has 3 infected neighbors, while V_5 has no infected neighbors. Since the total number of infected neighbors, m_i , of a node is upperbounded by the total number of neighbors, d_i , the infection rate depends on the underlying network topology.



(a) Agent V_3 becomes Infected



(b) Agent V_5 becomes Infected

Figure 3.2: Examples of Infection Transitions (Blue = Infected, Red = Healthy)

3.3 Continuous-Time Markov Process

The scaled SIS process, $\{X(t), t \geq 0\}$, is a stationary, homogenous, irreducible, finite-size continuous-time Markov process. Continuous-time Markov processes are reviewed in chapter 2.3. Each state of the Markov process corresponds to a configuration, \mathbf{x} . The state space of the scaled SIS process is the space of all possible configurations, \mathcal{X} . The possible transitions of the scaled SIS process are:

1. Consider the configuration

$$\mathbf{x} = [x_1, x_2, \dots, x_j = 1, x_k, \dots, x_N]^T.$$

Let \mathbf{x}' denote the configuration where the j th agent heals:

$$\mathbf{x}' = [x_1, x_2, \dots, x_j = 0, \dots, x_N]^T.$$

The transition rate of the scaled SIS process from \mathbf{x} to \mathbf{x}' is

$$q(\mathbf{x}, \mathbf{x}') = \mu. \quad (3.1)$$

2. Consider the configuration

$$\mathbf{x} = [x_1, x_2, \dots, x_j, x_k = 0, \dots, x_N]^T.$$

Let \mathbf{x}' denote the configuration where the k th agent becomes infected:

$$\mathbf{x}' = [x_1, x_2, \dots, x_k = 1, \dots, x_N]^T.$$

The transition rate of the scaled SIS process from \mathbf{x} to \mathbf{x}' is

$$q(\mathbf{x}, \mathbf{x}') = \lambda \gamma^{m_k}, \quad (3.2)$$

where

$$m_k = \sum_{j=1}^N \mathbb{1}(x_j = 1) A_{jk}$$

is the total number of infected neighbors of node k . The symbol $\mathbb{1}(\cdot)$ is the indicator function, and $\mathbf{A} = [A_{jk}]$ is the adjacency matrix of G .

3. The transition rate from a configuration \mathbf{x} to any other configuration \mathbf{x}' is 0 otherwise.

The dynamics of the scaled SIS process are summarized by the asymmetric, $2^N \times 2^N$ transition rate matrix, $\mathbf{Q} = [Q_{ij}]$. The i th row and j th column of \mathbf{Q} correspond to the decimal representations of different configurations, $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$, respectively. Depending on the type of the transition, the value of Q_{ij} , $i \neq j$ is determined by (3.1), (3.2), or is 0.

An alternative description of the scaled SIS process, using the transition rate matrix, is the master equation formulation

$$\frac{d}{dt} \mathbf{P}(t) = \mathbf{P}(t) \mathbf{Q}, \quad (3.3)$$

where the i th row and j th column of $\mathbf{P}(t)$ denote the probability that the scaled SIS process is in some configuration \mathbf{x}' at time $t > 0$ given that it started in some configuration \mathbf{x} at time $t = 0$

$$P(X(t) = \mathbf{x}' \mid X(0) = \mathbf{x}).$$

3.4 Equilibrium Distribution

The probability distribution for which

$$\frac{d}{dt}\mathbf{P}(t) = 0$$

is the equilibrium distribution. Adopting a convention from the Markov process literature, we denote the equilibrium distribution as $\pi(\mathbf{x})$, which is a probability mass function (PMF) over the configuration space \mathcal{X} . The equilibrium distribution is also the limiting distribution of the scaled SIS process [50]

$$\lim_{t \rightarrow \infty} \frac{d}{dt}\mathbf{P}(t) = \pi(\mathbf{x}).$$

The equilibrium distribution for a finite-size, continuous-time Markov process can be found, up to the normalization constant, by solving the eigenvalue-eigenvector problem

$$\pi\mathbf{Q} = 0.$$

The unnormalized distribution is the left eigenvector of the transition rate matrix corresponding to the null eigenvalue. However, for large networks, directly finding the equilibrium distribution of the scaled SIS process by this method is infeasible since the dimension of the \mathbf{Q} matrix is exponential in the size of the network. For $N = 100$, the transition rate matrix is $2^{100} \times 2^{100}$. Solving this eigenvalue-eigenvector problem is computationally intractable for all but small size networks. However, we can prove that:

Theorem 3.4.1. *The scaled SIS process, $\{X(t), t \geq 0\}$, is a reversible Markov process and the equilibrium distribution is*

$$\pi(\mathbf{x}) = \frac{1}{Z} \left(\frac{\lambda}{\mu}\right)^{1^T \mathbf{x}} \gamma^{\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2}}, \quad \mathbf{x} \in \mathcal{X} \quad (3.4)$$

where Z is the partition function defined as

$$Z = \sum_{\mathbf{x} \in \mathcal{X}} \left(\frac{\lambda}{\mu}\right)^{1^T \mathbf{x}} \gamma^{\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2}}. \quad (3.5)$$

According to Theorem 3.4.1, the equilibrium behavior of the scaled SIS process is completely described by 2 statistics:

$$1^T \mathbf{x} = \sum_{i=1}^N x_i,$$

the total number of infected nodes in the network configuration, and

$$\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2} = \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N A_{ij} x_i x_j,$$

the total number of edges where both end nodes are infected. Note that $0 \leq 1^T \mathbf{x} \leq N$ and $0 \leq \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2} \leq |E|$, where N is the total number of agents in the network and $|E|$ is the total number of edges.

The equilibrium distribution of the scaled SIS process is a Gibbs distribution [51]. An alternate formulation of (3.4) is

$$\pi(\mathbf{x}) = \frac{1}{Z} e^{H(\mathbf{x})} \quad (3.6)$$

with

$$H(\mathbf{x}) = 1^T \mathbf{x} \log\left(\frac{\lambda}{\mu}\right) + \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2} \log(\gamma) \quad (3.7)$$

$$= \sum_{i=1}^N x_i \log\left(\frac{\lambda}{\mu}\right) + \sum_{i=1}^N \sum_{j=1}^N A_{ij} x_i x_j \log(\gamma), \quad (3.8)$$

In statistical mechanics, $H(\mathbf{x})$ is called the **Hamiltonian** and is considered the energy, or potential, of the network configuration \mathbf{x} . By the Hammersley-Clifford theorem, the distribution $\pi(\mathbf{x})$, being a Gibbs distribution, is also a Markov random field (MRF, also known as undirected Markov network in the probabilistic graphical model literature [41]). Using MRF terminology, we will refer to $x_i \log\left(\frac{\lambda}{\mu}\right)$ as the **unary potential** and $A_{ij} x_i x_j \log(\gamma)$ as the **pairwise potential**.

Proof. Recall the definition for **reversible** Markov processes from chapter 2.3. The proof for Theorem 3.4.1 uses the following theorem regarding **reversible** Markov processes:

Theorem 3.4.2 (Theorem 1.3 in [50]). *A stationary Markov process is reversible if and only if there exists a collection of positive numbers $\pi(\mathbf{x})$, $\mathbf{x} \in \mathcal{X}$, summing to unity that satisfy the detailed balance equation (3.9). When there exists such a collection $\pi(\mathbf{x})$, $\mathbf{x} \in \mathcal{X}$, it is the equilibrium distribution of the process.*

The detailed balance equation is

$$\pi(\mathbf{x})q(\mathbf{x}, \mathbf{x}') = \pi(\mathbf{x}')q(\mathbf{x}', \mathbf{x}), \quad \mathbf{x}, \mathbf{x}' \in \mathcal{X} \quad (3.9)$$

We show now that equation (3.4) satisfies the detailed balance equation (3.9) and is therefore the equilibrium distribution.

1. Consider the scenario where the Markov process jumps from state \mathbf{x} , where the k th agent is susceptible, to the state \mathbf{x}' where it becomes an infective:

$$\mathbf{x}' = [x_1, x_2, \dots, x_k = 1, \dots, x_N]^T$$

By (3.2), $q(\mathbf{x}, \mathbf{x}') = \lambda \gamma^{m_k}$ where $m_k = \sum_{j=1}^N \mathbb{1}(x_j = 1) A_{jk}$ is the total number of infected neighbors of node k . Using (3.4), the LHS of (3.9) is

$$\begin{aligned} \pi(\mathbf{x})q(\mathbf{x}, \mathbf{x}') &= \frac{1}{Z} \left(\frac{\lambda}{\mu} \right)^{1^T \mathbf{x}} \gamma^{\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2}} (\lambda \gamma^{m_k}) \\ &= \frac{1}{Z} \left(\frac{\lambda^{(1^T \mathbf{x} + 1)}}{\mu^{1^T \mathbf{x}}} \right) \gamma^{\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2} + m_k} \end{aligned}$$

By (3.1), $q(\mathbf{x}', \mathbf{x}) = \mu$. We know that in state \mathbf{x}' there is one additional infected node than in state \mathbf{x} . Furthermore, transitioning from \mathbf{x} to \mathbf{x}' , we gain as many additional infected edges as the number of infected neighbors of agent k :

$$\frac{(\mathbf{x}')^T \mathbf{A} (\mathbf{x}')}{2} = \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2} + m_k$$

Therefore, the RHS of (3.9) is

$$\begin{aligned} \pi(\mathbf{x}')q(\mathbf{x}', \mathbf{x}) &= \frac{1}{Z} \left(\frac{\lambda}{\mu} \right)^{1^T \mathbf{x} + 1} \gamma^{\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2} + m_k} (\mu) \\ &= \frac{1}{Z} \left(\frac{\lambda^{(1^T \mathbf{x} + 1)}}{\mu^{1^T \mathbf{x}}} \right) \gamma^{\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2} + m_k} \end{aligned}$$

Since the RHS of (3.9) is the same as its LHS, the detailed balance condition (3.9) is satisfied.

2. Consider the scenario where the Markov process jumps from \mathbf{x} , where the j th agent is infective, to the state \mathbf{x}' where it is healed:

$$\mathbf{x}' = [x_1, x_2, \dots, x_j = 0, \dots, x_N]^T$$

We know that in state \mathbf{x}' there is one fewer infected node than in state \mathbf{x} . Furthermore, transitioning from \mathbf{x} to \mathbf{x}' , we lose as many infected edges as the number of infected neighbors of node j . Therefore, the LHS of (3.9) is

$$\begin{aligned}\pi(\mathbf{x})q(\mathbf{x}, \mathbf{x}') &= \frac{1}{Z} \left(\frac{\lambda}{\mu} \right)^{1^T \mathbf{x}} \gamma^{\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2}} (\mu) \\ &= \frac{1}{Z} \left(\frac{\lambda^{1^T \mathbf{x}}}{\mu^{(1^T \mathbf{x}-1)}} \right) \gamma^{\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2}}\end{aligned}$$

The RHS of (3.9) is

$$\begin{aligned}\pi(\mathbf{x}')q(\mathbf{x}', \mathbf{x}) &= \frac{1}{Z} \left(\frac{\lambda}{\mu} \right)^{1^T \mathbf{x}-1} \gamma^{\left(\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2} - m_j \right)} (\lambda \gamma^{m_j}) \\ &= \frac{1}{Z} \left(\frac{\lambda^{1^T \mathbf{x}}}{\mu^{(1^T \mathbf{x}-1)}} \right) \gamma^{\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2}}\end{aligned}$$

The detailed balance condition in (3.9) is again satisfied.

3. If $q(\mathbf{x}, \mathbf{x}') = 0$, then $q(\mathbf{x}', \mathbf{x}) = 0$, and the detailed balance condition is satisfied trivially.

By Theorem 3.4.2, equation (3.4) is the equilibrium distribution of the scaled SIS process. \square

3.5 Parameter Regimes

According to (3.4), the equilibrium behavior of the scaled SIS process, $\{X(t), t \geq 0\}$, depends on the underlying network, \mathbf{A} , and the dynamics parameters: $\frac{\lambda}{\mu}$ and γ . The parameter γ is the endogenous infection rate, also called the contagion rate. The parameter $\frac{\lambda}{\mu}$ is the ratio of the exogenous infection rate and the healing rate. Since the network topology affects only the γ term, we will also refer to the process controlled by the γ parameter as the **topology-dependent process** and the one controlled by $\frac{\lambda}{\mu}$ as the **topology-independent process**.

Depending on the value of the dynamics parameters, the scaled SIS process models different behaviors. The parameter space can be divided into 4 regimes, as shown in Table 3.1.

The parameter $\frac{\lambda}{\mu}$ models the preferences of individual agents.

- When $\frac{\lambda}{\mu} > 1$, on average, an agent is in the infected state longer than it is in the healthy state. When $\frac{\lambda}{\mu}$ falls in this parameter range, individual agents prefer the infected state ($x_i = 1$) to the healthy state ($x_i = 0$).

III) Exogenous Infection Dominant $\frac{\lambda}{\mu} > 1$ $0 < \gamma \leq 1$	IV) Infection Dominant $\frac{\lambda}{\mu} > 1$ $\gamma > 1$
I) Healing Dominant $0 < \frac{\lambda}{\mu} \leq 1$ $0 < \gamma \leq 1$	II) Endogenous Infection Dominant $0 < \frac{\lambda}{\mu} \leq 1$ $\gamma > 1$

Table 3.1: Parameter Regimes

- When $0 < \frac{\lambda}{\mu} < 1$, on average, an agent is in the healthy state longer than it is in the infected state. When $\frac{\lambda}{\mu}$ falls in this parameter range, individual agents *prefer* the healthy state ($x_i = 0$) to the infected state ($x_i = 1$).
- When $\frac{\lambda}{\mu} = 1$, individual agents have equal preference for the infected and healthy state.

The parameter γ models how an individual agent is affected by its neighborhood. This parameter couples the underlying network topology to the dynamics of the process. It determines if the network structure will facilitate or impede the spread of infection.

- When $\gamma > 1$, additional infected neighbors of a healthy agent will increase the infection rate, thereby making the agent more vulnerable to infection. As the number of infections increases, the population will be more vulnerable to the epidemics as a single infection may quickly lead to additional infections. In the context of opinion spreading in social networks, having more neighbors who are spreaders will positively influence an adopter agent to adjust its opinion to that of the spreaders.
- When $0 < \gamma < 1$, additional infected neighbors of a healthy agent will decrease the infection rate, thereby making the agent *less* vulnerable to infection. This means that additional infections will actually strengthen the population. For example, in a system with active countermeasures, increasing the number of infected neighbors is a signal to booster the susceptible agent's defense. Another way to interpret this scenario, in the context of opinion spread in social networks, is that having more spreaders as neighbors will negatively influence an adopter agent to switch its opinion to that of the spreaders.

- When $\gamma = 1$, the state of an agent is unaffected by the states of its neighbors. In this case, the underlying network topology does not affect the dynamics of the scaled SIS process.

The behavior modeled by the 4 regimes are

- I) **Healing Dominant:** $0 < \frac{\lambda}{\mu} \leq 1, 0 < \gamma \leq 1$. Agents prefer the healthy state. Agents adopt defensive mechanisms so that the infection rate decreases with increasing number of infected neighbors. The network also helps to impede the infection. In this regime, the topology-dependent process supports the topology-independent process.
- II) **Endogenous Infection Dominant:** $0 < \frac{\lambda}{\mu} \leq 1, \gamma > 1$. Agents prefer the healthy state. However, the infection rate increases with increasing number of infected neighbors. The network helps to spread the infection. This regime models the behaviors of traditional epidemics. In this regime, the topology-independent process opposes the topology-dependent process.
- III) **Exogenous Infection Dominant:** $\frac{\lambda}{\mu} > 1, 0 < \gamma \leq 1$. Agents prefer the infected state. However, the infection rate decreases with increasing number of infected neighbors. The network helps to impede the infection. In this regime, the topology-independent process opposes the topology-dependent process.
- IV) **Infection Dominant:** $\frac{\lambda}{\mu} > 1, \gamma > 1$. Agents prefer the infected state. The infection rate also increases with increasing number of infected neighbors. The network helps to spread the infection. In this regime, the topology-dependent process supports the topology-independent process.

3.6 Conclusion

This chapter introduced the scaled SIS process, which models an epidemics-like stochastic process over a static, undirected network. Unlike previous network process models, the equilibrium behavior of the scaled SIS process is analyzable for arbitrary network structure. We proved that the scaled SIS process is a reversible, continuous-time Markov process for which the equilibrium distribution has a simple, closed-form description. The equilibrium distribution is a Gibbs distribution. The sufficient statistics of the scaled SIS process at equilibrium are 1) the number of infected nodes and 2) the number of edges with both end nodes infected. The rest of the thesis studies the

equilibrium behavior of the scaled SIS process using the equilibrium distribution for the 4 different parameter regimes presented in this chapter, in section 3.5.

Most-Probable Configuration

4.1 Introduction

The equilibrium distribution of the scaled SIS process characterizes its behavior at equilibrium or as $t \rightarrow \infty$. We can infer about the effect of the network topology on the network process by studying various quantitative measures of the equilibrium distribution, $\pi(\mathbf{x})$. This chapter studies the mesoscopic vulnerability of the network by finding the configuration, out of 2^N possible configurations, with the maximum equilibrium probability; we call this the **most-probable configuration**. In the long-run, this would be the most often observed configuration. This configuration is known as the **ground state** in statistical mechanics, and in machine learning, this configuration is called **maximum a posteriori** (MAP) configuration. However, since each agent can only be in one of two states, the Most-Probable Configuration Problem is a combinatorial optimization problem.

Assuming that the infection and healing rates fall into regimes I) or IV), finding the most-probable configuration is trivial. In the other two regimes, the optimization is nontrivial. We show in section 4.4 that, in regime III), vulnerable structures in the network are related to the **maximum independent set**. Finding these structures is therefore NP-hard for an arbitrary network. It is only polynomial-time solvable for perfect graphs.

In section 4.5, we show that in regime II) the solution space for the Most-Probable Configuration Problem can be characterized exactly for several families of graphs: 1) k -regular, 2) complete multipartite, or 3) complete multipartite with k -regular islands. For these graphs, the solution space exhibits a single phase transition between $\mathbf{x}^0 = [0, 0, \dots, 0]^T$, the configuration where all the agents

are healthy, and $\mathbf{x}^N = [1, 1, \dots, 1]^T$, the configuration where all the agents are infected. The threshold can be determined exactly. We observe that, for other network topologies, the most-probable configuration may demonstrate multiple phase transitions; it may be a configuration where subsets of agents are infected. We call these types of solution to the Most-Probable Configuration Problem **non-degenerate** most-probable configuration.

4.2 Most-Probable Configuration Problem

The Most-Probable Configuration Problem solves for the configuration in \mathcal{X} with the maximum probability

$$\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{X}} \pi(\mathbf{x}) = \arg \max_{\mathbf{x} \in \mathcal{X}} \left\{ \left(\frac{\lambda}{\mu} \right)^{\mathbf{1}^T \mathbf{x}} \gamma^{\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2}} \right\}. \quad (4.1)$$

This is equivalent to solving for the maximizer of the Hamiltonian (3.7)

$$\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{X}} H(\mathbf{x}) = \arg \max_{\mathbf{x} \in \mathcal{X}} \left\{ \mathbf{1}^T \mathbf{x} \log \left(\frac{\lambda}{\mu} \right) + \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2} \log(\gamma) \right\}. \quad (4.2)$$

We call \mathbf{x}^* the **most-probable configuration**. In statistical mechanics, the most-probable configuration is called the *ground state* [51]; it is the configuration with minimum energy. Traditionally, the energy is defined as negative values in physics, but as positive values in the Markov random field literature; consequently, the most-probable configuration is the configuration with maximum energy in our case. In the Markov random field literature, the most-probable configuration is also called MAP (maximum a posteriori) as the distribution that is typically being maximized is the posteriori distribution.

The Most-Probable Configuration Problem (4.1) is a combinatorial optimization problem as agents can only be in one of two states. Its solution space depends on both the network topology, \mathbf{A} , and the infection and healing rates, λ, γ, μ . For large networks, it is infeasible to iterate through all 2^N configurations of the scaled SIS process to find the most-probable configuration.

Recall the 4 parameter regimes of the scaled SIS Process as defined in section 3.5. We analyzed the solution space of the Most-Probable Configuration Problem for these regimes in this chapter and chapter 5.

III) Exogenous Infection Dominant $\frac{\lambda}{\mu} > 1$ $0 < \gamma \leq 1$	IV) Infection Dominant $\frac{\lambda}{\mu} > 1$ $\gamma > 1$
I) Healing Dominant $0 < \frac{\lambda}{\mu} \leq 1$ $0 < \gamma \leq 1$	II) Endogenous Infection Dominant $0 < \frac{\lambda}{\mu} \leq 1$ $\gamma > 1$

Parameter Regimes

4.3 Regime I) Healing Dominant & IV) Infection Dominant

The Most-Probable Configuration Problem is trivial for regimes I) and IV) as the individual preferences of the agents (i.e., topology-independent process), controlled by $\frac{\lambda}{\mu}$, corroborates the network effect (i.e., topology-dependent process), controlled by γ .

In regime I) **Healing Dominant**, $0 < \frac{\lambda}{\mu} \leq 1, 0 < \gamma \leq 1$, as explained in chapter 3.5, both the topology-independent process and the topology-dependent process favor the healthy state. The most probable configuration in equilibrium is $\mathbf{x}^* = \mathbf{x}^0 = [0, 0, \dots, 0]^T$, the configuration where all the agents are healthy. This holds for any network topology, $G(V, E)$. The most-probable configuration has probability

$$\pi(\mathbf{x}^0) = \frac{1}{Z}.$$

In regime IV) **Infection Dominant**, $\frac{\lambda}{\mu} > 1, \gamma > 1$, as explained in chapter 3.5, both the topology-independent process and the topology-dependent process favor the infected state. The most probable configuration is $\mathbf{x}^* = \mathbf{x}^N = [1, 1, \dots, 1]^T$, the configuration where all the agents are infected. This holds for any network topology, $G(V, E)$. The most-probable configuration has probability

$$\pi(\mathbf{x}^N) = \frac{1}{Z} \left(\frac{\lambda}{\mu} \right)^N \gamma^{|E|}.$$

4.4 Regime III) Exogenous Infection Dominant

Unlike regime I) and IV), in regime III), the topology-independent process, controlled by $\frac{\lambda}{\mu}$, opposes the topology-dependent process, controlled by γ . This introduces an additional complexity in the analysis; unlike in regimes I) and IV), the solution space of the Most-Probable Configuration

Problem exhibits a **phase transition**; the solution of (4.1) changes depending on the parameter values.

With $\frac{\lambda}{\mu} > 1$, the exogenous infection rate, λ , is larger than the healing rate, μ . If the exogenous infection rate is ignored (i.e., we take $\gamma = 1$), then the most-probable network configuration would be $\mathbf{x}^* = \mathbf{x}^N = [1, 1, \dots, 1]^T$. However, since $0 < \gamma \leq 1$, additional infected nodes *decrease* the infection rate. This implies that healthy agents adopt a defense mechanism in response to having infected neighbors.

In regime III), the topology-dependent process wants to minimize the number of infected edges (i.e., edges where both end nodes are infected) whereas the topology-independent process wants to maximize the number of infected agents. In other words, these processes favor isolated infected nodes. Therefore, the solution space of the Most-Probable Configuration Problem is related to the graph theoretic concept of **independent sets**. An independent set is a subset of nodes such that the induced subgraph is composed entirely of isolated nodes. The **maximum independent set** is the largest possible independent set for a given graph [30]. The maximum independent set is also the largest maximal independent set (i.e., an independent set which is not a subset of any other independent set).

Intuitively, if $\frac{\lambda}{\mu} \gg 1$ and γ is close to 1, then the epidemics will be dominated by the topology-independent process; a good guess for \mathbf{x}^* is $\mathbf{x}^N = [1, 1, \dots, 1]^T$. When $0 < \gamma \ll 1$ and $\frac{\lambda}{\mu}$ is close to 1, the epidemics will be dominated by the topology-dependent process. However, as proved in Theorem 4.4.1, the most-probable configuration can not be $\mathbf{x}^0 = [0, 0, \dots, 0]^T$.

Theorem 4.4.1. *For any network topology, when $\frac{\lambda}{\mu} > 1, 0 < \gamma \leq 1$, then $\mathbf{x}^* \neq \mathbf{x}^0 = [0, 0, \dots, 0]^T$ for any feasible parameter values.*

Proof. We prove this by contradiction. Suppose that there is a network topology for which $\mathbf{x}^* = \mathbf{x}^0$. Therefore,

$$\pi(\mathbf{x}^0) = \frac{1}{Z} > \pi(\mathbf{x}), \forall \mathbf{x} \in \mathcal{X} \setminus \mathbf{x}^0$$

We know that, for any network configuration with 1 infected agent, \mathbf{x}^1 , the equilibrium is $\pi(\mathbf{x}^1) = \frac{1}{Z} \left(\frac{\lambda}{\mu} \right)$ regardless of the network topology. Since $\frac{\lambda}{\mu} > 1$,

$$\pi(\mathbf{x}^1) = \frac{1}{Z} \left(\frac{\lambda}{\mu} \right) > \frac{1}{Z} = \pi(\mathbf{x}^0)$$

which contradicts the premise that $\mathbf{x}^* = \mathbf{x}^0$. □

In regime III), according to Theorem 4.4.1, no matter the values of the parameters or the underlying network structure, it is not possible for $\mathbf{x}^0 = [0, 0, \dots, 0]^T$, the configuration where all the agents are healthy, to be a solution of the Most-Probable Configuration Problem. On the other hand, it is possible for $\mathbf{x}^N = [1, 1, \dots, 1]^T$ to be a most-probable configuration.

Theorem 4.4.2. *For any network topology and when $\frac{\lambda}{\mu} > 1, 0 < \gamma \leq 1$, sort the degrees so that $k_1 \geq k_2 \geq \dots \geq k_N$. If $\lambda\gamma^{k_1} > \mu$, then $\mathbf{x}^* = \mathbf{x}^N = [1, 1, \dots, 1]^T$.*

Proof. Suppose that there is a network topology for which $\mathbf{x}^* = \mathbf{x}', \mathbf{x}' \neq \mathbf{x}^N$. This means that $\pi(\mathbf{x}') > \pi(\mathbf{x}^N)$.

What are possible \mathbf{x}' ? Since $\frac{\lambda}{\mu} > 1, 0 < \gamma \leq 1$, we should consider configurations which either have more infected nodes or less infected edges than \mathbf{x}^N . However, only the latter option is possible. In addition, having less than $|E|$ infected edges means that we need to reduce the number of infected nodes in the configuration as well. Consider

$$\mathcal{X}' = \{\mathbf{x} \in \mathcal{X} \mid \mathbf{1}^T \mathbf{x} = N - 1, \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2} = |E| - k_1\}$$

Note that since

$$1 \leq \frac{\lambda}{\mu} \gamma^{k_1} \leq \frac{\lambda}{\mu} \gamma^{k_2} \leq \dots \leq \frac{\lambda}{\mu} \gamma^{k_N},$$

then

$$\pi(\mathbf{x}') \geq \pi(\mathbf{x}), \forall \mathbf{x} \in \mathcal{X} \setminus \mathbf{x}^N.$$

Realize that

$$\frac{\pi(\mathbf{x}^N)}{\pi(\mathbf{x}')} = \frac{\left(\frac{\lambda}{\mu}\right)^N \gamma^{|E|}}{\left(\frac{\lambda}{\mu}\right)^{N-1} \gamma^{|E|-k_1}} = \frac{\lambda}{\mu} \gamma^{k_1} > 1$$

This implies that $\pi(\mathbf{x}^N) > \pi(\mathbf{x}')$, which contradicts the premise that $\mathbf{x}^* = \mathbf{x}', \mathbf{x}' \neq \mathbf{x}^N$. □

Theorem 4.4.2 shows that the only consideration necessary for $\mathbf{x}^* = \mathbf{x}^N$ are the values of the dynamics parameters λ, γ, μ , and the maximum degree of the underlying network.

Theorem 4.4.3. For any network topology and when $\frac{\lambda}{\mu} > 1, 0 < \gamma \leq 1$, if in addition $\lambda\gamma < \mu$, then \mathbf{x}^* are the network configuration(s) with maximum number of infected nodes and 0 infected edges.

Proof. Let

$$\mathcal{X}' = \{\mathbf{x} \in \mathcal{X} \mid \mathbf{x}^T \mathbf{A} \mathbf{x} = 0, \text{ and } \mathbf{1}^T \mathbf{x} \text{ is maximum}\} \quad (4.3)$$

The set \mathcal{X}' is also known as the maximum independent set [30]. Since $\frac{\lambda}{\mu} > 1$, we know that

$$\pi(\mathbf{x}') = \frac{1}{Z} \left(\frac{\lambda}{\mu} \right)^{\mathbf{1}^T \mathbf{x}'} \geq \pi(\mathbf{x}), \forall \mathbf{x} \in \{\mathbf{x} \in \mathcal{X} \mid \mathbf{x}^T \mathbf{A} \mathbf{x} = 0\} \quad (4.4)$$

Consider $\mathbf{x} \in \mathcal{X}'$ and $\mathbf{x} \notin \mathcal{X}'$. Let $\bar{\mathcal{X}}_s$ represent the set of configurations in $\mathcal{X} \setminus \mathcal{X}'$ with s infected nodes.

1. If $\mathbf{1}^T \mathbf{x}' \geq \mathbf{1}^T \mathbf{x}$, then, since $\frac{\lambda}{\mu} > 1, 0 < \gamma \leq 1$, we know that $\pi(\mathbf{x}') > \pi(\mathbf{x})$.
2. If $\mathbf{1}^T \mathbf{x}' = m < \mathbf{1}^T \mathbf{x}$, then we know that $\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2} \geq 1$ by our definition of \mathcal{X}' . Consider the sets $\bar{\mathcal{X}}_{m+k}$ and $\bar{\mathcal{X}}_{m+k+1}$ where $1 \leq k \leq N - m - 1$. Suppose that all the configurations in $\bar{\mathcal{X}}_{m+k}$ have at least $q \geq 1$ edges. Condition (4.3) also implies that each additional infected node will result in at least one additional infected edge; therefore all the configurations in $\bar{\mathcal{X}}_{m+k+1}$ will have at least $q + 1$ infected edges.

Since $0 < \gamma \leq 1$, the configuration(s) with the maximum equilibrium probability in $\bar{\mathcal{X}}_{m+k}$ should have the fewest number of infected edges amongst all the configurations in $\bar{\mathcal{X}}_{m+k}$, similarly for $\bar{\mathcal{X}}_{m+k+1}$.

$$\max_{\mathbf{x} \in \bar{\mathcal{X}}_{m+k}} \pi(\mathbf{x}) = \frac{1}{Z} \left(\frac{\lambda}{\mu} \right)^{m+k} \gamma^q$$

and

$$\max_{\mathbf{x} \in \bar{\mathcal{X}}_{m+k+1}} \pi(\mathbf{x}) = \frac{1}{Z} \left(\frac{\lambda}{\mu} \right)^{m+k+1} \gamma^{q+1}$$

With $\lambda\gamma < \mu$, we see that

$$\max_{\mathbf{x} \in \bar{\mathcal{X}}_{m+k}} \pi(\mathbf{x}) > \max_{\mathbf{x} \in \bar{\mathcal{X}}_{m+k+1}} \pi(\mathbf{x}). \quad (4.5)$$

Equation (4.5) implies that the configuration with the largest equilibrium probability in $\{\bar{\mathcal{X}}_s \mid s = m + 1, m + 2, \dots, N\}$ is the configuration with $m + 1$ infected nodes and 1 infected

edge.

$$\max_{\mathbf{x} \in \{\bar{\mathcal{X}}_s | s=m+1, m+2, \dots, N\}} \pi(\mathbf{x}) = \frac{1}{Z} \left(\frac{\lambda}{\mu} \right)^{m+1} \gamma.$$

With $\lambda\gamma < \mu$, $\pi(\mathbf{x}') > \max_{\mathbf{x} \in \bar{\mathcal{X}}_{m+1}} \pi(\mathbf{x})$.

Since we showed for both cases 1) and 2), we conclude that $\mathbf{x}^* = \mathbf{x}'$. \square

Corollary 4.4.4 (Proof in Appendix 4.9). *For a complete graph, K_N , when $\frac{\lambda}{\mu} > 1, 0 < \gamma \leq 1$, if in addition $\lambda\gamma < \mu$, then*

$$\mathbf{x}^* \in \{\mathbf{x} \in \mathcal{X} \mid \mathbf{1}^T \mathbf{x} = 1\}.$$

Corollary 4.4.5 (Proof in Appendix 4.9). *For a bipartite graph with islands of size N_1 and N_2 , when $\frac{\lambda}{\mu} > 1, 0 < \gamma \leq 1$, if in addition $\lambda\gamma < \mu$, then*

$$\mathbf{x}^* \in \{\mathbf{x} \in \mathcal{X} \mid \mathbf{1}^T \mathbf{x} = \max\{N_1, N_2\}, \mathbf{x}^T \mathbf{A} \mathbf{x} = 0\}.$$

Maximum Independent Set Problem Theorem 4.4.3 shows that, in the parameter regime III), $\frac{\lambda}{\mu} > 1, 0 < \gamma \leq 1$, and $\lambda\gamma < \mu$, the most-probable configuration \mathbf{x}^* of the scaled SIS process at equilibrium is the maximum independent set. There exist polynomial-time algorithms to find this set for a special class of graphs called **perfect graphs** [52]. Well-known examples of perfect graphs are: complete graphs, bipartite graphs, and chordal graphs. Unfortunately, finding the maximum independent set is NP-hard for general graph topologies [30]. As a result, it is infeasible to analyze the equilibrium behavior for large, real-world networks when the dynamics parameters are in regime III).

4.5 Regime II) Endogenous Infection Dominant

In regime II), like regime III), the topology-independent process, controlled by the rate $\frac{\lambda}{\mu}$, opposes the topology dependent process, controlled by the contagion rate γ . Similarly, the solution space of the Most-Probable Configuration Problem also exhibits phase transition; the solution of (4.1) changes depending on the parameter values.

The behavior modeled in regime II) however, is the opposite of the behavior modeled in regime III). With $0 < \frac{\lambda}{\mu} \leq 1$, the exogenous infection rate, λ , is smaller than or equal to the healing rate, μ . If the exogenous infection rate is ignored (i.e., $\gamma = 1$), then the most-probable network

configuration would be $\mathbf{x}^0 = [0, 0, \dots, 0]^T$. However, since $\gamma > 1$, additional infected agents *increase* the infection rate. The process utilizes the network structure to spread the infection to healthy agents. Regime II) models the behavior of standard epidemics.

In regime II), the topology-dependent process wants to maximize the number of edges (i.e., edges where both end nodes are infected) whereas the topology independent process wants to minimize the number of infected agents. Intuitively, if $\frac{\lambda}{\mu}$ is very small and γ is close to 1, then the behavior of the process will be dictated by the topology-independent process; a good guess for \mathbf{x}^* is $\mathbf{x}^0 = [0, 0, \dots, 0]^T$. When $\gamma \gg 1$ and $\frac{\lambda}{\mu}$ is close to 1, the process behavior will be dictated by the topology-dependent process; a good guess for \mathbf{x}^* is $\mathbf{x}^N = [1, 1, \dots, 1]^T$.

We divide the analysis of the most-probable configuration in regime II) according to different classes of graphs: 1) structured networks and 2) unstructured networks. For structured graphs like k -regular, complete multipartite, and complete multipartite with k -regular island networks (see section 2.2.4), the solution space of (4.1) can be characterized exactly. Below some threshold β , the most-probable configuration is \mathbf{x}^0 , whereas above β $\mathbf{x}^* = \mathbf{x}^N$. The value of β depends on the dynamics parameters, λ, γ, μ and on the network topology. For general network topologies, the most-probable configuration may be other configurations; we call solutions \mathbf{x}^* to the Most-Probable Configuration Problem that are neither \mathbf{x}^0 nor \mathbf{x}^N **non-degenerate** configurations.

4.6 Structured Network Topology

4.6.1 k -Regular Graph

Consider the k -regular graph, $G(V, E)$. We can partition the set of all possible configurations, \mathcal{X} , as

$$\mathcal{X} = \mathcal{X}_0 \cup \mathcal{X}_1 \cup \dots \cup \mathcal{X}_N$$

where $\mathcal{X}_s = \{\mathbf{x} \in \mathcal{X} \mid \mathbf{1}^T \mathbf{x} = s\}$ is the set of configurations with s infected agents and $\mathbf{x}^s \in \mathcal{X}_s$. Note that $\mathcal{X}_0 = \{\mathbf{x}^0\}$ and $\mathcal{X}_N = \{\mathbf{x}^N\}$.

Lemma 4.6.1 (Proof in Appendix 4.9). *When $0 \leq \frac{\lambda}{\mu} \leq 1, \gamma > 1$ and for $s = 0, 1, \dots, N, \mathbf{x}^s \in \mathcal{X}_s$, the unnormalized equilibrium distribution for a k -regular graph is upperbounded by an exponential function:*

$$\pi(\mathbf{x}^s) \propto \left(\frac{\lambda}{\mu}\right)^s \gamma^{\frac{\mathbf{x}^s T A \mathbf{x}^s}{2}} \leq \left(\frac{\lambda}{\mu} \gamma^\beta\right)^s \quad (4.6)$$

where $\beta = \frac{k}{2}$. Furthermore, the relationship (4.6) holds with equality for \mathbf{x}^0 and \mathbf{x}^N .

Theorem 4.6.1. For a k -regular graph, when $0 \leq \frac{\lambda}{\mu} \leq 1, \gamma > 1$, a threshold exists for the most probable network configuration, \mathbf{x}^* . Let

$$\beta = \frac{k}{2},$$

then

1. $\lambda \gamma^\beta > \mu$ if and only if \mathbf{x}^* is unique and $\mathbf{x}^* = \mathbf{x}^N$.
2. $\lambda \gamma^\beta < \mu$ if and only if \mathbf{x}^* is unique and $\mathbf{x}^* = \mathbf{x}^0$.
3. $\lambda \gamma^\beta = \mu$ if and only if \mathbf{x}^* is no longer the unique maximizer, and \mathbf{x}^N and \mathbf{x}^0 are both maximizers.

Proof. 1. $\lambda \gamma^\beta > \mu$

Sufficiency: If $\lambda \gamma^\beta > \mu$, we show that \mathbf{x}^* is unique and $\mathbf{x}^* = \mathbf{x}^N$.

When $\frac{\lambda}{\mu} \gamma^\beta > 1$, the RHS of (4.6) is maximized when $s = N$. Since this is a growing exponential function, it is also the unique maximizer. By Lemma 4.6.1, $\mathbf{x}^* = \mathbf{x}^N$, and the relationship in (4.6) holds with equality.

Necessity: If \mathbf{x}^* is unique and $\mathbf{x}^* = \mathbf{x}^N$, we show that $\lambda \gamma^\beta > \mu$.

Since $\mathbf{x}^* = \mathbf{x}^N$, it follows from Lemma 4.6.1 that the bounding exponential function reaches a maximum at $s = N$, and this bound is met with equality. The RHS is monotonically increasing only when $\frac{\lambda}{\mu} \gamma^\beta > 1$.

2. $\lambda \gamma^\beta < \mu$

Sufficiency: If $\lambda \gamma^\beta < \mu$, we show that \mathbf{x}^* is unique and $\mathbf{x}^* = \mathbf{x}^0$.

When $\frac{\lambda}{\mu} \gamma^\beta < 1$, the RHS of (4.6) is maximized when $s = 0$. Since this is a decaying exponential function, it is also the unique maximizer. By Lemma 4.6.1, $\mathbf{x}^* = \mathbf{x}^0$, and the

relationship in (4.6) holds with equality.

Necessity: If \mathbf{x}^* is unique and $\mathbf{x}^* = \mathbf{x}^0$, we show that $\lambda\gamma^\beta < \mu$.

Since $\mathbf{x}^* = \mathbf{x}^0$, it follows from Lemma 4.6.1 that the bounding exponential function reaches a maximum at $s = 0$ and this bound is met with equality. The RHS is monotonically decreasing only when $\frac{\lambda}{\mu}\gamma^\beta < 1$.

3. $\lambda\gamma^\beta = \mu$

Sufficiency: If $\lambda\gamma^\beta = \mu$, we show that \mathbf{x}^* is no longer the unique maximizer, and \mathbf{x}^N and \mathbf{x}^0 are both maximizers.

When $\frac{\lambda}{\mu}\gamma^\beta = 1$, the RHS of (4.6) is 1 regardless of s . We know that this is satisfied with equality for \mathbf{x}^0 and \mathbf{x}^N . Therefore, \mathbf{x}^* is no longer the unique maximizer since \mathbf{x}^N and \mathbf{x}^0 are both maximizers.

Necessity: If \mathbf{x}^* is no longer the unique maximizer, and \mathbf{x}^N and \mathbf{x}^0 are both maximizers, we show that $\frac{\lambda}{\mu}\gamma^\beta = 1$.

For $s = 0$, the RHS of (4.6) equals to 1 and by Lemma 4.6.1 this bound is met with equality. If \mathbf{x}^N and \mathbf{x}^0 are both maximizers, this can only be achieved when $\frac{\lambda}{\mu}\gamma^\beta = 1$.

□

4.6.2 Complete Bipartite Network

Consider the complete bipartite graph, $G(V, E) = K_{N_1, N_2}$. We can partition the set of all possible configurations, \mathcal{X} , as

$$\mathcal{X} = \mathcal{X}_{0,0} \cup \mathcal{X}_{1,0} \cup \mathcal{X}_{0,1} \cup \mathcal{X}_{1,1} \cup \dots \cup \mathcal{X}_{N_1, N_2}$$

where

$$\mathcal{X}_{s_1, s_2} = \{\mathbf{x} \in \mathcal{X} \mid s_1 \text{ infected agents in } V_1, s_2 \text{ infected agents in } V_2\} \quad (4.7)$$

Let $\mathbf{x}^{s_1, s_2} \in \mathcal{X}_{s_1, s_2}$ denote a configuration with $0 \leq s_1 \leq N_1$ infected agents in set V_1 and $0 \leq s_2 \leq N_2$ infected agents in set V_2 .

Lemma 4.6.2 (Proof in Appendix 4.9). *For the complete bipartite graph, all configurations, \mathbf{x}^{s_1, s_2} , belonging to the same partition, \mathcal{X}_{s_1, s_2} , have the same equilibrium probability.*

Lemma 4.6.3 (Proof in Appendix 4.9). *When $0 \leq \frac{\lambda}{\mu} \leq 1, \gamma > 1$ and for $s_1 = 0, 1, 2, \dots, N_1, s_2 = 0, 1, 2, \dots, N_2$, the unnormalized equilibrium distribution for a complete bipartite graph is upperbounded by an exponential function:*

$$\pi(\mathbf{x}^{s_1, s_2}) \propto \left(\frac{\lambda}{\mu}\right)^{s_1 + s_2} \gamma^{s_1 s_2} \leq \left(\frac{\lambda}{\mu} \gamma^\beta\right)^{s_1 + s_2} \quad (4.8)$$

where $\beta = \frac{N_1 N_2}{N_1 + N_2}$. Furthermore, the relationship holds with equality for $\mathbf{x}^{s_1, s_2} = \mathbf{x}^N$ and $\mathbf{x}^{s_1, s_2} = \mathbf{x}^0$.

Theorem 4.6.2. *For a complete bipartite graph, when $0 \leq \frac{\lambda}{\mu} \leq 1, \gamma > 1$, a threshold exists for the most probable network configuration. Let*

$$\beta = \frac{N_1 N_2}{N_1 + N_2},$$

then

1. $\lambda \gamma^\beta > \mu$ if and only if \mathbf{x}^* is unique and $\mathbf{x}^* = \mathbf{x}^N$.
2. $\lambda \gamma^\beta < \mu$ if and only if \mathbf{x}^* is unique and $\mathbf{x}^* = \mathbf{x}^0$.
3. $\lambda \gamma^\beta = \mu$ if and only if \mathbf{x}^* is no longer the unique maximizer, $\mathbf{x}^* = \mathbf{x}^N, \mathbf{x}^* = \mathbf{x}^0$.

Proof. The proof follows that of Theorem 4.6.1. Note that when $N_1 = N_2$, the complete bipartite graph is also a k -regular graph, and $k = N_1$,

$$\beta = \frac{N_1 N_1}{N_1 + N_1} = \frac{N_1}{2}$$

□

4.6.3 Complete Multipartite

Consider the complete multipartite graph, $G(V, E) = K_{N_1, N_2, \dots, N_m}$, where the intra-island degree, $k^{\text{intra}} = 0$ for all nodes.

We can partition the set of all possible configurations, \mathcal{X} , as

$$\mathcal{X} = \mathcal{X}_{0,0,\dots,0} \cup \mathcal{X}_{1,0,\dots,0} \cup \mathcal{X}_{2,0,\dots,0} \cup \dots \cup \mathcal{X}_{N_1, N_2, \dots, N_m} \quad (4.9)$$

where $\mathcal{X}_{s_1, s_2, \dots, s_m}$ is the set of configurations with $0 \leq s_1 \leq N_1$ infected nodes in V_1 , $0 \leq s_2 \leq N_2$ infected nodes in V_2 , and so forth. Let $\mathbf{x}^{s_1, s_2, \dots, s_m}$ denote a particular configuration in $\mathcal{X}_{s_1, s_2, \dots, s_m}$.

For brevity, we state without proof the generalized version of Lemma 4.6.2. For the complete multipartite graph, all configurations, $\mathbf{x}^{s_1, s_2, \dots, s_m}$, belonging to the same partition, $\mathcal{X}_{s_1, s_2, \dots, s_m}$, have the same equilibrium probability.

Lemma 4.6.4 (Proof in Appendix 4.9). *When $0 \leq \frac{\lambda}{\mu} \leq 1$, $\gamma > 1$, and for $s_1 = 0, 1, 2, \dots, N_1$, $s_2 = 0, 1, 2, \dots, N_2, \dots, s_m = 0, 1, 2, \dots, N_m$, the unnormalized equilibrium distribution for a complete multipartite graph is upperbounded by an exponential function:*

$$\begin{aligned} \pi(\mathbf{x}^{s_1, s_2, \dots, s_m}) &\propto \left(\frac{\lambda}{\mu}\right)^{\sum_{p=1}^m s_p} \gamma^{\frac{(\mathbf{x}^{s_1, s_2, \dots, s_m})^T \mathbf{A} \mathbf{x}^{s_1, s_2, \dots, s_m}}{2}} \\ &\leq \left(\frac{\lambda}{\mu} \gamma^\beta\right)^{\sum_{p=1}^m s_p} \end{aligned} \quad (4.10)$$

where

$$\beta = \frac{\sum_{i=1}^m N_i (\sum_{j \neq i} \mathbf{1}(V_i \sim V_j) N_j)}{2 \sum_{p=1}^m N_p}.$$

Furthermore, the inequality holds with equality for $\mathbf{x}^{s_1, s_2, \dots, s_m} = \mathbf{x}^N$ and $\mathbf{x}^{s_1, s_2, \dots, s_m} = \mathbf{x}^0$.

Theorem 4.6.3. *For the complete multipartite graph, when $0 \leq \frac{\lambda}{\mu} \leq 1$, $\gamma > 1$, a threshold exists for the most probable network configuration. Let*

$$\beta = \frac{\sum_{i=1}^m N_i (\sum_{j \neq i} \mathbf{1}(V_i \sim V_j) N_j)}{2 \sum_{p=1}^m N_p}$$

Then

1. $\lambda \gamma^\beta > \mu$ if and only if \mathbf{x}^* is unique and $\mathbf{x}^* = \mathbf{x}^N$.
2. $\lambda \gamma^\beta < \mu$ if and only if \mathbf{x}^* is unique and $\mathbf{x}^* = \mathbf{x}^0$.
3. $\lambda \gamma^\beta = \mu$ if and only if \mathbf{x}^* is no longer the unique maximizer, $\mathbf{x}^* = \mathbf{x}^N$, $\mathbf{x}^* = \mathbf{x}^0$.

Proof. The proof follows that of Theorem 4.6.1. Note that the complete bipartite graph is a special case of the complete multipartite graph where $m = 2$ and $\sum_{j \neq 1} \mathbf{1}(V_1 \sim V_j) N_j = N_2$ and $\sum_{j \neq 2} \mathbf{1}(V_2 \sim V_j) N_j = N_1$. \square

4.6.4 Complete Multipartite with k -regular islands

Consider the complete multipartite graph with k -regular supernodes, $G(V, E)$, where the intra-island degree, $0 < k^{\text{intra}} \leq N_i, \forall i = 1, \dots, m$.

Lemma 4.6.5 (Proof in Appendix 4.9). *When $0 \leq \frac{\lambda}{\mu} \leq 1, \gamma > 1$, and for $s_1 = 0, 1, 2, \dots, N_1, s_2 = 0, 1, 2, \dots, N_2, \dots, s_m = 0, 1, 2, \dots, N_m$, the unnormalized equilibrium distribution for the complete multipartite graph with regular islands is upperbounded by an exponential function:*

$$\begin{aligned} \pi(\mathbf{x}^{s_1, s_2, \dots, s_m}) &\propto \left(\frac{\lambda}{\mu}\right)^{\sum_{p=1}^m s_p} \gamma^{\frac{(\mathbf{x}^{s_1, s_2, \dots, s_m})^T A \mathbf{x}^{s_1, s_2, \dots, s_m}}{2}} \\ &\leq \left(\frac{\lambda}{\mu} \gamma^\beta\right)^{\sum_{p=1}^m s_p} \end{aligned} \quad (4.11)$$

where

$$\beta = \sum_{i=1}^m \frac{k_i^{\text{intra}}}{2} + N_i \left(\frac{\sum_{j \neq i} \mathbb{1}(V_i \sim V_j) N_j}{2 \sum_{p=1}^m N_p} \right).$$

Furthermore, the relationship holds with equality for $\mathbf{x}^{s_1, s_2, \dots, s_m} = \mathbf{n}^N$ and $\mathbf{x}^{s_1, s_2, \dots, s_m} = \mathbf{x}^0$.

Theorem 4.6.4. *For the complete multipartite graph with k -regular islands, when $0 \leq \frac{\lambda}{\mu} \leq 1, \gamma > 1$, a threshold exists for the most probable network configuration. Let*

$$\beta = \sum_{i=1}^m \frac{k_i^{\text{intra}}}{2} + N_i \left(\frac{\sum_{j \neq i} \mathbb{1}(V_i \sim V_j) N_j}{2 \sum_{p=1}^m N_p} \right)$$

Then

1. $\lambda \gamma^\beta > \mu$ if and only if \mathbf{x}^* is unique and $\mathbf{x}^* = \mathbf{x}^N$.
2. $\lambda \gamma^\beta < \mu$ if and only if \mathbf{x}^* is unique and $\mathbf{x}^* = \mathbf{x}^0$.
3. $\lambda \gamma^\beta = \mu$ if and only if \mathbf{x}^* is no longer the unique maximizer, $\mathbf{x}^* = \mathbf{x}^N, \mathbf{x}^* = \mathbf{x}^0$.

Proof. The proof follows that of Theorem 4.6.1. □

4.6.5 Unstructured Networks

For networks that are not k -regular, complete multipartite, or complete multipartite with k -regular islands, in regime II), the most-probable configuration, \mathbf{x}^* , may be configurations other than \mathbf{x}^0 or \mathbf{x}^N . We call solutions of (4.1) that are neither \mathbf{x}^0 nor \mathbf{x}^N , **non-degenerate** most-probable

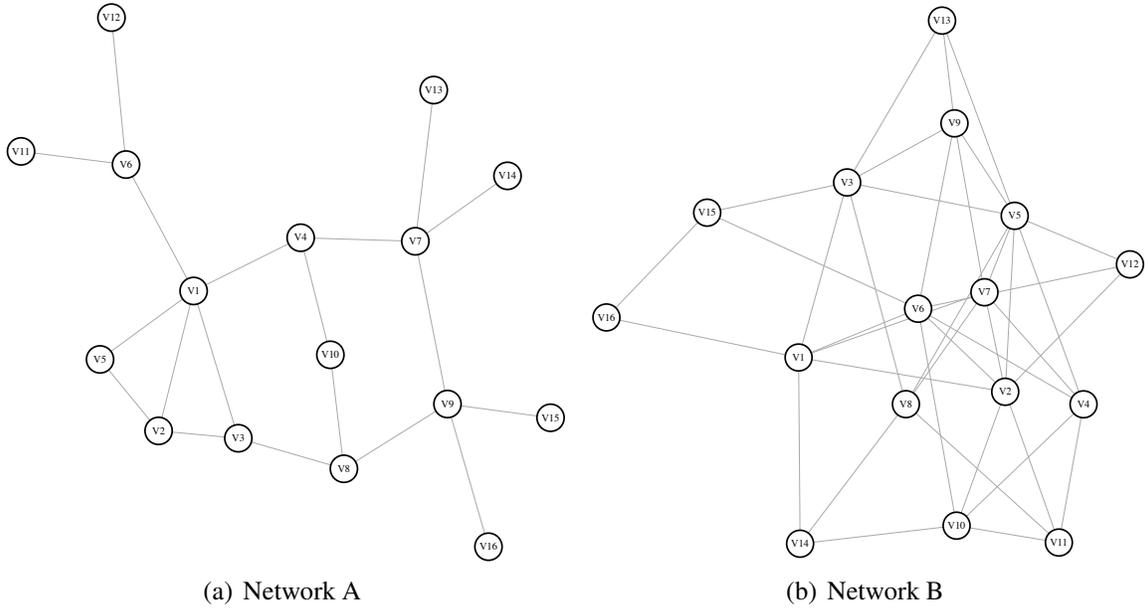


Figure 4.1: Example Network Topologies

configurations. Non-degenerate solutions mean that subsets of agents in the network are more vulnerable to infection than others.

For example, consider the two 16-node networks shown Figure 4.1(a) and Figure 4.1(b). Figure 4.2 shows the corresponding solution space of the Most-Probable Configuration Problem for $\frac{\lambda}{\mu} = 0.5$. The X-axis shows different γ values ranging from 1 to 4. On the Y-axis, we plot the total number of infected agents in \mathbf{x}^* . When the number of infected nodes is 0, $\mathbf{x}^* = \mathbf{x}^0$; when the number of infected nodes is 16, $\mathbf{x}^* = \mathbf{x}^N$. As γ increases, \mathbf{x}^* switches from \mathbf{x}^0 to \mathbf{x}^N as per our intuition.

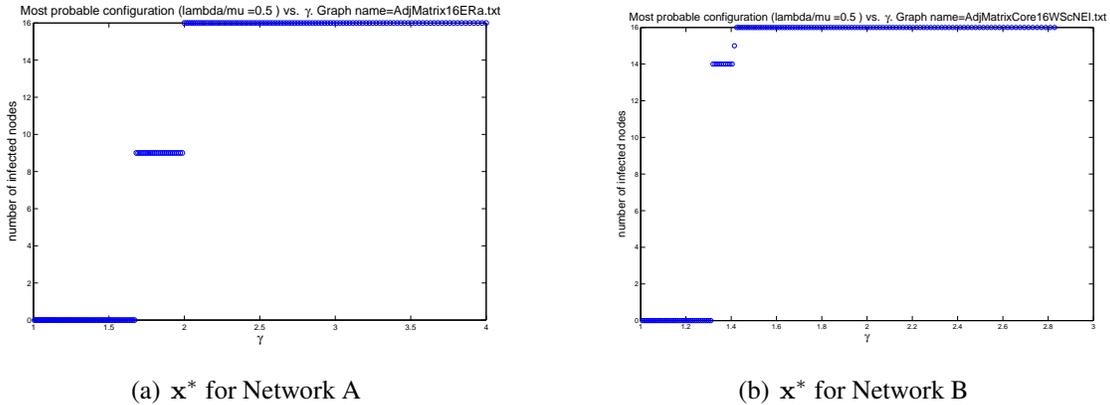


Figure 4.2: Solution Space of the Most-Probable Configuration Problem

However, for a small range of γ values the number of infected nodes in \mathbf{x}^* is neither 0 nor 16. These are non-degenerate solutions to the Most-Probable Configuration Problem (4.1). For network A, it is 9 agents out of the 16 total agents that are infected while the other 7 agents remain healthy. For network B, depending on the parameters, either 14 or 15 agents are more vulnerable than others. From Figure 4.3, we can see that there appears to be a structural dependence regarding which nodes in the network are more susceptible. In the next chapter, we will show that, surprisingly, these non-degenerate most-probable configurations can be found efficiently using polynomial time algorithms. We then prove the conditions for non-degenerate most-probable configurations and their relationships to the underlying network structure in the next chapter.

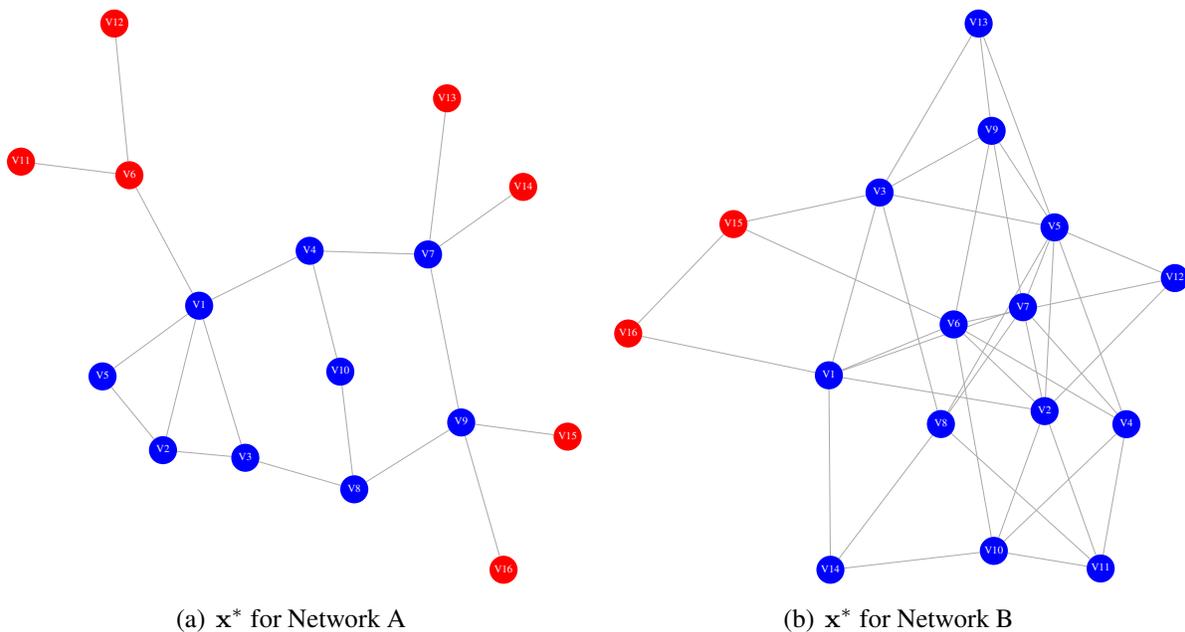


Figure 4.3: Non-Degenerate Most-Probable Configurations (Blue = Infected, Red = Healthy)

4.7 Ordering of Probability of \mathbf{x}^*

This section analyzes how $\pi(\mathbf{x}^*)$, the probability of the most-probable configuration, varies with parameter values in all 4 parameter regimes.

Lemma 4.7.1 (Proof in Appendix 4.9). *Consider two nonnegative vectors $\mathbf{x} = [x_1, \dots, x_n]$ and $\mathbf{y} = [y_1, \dots, y_n]$. We use $\mathbf{x}^\delta = [x_1^\delta, \dots, x_n^\delta]$ and $\mathbf{y}^\zeta = [y_1^\zeta, \dots, y_n^\zeta]$ to denote entrywise power. We*

denote entrywise product as

$$\mathbf{x} \circ \mathbf{y} = [x_1 y_1, \dots, x_n y_n].$$

If $\delta > 1$ and $\zeta > 1$, then

$$\left\| \frac{\mathbf{x}^\delta \circ \mathbf{y}^\zeta}{(\mathbf{x}^\delta)^T (\mathbf{y}^\zeta)} \right\|_\infty \geq \left\| \frac{\mathbf{x} \circ \mathbf{y}}{\mathbf{x}^T \mathbf{y}} \right\|_\infty$$

Theorem 4.7.1. For any network topology $G(V, E)$, consider two scaled SIS processes, $\{X(t), t \geq 0\}$ and $\{X'(t), t \geq 0\}$. Process $X(t)$ is parameterized by (λ, μ, γ) and process $X'(t)$ parameterized by $(\lambda', \mu', \gamma')$ with corresponding equilibrium distributions $\pi(\mathbf{x})$ and $\pi(\mathbf{x}')$ and most-probable network configurations, \mathbf{x}^* and \mathbf{x}'^* .

- **Regime I) Healing Dominant:** $0 < \frac{\lambda}{\mu} \leq 1, 0 < \gamma \leq 1$. If $\frac{\lambda'}{\mu'} < \frac{\lambda}{\mu}$ and/or $\gamma' < \gamma$, then $\pi(\mathbf{x}'^*) \geq \pi(\mathbf{x}^*)$.
- **Regime II) Endogenous Infection Dominant:** $0 < \frac{\lambda}{\mu} \leq 1, \gamma > 1$. If $\frac{\lambda'}{\mu'} < \frac{\lambda}{\mu}$ and/or $\gamma' > \gamma$, then $\pi(\mathbf{x}'^*) \geq \pi(\mathbf{x}^*)$.
- **Regime III) Exogenous Infection Dominant:** $\frac{\lambda}{\mu} > 1, 0 < \gamma \leq 1$. If $\frac{\lambda'}{\mu'} > \frac{\lambda}{\mu}$ and/or $\gamma' < \gamma$, then $\pi(\mathbf{x}'^*) \geq \pi(\mathbf{x}^*)$.
- **Regime IV) Infection Dominant:** $\frac{\lambda}{\mu} > 1, \gamma > 1$. If $\frac{\lambda'}{\mu'} > \frac{\lambda}{\mu}$ and/or $\gamma' > \gamma$, then $\pi(\mathbf{x}'^*) \geq \pi(\mathbf{x}^*)$.

Proof. We will prove Theorem 4.7.1 for regime IV; the other 3 regimes use similar arguments. We want to show that when $\frac{\lambda}{\mu} > 1, \gamma > 1$, if $\frac{\lambda'}{\mu'} > \frac{\lambda}{\mu}$ and/or $\gamma' > \gamma$, then $\pi(\mathbf{x}'^*) \geq \pi(\mathbf{x}^*)$.

We can express the parameters of $X'(t)$ as scaled versions of the parameters of $X(t)$ where $\lambda' = \lambda^\delta, \mu' = \mu^\delta, \gamma' = \gamma^\zeta$. Depending on if the parameters in $X'(t)$ are larger than, smaller than, or equal to the parameters of $X(t)$, we can set $\delta > 1$ or $0 < \delta \leq 1$, and likewise for ζ .

Alternatively, the equilibrium distribution of the scaled SIS process can be written as the Gibbs distribution. The equilibrium distribution of $X(t)$ is

$$\pi(\mathbf{x}) = \frac{1}{Z} e^{H(\mathbf{x})}, \tag{4.12}$$

where

$$H(\mathbf{x}) = \mathbf{1}^T \mathbf{x} \log \left(\frac{\lambda}{\mu} \right) + \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2} \log(\gamma).$$

The equilibrium distribution of $X'(t)$ is

$$\pi'(\mathbf{x}) = \frac{1}{Z'} e^{H'(\mathbf{x})}, \quad (4.13)$$

where

$$\begin{aligned} H'(\mathbf{x}) &= \mathbf{1}^T \mathbf{x} \log \left(\frac{\lambda'}{\mu'} \right) + \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2} \log(\gamma') \\ &= \delta \left(\mathbf{1}^T \mathbf{x} \log \left(\frac{\lambda}{\mu} \right) \right) + \zeta \left(\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2} \log(\gamma) \right). \end{aligned}$$

We can write $\pi'(\mathbf{x})$ as

$$\pi'(\mathbf{x}) = \frac{1}{Z'} \left(e^{\mathbf{1}^T \mathbf{x} \log \left(\frac{\lambda}{\mu} \right)} \right)^\delta \left(e^{\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2} \log(\gamma)} \right)^\zeta, \quad \mathbf{x} \in \mathcal{X}$$

Since the equilibrium is a PMF, there is a bijective mapping from $\pi(\mathbf{x})$ to the finite length vector $\frac{\mathbf{x} \circ \mathbf{y}}{\mathbf{x}^T \mathbf{y}}$ and from $\pi'(\mathbf{x})$ to the vector $\frac{\mathbf{x}^{\delta \circ \mathbf{y} \zeta}}{(\mathbf{x}^\delta)^T (\mathbf{y}^\zeta)}$. Using Lemma 4.7.1, we conclude that

$$\pi(\mathbf{x}^*) = \max_{\mathbf{x}' \in \mathcal{X}} \pi(\mathbf{x}') \geq \pi(\mathbf{x}^*) = \max_{\mathbf{x} \in \mathcal{X}} \pi(\mathbf{x})$$

□

4.8 Conclusion

This chapter studied the problem of solving for the configuration with the maximum equilibrium probability, \mathbf{x}^* . This inference problem is a combinatorial optimization problem. We showed that depending on the dynamics parameters, λ, γ, μ , the problem has different computational complexity. When $0 < \frac{\lambda}{\mu} \leq 1, 0 < \gamma \leq 1$ or $\frac{\lambda}{\mu} > 1, \gamma > 1$, the solution of the inference problem is uninteresting since, *regardless* of the underlying network structure, the most-probable configuration is either \mathbf{x}^0 , the configuration where all the agents are healthy, or \mathbf{x}^N , the configuration where all the agents are infected.

The most-probable configuration is more interesting in regime II) **endogenous infection dominant**: $0 < \frac{\lambda}{\mu} \leq 1, \gamma > 1$ and regime III) **exogenous infection dominant**: $\frac{\lambda}{\mu} > 1, 0 < \gamma \leq 1$ since the preference of the individual agents (i.e., controlled by $\frac{\lambda}{\mu}$) opposes the network effect (i.e.,

controlled by γ). Conflicting interests between individual preferences and network effect result in different solutions to the Most-Probable Configuration Problem depending on the parameter values. The solution of the inference problem in regime III) is related to the Maximum Independent Set Problem, a graph theoretic problem that is NP-hard to solve for all graph structures except perfect graphs. As a result, regime III), while interesting, is computationally infeasible to analyze for large networks.

Regime II) on the other hand, is more amenable for analysis. For k -regular, complete multipartite, and complete multipartite with k -regular islands, we can characterize the solution space exactly, including finding the threshold where the most-probable configuration transitions from \mathbf{x}^0 to \mathbf{x}^N . For other networks however, the most-probable configuration may be a non-degenerate solution where only a subset of agents is infected. This means that, at equilibrium, for these network structures, some agents are more vulnerable than others. These non-degenerate solutions are studied in more depth in the next chapter. Lastly, we showed that the probability of \mathbf{x}^* increases for extreme parameter values (i.e., either much larger than 1, or much smaller than 1).

4.9 Appendix

Proof of Corollary 4.4.4

Corollary. *For the complete graph, K_N , when $\frac{\lambda}{\mu} > 1, 0 < \gamma \leq 1$, if in addition $\lambda\gamma < \mu$, then*

$$\mathbf{x}^* \in \{\mathbf{x} \in \mathcal{X} \mid \mathbf{1}^T \mathbf{x} = 1\}.$$

Proof. For the complete graph, every node is connected to every other node. Therefore, if there are more than 2 infected nodes in a network configuration, at least 1 infected edge will be created. The maximum number of infected nodes that results in 0 infected edges is 1.

□

Proof of Corollary 4.4.5

Corollary. *For the bipartite graph, when $\frac{\lambda}{\mu} > 1, 0 < \gamma \leq 1$, if in addition $\lambda\gamma < \mu$, then*

$$\mathbf{x}^* \in \{\mathbf{x} \in \mathcal{X} \mid \mathbf{1}^T \mathbf{x} = \max\{N_1, N_2\}, \mathbf{x}^T \mathbf{A} \mathbf{x} = 0\}.$$

Proof. Without loss of generality, we assume that $N_1 \geq N_2$. We will prove that $\mathbf{x}^* \in \{\mathbf{x} \in \mathcal{X} \mid \mathbf{1}^T \mathbf{x} = N_1, \mathbf{x}^T \mathbf{A} \mathbf{x} = 0\}$ by contradiction. Suppose that \mathbf{x}' is a configuration where $\pi(\mathbf{x}') > \pi(\mathbf{x}^*)$.

Case 1: Consider when $\mathbf{1}^T \mathbf{x}' < N_1$. By the definition of a bipartite graph, we can obtain 0 infected edges by placing all the infected agents in V_1 . Since $\frac{\lambda}{\mu} > 1$, we know that $\pi(\mathbf{x}') < \pi(\mathbf{x}^*)$, which contradicts the supposition.

Case 2: Consider when $\mathbf{1}^T \mathbf{x}' = N_1 + d > N_1$, $d = 1, 2, \dots, N_2$. Since $|V_1| = N_1$, the minimum number of infected edges that will be created is d (i.e., we place all N_1 infected nodes in V_1 which generates no infected edges for a bipartite graph and the rest of the infected nodes in V_2).

Therefore

$$\pi(\mathbf{x}') = \left(\frac{\lambda}{\mu}\right)^{N_1+d} \gamma^d, \quad d = 1, 2, \dots, N_2$$

whereas

$$\pi(\mathbf{x}^*) = \left(\frac{\lambda}{\mu}\right)^{N_1}$$

Realize that

$$\frac{\pi(\mathbf{x}')}{\pi(\mathbf{x}^*)} = \left(\frac{\lambda}{\mu} \gamma\right)^d, \quad d = 1, 2, \dots, N_2$$

Since $\lambda\gamma < \mu$, $\pi(\mathbf{x}') < \pi(\mathbf{x}^*)$, which contradicts the supposition.

□

Proof of Lemma 4.6.1

Lemma. When $0 \leq \frac{\lambda}{\mu} \leq 1, \gamma > 1$ and for $s = 0, 1, \dots, N, \mathbf{x}^s \in \mathcal{X}_s$, the unnormalized equilibrium distribution for a k -regular graph is upperbounded by an exponential function:

$$\pi(\mathbf{x}^s) \propto \left(\frac{\lambda}{\mu}\right)^s \gamma^{\frac{\mathbf{x}^s T \mathbf{A} \mathbf{x}^s}{2}} \leq \left(\frac{\lambda}{\mu} \gamma^\beta\right)^s \quad (4.14)$$

where $\beta = \frac{k}{2}$. Furthermore, the relationship holds with equality for \mathbf{x}^0 and \mathbf{x}^N .

Proof. For the k -regular graph

$$\mathbf{A} \mathbf{x}^s = \begin{bmatrix} \sum_{i=1}^N A_{1i} x_i^s \\ \sum_{i=1}^N A_{2i} x_i^s \\ \vdots \\ \sum_{i=1}^N A_{Ni} x_i^s \end{bmatrix} \leq \begin{bmatrix} k \\ k \\ \vdots \\ k \end{bmatrix}, \quad \text{if } k \leq s$$

and

$$\mathbf{Ax}^s = \begin{bmatrix} \sum_{i=1}^N A_{1i}x_i^s \\ \sum_{i=1}^N A_{2i}x_i^s \\ \vdots \\ \sum_{i=1}^N A_{Ni}x_i^s \end{bmatrix} \leq \begin{bmatrix} s \\ s \\ \vdots \\ s \end{bmatrix}, \quad \text{if } k \geq s$$

Then

$$\begin{aligned} \mathbf{x}^{sT} \mathbf{Ax}^s &\leq sk, \text{ if } k \leq s \\ \mathbf{x}^{sT} \mathbf{Ax}^s &\leq s^2 \leq sk, \text{ if } k \geq s \end{aligned} \quad (4.15)$$

With $\gamma > 1$ and (4.15), we deduce that

$$\left(\frac{\lambda}{\mu}\right)^s \gamma^{\frac{\mathbf{x}^{sT} \mathbf{Ax}^s}{2}} \leq \left(\frac{\lambda}{\mu}\right)^s \gamma^{\frac{sk}{2}}, \forall s = 0, 1, \dots, N \quad (4.16)$$

When $s = 0$, equality is satisfied in (4.16). When $s = N$, $\frac{\mathbf{x}^{sT} \mathbf{Ax}^s}{2} = \frac{kN}{2}$; hence equality is also satisfied.

□

Proof of Lemma 4.6.2

Lemma. *For the complete bipartite graph, all configurations, \mathbf{x}^{s_1, s_2} , belonging to the same partition, \mathcal{X}_{s_1, s_2} , have the same equilibrium probability.*

Proof. Consider two configurations, $\mathbf{x}_1^{s_1, s_2}, \mathbf{x}_2^{s_1, s_2} \in \mathcal{X}_{s_1, s_2}$. By definition (4.7), $\mathbf{1}^T \mathbf{x}_1^{s_1, s_2} = s_1 + s_2$ and $\mathbf{1}^T \mathbf{x}_2^{s_1, s_2} = s_1 + s_2$. Furthermore, if the underlying topology is a complete bipartite graph, we know that $(\mathbf{x}_1^{s_1, s_2})^T \mathbf{Ax}_1^{s_1, s_2} = s_1 s_2$ and $(\mathbf{x}_2^{s_1, s_2})^T \mathbf{Ax}_2^{s_1, s_2} = s_1 s_2$. Since the number of infected nodes and the number of infected edges are the same, then $\pi(\mathbf{x}_1^{s_1, s_2}) = \pi(\mathbf{x}_2^{s_1, s_2})$.

□

Proof of Lemma 4.6.3

Lemma. *When $0 \leq \frac{\lambda}{\mu} \leq 1, \gamma > 1$ and for $s_1 = 0, 1, 2, \dots, N_1, s_2 = 0, 1, 2, \dots, N_2$, the unnormalized equilibrium distribution for a complete bipartite graph is upperbounded by an exponential function:*

$$\pi(\mathbf{x}^{s_1, s_2}) \propto \left(\frac{\lambda}{\mu}\right)^{s_1 + s_2} \gamma^{s_1 s_2} \leq \left(\frac{\lambda}{\mu} \gamma^\beta\right)^{s_1 + s_2} \quad (4.17)$$

where $\beta = \frac{N_1 N_2}{N_1 + N_2}$. Furthermore, the relationship holds with equality for $\mathbf{x}^{s_1, s_2} = \mathbf{x}^N$ and $\mathbf{x}^{s_1, s_2} = \mathbf{x}^0$.

Proof. Proving the inequality in (4.17) is equivalent to showing that

$$\frac{s_1 s_2}{s_1 + s_2} \leq \frac{N_1 N_2}{N_1 + N_2}, \text{ for } s_1 = 0, 1, \dots, N_1, \quad (4.18)$$

$$s_2 = 0, 1, \dots, N_2, \text{ excluding } s_1 = 0, s_2 = 0$$

Equation (4.18) implies that $\frac{N_1 N_2}{N_1 + N_2} - \frac{s_1 s_2}{s_1 + s_2} \geq 0$, which we can restate as

$$\frac{N_1 N_2 (s_1 + s_2) - s_1 s_2 (N_1 + N_2)}{(N_1 + N_2)(s_1 + s_2)} \geq 0 \quad (4.19)$$

The denominator of the LHS of (4.19) is always positive, so we only need to consider the numerator.

Realize that we can rewrite the numerator term of the LHS of (4.19) as $N_1 s_1 (N_2 - s_2) + N_2 s_2 (N_1 - s_1)$. When $s_1 = 0, 1, \dots, N_1 - 1, s_2 = 0, 1, \dots, N_2 - 1$,

$$N_1 s_1 (N_2 - s_2) + N_2 s_2 (N_1 - s_1) > 0 \quad (4.20)$$

When $s_1 = N_1, s_2 = N_2$

$$N_1 s_1 (N_2 - s_2) + N_2 s_2 (N_1 - s_1) = 0 \quad (4.21)$$

Therefore, (4.18) is satisfied. Since $\gamma > 1$, we can conclude that

$$\left(\frac{\lambda}{\mu} \gamma^{\frac{s_1 s_2}{s_1 + s_2}} \right)^{s_1 + s_2} \leq \left(\frac{\lambda}{\mu} \gamma^{\frac{N_1 N_2}{N_1 + N_2}} \right)^{s_1 + s_2} \quad (4.22)$$

for $s_1 = 0, 1, \dots, N_1, s_2 = 0, 1, \dots, N_2$. Equality is satisfied for $s_1 = N_1, s_2 = N_2$, which corresponds to the network configuration \mathbf{x}^N .

Consider the special case of $s_1 = 0, s_2 = 0$. If we rewrite $\left(\frac{\lambda}{\mu} \gamma^{\frac{s_1 s_2}{s_1 + s_2}} \right)^{s_1 + s_2}$ as $\left(\frac{\lambda}{\mu} \right)^{s_1 + s_2} \gamma^{s_1 s_2}$, then the relationship in (4.22) is also satisfied with equality.

□

Proof of Lemma 4.6.4

Lemma. When $0 \leq \frac{\lambda}{\mu} \leq 1, \gamma > 1$, and for $s_1 = 0, 1, 2, \dots, N_1, s_2 = 0, 1, 2, \dots, N_2, \dots, s_m = 0, 1, 2, \dots, N_m$, the unnormalized equilibrium distribution for a complete multipartite graph is

upperbounded by an exponential function:

$$\begin{aligned}\pi(\mathbf{x}^{s_1, s_2, \dots, s_m}) &\propto \left(\frac{\lambda}{\mu}\right)^{\sum_{p=1}^m s_p} \gamma^{\frac{(\mathbf{x}^{s_1, s_2, \dots, s_m})^T \mathbf{A} \mathbf{x}^{s_1, s_2, \dots, s_m}}{2}} \\ &\leq \left(\frac{\lambda}{\mu} \gamma^\beta\right)^{\sum_{p=1}^m s_p}\end{aligned}\quad (4.23)$$

where

$$\beta = \frac{\sum_{i=1}^m N_i (\sum_{j \neq i} \mathbb{1}(V_i \sim V_j) N_j)}{2 \sum_{p=1}^m N_p}.$$

Furthermore, the inequality holds with equality for $\mathbf{x}^{s_1, s_2, \dots, s_m} = \mathbf{x}^N$ and $\mathbf{x}^{s_1, s_2, \dots, s_m} = \mathbf{x}^0$.

Proof. This proof follows the same reasoning as the proof of Lemma 4.6.3 in appendix 4.9; certain steps are skipped for brevity. Proving the inequality in (4.10) is equivalent to showing that for $s_1 = 0, 1, 2, \dots, N_1, s_2 = 0, 1, 2, \dots, N_2, \dots, s_m = 0, 1, 2, \dots, N_m$ (excluding the case $s_1 = 0, s_2 = 0, \dots, s_m = 0$, which we will handle separately),

$$\beta - \frac{\sum_{i=1}^m s_i (\sum_{j \neq i} \mathbb{1}(V_i \sim V_j) s_j)}{2 \sum_{p=1}^m s_p} \geq 0 \quad (4.24)$$

Realize that the sign of the LHS of (4.24) is determined by

$$\begin{aligned}&\left(\sum_{i=1}^m N_i \left(\sum_{j \neq i} \mathbb{1}(V_i \sim V_j) N_j \right) \right) \left(\sum_{p=1}^m s_p \right) - \\ &\left(\sum_{i=1}^m s_i \left(\sum_{j \neq i} \mathbb{1}(V_i \sim V_j) s_j \right) \right) \left(\sum_{p=1}^m N_p \right)\end{aligned}\quad (4.25)$$

which we can rewrite as

$$\sum_{i=1}^m \sum_{p=1}^m N_i s_p \left(\sum_{j \neq i} \mathbb{1}(V_i \sim V_j) (N_j - s_j) \right) \quad (4.26)$$

Equation (4.26) is positive when $s_i = 0, 1, \dots, N_i, \forall i = 1, \dots, m$ and is 0 when $s_i = N_i, \forall i = 1, \dots, m$, which means that (4.24) is satisfied.

Consider the special case of $s_1 = 0, s_2 = 0, \dots, s_m = 0$. The inequality in (4.10) is satisfied with equality trivially.

□

Proof of Lemma 4.6.5

Lemma 4.9.1. *When $0 \leq \frac{\lambda}{\mu} \leq 1, \gamma > 1$, and for $s_1 = 0, 1, 2, \dots, N_1, s_2 = 0, 1, 2, \dots, N_2, \dots, s_m = 0, 1, 2, \dots, N_m$, the unnormalized equilibrium distribution for the complete multipartite graph with k -regular islands is upperbounded by an exponential function:*

$$\begin{aligned} \pi(\mathbf{x}^{s_1, s_2, \dots, s_m}) &\propto \left(\frac{\lambda}{\mu}\right)^{\sum_{p=1}^m s_p} \gamma^{\frac{(\mathbf{x}^{s_1, s_2, \dots, s_m})^T \mathbf{A} \mathbf{x}^{s_1, s_2, \dots, s_m}}{2}} \\ &\leq \left(\frac{\lambda}{\mu} \gamma^\beta\right)^{\sum_{p=1}^m s_p} \end{aligned} \quad (4.27)$$

where

$$\beta = \sum_{i=1}^m \frac{k_i^{\text{intra}}}{2} + N_i \left(\frac{\sum_{j \neq i} \mathbb{1}(V_i \sim V_j) N_j}{2 \sum_{p=1}^m N_p} \right).$$

Furthermore, the relationship holds with equality for $\mathbf{x}^{s_1, s_2, \dots, s_m} = \mathbf{x}^N$ and $\mathbf{x}^{s_1, s_2, \dots, s_m} = \mathbf{x}^0$.

Proof. Recall that N_i is the total number of nodes in island i . There are m islands so the total number of nodes in the entire network is $N = \sum_{i=1}^m N_i$. The number of infected nodes in the i th island is s_i . Without loss of generality, we can label the nodes in such a way that x_1, \dots, x_{N_1} refers to the state of all the nodes in island 1 and $x_{N_1+1}, \dots, x_{N_1+N_2}$ refers to the state of all the nodes in island 2, and so forth.

Proving (4.27) is equivalent to proving that for $s_1 = 0, 1, 2, \dots, N_1, s_2 = 0, 1, 2, \dots, N_2, \dots, s_m = 0, 1, 2, \dots, N_m$ (excluding the case $s_1 = 0, s_2 = 0, \dots, s_m = 0$, which we will handle separately),

$$\frac{(\mathbf{x}^{s_1, s_2, \dots, s_m})^T \mathbf{A} \mathbf{x}^{s_1, s_2, \dots, s_m}}{2 \sum_{p=1}^m s_p} \leq \beta \quad (4.28)$$

We will partition the rows of A into blocks corresponding to each island. Let $A_{[i]}$ denote the rows of the i th block where each block corresponds to different islands, $A = \{A_{[i]}\}_{i=1}^m$. Following the same reasoning as the proof of Lemma 4.6.1 in appendix 4.9, we can deduce the following bound for each island

$$A_{[i]} \cdot \mathbf{x}^{s_1, s_2, \dots, s_m} = \begin{cases} k_i^{\text{intra}} + \sum_{j \neq i} \mathbb{1}(V_i \sim V_j) s_j, & \text{if } k_i^{\text{intra}} < s_i \\ s_i + \sum_{j \neq i} \mathbb{1}(V_i \sim V_j) s_j, & \text{if } k_i^{\text{intra}} \geq s_i \end{cases}$$

Recognize that then

$$\begin{aligned} (\mathbf{x}^{s_1, s_2, \dots, s_m})^T \mathbf{A} \mathbf{x}^{s_1, s_2, \dots, s_m} &\leq \\ \sum_{i=1}^m s_i (k_i^{\text{intra}} + \sum_{i \neq j} \mathbb{1}(V_i \sim V_j) s_j) & \end{aligned} \quad (4.29)$$

We can rewrite the RHS of (4.29) as

$$\sum_{i=1}^m s_i(k_i^{\text{intra}}) + \sum_{i=1}^m s_i \left(\sum_{i \neq j} \mathbb{1}(V_i \sim V_j) s_j \right) \quad (4.30)$$

Divide the LHS and RHS of (4.29) by $2 \sum_{p=1}^m s_p$. We obtain the new relationship

$$\frac{(\mathbf{x}^{s_1, s_2, \dots, s_m})^T \mathbf{A} \mathbf{x}^{s_1, s_2, \dots, s_m}}{2 \sum_{p=1}^m s_p} \leq \frac{k_i^{\text{intra}} \sum_{i=1}^m s_i + \sum_{i=1}^m s_i \left(\sum_{i \neq j} \mathbb{1}(V_i \sim V_j) s_j \right)}{2 \sum_{p=1}^m s_p} \quad (4.31)$$

Since the largest possible s_i is N_i , we can upper bound the RHS of (4.31) with

$$\sum_{i=1}^m \frac{k_i^{\text{intra}}}{2} + \frac{\sum_{i=1}^m N_i \left(\sum_{i \neq j} \mathbb{1}(V_i \sim V_j) N_j \right)}{2 \sum_{p=1}^m N_p} \quad (4.32)$$

This means that

$$\frac{(\mathbf{x}^{s_1, s_2, \dots, s_m})^T \mathbf{A} \mathbf{x}^{s_1, s_2, \dots, s_m}}{2 \sum_{p=1}^m s_p} \leq \sum_{i=1}^m \frac{k_i^{\text{intra}}}{2} + \frac{\sum_{i=1}^m N_i \left(\sum_{i \neq j} \mathbb{1}(V_i \sim V_j) N_j \right)}{2 \sum_{p=1}^m N_p} \quad (4.33)$$

which is the relationship we want to prove in (4.28). Note that the first term of the RHS of equation (4.33) is the same threshold that we derived for k -regular graphs and the second term is the same term we derived for complete multipartite graphs.

□

Proof of Lemma 4.7.1

Lemma. Consider two nonnegative vectors $\mathbf{x} = [x_1, \dots, x_n]$ and $\mathbf{y} = [y_1, \dots, y_n]$. We use $\mathbf{x}^\delta = [x_1^\delta, \dots, x_n^\delta]$ and $\mathbf{y}^\zeta = [y_1^\zeta, \dots, y_n^\zeta]$ to denote entrywise power. We denote entrywise product as

$$\mathbf{x} \circ \mathbf{y} = [x_1 y_1, \dots, x_n y_n].$$

If $\delta > 1$ and $\zeta > 1$, then

$$\left\| \frac{\mathbf{x}^\delta \circ \mathbf{y}^\zeta}{(\mathbf{x}^\delta)^T (\mathbf{y}^\zeta)} \right\|_\infty \geq \left\| \frac{\mathbf{x} \circ \mathbf{y}}{\mathbf{x}^T \mathbf{y}} \right\|_\infty \quad (4.34)$$

Proof. Without loss of generality, assume that $x_1 = \|\mathbf{x}\|_\infty$ and $y_1 = \|\mathbf{y}\|_\infty$. Since $\delta > 1$ and $\zeta > 1$, then $x_1^\delta = \|\mathbf{x}^\delta\|_\infty$ and $y_1^\zeta = \|\mathbf{y}^\zeta\|_\infty$. Furthermore, as all the terms are nonnegative, $x_1 y_1 = \|\mathbf{x} \circ \mathbf{y}\|_\infty$ and $x_1^\delta y_1^\zeta = \|\mathbf{x}^\delta \circ \mathbf{y}^\zeta\|_\infty$.

To show (4.34) is equivalent as showing

$$\frac{x_1^\delta y_1^\zeta}{\sum_{i=1}^n x_i^\delta y_i^\zeta} - \frac{x_1 y_1}{\sum_{i=1}^n x_i y_i} \geq 0 \quad (4.35)$$

We can write the LHS of (4.35) as

$$x_1 y_1 \left(\frac{x_1^{\delta-1} y_1^{\zeta-1} (\sum_{i=1}^n x_i y_i) - (\sum_{i=1}^n x_i^\delta y_i^\zeta)}{(\sum_{i=1}^n x_i^\delta y_i^\zeta) (\sum_{i=1}^n x_i y_i)} \right) \quad (4.36)$$

The sign of (4.36) is determined by $x_1^{\delta-1} y_1^{\zeta-1} (\sum_{i=1}^n x_i y_i) - (\sum_{i=1}^n x_i^\delta y_i^\zeta)$, which we can rearrange into

$$\begin{aligned} & x_1^\delta y_1^m + x_1^{\delta-1} y_1^{\zeta-1} \left(\sum_{i=2}^n x_i y_i \right) - x_1^\delta y_1^\zeta - \left(\sum_{i=2}^n x_i^\delta y_i^\zeta \right) \\ &= x_1^{\delta-1} y_1^{\zeta-1} \left(\sum_{i=2}^n x_i y_i \right) - \left(\sum_{i=2}^n x_i^\delta y_i^\zeta \right) \\ &= \sum_{i=2}^n x_i y_i \left(x_1^{\delta-1} y_1^{\zeta-1} - x_i^{\delta-1} y_i^{\zeta-1} \right) \end{aligned} \quad (4.37)$$

Since all the terms are nonnegative and $x_1^\delta y_1^\zeta$ is the maximum, expression (4.37) is either greater than or equal to 0. This implies that

$$\frac{x_1^\delta y_1^\zeta}{\sum_{i=1}^n x_i^\delta y_i^\zeta} - \frac{x_1 y_1}{\sum_{i=1}^n x_i y_i} \geq 0$$

□

Most-Probable Configuration: Non-Degenerate Solutions

5.1 Introduction

When individual agents prefer the healthy state but network effects spread infection (i.e., $0 < \frac{\lambda}{\mu} \leq 1, \gamma > 1$), the solution to the Most-Probable Configuration Problem may be non-degenerate in that only subsets of agents in the network are infected while others remain healthy. This means that agents are *not* equally susceptible to the infection process. This chapter studies the non-degenerate solutions in regime II) and how it is related to the underlying network structure. First, we prove that the Most-Probable Configuration Problem, a combinatorial optimization problem, can be solved in polynomial-time when $\gamma \geq 1$. This is because it is finding the maximum of a supermodular function (i.e., solving for the minimum of a submodular function).

We then show an alternate formulation of the equilibrium distribution (3.4) using induced subgraphs and subgraph density. The existence of non-degenerate most-probable configurations is related to the existence of subgraphs in the network that are denser than the overall network. This means that agents that are more connected than the general population are more vulnerable to infection. Lastly, we discuss the uniqueness conditions of these non-degenerate solutions.

5.2 Submodular Function

This section proves that the Most-Probable Configuration Problem in regime II) can be solved exactly in polynomial-time for arbitrary network structure. We first present a brief summary on

the necessary concepts. Pseudo-Boolean functions are functions that map N binary variables to a real number [53]. Minimization of general pseudo-Boolean functions is NP-hard [54]. Grötschel, Lovász, and Schrijver, [55], proved that the minimization of a pseudo-Boolean function that is submodular could be solved in polynomial-time. If the function is supermodular, its maximization can be found in polynomial-time.

A pseudo-Boolean function, $f : \{0, 1\}^N \rightarrow \mathcal{R}$, is also a set function $g : \mathcal{P}(V) \rightarrow \mathcal{R}$ where $\mathcal{P}(V)$ is the power set of $V = \{1, 2, \dots, N\}$. There are many equivalent definitions of submodularity [56]. The one we use in this paper is the following:

Definition 5.2.1 ([53]). *A set function, $g : \mathcal{P}(V) \rightarrow \mathcal{R}$, is submodular if and only if for any $\alpha_1 \subseteq V, \alpha_2 \subseteq \alpha_1, i \in V \setminus \alpha_1$:*

$$g(\alpha_1 \cup \{i\}) - g(\alpha_1) \leq g(\alpha_2 \cup \{i\}) - g(\alpha_2).$$

For a submodular function, the incremental gain of adding an element to the set α_1 is less than or equal to the gain of adding the element to a smaller subset of α_1 . A supermodular function has the inequality in the opposite direction. Submodular functions are studied extensively in the field of combinatorial optimization.

5.2.1 Most-Probable Configuration: A Submodular Problem

The Most-Probable Configuration Problem (4.1) seeks the maximum of a pseudo-Boolean function that maps a 0-1 vector, the network configuration \mathbf{x} , to a scalar. The network configuration $\mathbf{x} \in \{0, 1\}^N$ is the characteristic vector or characteristic function of the set of infected agents: $\alpha_{\mathbf{x}} = \{i \mid i \in V, x_i = 1\}$. Let $h(\alpha_{\mathbf{x}})$ be the set of infected edges (i.e., edges where both end nodes are infected) in configuration \mathbf{x} : $h(\alpha_{\mathbf{x}}) = \{\{i, j\} \mid i, j \in V, A_{ij} = 1, x_i = 1, x_j = 1\}$.

The number of infected agents in configuration \mathbf{x} is $|\alpha_{\mathbf{x}}| = 1^T \mathbf{x}$. The number of infected edges is $|h(\alpha_{\mathbf{x}})| = \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2}$. The Most-Probable Configuration Problem is then to solve for the maximum argument of

$$g(\alpha_{\mathbf{x}}) = \left(\frac{\lambda}{\mu}\right)^{|\alpha_{\mathbf{x}}|} \gamma^{|h(\alpha_{\mathbf{x}})|}. \quad (5.1)$$

We will prove in Theorem 5.2.2 that $-\log(g(\alpha_{\mathbf{x}}))$ is a submodular function. Therefore, we can solve for its minimum argument in polynomial time. Lemma 5.2.1 sets up some basic conditions that makes proving Theorem 5.2.2 easier.

Lemma 5.2.1. Consider two sets of infected agents, $\alpha_1, \alpha_2 \subseteq V$ and $i \in V \setminus \alpha_1$. The cardinalities of α_1 and α_2 are $|\alpha_1| = n_1$ and $|\alpha_2| = n_2$, respectively; then $|\alpha_1 \cup \{i\}| = n_1 + 1$, and $|\alpha_2 \cup \{i\}| = n_2 + 1$. The numbers of infected edges induced by α_1 and α_2 are $|h(\alpha_1)| = e_1$ and $|h(\alpha_2)| = e_2$, respectively. Let $|h(\alpha_1 \cup \{i\})| = e_1 + m_1$ and $|h(\alpha_2 \cup \{i\})| = e_2 + m_2$; therefore m_1 is the number of additional infected edges created with the inclusion of agent i in α_1 and m_2 is the number of additional infected edges created with the inclusion of agent i in α_2 . Let $\alpha_2 \subseteq \alpha_1$. Then:

1. $n_1 \geq n_2$.
2. $e_1 \geq e_2$.
3. $m_1 \geq m_2$.

Proof. 1. When $\alpha_2 \subset \alpha_1$, α_2 must have strictly fewer infected agents than α_1 . When $\alpha_2 = \alpha_1$, then they contain the same number of infected agents. Hence, $n_1 \geq n_2$.

2. When $\alpha_2 \subset \alpha_1$, infected agents in α_2 can not induce more infected edges than the number of infected edges induced by the infected agents in α_1 . When $\alpha_2 = \alpha_1$, then the infected agents in α_1 and α_2 will induce the same number of infected edges. Hence, $e_1 \geq e_2$.

3. Every infected agent in α_2 is an infected agent in α_1 . Every new infected edge that is induced when adding infected agent i to α_2 is also a new infected edge when adding infected agent i to α_1 . Therefore, $m_1 \geq m_2$.

□

Theorem 5.2.2. Let $g(\alpha_x)$ be the set function given in (5.1). If $\lambda > 0, \mu > 0$ and $\gamma \geq 1$, then $-\log(g(\alpha_x))$ is a submodular function, where

$$-\log(g(\alpha_x)) = -|\alpha_x| \log\left(\frac{\lambda}{\mu}\right) - |h(\alpha_x)| \log(\gamma).$$

Proof. To prove submodularity of $-\log(g(\alpha_x))$, we need to show that

$$-\log(g(\alpha_1 \cup \{i\})) + \log(g(\alpha_1)) \leq -\log(g(\alpha_2 \cup \{i\})) + \log(g(\alpha_2)), \quad (5.2)$$

for any $\alpha_1 \subseteq V, \alpha_2 \subseteq \alpha_1, i \in V \setminus \alpha_1$.

The left-hand side (LHS) of (5.2) is

$$-(n_1 + 1) \log\left(\frac{\lambda}{\mu}\right) - (e_1 + m_1) \log(\gamma) + n_1 \log\left(\frac{\lambda}{\mu}\right) + e_1 \log(\gamma), \quad (5.3)$$

which reduces to

$$-\log\left(\frac{\lambda}{\mu}\right) - m_1 \log(\gamma). \quad (5.4)$$

The right-hand side (RHS) of (5.2) is

$$-(n_2 + 1) \log\left(\frac{\lambda}{\mu}\right) - (e_2 + m_2) \log(\gamma) + n_2 \log\left(\frac{\lambda}{\mu}\right) + e_2 \log(\gamma), \quad (5.5)$$

which reduces to

$$-\log\left(\frac{\lambda}{\mu}\right) - m_2 \log(\gamma). \quad (5.6)$$

Expression (5.2) reduces to

$$-\log\left(\frac{\lambda}{\mu}\right) - m_1 \log(\gamma) \leq -\log\left(\frac{\lambda}{\mu}\right) - m_2 \log(\gamma).$$

Since $\gamma \geq 1$, we know that $\log(\gamma) \geq 0$ and that $m_1 \geq m_2$ by Lemma 5.2.1. Therefore, the LHS of (5.2) is less than or equal to the RHS of (5.2) for any $\alpha_1 \subseteq V, \alpha_2 \subseteq \alpha_1, i \in V \setminus \alpha_1$. By definition, $-\log(g(\alpha_x))$ is a submodular function. \square

Theorem 5.2.2 proves that $-\log(g(\alpha_x))$ is submodular if $\lambda > 0, \mu > 0$, and $\gamma \geq 1$; this means that $\log(g(\alpha_x))$ is supermodular under the same condition. Since the logarithm function is a monotonic function, the maximum argument of $\log(g(\alpha_x))$ is also the maximum argument of $g(\alpha_x)$, which is the solution to the Most-Probable Configuration Problem. As regime II) **Endogenous Infection Dominant**: $0 < \frac{\lambda}{\mu} \leq 1, \gamma > 1$ satisfies the condition that $\gamma \geq 1$, using submodular minimization, the *exact* most-probable configuration of the scaled SIS process can be found for *arbitrary* network topology in polynomial-time.

Reference [57] shows an algorithm for using **Max-Flow/Min-Cut** algorithm for finding the minimum of submodular functions, which they called **regular** functions. This is a much more computationally efficient method for finding the most-probable configuration than using general submodular optimization solvers.

5.2.2 Relationship to Ground State

We showed in section 3.4 that the equilibrium distribution of the scaled SIS process is a Gibbs distribution, which means that it is also a Markov random field. An alternate interpretation of solving for the most-probable configuration of the scaled SIS process is that of finding the ground state of a Gibbs distribution; this is a well-studied problem in statistical mechanics especially relating to the **Ising model** [51, 58, 59]. However, models in statistical mechanics assume that the network configurations are $\{-1, 1\}^N$ whereas models in Markov random field assume configurations are $\{0, 1\}^N$. It is known in the field of statistical mechanics that the ground state of the **ferromagnetic** Ising model is submodular and can be found using polynomial-time algorithm, as we proved in Theorem 5.2.2 [51].

5.3 Examples Using Real-World Networks

The most-probable configuration changes depending on these parameters. When the healing rate dominates over the infection rates, $\mathbf{x}^* = \mathbf{x}^0$; this means that the epidemics is not severe. When the infection rates dominate over the healing rate, $\mathbf{x}^* = \mathbf{x}^N$; this means that the epidemics is severe. When \mathbf{x}^* is a non-degenerate configuration (i.e. $\mathbf{x}^* \neq \mathbf{x}^0, \mathbf{x}^N$), this indicates that sets of agents in the network are more vulnerable than others to the epidemics. We illustrate this by solving for the most-probable configuration under different $\left(\frac{\lambda}{\mu}, \gamma\right)$ parameters for 2 real-world networks (see chapter 2.2.4): a social network [60] (see Figure 2.4) and the Western United States power grid [12] (see Figure 2.5), obtained from [17].

The network shown in Fig. 5.1 is a 193 node, 273 edge social network of drug users in Hartford, CT. The network was determined through interviews. Reference [46] looked for influential agents in the network by considering it as a graph connectivity problem. However, they did not consider a dynamical model of influence. Assuming that we can model drug habits as an epidemics (i.e., there is a social contagion aspect to the behavior), we applied the scaled SIS process to this network and solved for the most-probable configuration under different parameters to find influential network structures.

We show the resultant most-probable configurations in Fig. 5.1(a), Fig. 5.1(b), Fig. 5.1(c), Fig. 5.1(d) as we change $\left(\frac{\lambda}{\mu}, \gamma\right)$. We can see from these results that there is a small community of users who are infected when others are healthy. The size of this community increases or decreases

depending on the parameters. If there is a social contagion component to drug usage, then these agents may be more vulnerable to the social contagion component of drug usage and therefore more likely to persist in their habit.

The network shown in Fig. 5.2 is the 4941 node, 6595 edge power grid network of the Western United States used by Watts and Strogatz. They showed through simulation of the SIR (susceptible-infected-removed) epidemics model on the western power grid that small-world networks like the western power grid are more conducive to spreading infection/failures than lattice networks. This is useful for explaining why failures propagate so quickly in a blackout. However, they can not identify *which* components in the power grid are more vulnerable to the epidemics with their approach. Here, we model the blackout as a SIS epidemics by assuming that failures and recoveries of grid components (e.g., power stations, substations, generators, switches, lines) are intermittent; a failed component may return to power, possibly failing again, as often happens in practice. Using the scaled SIS process, we can identify the most vulnerable substructures in the network.

Figure 5.2(a) and Fig. 5.2(b) show the most-probable configuration for the western US power grid when for the scaled SIS process parameterized $\left(\frac{\lambda}{\mu} = 0.33, \gamma = 2\right)$ and $\left(\frac{\lambda}{\mu} = 0.33, \gamma = 2.6\right)$, respectively. We can see that for the same $\frac{\lambda}{\mu}$, as γ increases, thereby increasing the infectiousness of cascading failures (i.e., epidemics), the number of vulnerable components increases. This is intuitive since, for large γ , the epidemics is severe, and the most-probable configuration is driven toward \mathbf{x}^N , the configuration where all the components are infected. Moreover, the most-probable configurations are both non-degenerate configurations. The components that are infected at equilibrium are more vulnerable to the cascading failures than components that remain healthy. By using submodular optimization, we can identify these more vulnerable components, by solving for the most-probable configuration out of 2^{4941} total possible configurations, exactly and in polynomial time.

5.3.1 Computation Time

Although the algorithms are both polynomial-time, there is a large difference in solving for the most-probable configuration using a general submodular optimizer or using the max-flow/min-cut algorithm.

Table 5.1 shows the computation time for the most-probable configuration for the correspond-

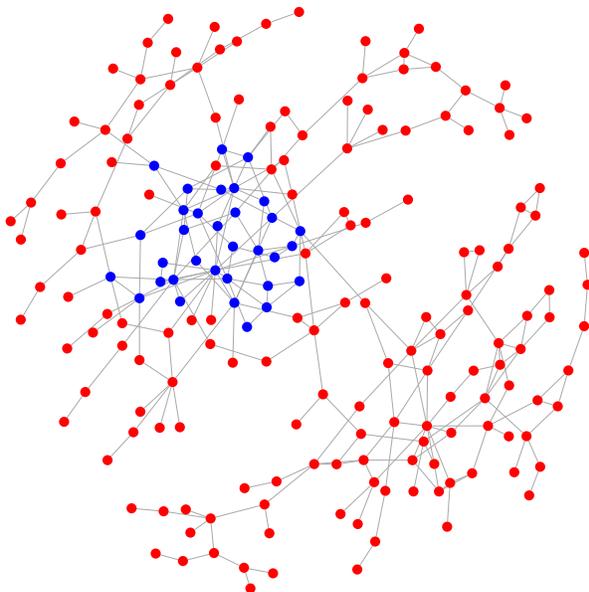
ing networks and parameter values using a desktop with 3.7 GHz Quad Core Xeon processor and 16GB of RAM. The solutions were obtained using two different methods: 1) general submodular optimizer [61] and 2) method described in [57] and the max-flow solver in [62]. The max-flow method is computationally much more efficient in comparison.

	Submodular Optimizer	Max-Flow/Min-Cut
Fig. 5.1(a)	1.54 sec	0.003 sec
Fig. 5.1(b)	0.96 sec	0.003 sec
Fig. 5.1(c)	0.14 sec	0.003 sec
Fig. 5.1(d)	1.76 sec	0.003 sec
Fig. 5.2(a)	1.29×10^4 sec	0.041 sec
Fig. 5.2(b)	3.71×10^3 sec	0.019 sec

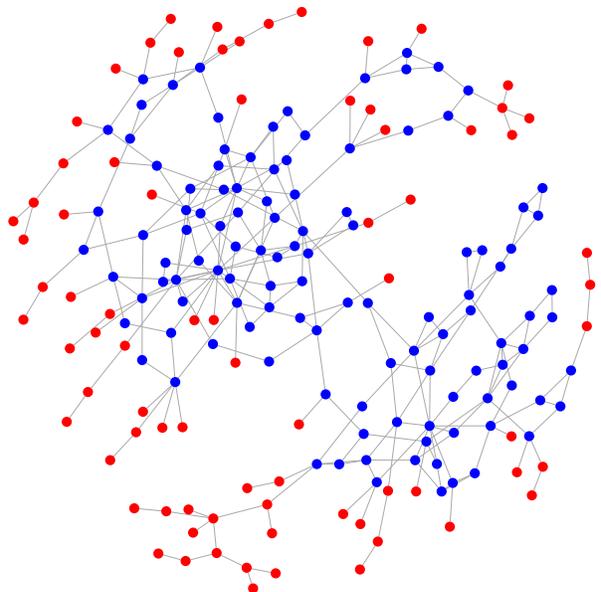
Table 5.1: Computation Time of the Most-Probable Configuration Problem

For large real-world networks, there are many non-degenerate solution to the Most-Probable Configuration Problem as seen in Figure 5.1 and Figure 5.2.

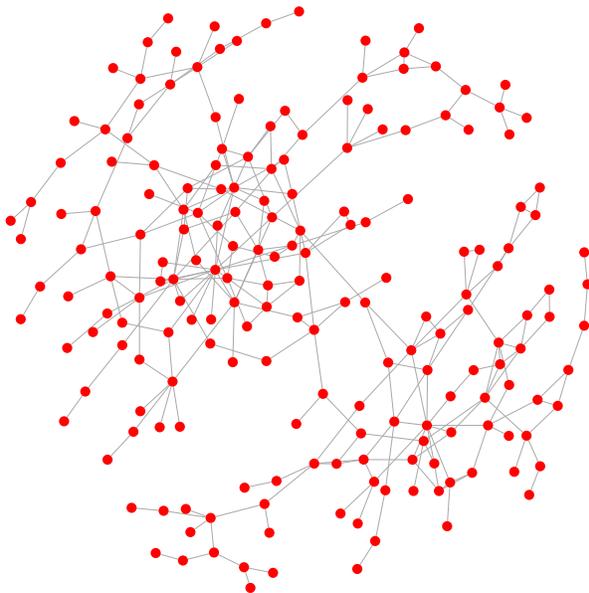
An important question is to relate the most-probable configuration to network structure. We will show in the next section that the most-probable configuration is related to subgraph density by rewriting the equilibrium distribution (3.4) in terms of induced subgraphs instead of network configurations.



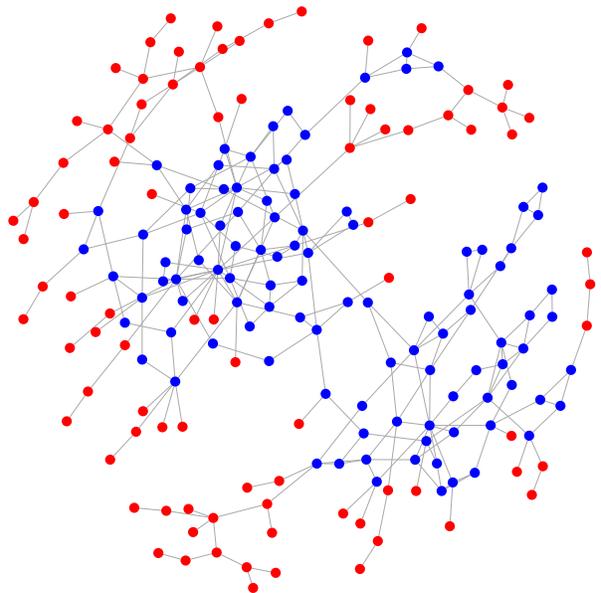
(a) $\frac{\lambda}{\mu} = 0.2, \gamma = 2.4$



(b) $\frac{\lambda}{\mu} = 0.267, \gamma = 3$

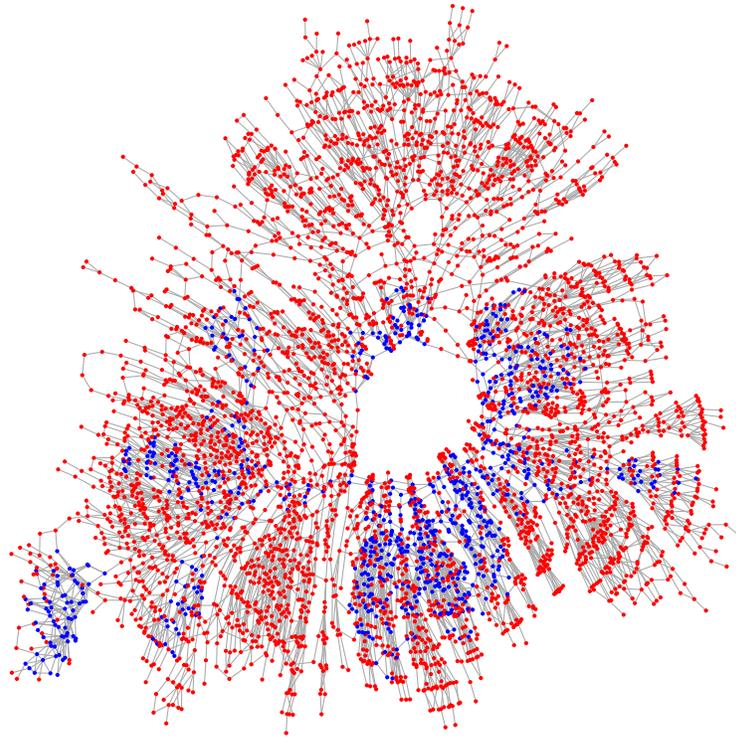


(c) $\frac{\lambda}{\mu} = 0.4, \gamma = 1.2$

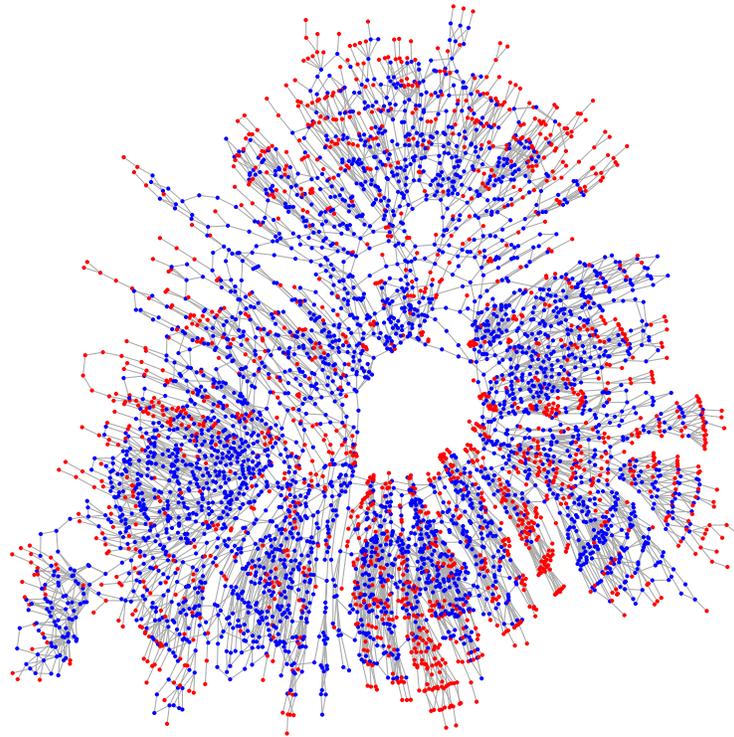


(d) $\frac{\lambda}{\mu} = 0.5, \gamma = 1.6$

Figure 5.1: Most-Probable Configuration \mathbf{x}^* under Different $\left(\frac{\lambda}{\mu}, \gamma\right)$ Parameters (Blue = Infected, Red = Healthy)



(a) $\frac{\lambda}{\mu} = 0.33, \gamma = 2$



(b) $\frac{\lambda}{\mu} = 0.33, \gamma = 2.6$

Figure 5.2: Most-Probable Configuration \mathbf{x}^* under Different $\left(\frac{\lambda}{\mu}, \gamma\right)$ Parameters (Blue = Infected, Red = Healthy)

5.4 Induced Subgraphs and Graph Density

In the previous section, we showed that we can exactly solve for the most-probable configuration with a polynomial time algorithm. The exact solution, however, does not give insight on how the most-probable configuration changes depending on the parameters $\left(\frac{\lambda}{\mu}, \gamma\right)$ and on the network topology. In this section, we draw the connection between the most-probable configuration and subgraphs in the network. As per our intuition for epidemics, densely connected network structures are more vulnerable to network epidemics; the scaled SIS process quantifies this intuition. First, we will define the graph theoretic terms used in this section.

Definition 5.4.1 (From [43]). *The graph F is an induced subgraph of G if two vertices in F are connected if and only if they are connected in G and the vertex set and edge set of F are subsets of the vertex set and edge set of G .*

$$V(F) \subseteq V(G), E(F) \subseteq E(G)$$

Definition 5.4.2. *The graph $F(\mathbf{x})$ is an induced subgraph of configuration $\mathbf{x} = [x_1, x_2, \dots, x_N]^T$ if the nodes/edges in the subgraph are the infected agents/edges in \mathbf{x} .*

$$V(F(\mathbf{x})) = \{v_i \in V(G) \mid x_i = 1\} \quad (5.7)$$

$$E(F(\mathbf{x})) = \{(i, j) \in E(G) \mid x_i = 1, x_j = 1\} \quad (5.8)$$

By definition, $|V(F(\mathbf{x}))| = \mathbf{1}^T \mathbf{x}$ and $|E(F(\mathbf{x}))| = \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2}$. Figure 5.3 and Figure 5.4 show two network configurations and their corresponding induced subgraphs. We proved in [35] that configurations whose induced subgraphs are isomorphic are equally probable. Unless we need to refer explicitly to the underlying network configuration \mathbf{x} , for notational simplicity, we will write F to denote an induced subgraph instead of writing $F(\mathbf{x})$.

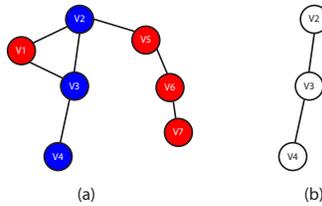


Figure 5.3: (a) Configuration $\mathbf{x}_1 = [0, 1, 1, 1, 0, 0, 0]^T$, (b) Induced Subgraph $F(\mathbf{x}_1) = F_1$

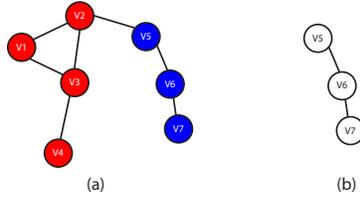


Figure 5.4: (a) Configuration $\mathbf{x}_2 = [0, 0, 0, 0, 1, 1, 1]^T$, (b) Induced Subgraph $F(\mathbf{x}_2) = F_2$

Definition 5.4.3. *The set of all possible induced subgraphs of G is $\mathcal{F} = \{F(\mathbf{x})\}$, $\forall \mathbf{x} \in \mathcal{X}$.*

The set \mathcal{H} includes the empty graph, which is induced by the configuration $\mathbf{x}^0 = [0, 0, \dots, 0]^T$, and G , which is the subgraph induced by the configuration $\mathbf{x}^N = [1, 1, \dots, 1]^T$.

Definition 5.4.4 (From [44]). *The density of the graph G is*

$$d(G) = \frac{|E(G)|}{|V(G)|}.$$

There is an alternative definition for graph density that is the number of edges divided by the total number of possible edges [45]. Unfortunately, these two definitions of density are not equivalent.

We will refer to the density of the entire network, $d(G) = d(F(\mathbf{x}^N))$, as the *network density*, and the density of an induced subgraph of G as the *subgraph density*. The density of the empty graph, $d(F(\mathbf{x}^0))$, is 0 by definition. The subgraphs in \mathcal{F} can be partially ordered by their density. There may be many subgraphs with the same density. A special induced subgraph in \mathcal{F} is the densest subgraph.

Definition 5.4.5. *Let \overline{F} be the densest subgraph in G . Then*

$$d(\overline{F}) \geq d(F), \quad \forall F \in \mathcal{F}.$$

Finding \overline{F} is known as the *Densest Subgraph Problem*. It is known that this problem can be solved in polynomial time exactly and in linear time in approximation for undirected graphs [44].

5.4.1 Equilibrium Distribution and Induced Subgraphs

Since there is a one-to-one relationship between the network configuration \mathbf{x} and its induced subgraph $F(\mathbf{x})$, we can rewrite the equilibrium distribution (3.4) of the scaled SIS process in terms

of the induced subgraph density and the size of the induced subgraph:

$$\pi(F) = \frac{1}{Z} \left(\left(\frac{\lambda}{\mu} \right) \gamma^{d(F)} \right)^{|V(F)|}, \quad F \in \mathcal{F}, \quad (5.9)$$

where $d(F)$ is the density of the induced subgraph F and Z is the partition function.

Stating the equilibrium distribution (3.4) as a function of induced subgraphs (5.9) allows us to see that when the induced subgraphs of two configurations are isomorphic, as configurations \mathbf{x}_1 in Figure 5.3 and \mathbf{x}_2 in Figure 5.4, then the configurations are equally probable at equilibrium for the scaled SIS process.

Theorem 5.4.6. *If the induced subgraphs of two network configurations, $\mathbf{x}_1 \in \mathcal{X}$ and $\mathbf{x}_2 \in \mathcal{X}$, are isomorphic, then $\pi(\mathbf{x}_1) = \pi(\mathbf{x}_2)$.*

Proof. By definition, $|V(F(\mathbf{x}))| = 1^T \mathbf{x}$ and $|E(F(\mathbf{x}))| = \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2}$. Consider that \mathbf{x}_1 and \mathbf{x}_2 induce two graphs $F(\mathbf{x}_1)$ and $F(\mathbf{x}_2)$. Since $F(\mathbf{x}_1)$ is isomorphic to $F(\mathbf{x}_2)$, the number of nodes and the number of edges are the same for $F(\mathbf{x}_1)$ and $F(\mathbf{x}_2)$ [43]. By construction, we know that the number of nodes in $F(\mathbf{x}_1) = 1^T \mathbf{x}_1 = 1^T \mathbf{x}_2$, and the number of edges in $F(\mathbf{x}_1) = \frac{\mathbf{x}_1^T \mathbf{A} \mathbf{x}_1}{2} = \frac{\mathbf{x}_2^T \mathbf{A} \mathbf{x}_2}{2}$. From (3.4), we can conclude that $\pi(\mathbf{x}_1) = \pi(\mathbf{x}_2)$. \square

5.4.2 Most-Probable Configuration and Induced Subgraphs

Using (5.9), the Most-Probable Configuration Problem (4.1) is then also an optimization problem over all the possible induced subgraphs in G :

$$F(\mathbf{x}^*) = \arg \max_{F \in \mathcal{F}} \left(\left(\frac{\lambda}{\mu} \right) \gamma^{d(F)} \right)^{|V(F)|}. \quad (5.10)$$

The subgraph induced by the most-probable configuration, $F(\mathbf{x}^*)$, is the *most-probable subgraph*, but this is *not* necessarily the same subgraph as the densest subgraph, \overline{F} .

Stating the equilibrium distribution in terms of the induced subgraph will allow us to derive several theorems regarding the most-probable configuration. For the theorems that follow, we make the following assumptions:

Assumption 1. The scaled SIS process operates in regime II) **Endogenous Infection Dominant**.

This limits the dynamics parameters to the range, $0 < \frac{\lambda}{\mu} \leq 1$ and $\gamma > 1$.

Assumption 2. The underlying network G is a simple, undirected, unweighted, and connected graph.

Theorem 5.4.7. [Proof in Appendix 5.8] *The most-probable configuration $\mathbf{x}^* \neq \mathbf{x}^0$ if and only if there exists at least one induced subgraph $F \in \mathcal{F}$ with density $d(F)$ for which $\lambda\gamma^{d(F)} > \mu$.*

Theorem 5.4.8. [Proof in Appendix 5.8] *The most-probable configuration $\mathbf{x}^* \neq \mathbf{x}^N$ if and only if there exists at least one induced subgraph $F \in \mathcal{F} \setminus G$ with density $d(F) = \frac{E'}{N'}$ for which*

$$\frac{\log(\frac{\lambda}{\mu}\gamma^{d(G)})}{\log(\frac{\lambda}{\mu}\gamma^{d(F)})} < \frac{N'}{N}. \quad (5.11)$$

Combining Theorem 5.4.7 and Theorem 5.4.8, we can obtain the following corollary regarding the non-degenerate most-probable configurations.

Corollary 5.4.9. [Proof in Appendix 5.8] *Let the density of the network be $d(G) = \frac{E}{N}$. Then, the most-probable configuration is a non-degenerate configuration, $\mathbf{x}^* \in \mathcal{X} \setminus \{\mathbf{x}^0, \mathbf{x}^N\}$, if and only if there exists at least one induced subgraph $F \in \mathcal{F}$ with density $d(F) = \frac{E'}{N'}$ for which $\lambda\gamma^{d(F)} > \mu$, and*

$$\frac{\log(\frac{\lambda}{\mu}\gamma^{d(G)})}{\log(\frac{\lambda}{\mu}\gamma^{d(F)})} < \frac{N'}{N}.$$

In regime II) individual agents have a preference for being healthy, but the epidemics might spread to other agents through neighbor-to-neighbor contagion. Under the scaled SIS process, the subgraph density $d(F)$ scales the exogenous infection rate γ , thereby affecting the overall infection rate. Theorem 5.4.7 states that, if the network contains *dense-enough* subgraphs, then even when the effective exogenous infection rate, $\frac{\lambda}{\mu}$, is small (i.e., $0 < \frac{\lambda}{\mu} \ll 1$), the exogenous infection rate, γ , can leverage dense subgraphs to spread the infection throughout the network.

On the other hand, if the endogenous infection rate, γ , is large (i.e., $\gamma \gg 1$), then most certainly the epidemics will spread throughout the entire network. Theorem 5.4.8 states when this does not happen. Furthermore, Theorem 5.4.8 shows that it is important to consider if the densest subgraph in the network is the entire network or a smaller subgraph. Corollary 5.4.9 proves that the existence of the non-degenerate configurations is related to the existence of subgraphs with density larger than the network density. The existence of these *denser-than G* subgraphs is crucial

to the existence of non-degenerate configurations (i.e., different from \mathbf{x}^0 and \mathbf{x}^N) as solutions to the Most-Probable Configuration Problem; when the most-probable configuration is a non-degenerate configuration, agents belonging to denser subgraphs are more vulnerable to the epidemics.

In network science, dense clusters of agents have often been identified as either the network *core* or *community* [9,21,63]. Solving for the non-degenerate configuration is an alternative method for determining these network structures. Previous works in core/community detection are algorithmic and do not consider the dynamical process on the network. The scaled SIS process, however, is a model for dynamical processes on networks and, therefore, what is considered a *community* changes depending on the parameters of the dynamical process: the most-probable configuration changes depending on the exogenous rates $\frac{\lambda}{\mu}$ and on the endogenous rates γ .

We illustrate Theorem 5.4.7 and 5.4.8 with two small 16 node examples, which we showed has non-degenerate solutions in regime II) in chapter 4.6.5; Network A shown in Fig. 5.5 and Network B in Fig. 5.6. For each network, we fix the effective exogenous infection rate, $\frac{\lambda}{\mu} = 0.5$. We then solve for the most-probable configuration for different γ , ranging from 1.2 to 3. As the endogenous infection rate, γ , changes, the most-probable configuration also changes. In Fig. 5.5(a) and Fig. 5.6(a), neither network supports dense enough subgraphs for the epidemics to be severe. But as γ increases, the infection starts to spread. In Network A, there is at least one subgraph denser than the network. The subgraph induced by $V1, V2, V3, V4, V5, V7, V8, V9, V10$ has a density of 1.33 whereas the density of the entire network is 1.19. In Fig. 5.5(b), the most-probable configuration has these 9 agents infected while the other 7 agents remain healthy. The 9 agents in the dense subgraph are more vulnerable to the epidemics when $\frac{\lambda}{\mu} = 0.5$ and $\gamma = 1.7$.

In Network B, there are at least two subgraphs denser than the network and they are induced by the set of infected agents of the most-probable configuration as shown in Fig. 5.6(b) and Fig. 5.6(c). We can see by solving for the most-probable configuration for different parameter values that, as the endogenous infection increases, the most-probable configuration goes toward \mathbf{x}^N as all agents become vulnerable to the epidemics.

It is easier for the infection to spread in Network B than in Network A, since, for the same rate parameters, $\mathbf{x}^* = \mathbf{x}^N$ for Network B while $\mathbf{x}^* \neq \mathbf{x}^N$ for Network A. This is because Network B is a denser graph ($d(G) = 2.4375$) than Graph A ($d(G) = 1.19$).

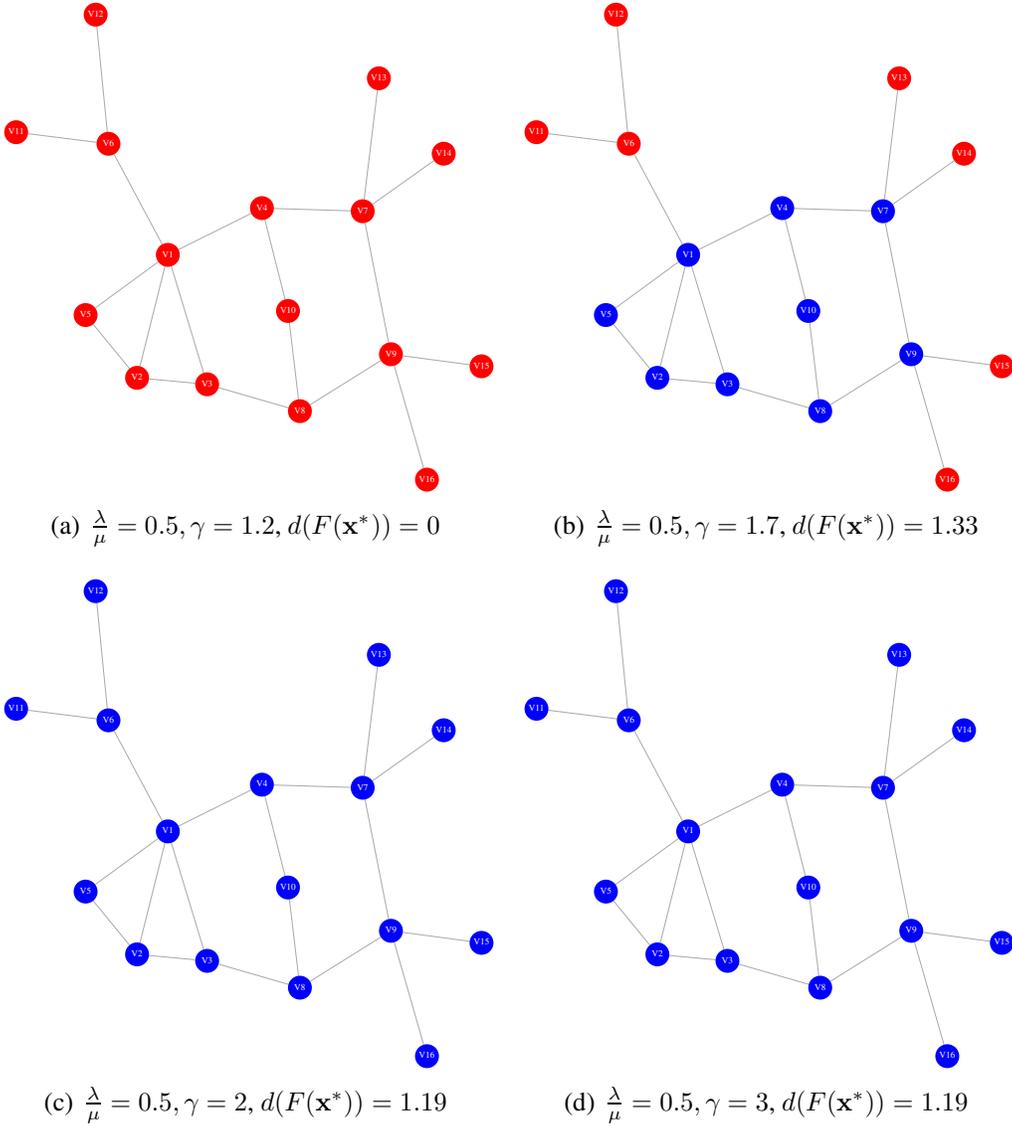


Figure 5.5: Most-Probable Configuration \mathbf{x}^* under Different $\left(\frac{\lambda}{\mu}, \gamma\right)$ Parameters (Blue = Infected, Red = Healthy)

5.4.3 Most-Probable Configuration and the Densest Subgraph

We showed that the most-probable configuration is related to the density of induced subgraphs in the network. The densest subgraph, \overline{H} , is a special induced subgraph. In this section, we focus specifically on the relationship between the most-probable configuration and the densest subgraph.

Corollary 5.4.10. *[Proof in Appendix 5.8] The most-probable configuration $\mathbf{x}^* = \mathbf{x}^0$ if and only if $\lambda\gamma^{d(\overline{F})} \leq \mu$.*

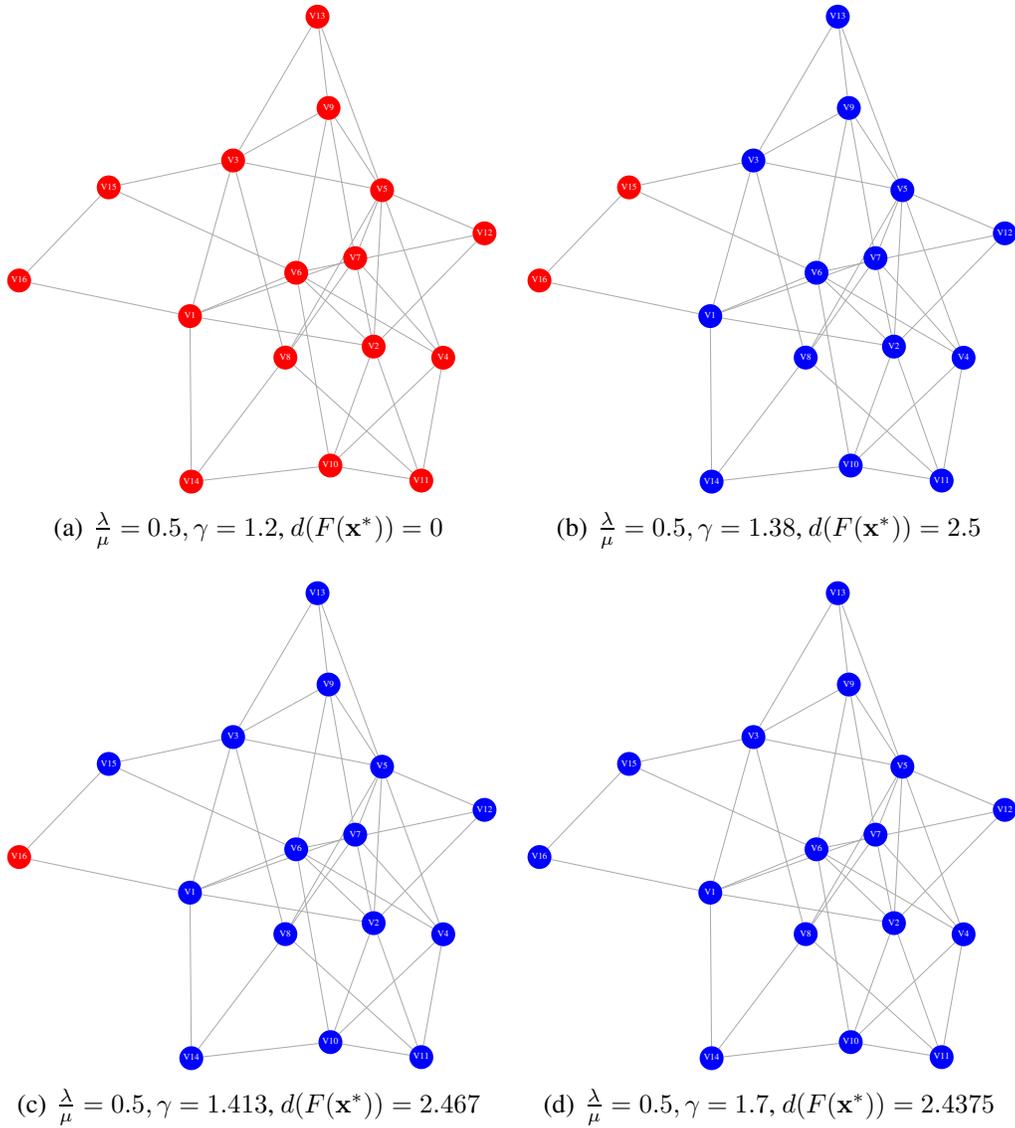


Figure 5.6: Most-Probable Configuration \mathbf{x}^* under Different $\left(\frac{\lambda}{\mu}, \gamma\right)$ Parameters (Blue = Infected, Red = Healthy)

Corollary 5.4.10 follows the result of Theorem 5.4.7. If the densest subgraph in the network is not *dense enough* to overcome individual preferences for being healthy, then the endogenous infection rate γ will not be able to drive the most-probable configuration away from \mathbf{x}^0 .

Lastly, because of the connection between the most-probable configuration of the scaled SIS process and the densest subgraph, we can prove a *general* statement regarding network structure using results from dynamical processes on networks.

Lemma 5.4.1. [Proof in Appendix 5.8] *If G is a k -regular, complete multipartite, or complete*

multipartite with k -regular islands network, then $\overline{F} = G$. That is, for these structured networks, the densest subgraph is the overall graph.

5.5 Uniqueness of the Most-Probable Configuration

The Most-Probable Configuration Problem (4.1) is a combinatorial optimization problem. Unlike convex optimization, for this class of optimization problem, solutions are not guaranteed to be unique.

We proved in section 4.4 that the most-probable configuration in regime III) is related to the maximum independent set problem. It is easy to see that for many network configurations, the maximum independent set is not unique. We also proved that in regime II), for k -regular, complete multipartite, and complete multipartite network with k -regular islands, the most-probable configuration is both \mathbf{x}^0 and \mathbf{x}^N at a particular configuration of the parameters $\frac{\lambda}{\mu}$ and γ .

Since the scaled SIS process is a dynamical process over a graph structure. The uniqueness of the ground state is further complicated by having to consider equivalence relationships of subgraph structures in the network.

Recall that the equilibrium distribution of the scaled SIS process can be rewritten as a function of induced subgraphs

$$\pi(F) = \frac{1}{Z} \left(\left(\frac{\lambda}{\mu} \right) \gamma^{d(F)} \right)^{|V(F)|}, \quad F \in \mathcal{F},$$

where $d(F)$ is the density of the induced subgraph F and Z is the partition function.

Definition 5.5.1. [42] *Two graphs, G and F , are equivalent if they are isomorphic: there is a bijection between the vertex sets of G and F , $f : V(G) \rightarrow V(F)$, such that any two vertices u and v of G are adjacent in G if and only if $f(u)$ and $f(v)$ are adjacent in F . This means that edge (u, v) is in G if and only if $(f(u), f(v))$ is in F .*

Theorem 5.4.6 states that two configurations have the same equilibrium probability if their induced subgraphs are isomorphic to each other. As a result, there are two types of uniqueness to be considered.

Definition 5.5.2. *A solution to the Most-Probable Configuration Problem (4.1), \mathbf{x}^* , is **subgraph unique** if there are no other subgraphs in the network isomorphic to the induced subgraph $F(\mathbf{x}^*)$.*

Definition 5.5.3. A solution to the Most-Probable Configuration Problem (4.1), \mathbf{x}^* , is **unique** if

1. it is subgraph unique
2. there is no other configurations \mathbf{x}' such that $\pi(\mathbf{x}^*) = \pi(\mathbf{x}')$ but $F(\mathbf{x}^*)$ and $F(\mathbf{x}')$ are not isomorphic.

When the most-probable configuration only satisfies the first condition, it is **subgraph unique**. When \mathbf{x}^* is $\mathbf{x}^0 = [0, 0, \dots, 0]^T$ or $\mathbf{x}^N = [1, 1, \dots, 1]^T$, the first condition is trivially satisfied since for any network, there are no other subgraphs isomorphic to $H(\mathbf{x}^0) = \emptyset$ or $H(\mathbf{x}^N) = G$.

For the most-probable configuration to be considered **unique** however, it must satisfy both conditions. Section 5.6.2 provides the conditions for \mathbf{x}^* to be unique.

5.6 Regime II) Endogenous Infection Dominant

We showed that in this regime, the Most-Probable Configuration Problem can be solved in polynomial-time and that it is related to dense subgraphs in the underlying network. This section discusses the uniqueness of solution with the assumptions that

Assumption 1. The scaled SIS process operates in regime II) **Endogenous Infection Dominant**.

This limits the dynamics parameters to the range, $0 < \frac{\lambda}{\mu} \leq 1$ and $\gamma > 1$.

Assumption 2. The underlying network G is a simple, undirected, unweighted, and connected graph.

5.6.1 Subgraph Uniqueness

This section proves that *all* the solutions of the Most-Probable Configuration Problem (4.1) are subgraph unique. This means that the induced subgraphs of all these configurations are unique subgraphs of the underlying network.

Lemma 5.6.1. [Proof in Appendix 5.8] Given an undirected graph G described by adjacency matrix \mathbf{A} and two different configurations \mathbf{x}_1 and \mathbf{x}_2 , which induce subgraphs $F(\mathbf{x}_1)$ and $F(\mathbf{x}_2)$, respectively. $F(\mathbf{x}_1)$ is isomorphic to $F(\mathbf{x}_2)$ if and only if $\mathbf{1}^T \mathbf{x}_1 = \mathbf{1}^T \mathbf{x}_2$ and $\mathbf{x}_1^T \mathbf{A} \mathbf{x}_1 = \mathbf{x}_2^T \mathbf{A} \mathbf{x}_2$.

Theorem 5.6.1. [Proof in Appendix 5.8] Consider Regime II) **Endogenous Infection Dominant:** $0 < \frac{\lambda}{\mu} \leq 1, \gamma > 1$ and network G . The most-probable configuration, \mathbf{x}^* , induces the subgraph $F(\mathbf{x}^*)$ with density $d(F(\mathbf{x}^*))$. If

$$\frac{\lambda}{\mu} \gamma^{d(F(\mathbf{x}^*))} > 1,$$

then \mathbf{x}^* is subgraph unique.

We showed that there can be no other configuration \mathbf{x}' that is equally probable to \mathbf{x}^* and its induced subgraph $F(\mathbf{x}')$ is isomorphic to the induced subgraph $F(\mathbf{x}^*)$. Next, we will consider the case where the configuration \mathbf{x}' is equally probable as \mathbf{x}^* but its induced subgraph $F(\mathbf{x}')$ is *not* isomorphic to the induced subgraph $F(\mathbf{x}^*)$

Corollary 5.6.2. If \mathbf{x}^* is the most-probable configuration in regime II) **Endogenous Infection Dominant:** $0 < \frac{\lambda}{\mu} \leq 1, \gamma > 1$, then it is subgraph unique.

Proof. Theorem 5.6.1 states the condition for which a configuration, \mathbf{x}^* , is subgraph unique. Theorem 5.4.7 states following necessary and sufficient conditions: The most-probable configuration $\mathbf{x}^* \neq \mathbf{x}^0$ if and only if there exists at least one induced subgraph $F \in \mathcal{F}$ with density $d(F)$ for which $\frac{\lambda}{\mu} \gamma^{d(F)} > 1$.

Since the induced subgraph of \mathbf{x}^* satisfies the condition that $\frac{\lambda}{\mu} \gamma^{d(F(\mathbf{x}^*))} > 1$, then $\mathbf{x}^* \neq \mathbf{x}^0$. However, since \mathbf{x}^0 induces the empty graph, a solution where $\mathbf{x}^* = \mathbf{x}^0$ is also subgraph unique. As a result, any solution to the Most-Probable Configuration in regime II) is subgraph unique.

□

Corollary 5.6.2 states that all the possible solutions of the Most Probable Configuration Problem in regime II) are subgraph unique. This means that the corresponding induced subgraphs are unique. Since the optimization problem can be solved efficiently for large networks, this is a method to identify unique subgraphs in large graphs.

5.6.2 Uniqueness

The condition for the uniqueness the most-probable configuration is *stricter* than subgraph uniqueness. First, we state the following result:

Theorem 5.6.3. [Proof in Appendix 5.8] Consider two scaled SIS processes, $\{X_1(t), t \geq 0\}$ and $\{X_2(t), t \geq 0\}$, with parameters $\left(\frac{\lambda_1}{\mu_1}, \gamma_1\right)$ and $\left(\frac{\lambda_2}{\mu_2}, \gamma_2\right)$ satisfying the relationship

$$\begin{bmatrix} \log\left(\frac{\lambda_1}{\mu_1}\right) \\ \log(\gamma_1) \end{bmatrix} \neq \alpha \begin{bmatrix} \log\left(\frac{\lambda_2}{\mu_2}\right) \\ \log(\gamma_2) \end{bmatrix} \text{ for any real-value } \alpha.$$

If two different configurations \mathbf{x}_1 and \mathbf{x}_2 are equally probable under both processes, then their induced subgraphs are isomorphic, $F(\mathbf{x}_1) \simeq F(\mathbf{x}_2)$.

Putting together Theorem 5.6.1 and Theorem 5.6.3, we can derive the sufficient conditions for which the solution to the most-probable configuration is unique.

Corollary 5.6.4. [Proof in Appendix 5.8] Consider two scaled SIS processes, $\{X_1(t), t \geq 0\}$ and $\{X_2(t), t \geq 0\}$, with parameters in regime II) **Endogenous Infection Dominant** that are not exponentially related. Therefore, $0 < \frac{\lambda_1}{\mu_1}, \frac{\lambda_2}{\mu_2} \leq 1, \gamma_1, \gamma_2 > 1$ and

$$\begin{bmatrix} \log\left(\frac{\lambda_1}{\mu_1}\right) \\ \log(\gamma_1) \end{bmatrix} \neq \alpha \begin{bmatrix} \log\left(\frac{\lambda_2}{\mu_2}\right) \\ \log(\gamma_2) \end{bmatrix} \text{ for any real-value } \alpha.$$

If \mathbf{x}^* is the most-probable configuration for both scaled SIS processes, $\{X_1(t), t \geq 0\}$ and $\{X_2(t), t \geq 0\}$, then \mathbf{x}^* is the unique maximizer of the Most-Probable Configuration Problem for both processes.

5.7 Conclusion

This chapter proved that the Most-Probable Configuration Problem can be solved exactly in regime II) **Endogenous Infection Dominant**: $0 < \frac{\lambda}{\mu} \leq 1, \gamma > 1$ in polynomial-time. For a network with 4941 nodes, the configuration with the highest equilibrium probability out of 2^{4941} , can be found on a standard desktop computer in less than 1 sec using Max-Flow/Min-Cut algorithm [57].

We then showed that the most-problem configuration for this regime, where individual agents prefer the healthy state but network effects spread the infection, is related to subgraphs in the network. In particular, non-degenerate solutions (i.e., most-probable configurations other than \mathbf{x}^0 and \mathbf{x}^N) are related to the existence of subgraphs that are denser than the overall network. These

non-degenerate solutions inform us, depending on the infection and healing rates, which subsets of agents in the network are vulnerable to infection. This chapter proves that these more vulnerable agents belong to subgraphs that are more densely connected than the overall network. We proved that the subgraphs corresponding to non-degenerate solutions of the Most-Probable Configuration Problem in regime II) are unique in the sense that there are no other isomorphic subgraphs in the network. The uniqueness of a solution to the Most-Probable Configuration Problem requires the additional condition that this solution is also the most-probable configuration of another scaled SIS process for which the dynamics parameters are not exponentially related.

5.8 Appendix

Proof for Theorem 5.4.7

Theorem. *The most-probable configuration $\mathbf{x}^* \neq \mathbf{x}^0$ if and only if there exists at least one induced subgraph $F \in \mathcal{F}$ with density $d(F)$ for which $\lambda\gamma^{d(F)} > \mu$.*

Proof. Sufficiency: If there exists at least one subgraph $F \in \mathcal{F}$ with density $d(F)$ for which $\lambda\gamma^{d(F)} > \mu$, then $\mathbf{x}^* \neq \mathbf{x}^0$.

Using the equilibrium distribution (3.4), $\pi(\mathbf{x}^0) = \frac{1}{Z}$. Let the subgraph $F \in \mathcal{F}$ be the subgraph induced by configuration $\mathbf{x}' \in \mathcal{X} \setminus \mathbf{x}^0$. The number of infected agents in configuration \mathbf{x}' is $1^T \mathbf{x}' = |V(F)| > 0$. Using (5.9), its equilibrium probability is

$$\pi(\mathbf{x}') = \pi(F) = \frac{1}{Z} \left(\left(\frac{\lambda}{\mu} \right) \gamma^{d(F)} \right)^{|V(F)|}$$

If $\left(\frac{\lambda}{\mu} \right) \gamma^{d(F)} > 1$, we know that $\pi(\mathbf{x}') > \pi(\mathbf{x}^0)$. Therefore, \mathbf{x}^0 can not be the most-probable configuration.

Necessity: If $\mathbf{x}^* \neq \mathbf{x}^0$, then there exist at least one subgraph $F \in \mathcal{F}$ with density $d(F)$ for which $\lambda\gamma^{d(F)} > \mu$.

If $\mathbf{x}^* \neq \mathbf{x}^0$, this means that there is some configuration \mathbf{x}' for which $\pi(\mathbf{x}') > \pi(\mathbf{x}^0)$. We know that $\pi(\mathbf{x}^0) = \frac{1}{Z}$. Using the equilibrium distribution in (5.9) and the fact that $1^T \mathbf{x} = |V(F)| > 0, \forall \mathbf{x} \in \mathcal{X} \setminus \mathbf{x}^0$, we can conclude that there must exist some induced subgraph whose density satisfies this condition $\left(\frac{\lambda}{\mu} \right) \gamma^{d(H(\mathbf{x}'))} > 1$. □

Proof for Theorem 5.4.8

Theorem. *The most-probable configuration $\mathbf{x}^* \neq \mathbf{x}^N$ if and only if there exists at least one induced subgraph $F \in \mathcal{F} \setminus G$ with density $d(F) = \frac{E'}{N'}$ for which*

$$\frac{\log\left(\frac{\lambda}{\mu}\gamma^{d(G)}\right)}{\log\left(\frac{\lambda}{\mu}\gamma^{d(F)}\right)} < \frac{N'}{N}. \quad (5.12)$$

Proof. Sufficiency: If there exists at least one induced subgraph $F \in \mathcal{F} \setminus G$ with density $d(F) = \frac{E'}{N'}$ such that $\frac{\log\left(\frac{\lambda}{\mu}\gamma^{d(G)}\right)}{\log\left(\frac{\lambda}{\mu}\gamma^{d(F)}\right)} < \frac{N'}{N}$, then $\mathbf{x}^* \neq \mathbf{x}^N$.

The subgraph H is induced by the configuration $\mathbf{x}' \in \mathcal{X}$. The log equilibrium probability according to (5.9) for \mathbf{x}' and \mathbf{x}^N , respectively, are:

$$\log(\pi(\mathbf{x}')) = \log\left(\frac{1}{Z}\right) + N' \log\left(\frac{\lambda}{\mu}\gamma^{d(F)}\right)$$

and

$$\log(\pi(\mathbf{x}^N)) = \log\left(\frac{1}{Z}\right) + N \log\left(\frac{\lambda}{\mu}\gamma^{d(G)}\right).$$

Condition $\frac{\log\left(\frac{\lambda}{\mu}\gamma^{d(G)}\right)}{\log\left(\frac{\lambda}{\mu}\gamma^{d(F)}\right)} < \frac{N'}{N}$ implies that $N \log\left(\frac{\lambda}{\mu}\gamma^{d(G)}\right) < N' \log\left(\frac{\lambda}{\mu}\gamma^{d(F)}\right)$. Therefore, $\log(\pi(\mathbf{x}')) > \log(\pi(\mathbf{x}^N))$. Since the logarithm is a monotonic function, we can conclude that $\mathbf{x}^* \neq \mathbf{x}^N$.

Necessity: If $\mathbf{x}^* \neq \mathbf{x}^N$, then there exists at least one induced subgraph $F \in \mathcal{F}$ such that $\frac{\log\left(\frac{\lambda}{\mu}\gamma^{d(G)}\right)}{\log\left(\frac{\lambda}{\mu}\gamma^{d(F)}\right)} < \frac{N'}{N}$.

Let $\mathbf{x}^* = \mathbf{x}'$, which induces a subgraph $F \in \mathcal{F}$ with density $d(F)$. Using (5.9),

$$\pi(\mathbf{x}') = \log\left(\frac{1}{Z}\right) + N' \log\left(\frac{\lambda}{\mu}\gamma^{d(F)}\right)$$

$$\pi(\mathbf{x}^N) = \log\left(\frac{1}{Z}\right) + N \log\left(\frac{\lambda}{\mu}\gamma^{d(G)}\right).$$

This means $\pi(\mathbf{x}') - \pi(\mathbf{x}^N) > 0$, which implies

$$N' \log\left(\frac{\lambda}{\mu}\gamma^{d(F)}\right) - N \log\left(\frac{\lambda}{\mu}\gamma^{d(G)}\right) > 0$$

This reduces to the condition that

$$\frac{\log\left(\frac{\lambda}{\mu}\gamma^{d(G)}\right)}{\log\left(\frac{\lambda}{\mu}\gamma^{d(F)}\right)} < \frac{N'}{N}.$$

□

Proof for Corollary 5.4.9

Corollary. Let the density of the network be $d(G) = \frac{E}{N}$. Then, the most-probable configuration is a non-degenerate configuration, $\mathbf{x}^* \in \mathcal{X} \setminus \{\mathbf{x}^0, \mathbf{x}^N\}$, if and only if there exists at least one induced subgraph $F \in \mathcal{F}$ with density $d(F) = \frac{E'}{N'}$ for which $\lambda\gamma^{d(F)} > \mu$, and

$$\frac{\log\left(\frac{\lambda}{\mu}\gamma^{d(G)}\right)}{\log\left(\frac{\lambda}{\mu}\gamma^{d(F)}\right)} < \frac{N'}{N}.$$

Proof. We want to determine the necessary and sufficient conditions such that $\mathbf{x}^* = \mathbf{x}'$, which induces subgraph H , such that we have both $\mathbf{x}' \neq \mathbf{x}^0$ and $\mathbf{x}' \neq \mathbf{x}^N$. This is equivalent to showing

$$\pi(\mathbf{x}') > \pi(\mathbf{x}^0) \quad (5.13)$$

and

$$\pi(\mathbf{x}') > \pi(\mathbf{x}^N). \quad (5.14)$$

Condition (5.13) holds if and only if $\lambda\gamma^{d(F)} > \mu$ by Theorem 5.4.7. Condition (5.14) holds if and only if

$$\frac{\log\left(\frac{\lambda}{\mu}\gamma^{d(G)}\right)}{\log\left(\frac{\lambda}{\mu}\gamma^{d(F)}\right)} < \frac{N'}{N}.$$

by Theorem 5.4.8. This proves the corollary. \square

Proof for Corollary 5.4.10

Corollary. The most-probable configuration $\mathbf{x}^* = \mathbf{x}^0$ if and only if $\lambda\gamma^{d(\overline{H})} \leq \mu$.

Proof. Sufficiency: If $\lambda\gamma^{d(\overline{F})} \leq \mu$, then $\mathbf{x}^* = \mathbf{x}^0$.

Recall the definition of the densest subgraph 5.4.5. With $\gamma > 1$, $\lambda\gamma^{d(F(\mathbf{x}))} \leq \lambda\gamma^{d(\overline{F}(\mathbf{x}))} \leq \mu$ for all possible induced subgraphs in G . This means that there is no subgraph, $F \in \mathcal{F}$, for which $\lambda\gamma^{d(F)} > \mu$. We can conclude that $\mathbf{x}^* = \mathbf{x}^0$ using the contrapositive of Theorem 5.4.7: If there is no subgraph $F \in \mathcal{F}$ with density $d(F)$ for which $\lambda\gamma^{d(F)} > \mu$, then $\mathbf{x}^* = \mathbf{x}^0$.

Necessity: If $\mathbf{x}^* = \mathbf{x}^0$, then $\lambda\gamma^{d(\overline{F})} \leq \mu$.

The result follows from the contrapositive of Theorem 5.4.7: If $\mathbf{x}^* = \mathbf{x}^0$, then there is no subgraph $F \in \mathcal{F}$ with density $d(F)$ for which $\lambda\gamma^{d(F)} > \mu$. Therefore, all induced subgraphs, including the densest subgraph have density for which $\lambda\gamma^{d(F)} \leq \mu$. \square

Proof for Lemma 5.4.1

Lemma. *If G is a k -regular, complete multipartite, or complete multipartite with k -regular islands network, then $\overline{F} = G$. That is, for these structured networks, the densest subgraph is the overall graph.*

Proof. Theorems 4.6.1, 4.6.3, 4.6.4 state that the solution of the Most-Probable Configuration Problem for any parameters $\left(\frac{\lambda}{\mu}, \gamma\right)$ in regime II) **Endogenous Infection Dominant:** $0 < \frac{\lambda}{\mu} \leq 1, \gamma > 1$, over k -regular, complete multipartite, complete multipartite with k -regular islands networks is either \mathbf{x}^0 and/or \mathbf{x}^N ; the solution to the Most-Probable Configuration Problem for these networks is not a non-degenerate configuration in regime II). We will use this and Corollary 5.4.9 to prove this corollary.

Consider the contrapositive of Corollary 5.4.9: Let the density of the network be $d(G) = \frac{E}{N}$. Then, the most-probable configuration is not a non-degenerate configuration, $\mathbf{x}^* \in \{\mathbf{x}^0, \mathbf{x}^N\}$, if and only if there does not exist any subgraph $F \in \mathcal{F}$ with density $d(F) = \frac{E'}{N'}$ for which $\lambda\gamma^{d(F)} > \mu$, or

$$\frac{\log\left(\frac{\lambda}{\mu}\gamma^{d(G)}\right)}{\log\left(\frac{\lambda}{\mu}\gamma^{d(F)}\right)} < \frac{N'}{N}.$$

This implies that all the induced subgraphs, $F \in \mathcal{F}$, in networks whose solution to the Most-Probable Configuration Problem is not a non-degenerate configuration in regime II), satisfy the condition that $\lambda\gamma^{d(F)} \leq \mu$ or

$$\frac{\log\left(\frac{\lambda}{\mu}\gamma^{d(G)}\right)}{\log\left(\frac{\lambda}{\mu}\gamma^{d(F)}\right)} \geq \frac{N'}{N},$$

for all $0 < \frac{\lambda}{\mu} \leq 1, \gamma > 1$.

Depending on the effective infection rate and the endogenous infection rate, $\left(\frac{\lambda}{\mu}, \gamma\right)$, the first condition $\lambda\gamma^{d(F)} \leq \mu$ may not be satisfied. However, since $\frac{N'}{N} \leq 1$ regardless of the parameters and the underlying network, the second condition is satisfied if

$$\frac{\log\left(\frac{\lambda}{\mu}\gamma^{d(G)}\right)}{\log\left(\frac{\lambda}{\mu}\gamma^{d(F)}\right)} \geq 1, \quad \forall F \in \mathcal{F}.$$

Since $\gamma > 1$, this means that $d(F) \leq d(G)$ for all possible induced subgraph. As this only depend on the structure of the underlying network, we can conclude that $d(F) \leq d(G)$ for networks whose most-probable configuration can only be \mathbf{x}^0 and/or \mathbf{x}^N .

□

Proof for Lemma 5.6.1

Lemma. Given an undirected graph G described by adjacency matrix \mathbf{A} and two different configurations \mathbf{x}_1 and \mathbf{x}_2 , which induce subgraphs $F(\mathbf{x}_1)$ and $F(\mathbf{x}_2)$, respectively. $F(\mathbf{x}_1)$ is isomorphic to $F(\mathbf{x}_2)$ if and only if $1^T \mathbf{x}_1 = 1^T \mathbf{x}_2$ and $\mathbf{x}_1^T \mathbf{A} \mathbf{x}_1 = \mathbf{x}_2^T \mathbf{A} \mathbf{x}_2$.

Proof. Recall that $1^T \mathbf{x}_1$ and $1^T \mathbf{x}_2$ are the number of infected nodes in the two configurations. Therefore, they are the number of nodes in each induced subgraph, $F(\mathbf{x}_1)$ and $F(\mathbf{x}_2)$. The number of infected edges in the two configurations are, $\frac{\mathbf{x}_1^T \mathbf{A} \mathbf{x}_1}{2}$ and $\frac{\mathbf{x}_2^T \mathbf{A} \mathbf{x}_2}{2}$, respectively. By the definition of the induced subgraphs, they are equal to twice the edge in each subgraph.

Necessity

If $F(\mathbf{x}_1)$ is isomorphic to $F(\mathbf{x}_2)$, then $1^T \mathbf{x}_1 = 1^T \mathbf{x}_2$ and $\mathbf{x}_1^T \mathbf{A} \mathbf{x}_1 = \mathbf{x}_2^T \mathbf{A} \mathbf{x}_2$.

This follows from the definition of isomorphism.

Sufficiency

If $1^T \mathbf{x}_1 = 1^T \mathbf{x}_2$ and $\mathbf{x}_1^T \mathbf{A} \mathbf{x}_1 = \mathbf{x}_2^T \mathbf{A} \mathbf{x}_2$, then $F(\mathbf{x}_1)$ is isomorphic to $F(\mathbf{x}_2)$.

We prove this by contrapositive. We need to prove that if $F(\mathbf{x}_1)$ is not isomorphic to $F(\mathbf{x}_2)$, then $1^T \mathbf{x}_1 \neq 1^T \mathbf{x}_2$ or $\mathbf{x}_1^T \mathbf{A} \mathbf{x}_1 \neq \mathbf{x}_2^T \mathbf{A} \mathbf{x}_2$.

There are two ways $F(\mathbf{x}_1)$ is not isomorphic to $F(\mathbf{x}_2)$: 1) There is no bijective function f or 2) There is a bijective function f but two vertices adjacent in $F(\mathbf{x}_1)$ is not adjacent in $F(\mathbf{x}_2)$. The bijection function f does not exist if the induced subgraphs, $F(\mathbf{x}_1)$ and $F(\mathbf{x}_2)$, have different number of nodes; this mean that $1^T \mathbf{x}_1 \neq 1^T \mathbf{x}_2$. 2) There is a bijection function, f , but u and v are adjacent in $F(\mathbf{x}_1)$ but are not adjacent in $F(\mathbf{x}_2)$. This would mean that $\mathbf{x}_1^T \mathbf{A} \mathbf{x}_1 \neq \mathbf{x}_2^T \mathbf{A} \mathbf{x}_2$. This is impossible by the definition of induced subgraph.

□

Proof for Theorem 5.6.1

Theorem. Consider Regime II) *Endogenous Infection Dominant*: $0 < \frac{\lambda}{\mu} \leq 1, \gamma > 1$ and network G . The most-probable configuration, \mathbf{x}^* , induces the subgraph $F(\mathbf{x}^*)$ with density $d(F(\mathbf{x}^*))$. If

$$\frac{\lambda}{\mu} \gamma^{d(F(\mathbf{x}^*))} > 1,$$

then \mathbf{x}^* is subgraph unique.

Proof. We will prove by contradiction. We can show that when the solution to the Most-Probable Configuration Problem are two equally probable configurations, \mathbf{x}_1^* and \mathbf{x}_2^* , whose induced subgraphs, $F_1 = F(\mathbf{x}_1^*)$ and $F_2 = F(\mathbf{x}_2^*)$, are isomorphic, we can always create another subgraph by combining F_1 and F_2 such that the configuration that induces this third subgraph will have a higher probability than \mathbf{x}_1^* and \mathbf{x}_2^* .

Suppose that the solution to the Most-Probable Configuration Problem are two equally probable configurations, \mathbf{x}_1^* and \mathbf{x}_2^* , whose induced subgraphs, $F_1 = F(\mathbf{x}_1^*)$ and $F_2 = F(\mathbf{x}_2^*)$, are isomorphic. From Lemma 5.6.1, this means that $|V(F_1)| = |V(F_2)| = N_1$, $|E(F_1)| = |E(F_2)| = E_1$, $d(F_1) = d(F_2)$. Additionally, we know that $\frac{\lambda}{\mu} \gamma^{d(F_1)} = \frac{\lambda}{\mu} \gamma^{d(F_2)} > 1$.

We now consider two cases: 1) $F_1 \cap F_2 = \emptyset$ and 2) $F_1 \cap F_2 \neq \emptyset$.

$$F_1 \cap F_2 = \emptyset$$

The induced subgraphs F_1 and F_2 are disconnected.

Define a new subgraph

$$\tilde{F} = F_1 \cup F_2.$$

We know that $|V(\tilde{F})| = 2N_1$ and $|E(\tilde{F})| = 2E_1$, while $d(\tilde{F}) = d(F_1) = d(F_2)$. The subgraph \tilde{F} has the same density as F_1 , which means

$$\frac{\lambda}{\mu} \gamma^{d(\tilde{F})} = \frac{\lambda}{\mu} \gamma^{d(F_1)} > 1.$$

Additionally $\tilde{N} > N_1$. Then as we are in regime II), by (5.9), the configuration that induces \tilde{F} has a larger equilibrium probability than \mathbf{x}_1^* , \mathbf{x}_2^* , thereby contradicting the premise that \mathbf{x}_1^* , \mathbf{x}_2^* are the most-probable configuration.

$$F_1 \cap F_2 \neq \emptyset$$

The induced subgraphs F_1 and F_2 are not disconnected.

Define a new subgraph

$$\hat{F} = F_1 \cap F_2. \tag{5.15}$$

We know then that $|V(\hat{F})| = \hat{N} < N_1$, $|E(\hat{F})| = \hat{E} < E_1$, with density $d(\hat{F}) = \frac{\hat{E}}{\hat{N}}$.

We have 3 cases to consider: 1) $d(\hat{F}) = d(F_1)$, 2) $d(\hat{F}) < d(F_1)$, 3) $d(\hat{F}) > d(F_1)$

1. $d(\hat{F}) = d(F_1)$. The intersection of F_1 and F_2 is a subgraph with the same density as F_1 and F_2 .

Define a new subgraph

$$\tilde{F} = (F_1 \cup F_2).$$

Let $|V(\tilde{F})| = \tilde{N}$ and $|E(\tilde{F})| = \tilde{E}$. Since F_1 and F_2 are not disconnected subgraphs, we know that $\tilde{N} = 2N_1 - \hat{N} > N_1$ and $\tilde{E} = 2E_1 - \hat{E} > E_1$, where \hat{N} and \hat{E} are the number of nodes and edges in subgraph \hat{F} defined in (5.15). First, we claim that $d(\tilde{F}) = d(\hat{F}) = d(F_1)$. This means that

$$d(\tilde{F}) = \frac{\tilde{E}}{\tilde{N}} = \frac{2E_1 - \hat{E}}{2N_1 - \hat{N}} = \frac{E_1}{N_1} \quad (5.16)$$

is a true statement. Using algebra, we see the above condition is equivalent to stating that

$$(2E_1 - \hat{E})N_1 = (2N_1 - \hat{N})E_1 \quad (5.17)$$

$$2E_1N_1 - \hat{E}N_1 = 2N_1E_1 - \hat{N}E_1 \quad (5.18)$$

$$\frac{\hat{E}}{\hat{N}} = \frac{E_1}{N_1}, \quad (5.19)$$

which is true under the case condition $d(\hat{F}) = d(F_1)$. The subgraph \tilde{F} has the same density as F_1 , which means

$$\frac{\lambda}{\mu} \gamma^{d(\tilde{F})} = \frac{\lambda}{\mu} \gamma^{d(F_1)} > 1.$$

Additionally $\tilde{N} > N_1$. Then as we are in regime II), using (5.9), the equilibrium probability of the configuration that induces \tilde{F} is larger than the equilibrium probability of $\mathbf{x}_1^*, \mathbf{x}_2^*$, thereby contradicting the premise that $\mathbf{x}_1^*, \mathbf{x}_2^*$ are the most-probable configuration.

2. $d(\hat{F}) < d(F_1)$. The intersection of F_1 and F_2 is a sparser subgraph than F_1 and F_2 .

Define a new subgraph

$$\tilde{F} = (F_1 \cup F_2).$$

Let $|V(\tilde{F})| = \tilde{N} = 2N_1 - \hat{N} > N_1$ and $|E(\tilde{F})| = \tilde{E} = 2E_1 - \hat{E} > E_1$, where \hat{N} and \hat{E} are the number of nodes and edges in subgraph \hat{F} defined in (5.15). First, we claim

that $d(\tilde{F}) > d(F_1)$. This means that

$$d(\tilde{F}) = \frac{\tilde{E}}{\tilde{N}} = \frac{2E_1 - \hat{E}}{2N_1 - \hat{N}} > \frac{E_1}{N_1} \quad (5.20)$$

is a true statement. Using algebra, this relationship is equivalently stating that

$$(2E_1 - \hat{E})N_1 > (2N_1 - \hat{N})E_1 \quad (5.21)$$

$$2E_1N_1 - \hat{E}N_1 > 2N_1E_1 - \hat{N}E_1 \quad (5.22)$$

$$\frac{\hat{E}}{\hat{N}} < \frac{E_1}{N_1}, \quad (5.23)$$

which is true under the case condition $d(\hat{F}) < d(F_1)$. The subgraph \tilde{F} is a denser subgraph than F_1 . Therefore,

$$\frac{\lambda}{\mu} \gamma^{d(\tilde{F})} > \frac{\lambda}{\mu} \gamma^{d(F_1)} > 1$$

and $\tilde{N} > N_1$. Then as we are in regime II), using (5.9), the equilibrium probability of the configuration that induces \tilde{F} is larger than the equilibrium probability of $\mathbf{x}_1^*, \mathbf{x}_2^*$, thereby contradicting the premise that $\mathbf{x}_1^*, \mathbf{x}_2^*$ are the most-probable configuration.

3. $d(\hat{F}) > d(F_1)$. The intersection of F_1 and F_2 is a denser subgraph than F_1 and F_2 .

Define a new subgraph

$$\tilde{F} = (F_1 \cup F_2).$$

Let $|V(\tilde{F})| = \tilde{N} = 2N_1 - \hat{N} > N_1$ and $|E(\tilde{F})| = \tilde{E} = 2E_1 - \hat{E} > E_1$, where \hat{N} and \hat{E} are the number of nodes and edges in subgraph \hat{F} defined in (5.15).

This is the most complicated case. We will show that the equilibrium probability of $\mathbf{x}_1^*, \mathbf{x}_2^*$ can not simultaneously be larger than the equilibrium probability of the configuration, $\hat{\mathbf{x}}$, which induces the subgraph \hat{F} and the equilibrium probability of the configuration, $\tilde{\mathbf{x}}$, which induces the subgraph \tilde{F} .

Suppose that $\pi(\mathbf{x}_1^*) = \pi(\mathbf{x}_2^*) > \pi(\hat{\mathbf{x}})$. This means that

$$\left(\frac{\lambda}{\mu} \gamma^{\frac{E_1}{N_1}}\right)^{N_1} > \left(\frac{\lambda}{\mu} \gamma^{\frac{\hat{E}}{\hat{N}}}\right)^{\hat{N}}.$$

This equivalently reduces to the relationship

$$\left(\frac{\lambda}{\mu}\right)^{N_1} \gamma^{E_1} > \left(\frac{\lambda}{\mu}\right)^{\hat{N}} \gamma^{\hat{E}}. \quad (5.24)$$

Suppose also that $\pi(\mathbf{x}_1^*) = \pi(\mathbf{x}_2^*) > \pi(\tilde{\mathbf{x}})$. This means that

$$\left(\frac{\lambda}{\mu}\gamma^{\frac{E_1}{N_1}}\right)^{N_1} > \left(\frac{\lambda}{\mu}\gamma^{\frac{\hat{E}}{N}}\right)^{\tilde{N}} = \left(\frac{\lambda}{\mu}\gamma^{\frac{2E_1 - \hat{E}}{2N_1 - \tilde{N}}}\right)^{2N_1 - \tilde{N}}.$$

This equivalently reduces to the relationship

$$\left(\frac{\lambda}{\mu}\right)^{N_1} \gamma^{E_1} > \left(\frac{\lambda}{\mu}\right)^{2N_1 - \tilde{N}} \gamma^{2E_1 - \hat{E}}. \quad (5.25)$$

By algebra, (5.25) reduce further to

$$\left(\frac{\lambda}{\mu}\right)^{N_1} \gamma^{E_1} < \left(\frac{\lambda}{\mu}\right)^{\tilde{N}} \gamma^{\hat{E}},$$

which contradicts the condition in (5.24). Therefore, $\pi(\mathbf{x}_1^*) = \pi(\mathbf{x}_2^*)$ can not be larger than both $\pi(\tilde{\mathbf{x}})$ and $\pi(\hat{\mathbf{x}})$ simultaneously; hence \mathbf{x}_1^* is not the most-probable configuration.

□

Proof for Theorem 5.6.3

Theorem. Consider two scaled SIS processes, $\{X_1(t), t \geq 0\}$ and $\{X_2(t), t \geq 0\}$, with parameters $\left(\frac{\lambda_1}{\mu_1}, \gamma_1\right)$ and $\left(\frac{\lambda_2}{\mu_2}, \gamma_2\right)$ satisfying the relationship

$$\begin{bmatrix} \log\left(\frac{\lambda_1}{\mu_1}\right) \\ \log(\gamma_1) \end{bmatrix} \neq \alpha \begin{bmatrix} \log\left(\frac{\lambda_2}{\mu_2}\right) \\ \log(\gamma_2) \end{bmatrix} \text{ for any real-value } \alpha.$$

If two different configurations \mathbf{x}_1 and \mathbf{x}_2 are equally probable under both processes, then their induced subgraphs are isomorphic, $F(\mathbf{x}_1) \simeq F(\mathbf{x}_2)$.

Proof. Let $\pi_1(\mathbf{x}_1)$ and $\pi_1(\mathbf{x}_2)$ be the equilibrium probability of configuration \mathbf{x}_1 and \mathbf{x}_2 under the first scaled SIS process with rates $\left(\frac{\lambda_1}{\mu_1}, \gamma_1\right)$ and $\pi_2(\mathbf{x}_1)$ and $\pi_2(\mathbf{x}_2)$ be the equilibrium probability of the configurations under the second scaled SIS process with rates $\left(\frac{\lambda_2}{\mu_2}, \gamma_2\right)$.

Since $\pi_1(\mathbf{x}_1) = \pi_1(\mathbf{x}_2)$, this is equivalent to stating that $\log(\pi_1(\mathbf{x}_1)) = \log(\pi_1(\mathbf{x}_2))$, which means

$$(1^T \mathbf{x}_1 - 1^T \mathbf{x}_2) \log\left(\frac{\lambda_1}{\mu_1}\right) + \left(\frac{\mathbf{x}_1^T \mathbf{A} \mathbf{x}_1}{2} - \frac{\mathbf{x}_2^T \mathbf{A} \mathbf{x}_2}{2}\right) \log(\gamma_1) = 0. \quad (5.26)$$

And $\pi_2(\mathbf{x}_1) = \pi_2(\mathbf{x}_2)$ is equivalent to

$$(1^T \mathbf{x}_1 - 1^T \mathbf{x}_2) \log \left(\frac{\lambda_2}{\mu_2} \right) + \left(\frac{\mathbf{x}_1^T \mathbf{A} \mathbf{x}_1}{2} - \frac{\mathbf{x}_2^T \mathbf{A} \mathbf{x}_2}{2} \right) \log(\gamma_2) = 0. \quad (5.27)$$

We can express (5.26) and (5.27) as a system of linear equations

$$\begin{bmatrix} (1^T \mathbf{x}_1 - 1^T \mathbf{x}_2) & \left(\frac{\mathbf{x}_1^T \mathbf{A} \mathbf{x}_1}{2} - \frac{\mathbf{x}_2^T \mathbf{A} \mathbf{x}_2}{2} \right) \end{bmatrix} \begin{bmatrix} \log \left(\frac{\lambda_1}{\mu_1} \right) & \log \left(\frac{\lambda_2}{\mu_2} \right) \\ \log(\gamma_1) & \log(\gamma_2) \end{bmatrix} = \begin{bmatrix} 0 & 0 \end{bmatrix}. \quad (5.28)$$

From linear algebra theory, we know that if

$$\begin{bmatrix} \log \left(\frac{\lambda_1}{\mu_1} \right) \\ \log(\gamma_1) \end{bmatrix} \neq \alpha \begin{bmatrix} \log \left(\frac{\lambda_2}{\mu_2} \right) \\ \log(\gamma_2) \end{bmatrix} \text{ for any real-value } \alpha, \text{ then the matrix}$$

$$\begin{bmatrix} \log \left(\frac{\lambda_1}{\mu_1} \right) & \log \left(\frac{\lambda_2}{\mu_2} \right) \\ \log(\gamma_1) & \log(\gamma_2) \end{bmatrix}$$

is full-rank. Then condition (5.28) is satisfied only when $1^T \mathbf{x}_1 - 1^T \mathbf{x}_2 = 0$ and $\mathbf{x}_1^T \mathbf{A} \mathbf{x}_1 - \mathbf{x}_2^T \mathbf{A} \mathbf{x}_2 = 0$ [64]. By Lemma 5.6.1, this is equivalent to stating that the induced subgraphs $F(\mathbf{x}_1)$ and $F(\mathbf{x}_2)$ are isomorphic. □

Proof for Corollary 5.6.4

Corollary. Consider two scaled SIS processes, $\{X_1(t), t \geq 0\}$ and $\{X_2(t), t \geq 0\}$, with parameters in regime II) **Endogenous Infection Dominant** that are not exponentially related. Therefore,

$0 < \frac{\lambda_1}{\mu_1}, \frac{\lambda_2}{\mu_2} \leq 1, \gamma_1, \gamma_2 > 1$ and

$$\begin{bmatrix} \log \left(\frac{\lambda_1}{\mu_1} \right) \\ \log(\gamma_1) \end{bmatrix} \neq \alpha \begin{bmatrix} \log \left(\frac{\lambda_2}{\mu_2} \right) \\ \log(\gamma_2) \end{bmatrix} \text{ for any real-value } \alpha.$$

If \mathbf{x}^* is the most-probable configuration for both scaled SIS processes, $\{X_1(t), t \geq 0\}$ and $\{X_2(t), t \geq 0\}$, then \mathbf{x}^* is the unique maximizer of the Most-Probable Configuration Problem for both processes.

Proof. Let $\pi_1(\mathbf{x}_1^*)$ be the equilibrium probability of configuration \mathbf{x}_1^* of the process with parameters $\lambda_1, \gamma_1, \mu_1$ and $\pi_2(\mathbf{x}_1^*)$ be the equilibrium probability of configuration \mathbf{x}_1^* of the process with parameters $\lambda_2, \gamma_2, \mu_2$. We will prove this by contradiction. Suppose that the most-probable configuration is not unique; there are two equally probable most-probable configurations, \mathbf{x}_1^* and \mathbf{x}_2^* .

$$\pi_1(\mathbf{x}_1^*) = \pi_1(\mathbf{x}_2^*) = \arg \max_{\mathbf{x} \in \mathcal{X}} \pi_1(\mathbf{x})$$

$$\pi_2(\mathbf{x}_1^*) = \pi_2(\mathbf{x}_2^*) = \arg \max_{\mathbf{x} \in \mathcal{X}} \pi_2(\mathbf{x})$$

From Theorem 5.6.3, we can then conclude that $F(\mathbf{x}_1^*)$ is isomorphic to $F(\mathbf{x}_2^*)$. However, this contradicts Theorem 5.6.1 which states that if \mathbf{x}_1^* is the most-probable configuration and $\frac{\lambda_1}{\mu_1} \gamma_1^{d(F(\mathbf{x}_1^*))} > 1$, there are no other induced subgraphs $F(\mathbf{x}')$ isomorphic to $F(\mathbf{x}_1^*)$. Similarly, if \mathbf{x}_1^* is the most-probable configuration and $\frac{\lambda_2}{\mu_2} \gamma_2^{d(F(\mathbf{x}_2^*))} > 1$, there are no other induced subgraphs $F(\mathbf{x}')$ isomorphic to $F(\mathbf{x}_2^*)$.

□

Characterizing Individual Susceptibility in Polynomial-Time

6.1 Introduction

Which agents in the population are more vulnerable to infection? Such microscale analysis can only be addressed by network processes, which consider interactions between individual agents. Identifying the agents most vulnerable to infection can impact control and treatment measures. The scaled SIS process models the interactions of individual agents in a population. It is therefore theoretically possible to study the susceptibility of individual agents to infection at equilibrium using the equilibrium distribution (3.4). In practice, this is intractable for large networks since solving for the marginal probability of infection of an agent, $P(x_i = 1)$, requires knowing the value of the partition function Z . This chapter shows that it is possible to efficiently and accurately approximate $P(x_i = 1)$ in regime II) **Endogenous Infection Dominant**: $0 < \frac{\lambda}{\mu} \leq 1, \gamma > 1$ using the Perturb-and-MAP sampling method [40], a Monte Carlo sampling technique.

Section 6.2 defines the marginal probability of infection. Section 6.3 reviews the Perturb-and-MAP sampling method and shows that an estimate, $\hat{P}(x_i = 1)$, can be found efficiently and with high accuracy approximates the true probability of infection, $P(x_i = 1)$. Using both small artificial networks, for which we know the exact marginal probability of infection, and large real-world networks, section 6.4 analyzes the dependence of individual agents susceptibility to infection on both the healing and infection rates, λ, γ, μ and on the topological location of the agents in the network.

6.2 Marginal Probability of Infection

The equilibrium distribution of the scaled SIS process (3.4), $\pi(\mathbf{x})$, is a joint probability distribution. It characterizes exactly the probabilistic relationships between all the agents in the network at the equilibrium of the scaled SIS process, $\{X(t), t \geq 0\}$. The susceptibility of individual agents to the epidemics can be characterized by its marginal probability. The singleton marginal probability of $\pi(\mathbf{x})$ is the probability that the i th agent is infected at equilibrium,

$$\begin{aligned} P(x_i = 1) &= \sum_{\mathbf{x} \in \mathcal{X}: x_i=1} \pi(\mathbf{x}) \\ &= \frac{1}{Z} \sum_{\mathbf{x} \in \mathcal{X}: x_i=1} \left(\frac{\lambda}{\mu}\right)^{1^T \mathbf{x}} \gamma^{\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2}}. \end{aligned} \quad (6.1)$$

The probability that the i th agent is healthy is $P(x_i = 0) = 1 - P(x_i = 1)$. Solving for the *exact* singleton marginal probability is intractable for large networks. It requires knowing the partition function, Z , which belongs to the class of #P counting problems [65]. For the scaled SIS process, solving for the partition function means summing over 2^N configurations in \mathcal{X} .

Monte Carlo methods can be applied to approximate intractable inference problems [41]. These methods consist of generating large number of random samples to obtain numerical approximations. Let $k = 1, 2, \dots, m$ be the number of independent samples drawn from the equilibrium distribution $\pi(\mathbf{x})$. The k th sample is $\mathbf{x}_k = [(\mathbf{x}_k)_1, (\mathbf{x}_k)_2, \dots, (\mathbf{x}_k)_N]^T$, which is the N-tuple collection of all the agent states in the network. The estimated marginal probability that the i th agent is infected at equilibrium, $\widehat{P}(x_i = 1)$, is the sample mean of the number of times agent i is infected

$$\widehat{P}(x_i = 1) = \frac{1}{m} \sum_{k=1}^m \mathbb{1}((\mathbf{x}_k)_i = 1), \quad (6.2)$$

where $\mathbb{1}(\cdot)$ is the indicator function.

The challenge of this approach is how to efficiently generate independent and identically distributed (i.i.d.) samples from the equilibrium distribution of the scaled SIS process.

6.3 Approximating Marginal Probability of Infection Using Sampling

The equilibrium distribution of the scaled SIS process is a Gibbs distribution. Reference [40] proposed an efficient method for generating independent, unbiased samples from a subclass of

the Gibbs distributions. This method is based on finding the MAP (maximum a posteriori) of the perturbed energy function (i.e., the Hamiltonian); it is called **Perturb-and-MAP** and is suitable for sampling from Gibbs distributions for which the MAP can be solved efficiently. We proved in chapter 5.2 that the most-probable configuration of the equilibrium distribution of the scaled SIS process can be found efficiently in regime II) **Endogenous Infection Dominant**: $0 < \frac{\lambda}{\mu} \leq 1, \gamma > 1$. We showed using an example network that the configuration, out of 2^{4941} configurations, with the maximum equilibrium probability can be found in 0.0041 secs on a standard desktop.

Perturb-and-MAP is based on the following properties of the Gumbel distribution. Consider a Gibbs distribution

$$P(\mathbf{x}) = \frac{1}{Z} e^{H(\mathbf{x})}, \mathbf{x} \in \mathcal{X}$$

where $H(\mathbf{x})$ is the Hamiltonian (i.e., energy function) of the distribution. If

$$G(\mathbf{x}) = [g_1, g_2, \dots, g_N]^T,$$

where g_1, g_2, \dots, g_N are i.i.d. samples from Gumbel(0,1), then

$$\arg \max_{\mathbf{x} \in \mathcal{X}} H(\mathbf{x}) + G(\mathbf{x}) \sim \frac{e^{H(\mathbf{x})}}{Z}.$$

The term $H(\mathbf{x}) + G(\mathbf{x})$ is a randomly perturbed version of the original Hamiltonian, $H(\mathbf{x})$. The maximizer of $H(\mathbf{x}) + G(\mathbf{x})$ is a sample from the original Gibbs distribution, $P(\mathbf{x})$. Samples can be obtained from $P(\mathbf{x})$ by solving for the maximizer of the randomly perturbed Hamiltonian multiple times. However, this method requires generating $|\mathcal{X}|$ number of random Gumbel perturbations $G(\mathbf{x})$. In practice, this is often not practical since the configuration space \mathcal{X} can be very large. For instance for the scaled SIS process, the size of the configuration space is exponential in the number of nodes in the network, 2^N .

Instead of sampling directly from the Gibbs distribution $P(\mathbf{x})$, reference [40] proposed two variations of the Perturb-and-MAP method, which sample from low dimensional approximations of the Gibbs distribution: order-1 and order-2. Low dimension approximations require generating only $2N$ number of i.i.d. Gumbel(0,1) random variables. For the scaled SIS process, the number of perturbations needed to sample from an approximation of the equilibrium distribution would be linear in the size of the network, N , instead of exponential.

Recall that the equilibrium distribution of the scaled SIS process, written as a Gibbs distribution, has the following form

$$\pi(\mathbf{x}) = \frac{1}{Z} e^{H(\mathbf{x})}$$

with

$$H(\mathbf{x}) = \sum_{i=1}^N x_i \log\left(\frac{\lambda}{\mu}\right) + \sum_{i=1}^N \sum_{j=1}^N A_{ij} x_i x_j \log(\gamma),$$

The term $x_i \log\left(\frac{\lambda}{\mu}\right)$ is the unary potential and the term $A_{ij} x_i x_j \log(\gamma)$ is the pairwise potential.

Order-1 approximation adds independent Gumbel(0,1) perturbations only to the unary potential. We proved in chapter 5.2, that the supermodular structure of the equilibrium distribution depends only on γ . As a result, the order-1 approximation is guaranteed to preserve the supermodular structure of the perturbed Hamiltonian $H(\mathbf{x}) + G(\mathbf{x})$. Order-2 approximation adds random perturbations to both the unary and pairwise potentials. It would generate samples from a more accurate approximation of the equilibrium distribution. However, it is not guaranteed to preserve supermodularity and is impractical as a method to estimate the marginal probabilities of large networks. Therefore, we only consider order-1 Perturb-and-MAP in this thesis.

6.3.1 Order-1 Perturb-and-MAP Accuracy

This section compares the approximate singleton marginal probabilities, $\hat{P}(x_i), i = 1, \dots, N$ (6.2), to the true singleton marginal probabilities, $P(x_i), i = 1, \dots, N$ (6.1). The true marginal probability can only be found for small networks. We consider the networks shown in Figure 6.1.

To quantify the accuracy of the estimator, we compare the true marginal probabilities with the estimate marginal probabilities using the following metric:

$$\text{Error} = \frac{1}{N} \sum_{i=1}^N |P(x_i = 1) - \hat{P}(x_i = 1)|. \quad (6.3)$$

The maximum possible error is 1.

To demonstrate the convergence rate and accuracy of the estimate, we plot the error (6.3) on the Y-axis as a function of the number of samples, m (from 10 to 10510 with a step size of 500) on the X-axis for network A, B, and C (see Figure 6.1) in Figures 6.2, 6.3, and 6.4, respectively.

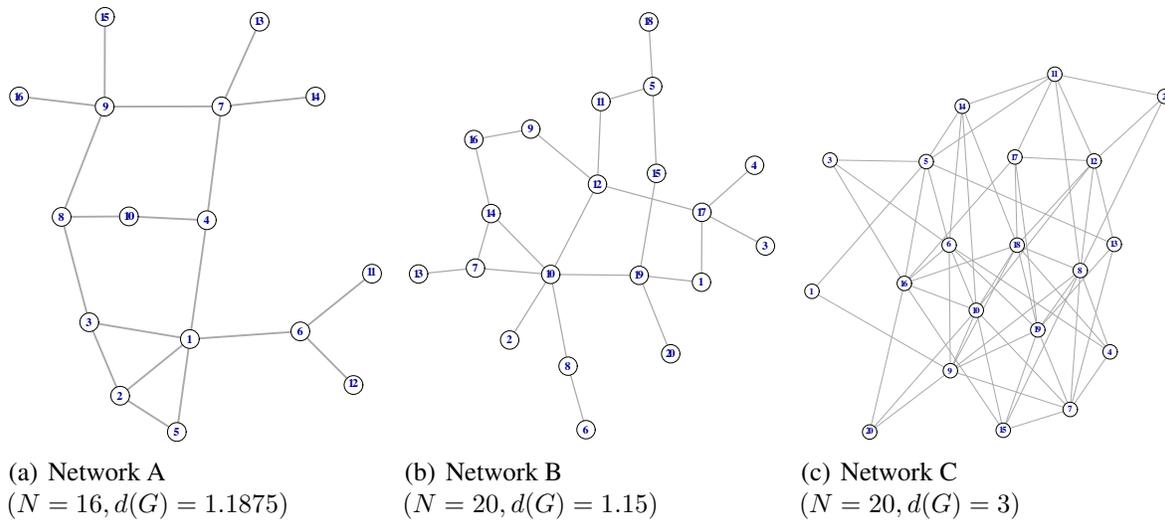
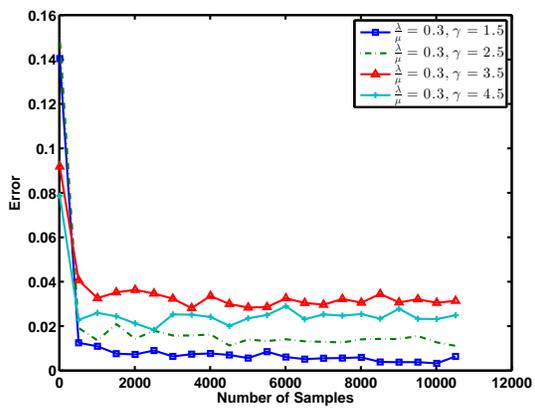
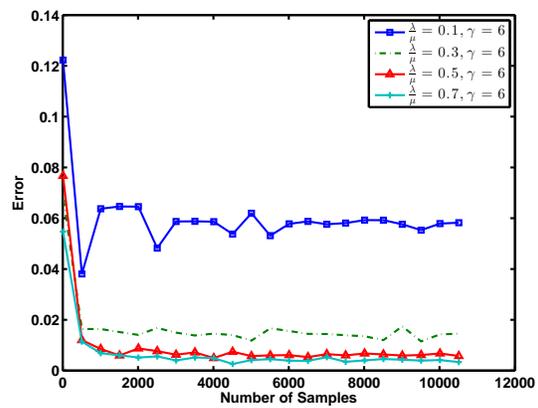


Figure 6.1: Small Networks



(a) Network A: Constant $\frac{\Delta}{\mu}$, Different γ



(b) Network A: Different $\frac{\Delta}{\mu}$, Constant γ

Figure 6.2: Network A: Error vs. Number of Samples

Different lines in Figures 6.2, 6.3, and 6.4 correspond to scaled SIS processes with different $\frac{\lambda}{\mu}, \gamma$ parameters. All the parameters are in the range of regime II) **Endogenous Infection Dominant**.

From Figures 6.2, 6.3, and 6.4, the error converges rapidly with a small number of samples, m , for all three example networks. Due to computational constraints, we have to use the order-1 Perturb-and-MAP method. This means that we are sampling from an *approximation* of the true equilibrium distribution, $\pi(\mathbf{x})$ (3.4). As a result, the error will depend on the parameters $\frac{\lambda}{\mu}, \gamma$, which determines ‘raggedness of landscape of the equilibrium distribution and consequently how well i.i.d. samples from the low-dimensional Perturb-and-MAP approximates i.i.d. samples from the equilibrium distribution.

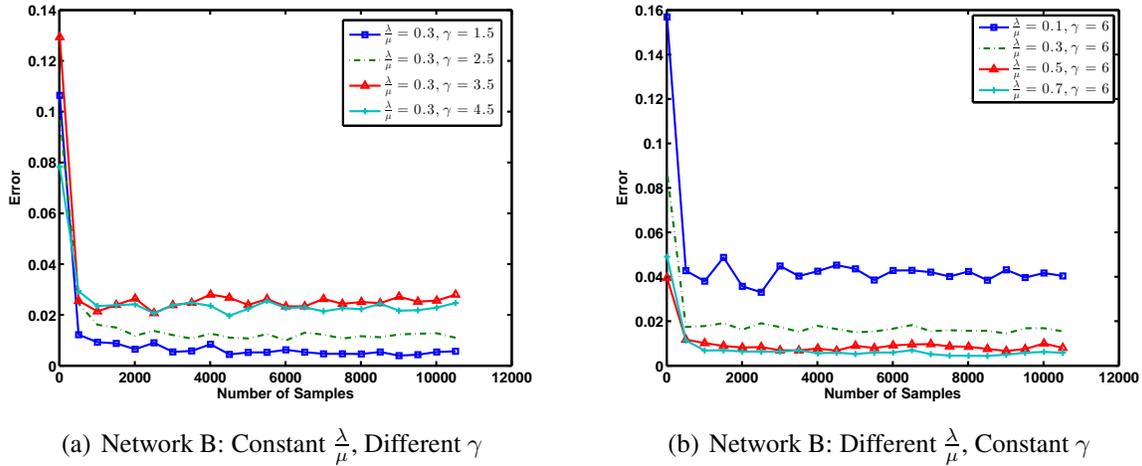
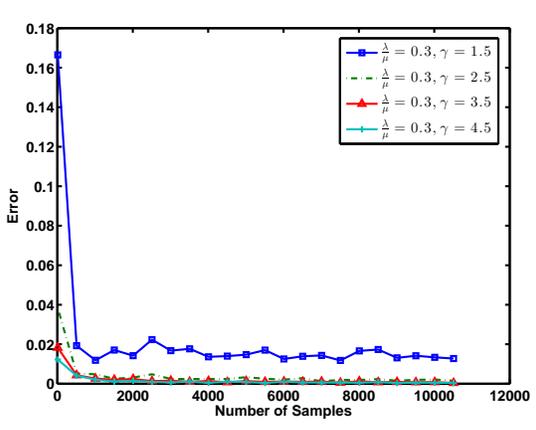


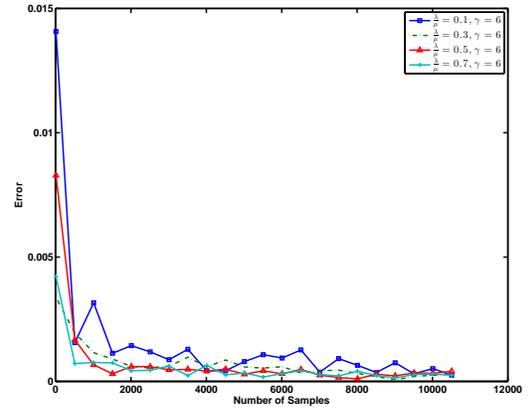
Figure 6.3: Network B: Error vs. Number of Samples

To investigate the dependence of the error on the dynamical parameters of the process, Figures 6.5, 6.6, and 6.7 show the error (6.3), represented by different color, as a function of $\frac{\lambda}{\mu}$ (from 0.1 to 0.9 with step size = 0.05) on the Y-axis and γ (from 1.5 to 6 with step size = 0.5) on the X-axis. The number of samples is fixed to $m = 6010$ samples, which is a very small fraction of the 2^N configuration space.

The error is less than 7% for all the networks. From the plot, the error appears to be larger for processes, $\{X(t), t \geq 0\}$, with small $\frac{\lambda}{\mu}$ and large γ value. It is known from Markov random field literature that Gibbs distributions with strong unary potential (i.e., $0 < \frac{\lambda}{\mu} \ll 1$) and strong pairwise potential (i.e., $\gamma \gg 1$) have more ragged energy landscape and are more challenging computationally [65]. For network C, because it is a more densely connected network than net-



(a) Network C: Constant $\frac{\lambda}{\mu}$, Different γ



(b) Network C: Different $\frac{\lambda}{\mu}$, Constant γ

Figure 6.4: Network C: Error vs. Number of Samples

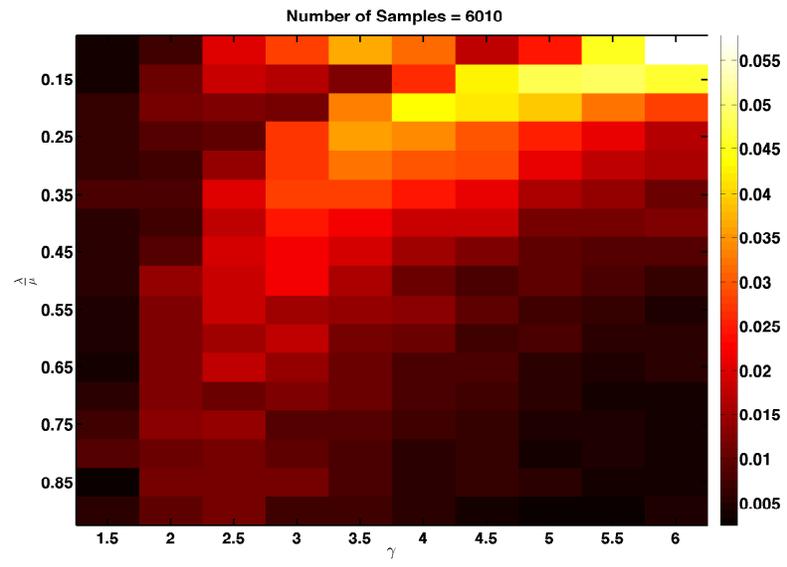


Figure 6.5: Network A: Error vs. Dynamics Parameters (Y-axis = $\frac{\lambda}{\mu}$, X-axis = γ)

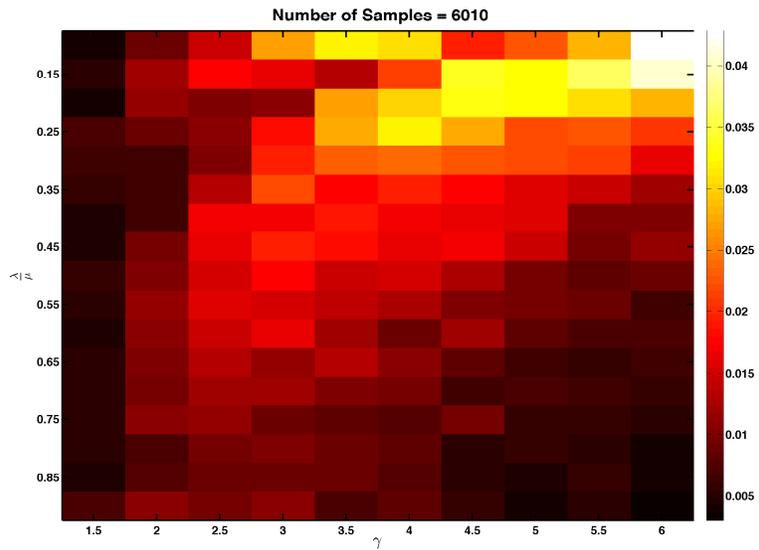


Figure 6.6: Network B: Error vs. Dynamics Parameters (Y-axis = $\frac{\lambda}{\mu}$, X-axis = γ)

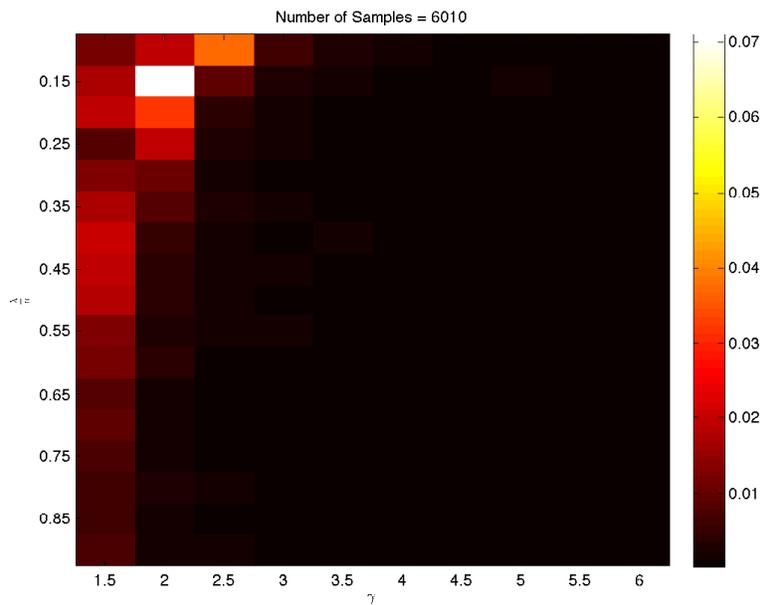


Figure 6.7: Network C: Error vs. Dynamics Parameters (Y-axis = $\frac{\lambda}{\mu}$, X-axis = γ)

works A and B, a strong pairwise potential corresponds to a smaller γ value than for networks A and B.

6.3.2 Comparing Perturb-and-MAP with Loopy Belief Propagation

Belief propagation is a popular algorithm for finding the marginal probabilities and solving other types of inference problems for Markov random fields [41]. It is known to be exact for networks that are trees but only approximate for networks with cycles. We used the UGM toolbox [66] to find the marginal probabilities for the equilibrium distribution of the scaled SIS process on networks A, B, and C using LBP (loopy belief propagation). For the equilibrium distribution of the scaled SIS process, the approximate marginal probabilities obtained by the order-1 Perturb-and-MAP method tracks the true marginal probabilities better than LBP.

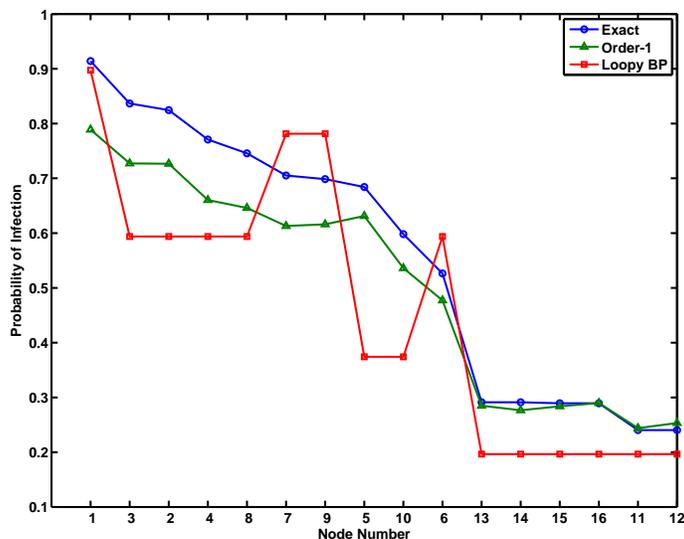


Figure 6.8: Network A: $\frac{\lambda}{\mu} = 0.1, \gamma = 6$, error = 0.06

We compare the performance between order-1 Perturb-and-MAP and LBP when the error is the largest. For network A, deviation between the true marginal probability, $P(x_i = 1), i = 1, \dots, N$, and the estimate marginal probability, $\hat{P}(x_i = 1), i = 1, \dots, N$, is the largest when $\frac{\lambda}{\mu} = 0.1, \gamma = 6$ (see Figure 6.5). Figure 6.8 plots the true marginal probability of infection, $P(x_i = 1), i = 1, \dots, N$ on the Y-axis and the corresponding node ID on the X-axis for Network A in blue. The nodes are ordered such that the marginal probability of infection decreases. The

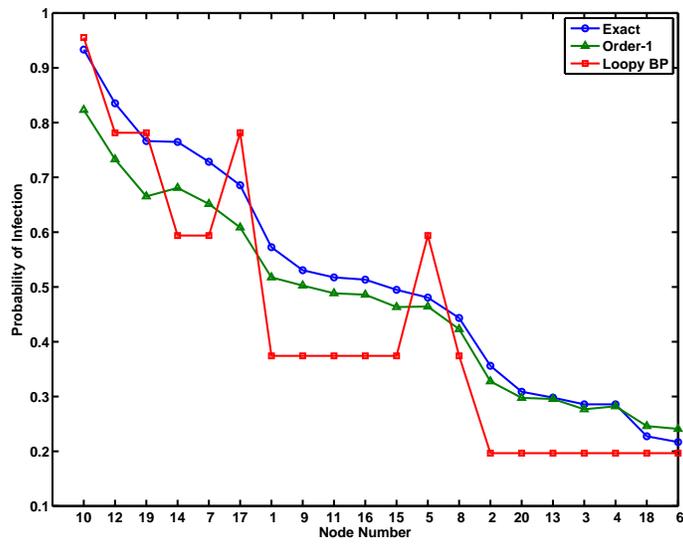


Figure 6.9: Network B: $\frac{\lambda}{\mu} = 0.1, \gamma = 6, \text{error} = 0.05$

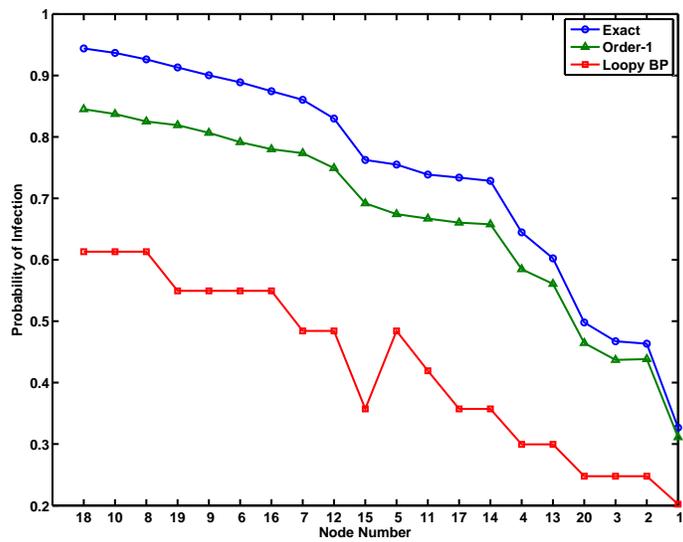


Figure 6.10: Network C: $\frac{\lambda}{\mu} = 0.15, \gamma = 2, \text{error} = 0.07$

green line corresponds to the estimates of $\hat{P}(x_i = 1), i = 1, \dots, N$ derived using order-1 Perturb-and-MAP, and the red line corresponds to the results of LBP. Figure 6.9 and Figure 6.10 show similar comparison for networks B and C, respectively.

All three figures show scenarios that have strong unary potential and strong pairwise potential; the energy landscape of the equilibrium distribution is ragged, which means that approximation methods such as Perturb-and-MAP and LBP perform poorly. We can see that Perturb-and-MAP tracks the true marginal probabilities much better than LBP in these cases.

6.4 Dependence of $P(x_i = 1)$ on Parameters and Topology

The susceptibility of an individual agent to infection depends on the infection and healing rates of the epidemics and on the location of the agent in the network; this can be analyzed and quantized by the marginal probability of the equilibrium distribution of the scaled SIS process $\{X(t), t \geq 0\}$. The previous section showed that the marginal probabilities, $P(x_i = 1)$, could be approximated efficiently using sampling techniques when $0 < \frac{\lambda}{\mu} \leq 1, \gamma > 1$. This section explores the dependence of the marginal probability on the dynamics parameters λ, γ, μ and on the network topology.

The marginal probability of each agent is a Bernoulli random variable that measures the probability the agent will be infected or healthy at equilibrium. At one extreme, we expect $P(x_i = 1) \approx 0, \forall i = 1, \dots, N$ for processes with low infection rates (i.e., $\gamma \approx 1$) and high healing rate (i.e., $0 < \frac{\lambda}{\mu} \ll 1$), and $P(x_i = 1) \approx 1, \forall i = 1, \dots, N$ for processes with high infection rates (i.e., $\gamma \gg 1$) and low healing rate (i.e., $\frac{\lambda}{\mu} \approx 1$). What we want to understand is the behavior of the marginal probabilities in between. This section shows the complex, nonlinear relationships between $P(x_i = 1)$ and the parameters and topology. First, we consider the simplest scenario of when $\gamma = 1$.

6.4.1 Simplest Scenario: $\gamma = 1$

When $\gamma = 1$, the number of infected neighbors does not affect the infection rate. Therefore, the structure of the network topology does not have an effect on the dynamics of the epidemics. We can find, in closed-form, all the statistics of the equilibrium distribution of the scaled SIS process

with

$$\pi(\mathbf{x}) = \frac{1}{Z} e^{\log\left(\frac{\lambda}{\mu}\right) \mathbf{1}^T \mathbf{x}},$$

where

$$\begin{aligned} Z &= \sum_{\mathbf{x} \in \mathcal{X}} e^{\log\left(\frac{\lambda}{\mu}\right) \mathbf{1}^T \mathbf{x}} \\ &= \sum_{k=0}^N e^{k \log\left(\frac{\lambda}{\mu}\right)} \binom{N}{k}. \end{aligned}$$

Using (6.1), regardless of the underlying network topology, the marginal probability of infection is the *same* for all the agents in the network and is

$$\begin{aligned} P(x_1 = 1) &= P(x_2 = 1) = \dots = P(x_N = 1) \\ &= \frac{\sum_{k=1}^N e^{k \log\left(\frac{\lambda}{\mu}\right)} \binom{N-1}{k-1}}{\sum_{k=0}^N e^{k \log\left(\frac{\lambda}{\mu}\right)} \binom{N}{k}}. \end{aligned} \tag{6.4}$$

Figure 6.11 plots the exact dependence of the marginal probability on $\frac{\lambda}{\mu}$ when $\gamma = 1$ using equation 6.4. The Y-axis shows the marginal probability of infection and the X-axis shows varying $\frac{\lambda}{\mu}$ values. Note that the relationship is nonlinear. In regime II), the parameter $\frac{\lambda}{\mu}$ is restricted to values between 0 and 1. This means that the healing rate, μ , is larger than or equal to the exogenous infection rate, λ . It is more likely that agents are in the healthy state than the infected state. We can see that the marginal probability of infection for all the agents is less than 0.5, $P(x_i = 1) < 0.5$.

6.4.2 Complex Scenario: $\gamma > 1$

When $\gamma > 1$, the additional number of infected neighbors increases the infection rate. As a result, the structure of the network topology may have an effect on the dynamics of the epidemics. Since the infection rate depends on the number of infected neighbors, does this mean that high degree agents are more susceptible to infection? In this section, we will show that this assumption holds *only* for some scaled SIS processes, $\{X(t), t \geq 0\}$. Which agent is more vulnerable to the epidemics depends on both the dynamics parameters and on the location of the agent in the network. When $0 < \frac{\lambda}{\mu} \leq 1, \gamma > 1$, the exact marginal probability of infection for large networks can not be found exactly but can be approximated efficiently with $\hat{P}(x_i = 1)$ (6.2) using order-1 Perturb-and-MAP sampling. We will use $P(x_i = 1)$ for small networks and $\hat{P}(x_i = 1)$ for large networks.

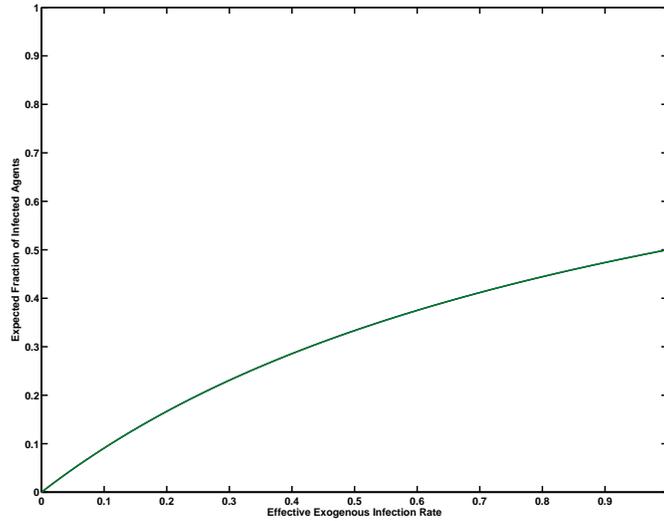
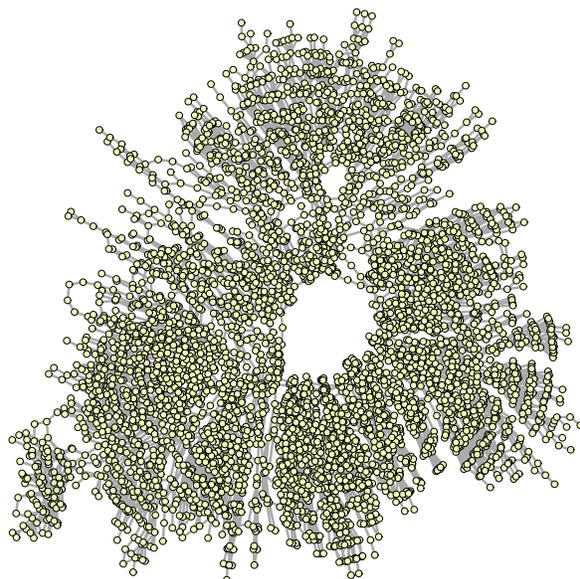


Figure 6.11: $P(x_i = 1), \forall i = 1, \dots, N$ vs $0 < \frac{\lambda}{\mu} \leq 1$ when $\gamma = 1$

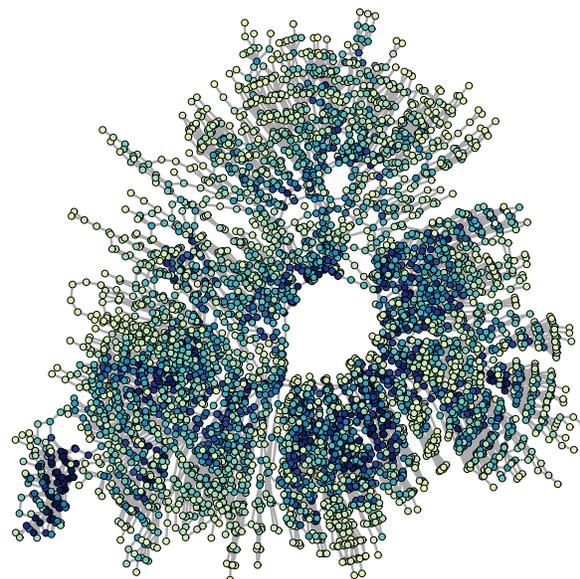
Figures 6.12 and 6.13 show the $\hat{P}(x_i = 1)$ under different dynamics parameters for the 4091 node US Western Power Grid. As expected, when the healing rate is much higher than the exogenous infection rate (i.e., $0 < \frac{\lambda}{\mu} \ll 1$) and the endogenous infection rate is low (i.e., $\gamma \approx 1$), the singleton marginal probability of infection is close to 0 for *all* the agents. When the healing rate is comparable with the exogenous infection rate (i.e., $0 < \frac{\lambda}{\mu} \approx 1$) and the endogenous infection rate is high (i.e., $\gamma \gg 1$), the singleton marginal probability of infection is close to 1 for *all* the agents. In between these extremes however, different agents have different probability of infection.

For the small 16 node Network A, we can consider how the individual marginal probability of infection changes depending on the parameters. Consider agent 5, 6, 9, and 10 in Network A (see Figure 6.1(a)). Node 5 has 2 neighbors; node 6 has 3 neighbors; node 9 has 4 neighbors, and node 10 has 2 neighbors. Since the infection rate depends on the number of infected neighbors, it is intuitive that an agent with more neighbors should have higher probability of infection. Surprisingly, this is not always the case.

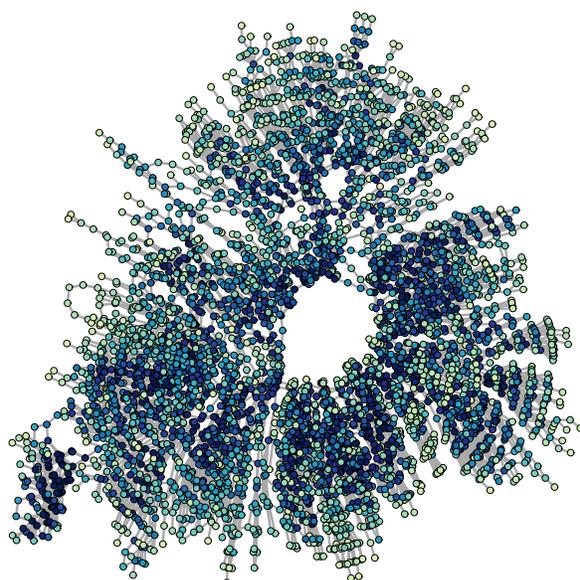
Figure 6.14, Figure 6.15, and Figure 6.16 plot $P(x_5 = 1)$, $P(x_6 = 1)$, $P(x_9 = 1)$, and $P(x_{10} = 1)$ as a function of γ for different $\frac{\lambda}{\mu}$ values in Network A. We choose agents 5, 6, 9 and 10 in Network A because they demonstrate the most surprising behavior. As γ increases, the topology-dependent infection process begins to dominate the healing process; as a result, the



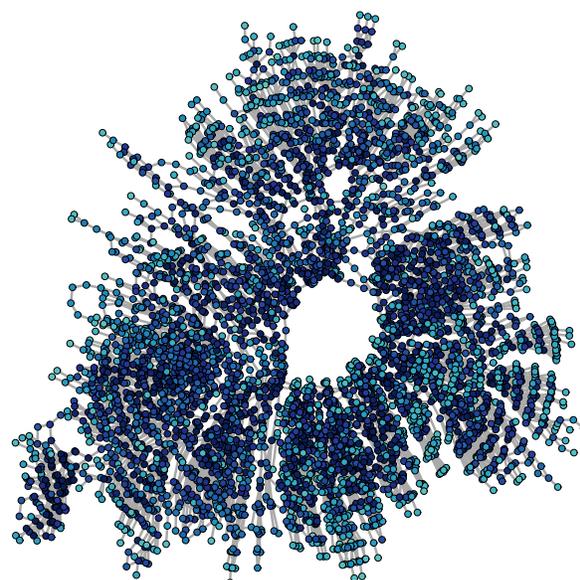
(a) $\left(\frac{\lambda}{\mu} = 0.1, \gamma = 1.5\right)$



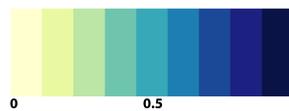
(b) $\left(\frac{\lambda}{\mu} = 0.1, \gamma = 4.5\right)$



(c) $\left(\frac{\lambda}{\mu} = 0.1, \gamma = 6\right)$



(d) $\left(\frac{\lambda}{\mu} = 0.3, \gamma = 4.5\right)$



(e) Colormap

Figure 6.12: $\widehat{P}(x_i = 1)$ of the US Western Power Grid

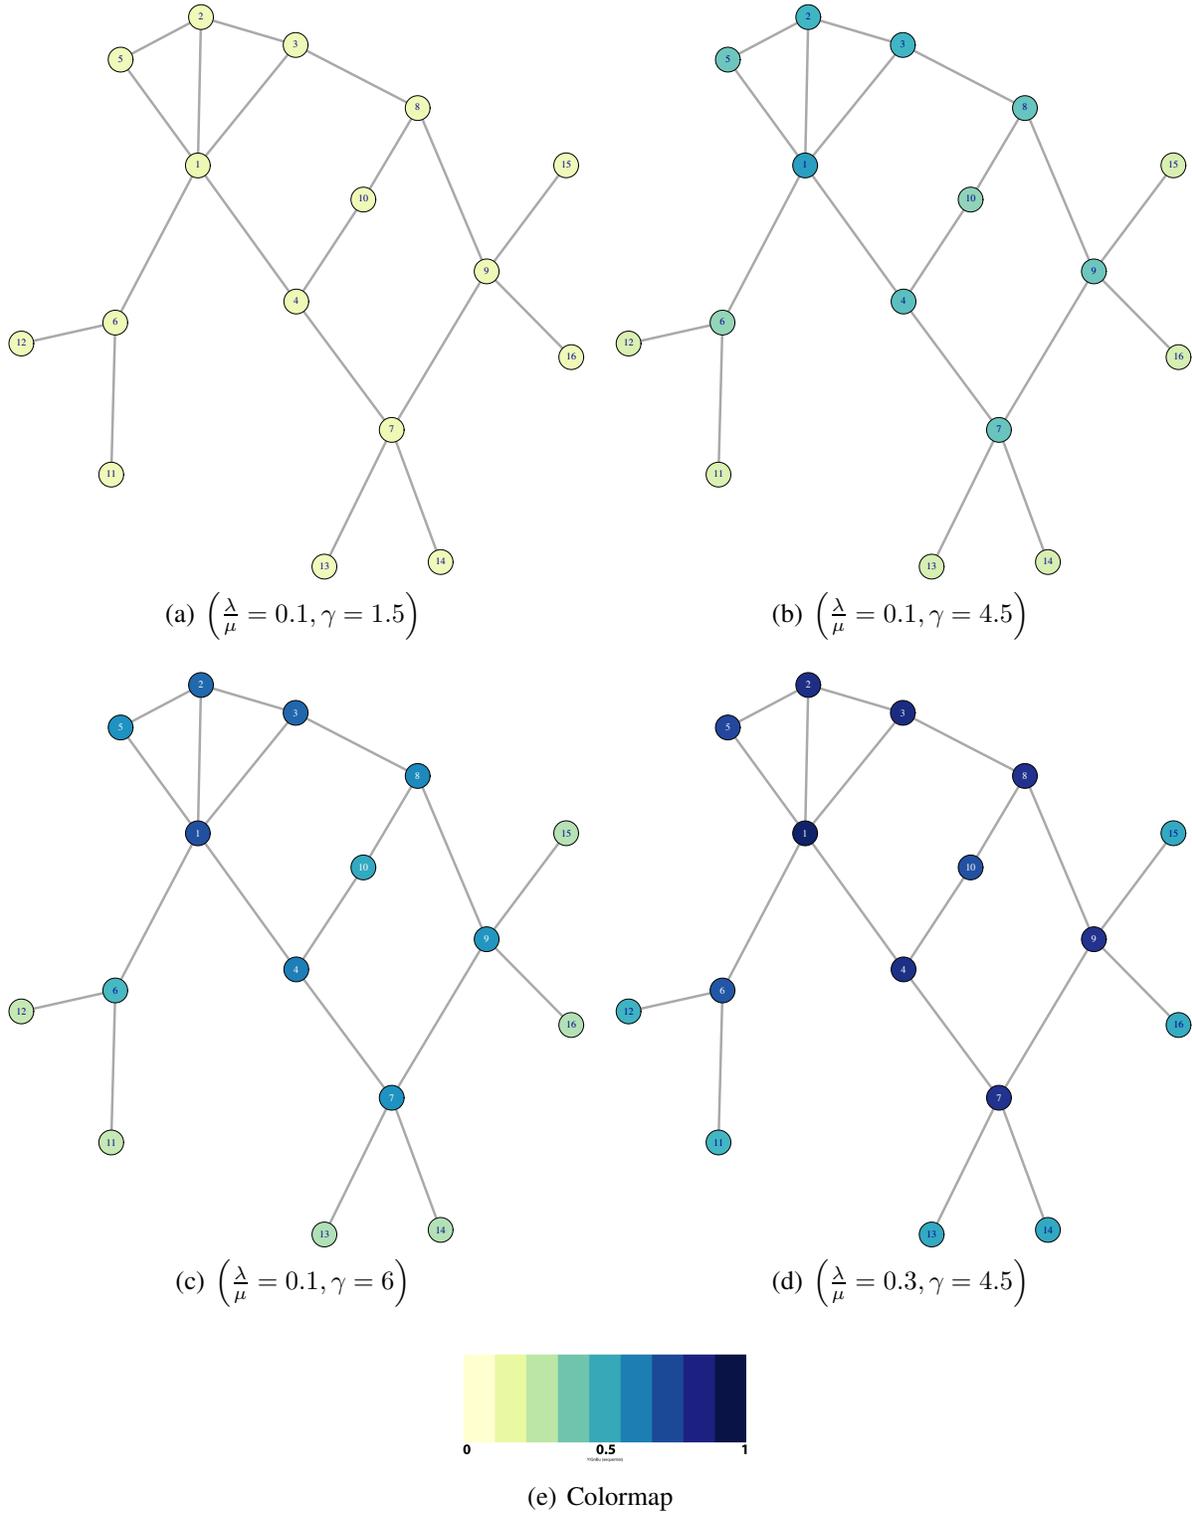


Figure 6.13: $P(x_i = 1)$ of Network A

marginal probability of infection for all agents in the network increases from 0 towards 1. This increase is nonlinear. Furthermore, the qualitative comparison of agent vulnerability changes depending on the dynamics parameters.

For example, in Figure 6.14, the marginal probability of infection of agent 9 is the largest for γ from 1.1 to 10, followed by agents 6, 5, and 10, respectively. In Figure 6.15 however, agent 5 has a higher probability of infection than agent 6. There is also a transition point where $P(x_{10} = 1) > P(x_6 = 1)$. In Figure 6.16, agent 5 is the node with the highest marginal probability of infection for γ from 1.5 to 1200, followed by agents 10, 9, and 6, respectively.

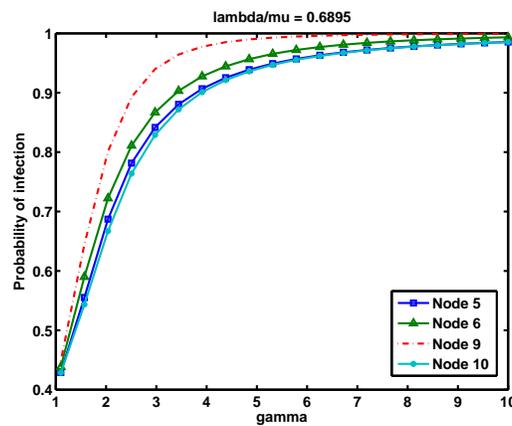


Figure 6.14: Network A: $P(x_i = 1)$ as a Function of γ , $\left(\frac{\lambda}{\mu} = 0.6895\right)$

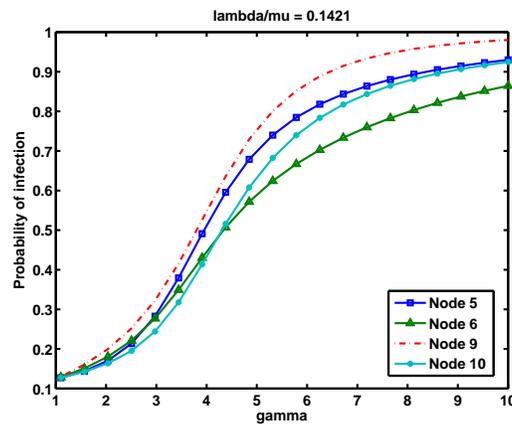


Figure 6.15: Network A: $P(x_i = 1)$ as a Function of γ , $\left(\frac{\lambda}{\mu} = 0.1421\right)$

Figure 6.17, Figure 6.18, and Figure 6.19 plot $P(x_5 = 1)$, $P(x_6 = 1)$, $P(x_9 = 1)$, and $P(x_{10} = 1)$ as a function of $\frac{\lambda}{\mu}$ for different γ values in network A. The marginal probability of

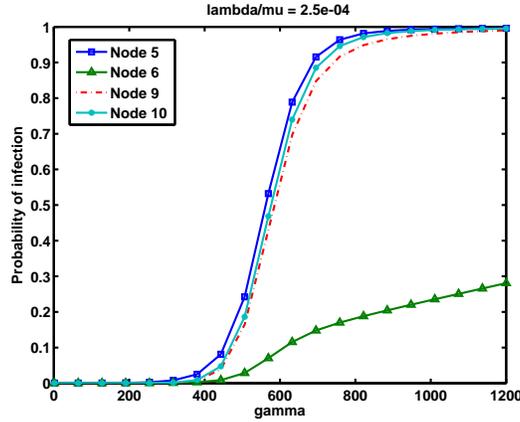


Figure 6.16: Network A: $P(x_i = 1)$ as a Function of γ , $\left(\frac{\lambda}{\mu} = 0.00025\right)$

infection also increases nonlinearly as $\frac{\lambda}{\mu}$ increases. Similarly, the marginal probability of agent 5 and 10 is larger than the marginal probability of infection of agent 6 in some range of parameters but is smaller than $P(x_6 = 1)$ in other range of parameter values.

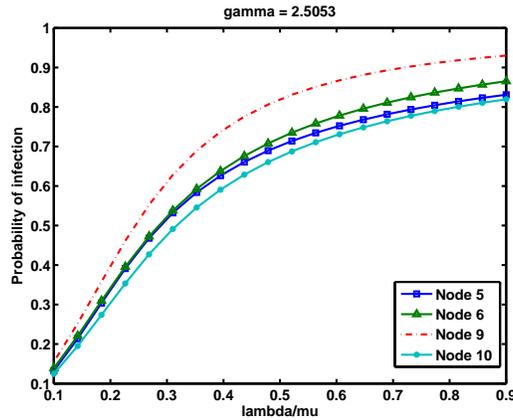


Figure 6.17: Network A: $P(x_i = 1)$ as a Function of $\frac{\lambda}{\mu}$ ($\gamma = 2.5053$)

Figure 6.20 shows the probability of infection of agent 6, $P(x_6 = 1)$, and the probability of infection of agent 10, $P(x_{10} = 1)$, as a function of both $\frac{\lambda}{\mu}$ (X-axis) and γ (Y-axis). For much of the parameter space, agent 6 is more likely to be infected at equilibrium than agent 10. However, there is a region for which agent 6 is *less* likely to be infected at equilibrium than agent 10.

This example illustrates that even though agent 6 is more connected in the underlying network (i.e., it has more neighbors), it is not necessarily more vulnerable than agent 10, which has fewer number of neighbors. One explanation as to why this may be the case comes from the solution of

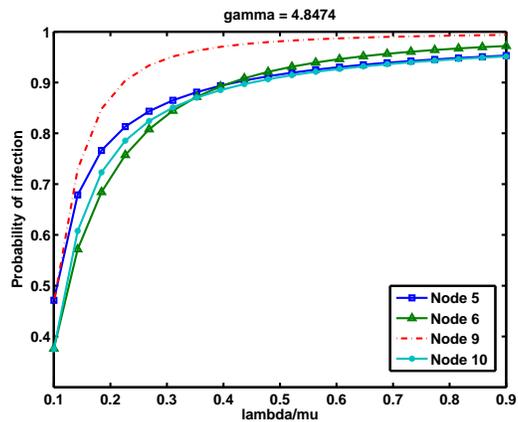


Figure 6.18: Network A: $P(x_i = 1)$ as a Function of $\frac{\lambda}{\mu}$, ($\gamma = 4.8474$)

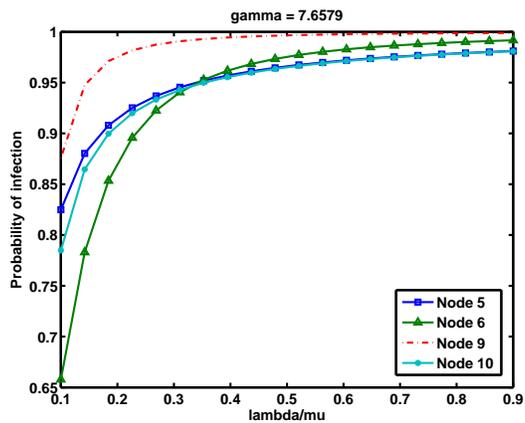


Figure 6.19: Network A: $P(x_i = 1)$ as a Function of $\frac{\lambda}{\mu}$, ($\gamma = 7.6579$)

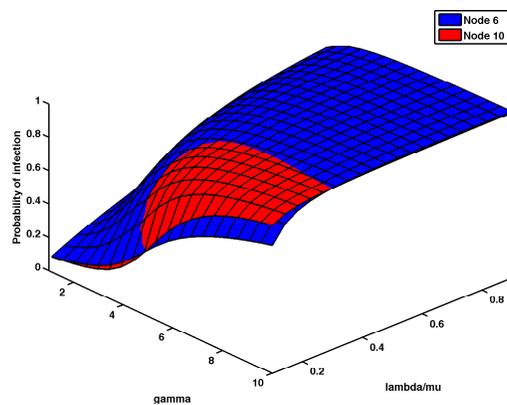


Figure 6.20: Network A: $P(x_6 = 1)$ and $P(x_{10} = 1)$

the Most-Probable Configuration Problem (4.1). We showed in Figure 5.5(c) that node 10 belongs to a subgraph that is denser than the overall network whereas node 6 does not. For a certain range of dynamics parameters, λ, γ, μ , local characterizations of the network topology such as node degree are sufficient to indicate which agent may be more vulnerable to infection. However, for other ranges of dynamic parameters, a coarser characterization such as membership in dense subgraphs is a better indicator of susceptibility.

The next section studies how the qualitative behavior of nodal susceptibility depends on the underlying network topology by measuring the correlation between the marginal probability of infection and topological measures such as degree centrality and eigenvector centrality.

6.4.3 Relating Marginal Probabilities to Centrality Measures

Before we explore the effect of network topology on the marginal probability of infection, we prove that, for a special subclass of networks, the network topology does not matter. We introduce the definition of vertex transitive graph.

Definition 6.4.1. *A graph $G(V, E)$ is **vertex transitive** if for any pair of vertices in V , there is an automorphism mapping from one to the other. A vertex transitive is always k -regular, but a k -regular graph may not be vertex transitive [43].*

Theorem 6.4.2. *If $G(V, E)$ is a vertex transitive graph, then $P(x_1 = 1) = P(x_2 = 1) = \dots = P(x_N = 1)$ for any dynamics parameters, λ, γ, μ .*

Proof. Under one vertex labeling, agent i is the i th node in graph $G(V, E)$. Since the graph is vertex transitive, there is an automorphism, f , such that agent i is the $f(i) = j \neq i$ in another network, $G_2(V, E)$. Since $G(V, E)$ and $G_2(V, E)$ are the same network but with different vertex labelings (i.e., isomorphic), then the equilibrium behavior of the scaled SIS process over both networks is the same. Therefore, $P(x_i = 1) = P(x_j = 1), \forall, i, j = 1, \dots, N$. \square

A vertex transitive graph is a very restrictive structure. Real-world networks such as Figure 2.4, Figure 2.5, Figure 2.6, and Figure 2.7 are not vertex transitive.

Nodal centrality measures the importance of a node within the network structure. They are based solely on the network structure. Common nodal centrality measures are [18]:

1. Degree centrality (DC) of the i th node is

$$\frac{d_i}{N - 1}, \quad (6.5)$$

where d_i is the number of neighbors of node i and N is the total number of nodes in the network. Degree centrality characterizes the local neighborhood of a node. The contagion rate of the scaled SIS process is dependent on the number of infected neighbors of a susceptible node; it is therefore dependent on the degree of the node.

2. Betweenness centrality (BC) of the i th node is

$$\sum_{k \neq j: i \notin \{k, j\}} \frac{2P_i(kj)}{P(kj)(N - 1)(N - 2)}, \quad (6.6)$$

where $P_i(kj)$ is the number of shortest paths between node k and node j that go through node i , and $P(kj)$ is the number of shortest paths between node k and node j . Betweenness centrality of the i th node quantifies how important the i th node is for connecting any pair of nodes k and j in the network

3. Eigenvector centrality (EC) is the eigenvector associated with the largest eigenvalue of the adjacency matrix, A , of the network G . The centrality of the i th node is proportional to the sum of the eigenvector centralities of its neighbors.

4. Subgraph centrality (SC) of the i th node is

$$\sum_{k=0}^{\infty} \frac{(A^k)_{ii}}{k!}, \quad (6.7)$$

where $(A^k)_{ii}$ is the i th diagonal entry of the adjacency matrix raised to the k th power. Subgraph centrality counts the number of closed walks starting and ending at the i th node. Subgraph centrality can only be calculated *exactly* for small size networks since it requires raising the adjacency matrix to many different powers [62]

To gain insight to how network topology affects agent vulnerability, we need to quantitatively measure the relationship between the probability of infection, $P(x_i = 1)$, which depend on both the underlying network structure and the dynamics parameters, λ, γ, μ and nodal centralities, $C_i, i =$

$1, 2, \dots, N$, which depends on the network structure. We choose to use the correlation coefficient [67]:

$$r = \frac{\sum_{i=1}^N (P(x_i = 1) - \overline{P(x_i = 1)})(C_i - \overline{C_i})}{\sqrt{\sum_{i=1}^N (P(x_i = 1) - \overline{P(x_i = 1)})^2 (C_i - \overline{C_i})^2}}, \quad (6.8)$$

where

$$\overline{P(x_i = 1)} = \frac{1}{N} \sum_{i=1}^N P(x_i = 1), \text{ and}$$

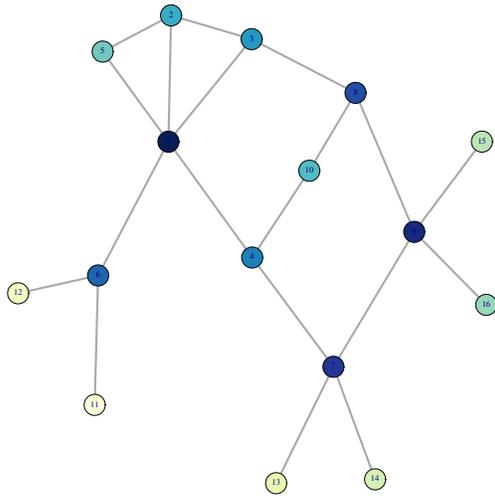
$$\overline{C_i} = \frac{1}{N} \sum_{i=1}^N C_i.$$

For small size networks, we can find $P(x_i = 1)$ exactly for all agents; for large networks, we can approximate the value with $\hat{P}(x_i = 1)$ using order-1 Peturb-and-MAP sampling. The correlation coefficient measures the linear statistical dependence between the marginal probability of infection and the centrality measure. It is between -1 and to 1 . When $r = 1$, the two values are completely, positively correlated; this means that a node with high centrality will also have large marginal probability of infection. When $r = 0$, the centrality gives no indication of the marginal probability of infection.

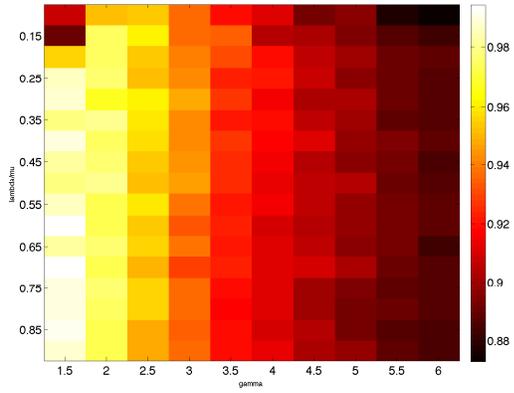
Figure 6.21 and Figure 6.22 show on the left the degree centrality, eigenvector centrality, betweenness centrality, and subgraph centrality of the nodes in Network A. The correlation coefficients (6.8) between the marginal probability of infection, $P(x_i = 1)$, for different parameters, λ, γ, μ , and the centrality measures are plotted on the right of these figures. The X-axis corresponds to different γ values ranging from 1.5 to 6 with a step size of 0.5. The Y-axis corresponds to different $\frac{\lambda}{\mu}$ ranging from 0.1 to 0.9 with a step size of 0.05. The correlation coefficients are similar for all the centrality measures. It has been shown that strong correlations exist between the different centralities [68].

As shown in Figure 6.13(a), when $0 < \frac{\lambda}{\mu} \ll 1$ and $\gamma \approx 1$, then $P(x_i = 1) \approx 0$ for all the agents. When $\frac{\lambda}{\mu} \approx 1$ and $\gamma \gg 1$, then $P(x_i = 1) \approx 1$ for all the agents. In these extreme cases, the marginal probability of infection has low correlation with the various centrality measures.

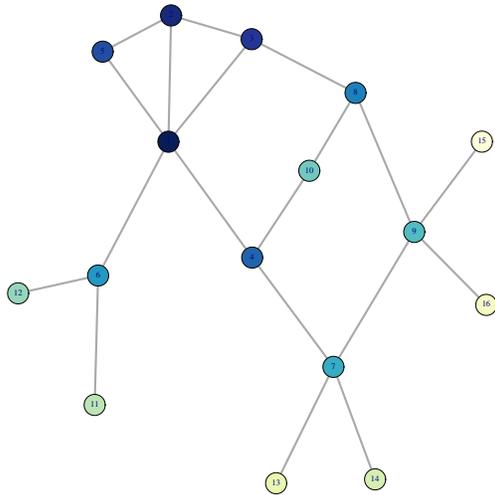
When $\frac{\lambda}{\mu}$ is large and γ is small, degree centrality is highly correlated with the marginal probability of infection. We can see from Figure 6.21(b) that the correlation *decreases* as γ increases. When γ is small, the topology-independent process, controlled by $\frac{\lambda}{\mu}$, dominates the topology-dependent process, controlled by γ ; the local characteristics of a node (i.e., its degree) is highly indicative of



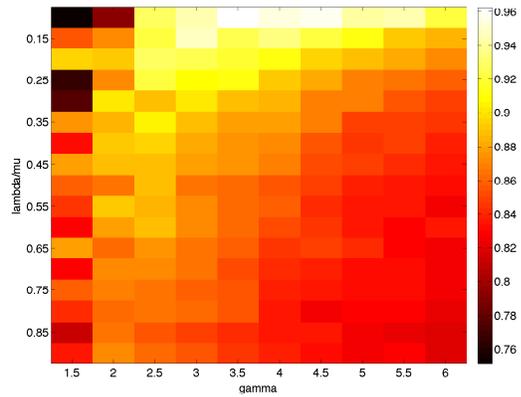
(a) Degree Centrality (Light = Small, Dark = Large)



(b) Correlation Coefficient between $P(x_i = 1)$ and Degree Centrality

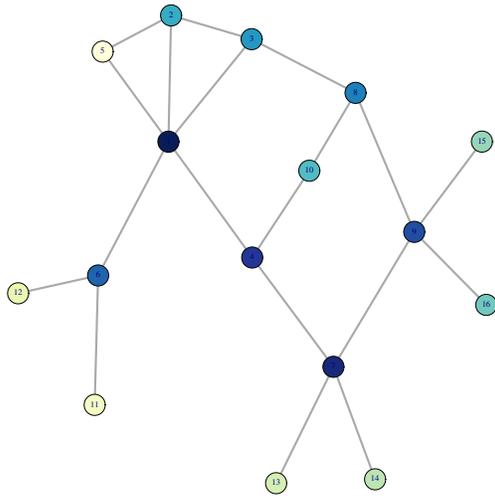


(c) Eigenvector Centrality (Light = Small, Dark = Large)

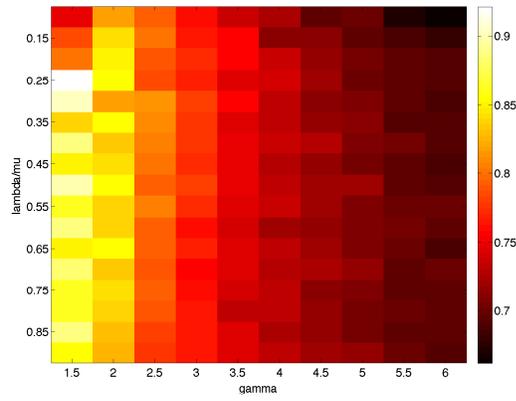


(d) Correlation Coefficient between $P(x_i = 1)$ and Eigenvector Centrality

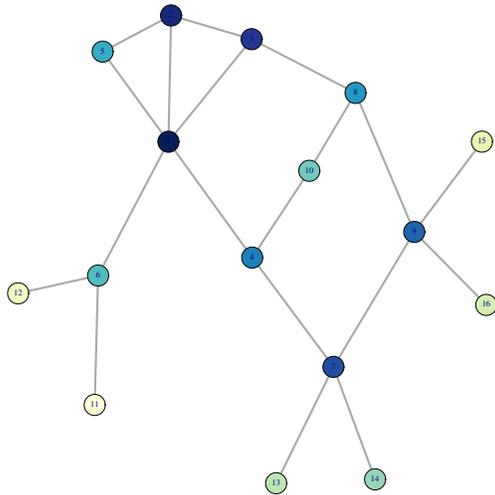
Figure 6.21: Network A: Centrality and Correlation



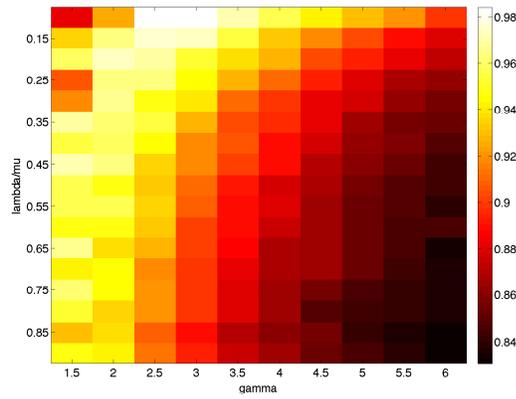
(a) Betweenness Centrality (Light = Small, Dark = Large)



(b) Correlation Coefficient between $P(x_i = 1)$ and Betweenness Centrality



(c) Subgraph Centrality (Light = Small, Dark = Large)



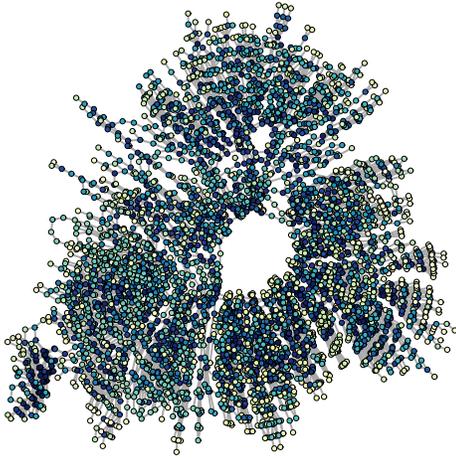
(d) Correlation Coefficient between $P(x_i = 1)$ and Subgraph Centrality

Figure 6.22: Network A: Centrality and Correlation

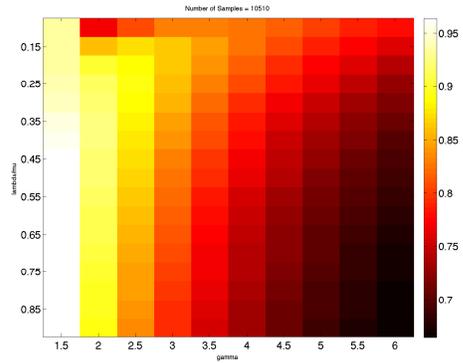
its vulnerability. However, as γ increases, the topology-dependent process starts to dominate the topology-independent process. Therefore, the degree centrality of an agent is less correlated with its marginal probability of infection.

Figure 6.23 shows the degree centrality, eigenvector centrality, and the betweenness centrality of the 4091 node US Western power grid and the corresponding correlation coefficients. Due to the size of the network, it is computationally intractable to find the subgraph centrality. The degree centrality best correlates with the marginal probability of infection; the dependence on the dynamics parameters is similar to that of Network A. However, unlike for the small 16 node Network A, the eigenvector and betweenness centrality are poor indicators of $P(x_i = 1)$.

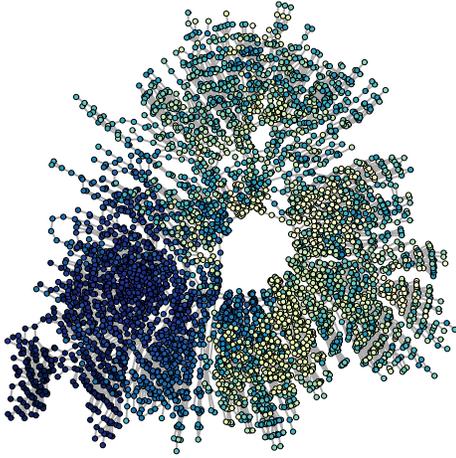
For the US Western power grid examples we considered, correlation between $P(x_i = 1)$ and the degree centrality is the smallest when $\frac{\lambda}{\mu} = 0.1$ and $\gamma = 6$. Figure 6.24(a) plots the marginal probability of infection. Figure 6.24(c) indicates the set of agents with marginal probability of infection greater than 0.5 as blue and the set of agents with marginal probability less than or equal to 0.5 as white. Figure 6.24(d) shows the most-probable configuration, \mathbf{x}^* , which is the configuration with the maximum equilibrium probability. We argued in chapter 5.4 that the nodes infected in \mathbf{x}^* belong to dense subgraphs; the marginal probability of infection of these nodes are high when $\frac{\lambda}{\mu} = 0.1$ and $\gamma = 6$. This indicates that a more global structural characterization such as dense subgraphs, as we consider here, is a better qualifier of nodal vulnerability. We can see that many of the nodes with marginal probability of infection larger than 0.5 in Figure 6.24(c) are also the infected nodes in the most-probable configuration in Figure 6.24(d). This is consistent with Theorem 4.7.1, which shows that the most-probable configuration dominates the configuration space for scaled SIS processes with small $\frac{\lambda}{\mu}$ and large γ in regime II).



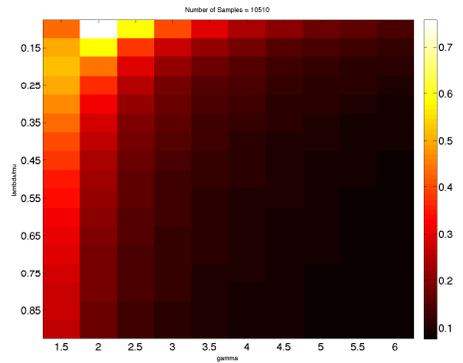
(a) Degree Centrality (Light = Small, Dark = Large)



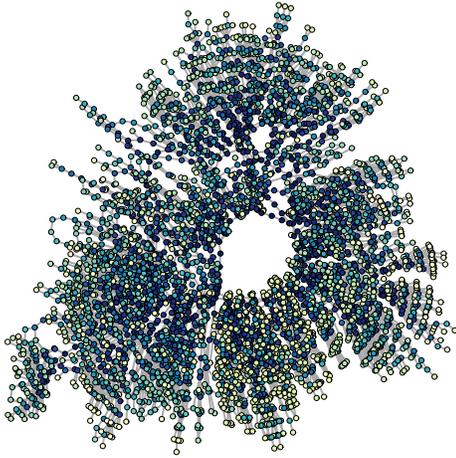
(b) Correlation Coefficient between $\hat{P}(x_i = 1)$ and Degree Centrality



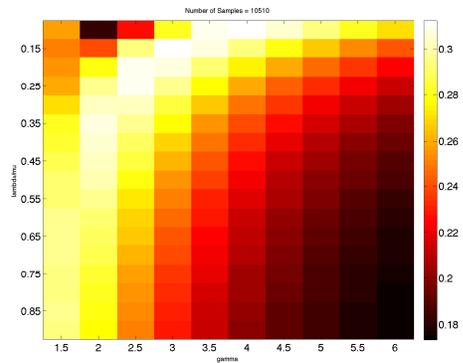
(c) Eigenvector Centrality (Light = Small, Dark = Large)



(d) Correlation Coefficient $\hat{P}(x_i = 1)$ and Eigenvector Centrality

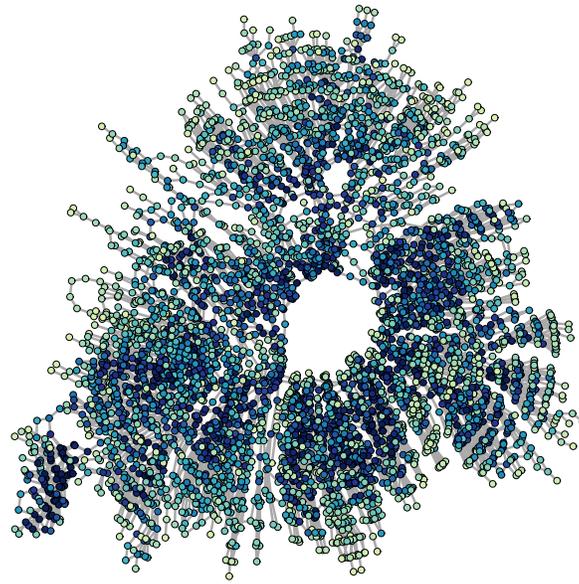


(e) Betweenness Centrality (Light = Small, Dark = Large)

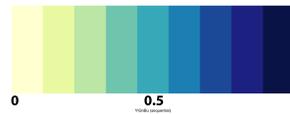


(f) Correlation Coefficient $\hat{P}(x_i = 1)$ and Degree Centrality

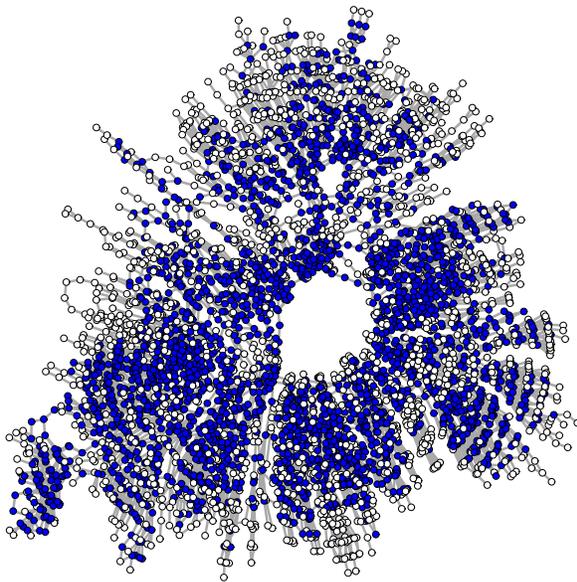
Figure 6.23: US Western Power Grid: Centrality and Correlation



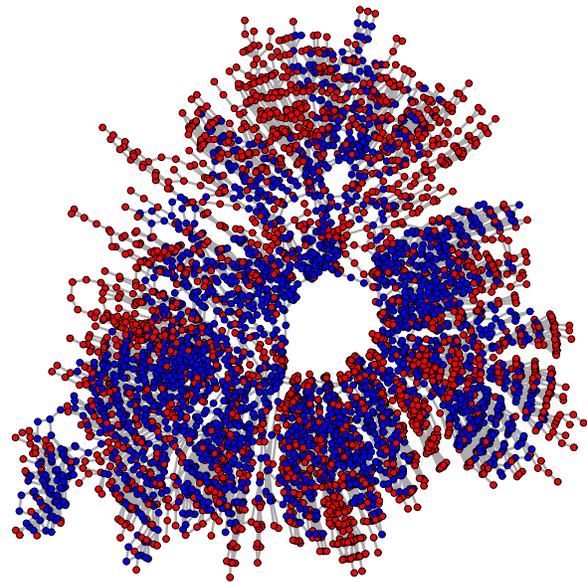
(a) $\hat{P}(x_i = 1)$ when $\left(\frac{\lambda}{\mu} = 0.1, \gamma = 6\right)$



(b) Colormap



(c) Blue = $\hat{P}(x_i = 1) > 0.5$, White = $\hat{P}(x_i = 1) \leq 0.5$



(d) Most-Probable Configuration when $\left(\frac{\lambda}{\mu} = 0.1, \gamma = 6\right)$ (Blue = Infected, Red = Healthy)

Figure 6.24: US Western Power Grid $P(x_i = 1)$ and Most-Probable Configuration when $\left(\frac{\lambda}{\mu} = 0.1, \gamma = 6\right)$

6.5 Conclusion

This chapter showed that the marginal probability of infection of individual nodes in a network can be efficiently approximated using Monte Carlo sampling techniques when the dynamics parameters of the scaled SIS process are in the regime where $0 < \frac{\lambda}{\mu} \leq 1, \gamma > 1$. This is due to the fact that the most-probable configuration can be solved in polynomial time in this regime. The error of the approximation is affected by the energy ‘landscape’ of the equilibrium distribution. It is known from Markov random field literature that Gibbs distributions with strong unary potential, corresponding to $\frac{\lambda}{\mu} \approx 0$ in regime II), and strong pairwise potential, corresponding to $\gamma \gg 1$ in regime II), result in more ragged probability distributions. Consequently, these are the parameter values that result in the largest discrepancy between the true marginal probabilities of infection, $P(x_i = 1)$, and the estimates, $\hat{P}(x_i = 1)$. We show that estimates using Perturb-and-MAP outperforms Loopy Belief Propagation in these cases.

Using simulation on both small, artificial graphs and large, real-world networks, our analysis show that, depending on the dynamics parameters of the epidemics, vulnerability may or may not be characterized by the total number of neighbors an agent has. The marginal probability of infection, $P(x_i = 1)$, and its estimate, $\hat{P}(x_i = 1)$, are *nonlinearly* dependent on the dynamics parameters. Furthermore, node A may have higher marginal probability of infection than node B for some values of $\lambda_1, \gamma_1, \mu_1$, but have lower probability of infection than node B for other values $\lambda_2, \gamma_2, \mu_2$. We explored this by studying the correlation between $P(x_i = 1)$ and nodal centrality measures such as degree centrality, eigenvector centrality, and betweenness centrality. When $\gamma \approx 1$, then $P(x_i = 1)$ is highly correlated with degree; the vulnerability of individual agents depends on their number of neighbors. As γ increases however, local structural characterization like node degree is no longer a good predictor of vulnerability. Instead, vulnerability depends on more global structural characterization such as if the agent belongs to a dense subgraph.

Expected Fraction of Infected Agents

7.1 Introduction

The susceptibility of the entire network population to infection can be characterized by macroscopic statistics such as the expected fraction of infected agents, $E[Y]$. A weak epidemics (i.e., high healing rates, low infection rates) means that $E[Y] \approx 0$, whereas a virulent epidemics (i.e., low healing rates, high infection rates) means that $E[Y] \approx 1$. However, solving for the expected fraction of infected agents of the equilibrium distribution of the scaled SIS process is intractable for large networks because it requires 2^N summations, where N is the total number of agents. First, we prove that $E[Y]$ is the average of the marginal probabilities of infection of all the agents. Therefore, we can approximate the value of $E[Y]$ using the Perturb-and-MAP sampling technique described in chapter 6 where we showed we could accurately compute an estimate of the marginal probabilities of infection, $\hat{P}(x_i = 1)$.

Section 7.2 defines the expected fraction of infected agents. Section 7.3 analyzes the dependence of $E[Y]$ on the dynamics parameters and the underlying network topology. The expected fraction of infected agents demonstrates nonlinear, sigmoidal behavior (i.e., S-curve) with respect to the dynamics parameters. Networks that are more densely connected have larger $E[Y]$ than sparsely connected networks for the same dynamic parameters.

7.2 Relating Expected Fraction of Infected Agents to Marginal Probabilities

The vulnerability of the entire population to infection can be characterized by the expected fraction of infected agents,

$$\begin{aligned} E[Y] &= E \left[\frac{1^T \mathbf{x}}{N} \right] = \sum_{\mathbf{x} \in \mathcal{X}} \frac{1^T \mathbf{x}}{N} \pi(\mathbf{x}) \\ &= \frac{1}{Z} \sum_{\mathbf{x} \in \mathcal{X}} \frac{1^T \mathbf{x}}{N} \left(\frac{\lambda}{\mu} \right)^{1^T \mathbf{x}} \gamma^{\frac{\mathbf{x}^T A \mathbf{x}}{2}}. \end{aligned} \tag{7.1}$$

Macroscopic characterizations relate to microscopic characterizations through the following theorem:

Theorem 7.2.1. *The expected fraction of infected agents is the average of the marginal probability of infection*

$$E[Y] = \frac{1}{N} \sum_{i=1}^N P(x_i = 1).$$

Proof. The marginal probability of infection of the i th agent is

$$P(x_i = 1) = \sum_{\mathbf{x} \in \mathcal{X}: x_i = 1} \pi(\mathbf{x}).$$

As $x_i \in \{0, 1\}$, the marginal probability can also be written as

$$P(x_i = 1) = \sum_{\mathbf{x} \in \mathcal{X}} x_i \pi(\mathbf{x}).$$

From (7.1),

$$\begin{aligned} E[Y] &= \frac{1}{N} \sum_{\mathbf{x} \in \mathcal{X}} (x_1 + x_2 + \dots + x_N) \pi(\mathbf{x}) \\ &= \frac{1}{N} \sum_{i=1}^N P(x_i = 1) \end{aligned}$$

□

Computing the exact expected fraction of infected agents in large networks, like finding the marginal probability of infection, is intractable. However, it can also be approximated using Monte

Carlo methods. Let $k = 1, 2, \dots, m$ be the number of independent samples drawn from the equilibrium distribution $\pi(\mathbf{x})$. The k th sample is $\mathbf{x}_k = [(\mathbf{x}_k)_1, (\mathbf{x}_k)_2, \dots, (\mathbf{x}_k)_N]^T$, which is the N -tuple collection of all the agent states in the network. The estimated expected fraction of infected agents at equilibrium, $\widehat{E}[Y]$, is

$$\widehat{E}[Y] = \frac{1}{m} \sum_{k=1}^m \frac{1^T \mathbf{x}^k}{N}, \quad (7.2)$$

where $1^T \mathbf{x}^k$ is the total number of infected agents in the k th sample.

From Theorem 7.2.1, the estimated fraction infected agents is also the average of the estimated marginal probability of infection

$$\widehat{E}[Y] = \frac{1}{N} \sum_{i=1}^N \widehat{P}(x_i = 1).$$

The samples can be generated efficiently when $0 < \frac{\lambda}{\mu} \leq 1$ and $\gamma > 1$ using order-1 Perturb-and-MAP, see section 6.3 in chapter 6. Since $\widehat{E}[Y]$ can be computed with $\widehat{P}(x_i = 1)$, the accuracy of $\widehat{E}[Y]$ is related to how well $\widehat{P}(x_i = 1)$ approximates $P(x_i = 1)$.

7.3 Dependence of $E[Y]$ on Parameters and Topology

The mean fraction of infected agents, $E[Y]$, depends on both the dynamics parameters, λ, γ, μ , and on the underlying network topology.

Corollary 7.3.1. *If $\gamma = 1$, then*

$$\begin{aligned} E[Y] &= P(x_i = 1), \forall i = 1, \dots, N \\ &= \frac{\sum_{k=1}^N e^{k \log(\frac{\lambda}{\mu})} \binom{N-1}{k-1}}{\sum_{k=0}^N e^{k \log(\frac{\lambda}{\mu})} \binom{N}{k}}. \end{aligned} \quad (7.3)$$

Proof. Equation 6.4 shows that when $\gamma = 1$,

$$\begin{aligned} P(x_1 = 1) &= P(x_2 = 1) = \dots = P(x_N = 1) \\ &= \frac{\sum_{k=1}^N e^{k \log(\frac{\lambda}{\mu})} \binom{N-1}{k-1}}{\sum_{k=0}^N e^{k \log(\frac{\lambda}{\mu})} \binom{N}{k}}. \end{aligned} \quad (7.4)$$

According to Theorem 7.2.1, $E[Y]$ is the average of the marginal probability of infection. This means that when $\gamma = 1$,

$$E[Y] = P(x_i = 1), \forall i = 1, \dots, N.$$

□

The network topology can also affect the mean fraction of infected agents. Recall the definition of vertex transitive graphs (see Definition 6.4.1).

Corollary 7.3.2. *If $G(V, E)$ is a vertex transitive graph [43], then*

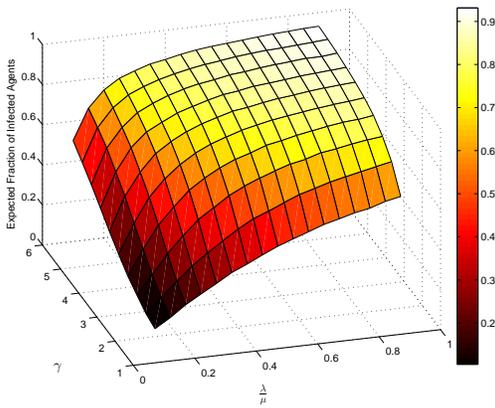
$$E[Y] = P(x_i = 1), \forall i = 1, \dots, N.$$

Proof. Theorem 6.4.2 states that if $G(V, E)$ is a vertex transitive graph, then $P(x_1 = 1) = P(x_2 = 1) = \dots = P(x_N = 1)$ for any dynamics parameters, λ, γ, μ . Consequently, this means that $E[Y]$ is equal to the average of the marginal probability of infection, which is the same as $P(x_i = 1)$ for any node in the network. □

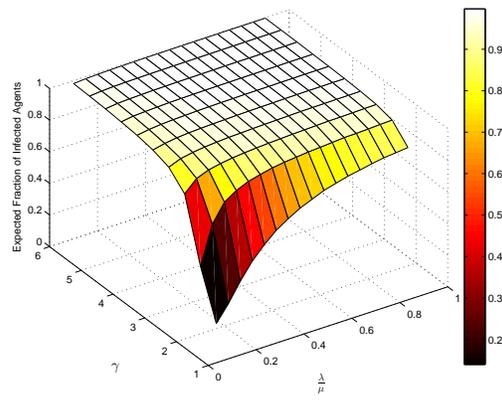
In the special scenario of $\gamma = 1$, or the underlying network being a vertex transitive graph, the susceptibility of individual agents at equilibrium under the scaled SIS process (i.e., the marginal probability of infection) can be inferred from the susceptibility of the entire population (i.e., the expected fraction of infected agents). Under other scenarios, the macroscopic characterization of the entire network is *not* indicative of the microscopic characterization.

Under more realistic scenarios, the expected fraction of infected agents can be approximated effectively in regime II): **Endogenous Infection Dominant** $0 < \frac{\lambda}{\mu} \leq 1, \gamma > 1$. We expect $E[Y] \approx 0$ for processes with low infection rates (i.e., $\gamma \approx 1$) and high healing rate (i.e., $0 < \frac{\lambda}{\mu} \ll 1$), and expect $E[Y] \approx 1$ for processes with high infection rates (i.e., $\gamma \gg 1$) and low healing rate (i.e., $\frac{\lambda}{\mu} \approx 1$).

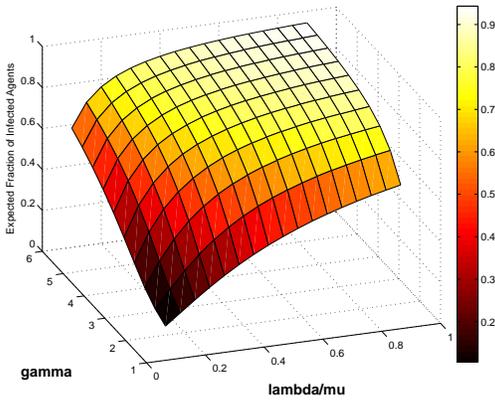
Figure 7.1 plots how $E[Y]$ depends on $\frac{\lambda}{\mu}$ and γ for four networks: Network A (see Figure 6.1(a)); Network C (see Figure 6.1(c)); the US Western power grid (see Figure 2.5); and a sample of the Facebook network (see Figure 2.7). The expected fraction of infected agents, $E[Y]$, is plotted on the z-axis. The parameter γ goes from 1.5 to 6 with step size = 0.5 along the Y-axis and $\frac{\lambda}{\mu}$ goes from 0.1 to 0.9 with step size = 0.05 along the X-axis. As expected, $E[Y]$ is close to 0 for small γ and small $\frac{\lambda}{\mu}$. The expected fraction of infected agents, $E[Y]$, increases towards 1 as γ and $\frac{\lambda}{\mu}$ increase. We can see from Figure 7.1 that the change is highly *nonlinear*. The rate of increase depends on the network topology; the increase is much steeper for Network C.



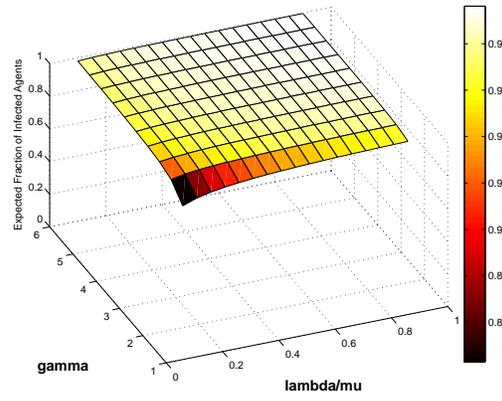
(a) Network A ($N = 16, |E| = 19$)



(b) Network C ($N = 20, |E| = 60$)



(c) US Western Power Grid ($N = 4941, |E| = 6595$)



(d) Sample of Facebook Network ($N = 4039, |E| = 88234$)

Figure 7.1: $E[Y]$ for Different $\frac{\lambda}{\mu}$ and γ for Four Networks

Figures 7.2 and 7.3 show the projection of the plots in Figure 7.1 onto the $\frac{\lambda}{\mu}$ axis for four different networks: network A, network C, US Western power grid, and sample of the Facebook network.

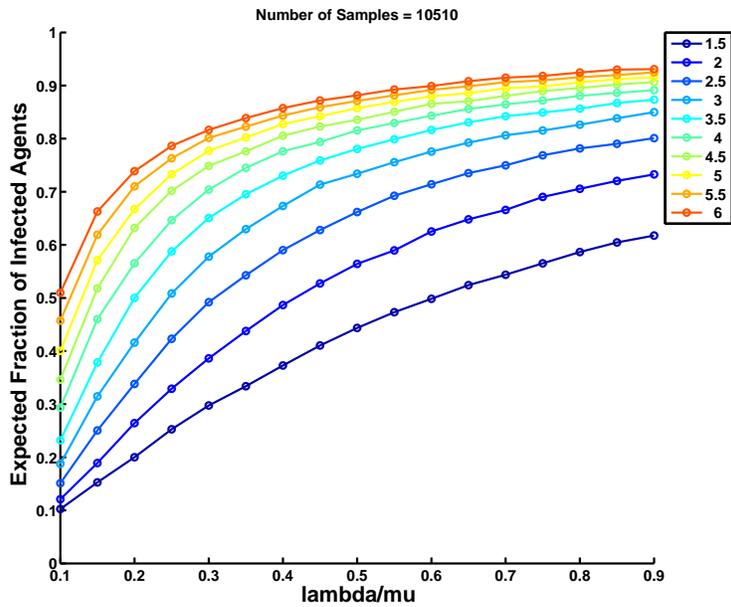
Different curves correspond to scaled SIS processes with varying γ value from 1.5 to 6 with step size = 0.5. As γ increases, the dependence of $E[Y]$ on the exogenous infection rate and the healing rate, $\frac{\lambda}{\mu}$, becomes more nonlinear until the saturation point where $E[Y] \approx 1$ for all the possible $\frac{\lambda}{\mu}$ from 0.1 to 0.9. The parameter $\frac{\lambda}{\mu}$ models the individual preferences of the agents; consequently this means that as γ increases, the dependence of the macroscopic behavior of the scaled SIS process on individual preferences becomes more and more nonlinear. This is intuitive because γ controls the

topology-dependent process. As γ increases, the topology-dependent process begins to dominate the topology-independent process, controlled by $\frac{\lambda}{\mu}$. Therefore, the dynamics of individual agents is less indicative of the network as a whole.

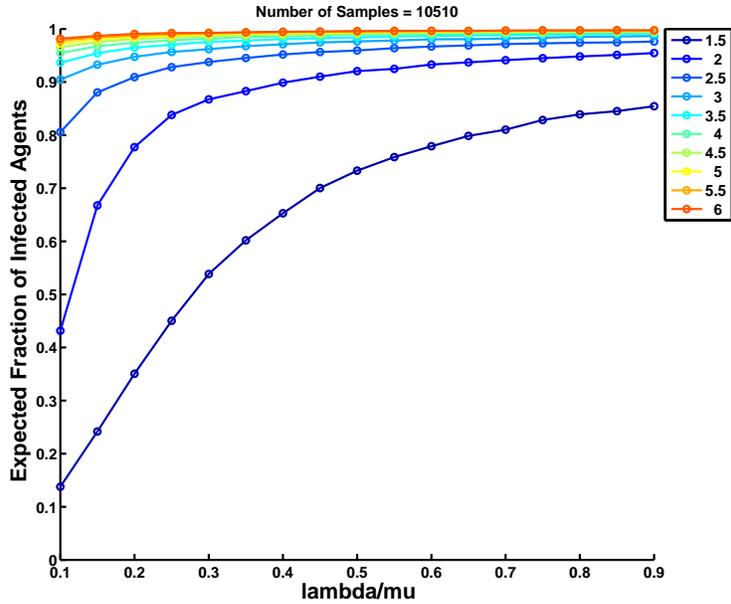
Topology also influences the expected fraction of infected agents. Network A and the US Western power grid are both sparser networks than Network C and the Facebook network. For the same dynamics parameters, the value of $E[Y]$ in Network C (see Figure 7.2(b)) and the Facebook network (see Figure 7.3(b)) is larger than the expected fraction of infected agents in Network A (see Figure 7.2(a)) and in the US Western power grid (see Figure 7.3(a)).

Figure 7.4 shows the projection of the surface plot in Figure 7.1 onto the γ axis (x-axis) for the four different networks: network A, network C, US Western power grid, and sample of the Facebook network. Different curves correspond to scaled SIS processes with varying $\frac{\lambda}{\mu}$ from 0.1 to 0.9 with step size = 0.05. For different $\frac{\lambda}{\mu}$, the dependence of $E[Y]$ on the γ parameter looks like different parts of a S-curve (i.e., sigmoid function); when $\frac{\lambda}{\mu}$ is small, dependences correspond more to the convex side of the S-curve (i.e., to the left of the inflection point) whereas when $\frac{\lambda}{\mu}$ is large, dependences correspond more to the concave side of the S-curve (i.e., to the right of the inflection point). This shows the interdependence between the network topology, the topology-independent process, controlled by $\frac{\lambda}{\mu}$, and the topology-dependent process, controlled by γ .

This behavior is intuitive in that when $\frac{\lambda}{\mu}$ is very small, the topology-independent process dominates the topology-dependent process, and individual agents prefer the healthy state; consequently, $E[Y]$ is small and does not change much with γ . The point where the topology-independent process dominates the topology-dependent process is an inflection point. As $\frac{\lambda}{\mu}$ increases towards 1, the value of $E[Y]$ begins to saturate at 1 regardless of γ . We can see from Figure 7.5(b) and Figure 7.5(a) that, on average, for half of the population to be healthy and for the same γ parameter, a densely connected network requires a smaller $\frac{\lambda}{\mu}$ than a sparsely connected network.

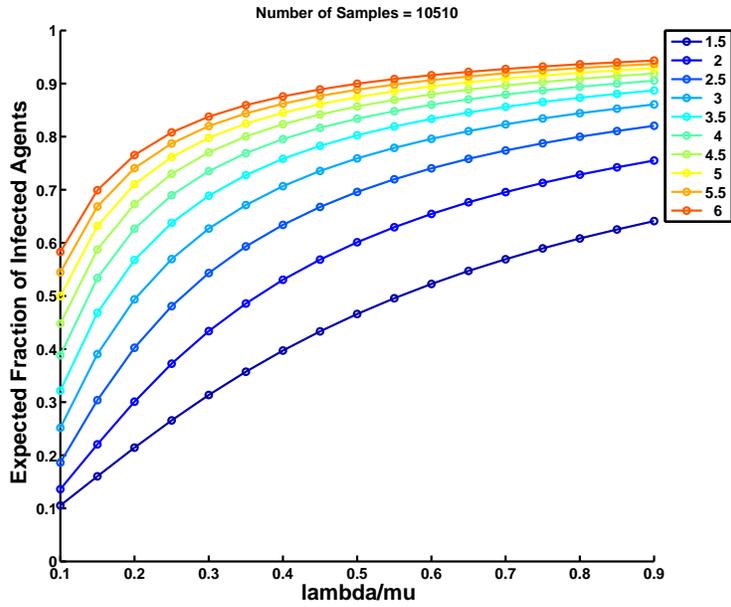


(a) Network A ($N = 16, |E| = 19, d(G) = 1.1875$)

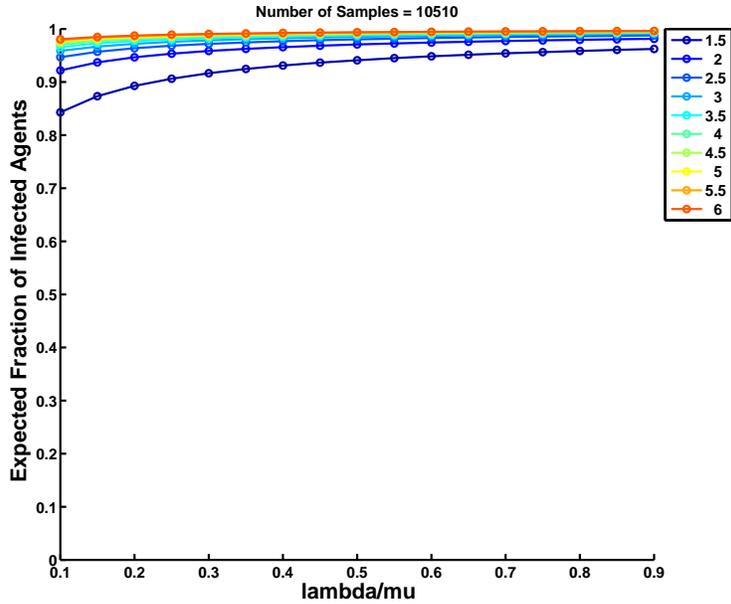


(b) Network C ($N = 20, |E| = 60, d(G) = 3$)

Figure 7.2: Projection of $E[Y]$ Onto $\frac{\lambda}{\mu}$ -Axis

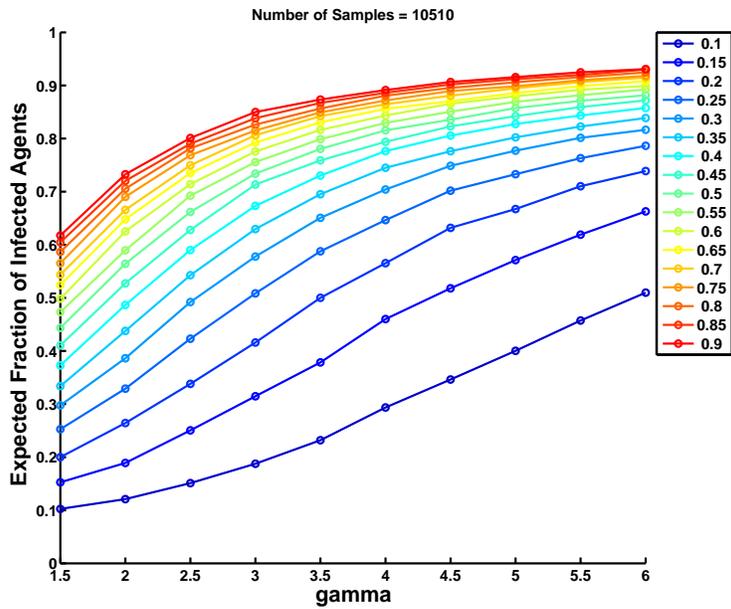


(a) US Western Power Grid ($N = 4941, |E| = 6595, d(G) = 1.335$)

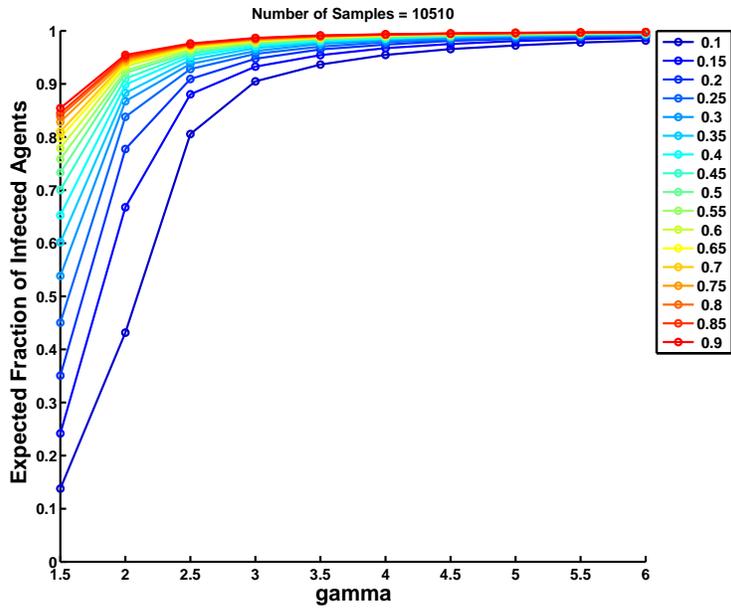


(b) Sample of Facebook Network ($N = 4039, |E| = 88234, d(G) = 21.846$)

Figure 7.3: Projection of $E[Y]$ Onto $\frac{\lambda}{\mu}$ -Axis

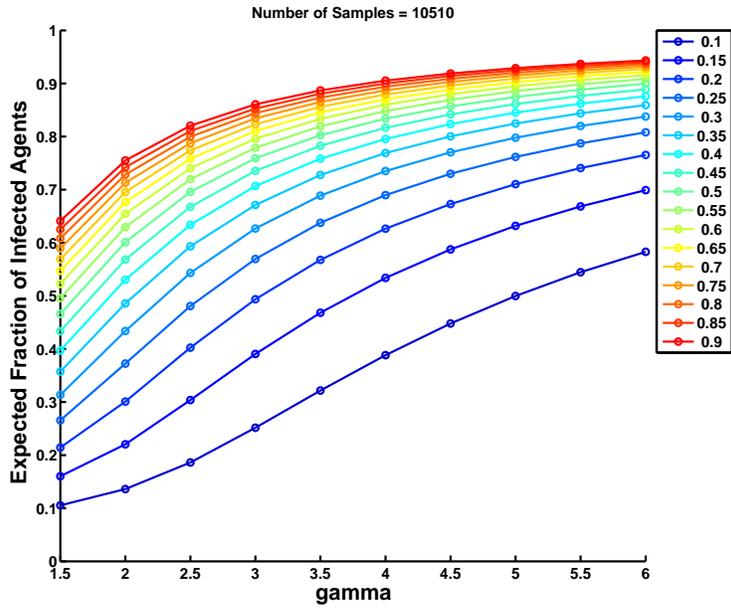


(a) Network A ($N = 16, |E| = 19, d(G) = 1.1875$)

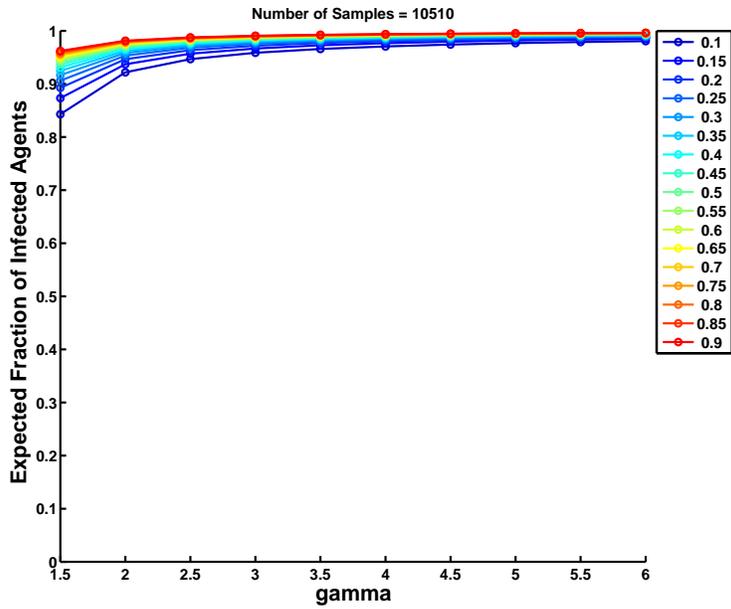


(b) Network C ($N = 20, |E| = 60, d(G) = 3$)

Figure 7.4: Projection of $E[Y]$ Onto γ -Axis



(a) US Western Power Grid ($N = 4941, |E| = 6595, d(G) = 1.335$)



(b) Sample of Facebook Network ($N = 4039, |E| = 88234, d(G) = 21.846$)

Figure 7.5: Projection of $E[Y]$ Onto γ -Axis

7.4 Conclusion

We proved that the expected fraction of infected agents, $E[Y]$, is equal to the average of the marginal probability of infections, $P(x_i = 1)$, of all the agents in the network. In the special scenario where $\gamma = 1$, or the underlying network is a vertex transitive graph, the susceptibility of individual agents at equilibrium under the scaled SIS process (i.e., the marginal probability of infection) can be inferred from the susceptibility of the entire population (i.e., the expected fraction of infected agents). For scenarios where $\gamma \neq 1$ and $G(V, E)$ is not a vertex transitive graph, the macroscopic characterization of the entire network is *not* indicative of the microscopic characterization. The approximation $\widehat{E[Y]}$ can be found by solving for the average of the estimated marginal probability of infections $\widehat{P}(x_i = 1)$, which we showed can be computed efficiently for large networks using order-1 Perturb-and-MAP sampling when the parameters are in regime II).

We showed with examples that, in this regime, the expected fraction of infected agents, $E[Y]$, demonstrates complex dependence on the dynamics parameters and topology. For the same dynamics parameters, networks that are more densely connected have larger $E[Y]$. We also examined $E[Y]$ as a function of $\frac{\lambda}{\mu}$ for different γ . As γ increases, the dependence becomes more nonlinear. Then we examined $E[Y]$ as a function of γ for different $\frac{\lambda}{\mu}$. The dependence is sigmoidal in shape since when $\frac{\lambda}{\mu}$ is at the extremes, the topology-independent process dominates the scaled SIS process, and $E[Y]$ changes little for different γ values. The point where the topology-independent process (i.e., preference for the healthy state) dominates the topology-dependent process (i.e., preference for the infected state) depends on the density of the underlying network; a densely connected network requires a much smaller $\frac{\lambda}{\mu}$ than a sparsely connected network.

Scaled SIS Process and the Contact Process

(ϵ -SIS process)

8.1 Introduction

This chapter shows how the scaled SIS process can give insight on the equilibrium behavior of another network-based epidemics model, the **contact process**. The contact process [5, 6] and its extension, the ϵ -SIS (susceptible-infected-susceptible) model [7], which we refer to as the extended contact process, are models widely considered for describing the propagation dynamics of failures or epidemics in complex networks. The contact process differs from the scaled SIS process in that it assumes that the infection rate of a susceptible agent is *linearly* dependent on its number of infected neighbors; the extended contact process further assumes a non-zero exogenous infection rate, λ .

The equilibrium distribution of the extended contact process can only be found numerically, which means solving for the left eigenvector corresponding to the zero eigenvalue of the transition rate matrix; this is an infeasible computation problem for networks with more than a few agents. For a subclass of extended contact processes, the equilibrium distribution can be approximated by that of the scaled SIS process. Consequently, analysis of the scaled SIS process at the microscopic, mesoscopic, and macroscopic scale are applicable to the extended contact process.

In section 8.2, we introduce the contact process and the extended contact process. We will point out the similarities and differences between the scaled SIS process and the extended contact process. We then prove and demonstrate with numerical experiments with small networks, the

conditions under which the contact process has the same equilibrium behavior as the scaled SIS process in sections 8.3 and 8.4. Section 8.5 considers the solution space of the Most-Probable Configuration Problem for the extended contact process.

8.2 Contact Process

The contact process, like the scaled SIS process, is a binary state, irreducible, continuous-time Markov process on a static, simple, connected, undirected network G [5, 6]. Each node in the network is an agent in the population. Each node can be in one of two states, $\{0, 1\}$, representing, for example, a healthy or an infected state. For a system with N nodes, the microscopic network configuration is

$$\mathbf{x} = [x_1, x_2, \dots, x_N]^T, \text{ where } x_i = \{0, 1\}.$$

As a result, there are 2^N possible configurations.

The contact process also models SIS (susceptible-infected-susceptible) epidemics on networks. There are two types of state transitions representing 1) healing of infected agents and 2) infection of susceptible agents.

1. Consider the configuration $\mathbf{x} = [x_1, x_2, \dots, x_j = 1, x_k, \dots, x_N]^T$. Let \mathbf{x}' denote the configuration where the j th agent heals:

$$\mathbf{x}' = [x_1, x_2, \dots, x_j = 0, x_k, \dots, x_N]^T.$$

The contact process transitions from \mathbf{x} to \mathbf{x}' in an exponentially distributed random amount of time with transition rate

$$q(\mathbf{x}, \mathbf{x}') = \mu. \tag{8.1}$$

Parameter μ is the **healing rate**. The healing rate of the contact process is the same as the healing rate of the scaled SIS process (see chapter 3). Without loss of generality, typically $\mu \equiv 1$.

2. Consider the configuration $\mathbf{x} = [x_1, x_2, \dots, x_j, x_k = 0, \dots, x_N]^T$. Let \mathbf{x}' be the configuration where the k th agent becomes infected:

$$\mathbf{x}' = [x_1, x_2, \dots, x_j, x_k = 1, \dots, x_N]^T.$$

The contact process transitions from \mathbf{x} to \mathbf{x}' in an exponentially distributed random amount of time with transition rate

$$q(\mathbf{x}, \mathbf{x}') = \gamma_e m_k, \quad (8.2)$$

where $m_k = \sum_{j=1}^N \mathbb{1}(x_j = 1) A_{jk}$ is the total number of infected neighbors of node k . The symbol $\mathbb{1}(\cdot)$ is the indicator function, and $\mathbf{A} = [A_{jk}]$ is the adjacency matrix of G . The parameter $\gamma_e > 0$ is the **endogenous infection rate**. The contact process assumes that the infection rate of the k th agent is *linearly* dependent on its number of infected neighbors m_k , whereas the scaled SIS process assumes that the infection rate is exponentially dependent on m_k .

In the contact process, when all the agents in the network are healthy, the process dies out. The configuration where all the agents are healthy ($\mathbf{x}^0 = [0, 0, \dots, 0]^T$) is an absorbing state of the Markov process. For networks with N agents, the contact process will eventually reach the configuration \mathbf{x}^0 and remain there indefinitely. Thus, the equilibrium distribution is trivial for contact processes on finite-size networks [6]. Because of this, we consider the extended contact process.

8.2.1 Extended Contact Process

In the contact process, a healthy agent can only become infected through contagion from an infected neighbor. It may be the case that a healthy agent (or working component) may also become infected (or fail) due to an exogenous (i.e., outside of the network) source—the agent is infected spontaneously [7, 25, 31]. For SIS epidemics, this is captured by a non-zero **exogenous infection rate**, $\lambda > 0$. The transition rate of the extended contact process from a configuration to another configuration where the k th agent becomes infected is

$$q(\mathbf{x}, \mathbf{x}') = \lambda + \gamma_e m_k, \quad (8.3)$$

where $m_k = \sum_{j=1}^N \mathbb{1}(x_j = 1) A_{jk}$ is the total number of infected neighbors of node k . The healing rate is the same as (8.1). We call this modified model the **extended contact process**, whereas [7] referred to it as the ϵ -SIS model. When agent k has 0 infected neighbors, the rate at which agent k becomes infected is the exogenous infection rate. For a system where spontaneous infection is rare,

the exogenous infection rate can be made arbitrarily small, but for the extended contact process, it has to remain greater than zero.

The configuration where all the agents are healthy ($\mathbf{x}^0 = [0, 0, \dots, 0]^T$) is no longer an absorbing state in the Markov process, since susceptible agents can spontaneously become infected. As a result, the equilibrium distribution of the Markov process is no longer trivial. There is currently no known tractable analytical results regarding this equilibrium distribution for the extended contact process for *arbitrary* network topologies; reference [7] provided the exact equilibrium distribution only for the complete graph.

As stated in chapter 2.3, the equilibrium distribution for a continuous-time Markov process can be found numerically. However, this approach is infeasible for large networks. In the case of an irreducible, continuous-time Markov process, the equilibrium distribution, π , is the left eigenvector of the transition rate matrix, \mathbf{Q}_e , corresponding to the 0 eigenvalue. However, the transition rate matrix is a $2^N \times 2^N$ matrix, where N is the size of the network. Solving for the equilibrium distribution of the extended contact process over a 200-node network with arbitrary topology means finding the eigenvector of a $2^{200} \times 2^{200}$ matrix; even taking into account sparsity, such computation is intractable. Instead, we can obtain a closed-form approximation of the equilibrium distribution over arbitrary network topologies for a subset of extended contact processes using the scaled SIS process.

8.2.2 Scaled SIS Process vs. Extended Contact Process

The infection rate of a susceptible agent in both the extended contact process and the scaled SIS process depends on its number of infected neighbors. The two models make different assumptions regarding the underlying mechanism of the contagion process:

Extended Contact Process

The extended contact process is parameterized by the exogenous infection rate, λ , the healing rate, μ , and the endogenous infection rate γ_e . Consider the scenario in Figure 8.1. Let T_3 be the random amount of time it takes for agent V_3 to become infected. Each infected neighbors of agent V_3 (i.e., V_1, V_2, V_4) and the exogenous (i.e., external) source may infect V_3 in an exponentially distributed amount of time $T_3^i \sim \exp(\gamma_e), i = 1, 2, 4$, and $T_3^e \sim \exp(\lambda)$, respectively. Let $T_3 = \min\{T_3^1, T_3^2, T_3^4, T_3^e\}$. Assuming that these sources act independently,

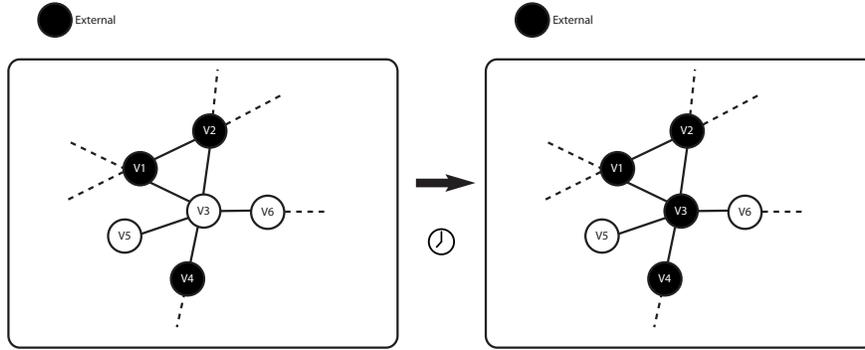


Figure 8.1: Agent V_3 Becomes Infected

then $T_3 \sim \exp(\lambda + 3\gamma_e)$. As the number of infected neighbors of V_3 increases, its infection rate also increases. The extended contact process models a *distributed* contagion scenario where all the infection sources compete to be the first to infect a healthy agent.

Scaled SIS Process

The scaled SIS process is parameterized by the exogenous infection rate, λ , healing rate, μ , and the endogenous infection rate γ_s . Consider the scenario in Figure 8.1. Let T_3 be the random amount of time it takes for agent V_3 to become infected. As agent V_3 has three infected neighbors (i.e., V_1, V_2, V_4), the scaled SIS process assumes that $T_3 = \frac{1}{(\gamma_s)^3} T \sim \exp(\lambda(\gamma_s)^3)$, where $T \sim \exp(\lambda)$ is the random amount of time a healthy agent becomes infected when it has no infected neighbors. When $\gamma_s > 1$, as the number of infected neighbors of V_3 increases, its infection rate also increases.

8.3 Equilibrium Behavior of the Extended Contact Process

For finite-size networks, unlike the contact process, the equilibrium distribution (also the limiting distribution of the process as $t \rightarrow \infty$) of the extended contact process is nontrivial. In this section, we show that, for a subclass of extended contact processes over *arbitrary* network topology, this equilibrium distribution is well approximated by the equilibrium distribution of the scaled SIS process; for these processes, the time-asymptotic behaviors of both processes are similar.

Lemma 8.3.1. [Proof in Appendix 8.8] For any nonnegative integer m from 0 to d_{\max} , if

$$\Delta^2 \ll \frac{2}{d_{\max}(d_{\max} - 1)},$$

then

$$\frac{\lambda}{\mu}(1 + \Delta)^m \approx \frac{\lambda}{\mu} + \frac{\lambda}{\mu}\Delta m.$$

Using Lemma 8.3.1, we can prove the following theorem.

Theorem 8.3.1. [Proof in Appendix 8.8] Consider the extended contact process with exogenous infection rate λ , healing rate μ , and endogenous infection rate γ_e , over a static, simple, connected, undirected network of arbitrary topology, G , with maximum degree d_{\max} . Let $\gamma_e = \frac{\lambda}{\mu}\Delta$. If

$$\Delta^2 \ll \frac{2}{d_{\max}(d_{\max} - 1)},$$

then the equilibrium distribution of the extended contact process is well approximated by

$$\pi_{\text{approx}}(\mathbf{x}) = \frac{1}{Z} \left(\frac{\lambda}{\mu} \right)^{1^T \mathbf{x}} (1 + \Delta)^{\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2}}, \quad \mathbf{x} \in \mathcal{X}, \quad (8.4)$$

where \mathbf{A} is the adjacency matrix of the network G , and Z is the partition function. The approximate distribution, $\pi_{\text{approx}}(\mathbf{x})$, is approximated by the equilibrium distribution (8.4) of a scaled SIS process over the same network G with exogenous infection rate λ , healing rate μ , and endogenous infection rate $\gamma_s = 1 + \Delta$.

Theorem 8.3.1 gives an upperbound on the factor, Δ , between the endogenous infection rate, γ_e , and the ratio, $\frac{\lambda}{\mu}$, of the exogenous infection rate, λ , and the healing rate, μ . This bound depends on the maximum degree of the underlying network topology. When γ_e is much smaller than $\frac{\lambda}{\mu}$, then the equilibrium distribution, $\pi_e(\mathbf{x})$, of the extended contact process is well approximated by that of an equivalent scaled SIS process. What does this imply about the extended contact process?

Recall that, for the extended contact process, all infection sources are *independent*. Suppose that susceptible agent i has one infected neighbor. Let $T_i^1 \sim \exp(\gamma_e)$ be the random amount of time it takes for susceptible agent i to be infected by this infected neighbor, and $T_i^e \sim \exp(\lambda)$ be the random amount of time it takes for susceptible agent i to become infected by an exogenous source. The probability that the agent i is infected by the exogenous source rather than by its infected neighbors is

$$P(T_i^e \leq T_i^1) = \frac{\lambda}{\gamma_e + \lambda} = \frac{1}{\frac{\Delta}{\mu} + 1}, \quad (8.5)$$

since $\gamma_e = \frac{\lambda}{\mu}\Delta$. (See Appendix 8.8 for a review regarding functions of exponentially distributed random variables.) Suppose that susceptible agent i has multiple (i.e., $m > 1$) infected neighbors. The probability that agent i is infected by the exogenous source rather than by its infected neighbors is

$$P(T_i^e \leq \min\{T_i^1, \dots, T_i^m\}) = \frac{\lambda}{m\gamma_e + \lambda} = \frac{1}{\frac{m\Delta}{\mu} + 1}. \quad (8.6)$$

Without loss of generality, let $\mu = 1$. According to Theorem 8.3.1, the scaled SIS process is a valid approximation for the extended contact process when Δ is small. In this case, according to (8.5) and (8.6), the probability that the source of infection is exogenous rather than endogenous is high; infection due to contagion from infected neighbor is rare but not impossible.

8.4 Experimental Simulations

We showed when the extended contact process can be well approximated by the scaled SIS process. We confirm Theorem 8.3.1 with numerical simulations. Further, we show that this upperbound is conservative. Below it, the equilibrium distribution of the extended contact process, $\pi_e(\mathbf{x})$, for arbitrary network topology is well approximated by the equilibrium distribution, $\pi_{\text{approx}}(\mathbf{x})$, of a scaled SIS process. However, depending on the underlying network topology, the approximation may still remain accurate (< 0.1 deviation) for extended contact processes with parameters *away from* the bound.

8.4.1 Setup

We will compare the true equilibrium distribution, $\pi_e(\mathbf{x})$, of the extended contact process, with infection and healing rates $\left(\lambda, \mu, \gamma_e = \frac{\lambda}{\mu}\Delta\right)$ over network G , with the approximation distribution, $\pi_{\text{approx}}(\mathbf{x})$. The true distribution, $\pi_e(\mathbf{x})$, is found numerically by forming the transition rate matrix, \mathbf{Q}_e , according to (8.1) and (8.3) and solving for the left eigenvector of \mathbf{Q}_e corresponding to eigenvalue 0. The approximate distribution, $\pi_{\text{approx}}(\mathbf{x})$, is obtained from the closed-form equation according to Theorem 8.3.1

$$\pi_{\text{approx}}(\mathbf{x}) = \frac{1}{Z} \left(\frac{\lambda}{\mu}\right)^{1^T \mathbf{x}} (1 + \Delta)^{\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2}}, \quad \mathbf{x} \in \mathcal{X}.$$

We solve for $\pi_e(\mathbf{x})$ and $\pi_{\text{approx}}(\mathbf{x})$ for different $\frac{\lambda}{\mu}$ values and different Δ values, both below and above the upperbound,

$$\Delta_u = \sqrt{\frac{2}{d_{\max}(d_{\max} - 1)}}.$$

To quantify the difference between the exact and the approximation equilibrium distribution, $\pi_e(\mathbf{x})$ and $\pi_{\text{approx}}(\mathbf{x})$, we use the total variation distance (TVD) [69]:

$$\text{TVD}(\pi_e, \pi_{\text{approx}}) = \frac{1}{2} \sum_{\mathbf{x} \in \mathcal{X}} |\pi_e(\mathbf{x}) - \pi_{\text{approx}}(\mathbf{x})|. \quad (8.7)$$

When the two distributions are equal, TVD is 0. The maximum TVD between any two probability distributions over the same support is 1.

As the true distribution of the extended contact process, $\pi_e(\mathbf{x})$, is obtained by solving the zero eigenvalue-eigenvector problem of \mathbf{Q}_e , which is a $2^N \times 2^N$ matrix, we are restricted to simulating examples with small networks of size N . We consider six 16-node networks (see Figure 8.2) with different maximum degree, d_{\max} , corresponding to different upperbounds Δ_u . Networks A and B have the smallest possible maximum degree of any connected graph ($d_{\max} = 2$); they have the largest possible upperbound ($\Delta_u = 1$). Network F has the largest maximum degree of the networks studied ($d_{\max} = 15$) and has the smallest upperbound ($\Delta_u = 0.098$).

In Matlab, on a Microsoft Azure cloud virtual machine with 2.6GHz Intel Xeon E5-2670 and 56GB of RAM, for a 16-node network, it takes approximately 2 secs to generate the sparse transition rate matrix \mathbf{Q}_e and 460 secs to solve for the eigenvector corresponding to the 0 eigenvalue. For a 20-node network, it takes approximately 30 secs to generate the transition rate matrix \mathbf{Q}_e ; we receive an OUT-OF-MEMORY error when computing the eigenvector.

8.4.2 Results: $\pi_e(\mathbf{x})$ and $\pi_{\text{approx}}(\mathbf{x})$

To provide intuition on the quality of the approximations for different TVDs, we plot in Figure 8.3 the true equilibrium distribution, $\pi_e(\mathbf{x})$, of the extended contact process together with the approximate equilibrium distribution, $\pi_{\text{approx}}(\mathbf{x})$. The Y-axis displays both equilibrium distributions; we use log scaling for better visualization. The 2^{16} network configurations are on the X-axis. The configurations are ordered such that high probability configurations in $\pi_e(\mathbf{x})$ are in the center.

Figure 8.3 shows $\pi_e(\mathbf{x})$ and $\pi_{\text{approx}}(\mathbf{x})$ and their corresponding TVD, see (8.7), for the six different network topologies (see Figure 8.2) for parameters $\frac{\lambda}{\mu} = 0.7$, $\Delta = 0.0023$. This value of

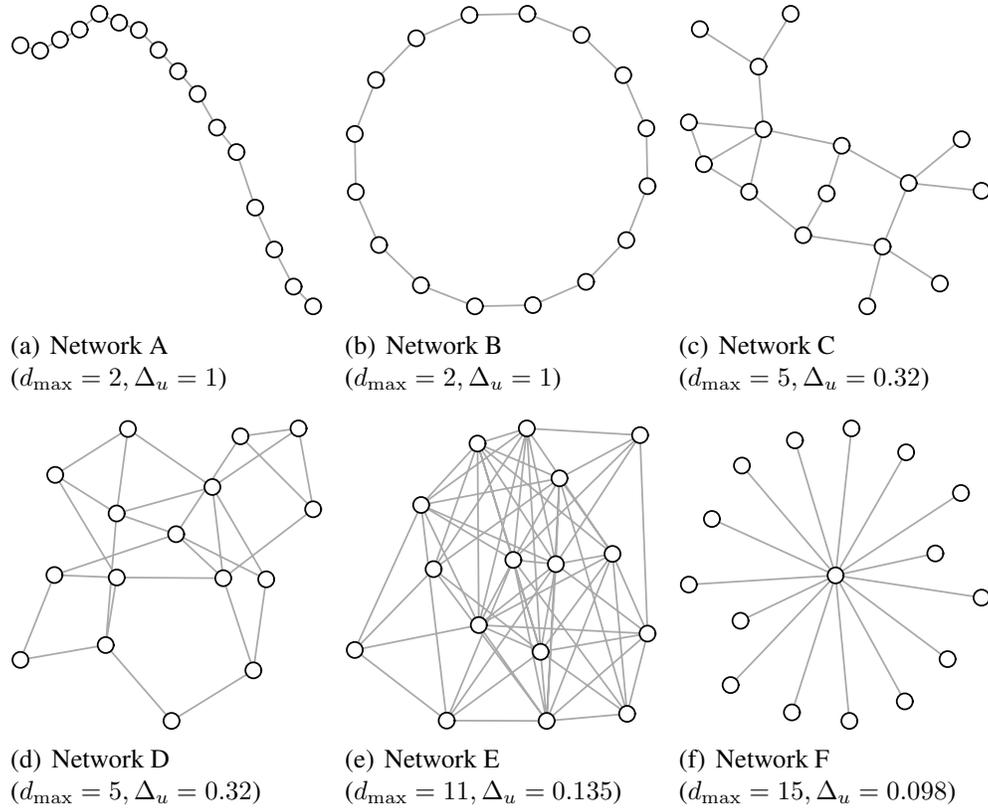
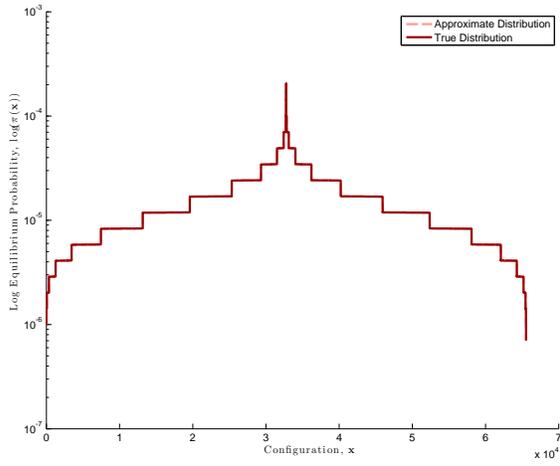


Figure 8.2: Different Network Topologies with Different Maximum Degree

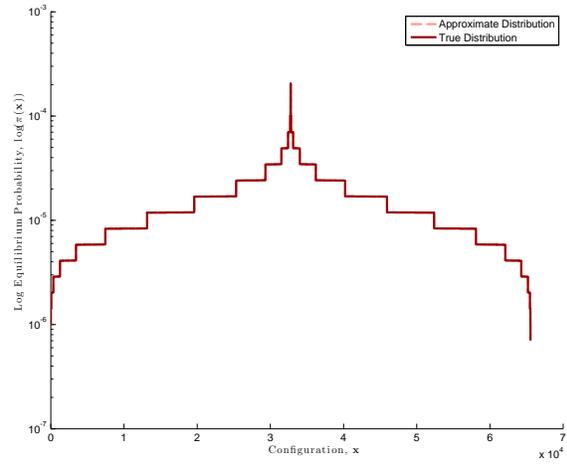
Δ is much smaller than the upperbound, Δ_u , for all the networks. The equilibrium distribution, $\pi_e(\mathbf{x})$, of the extended contact process is well approximated (i.e., the TVD is on the order of 10^{-5} or smaller) by equilibrium distribution, $\pi_{\text{approx}}(\mathbf{x})$, of the scaled SIS process. Note that this value of TVD is over 2^{16} configurations; so the actual divergence for any configuration is very small. The two distributions are almost identical for all the networks.

We also considered the case when the condition of Theorem 8.3.1 is not satisfied. Figure 8.4 shows $\pi_e(\mathbf{x})$ and $\pi_{\text{approx}}(\mathbf{x})$ and their corresponding TVD for parameters $\frac{\lambda}{\mu} = 0.7, \Delta = 1.0496$. In this case, the value of Δ is above the upperbound Δ_u for all the networks in Figure 8.2. As we expect, TVD is larger when compared to the TVD for processes with Δ well below Δ_u (see Figure 8.3). Again, for the same infection and healing rates, different networks have different TVD values.

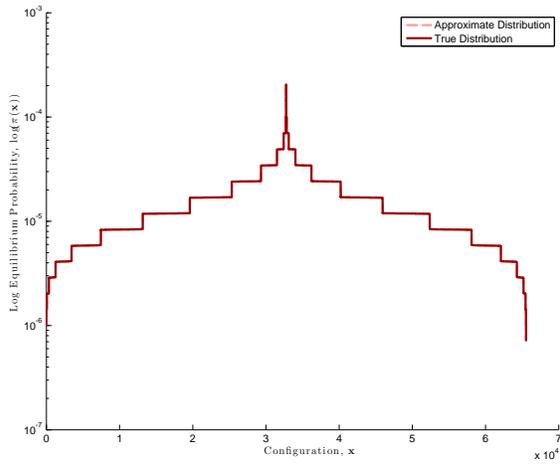
For Networks A and B, the deviation between the true and approximate equilibrium distribution, 0.1073 and 0.1186, respectively, is relatively small. We see from Figure 8.4(a) and Figure 8.4(b)



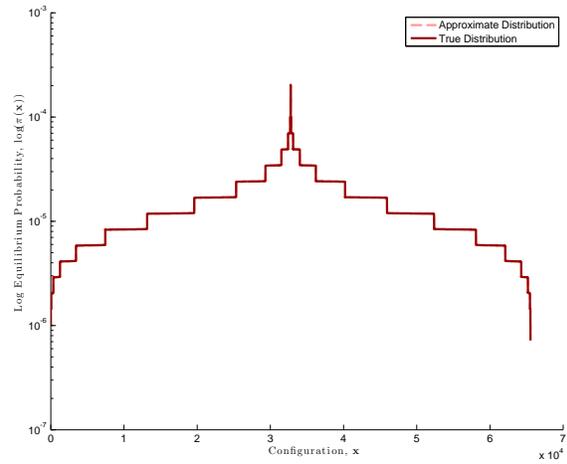
(a) Network A (TVD = 1.0384×10^{-6})



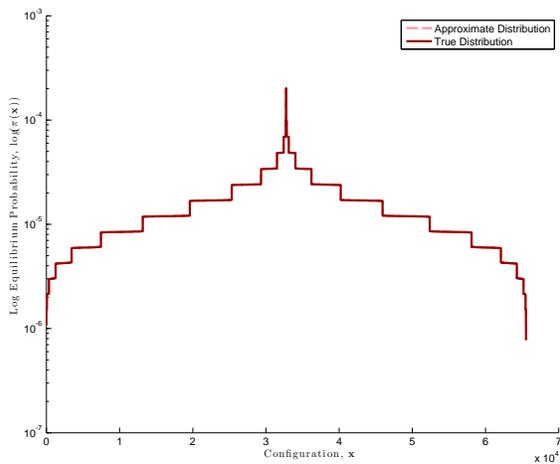
(b) Network B (TVD = 1.1236×10^{-6})



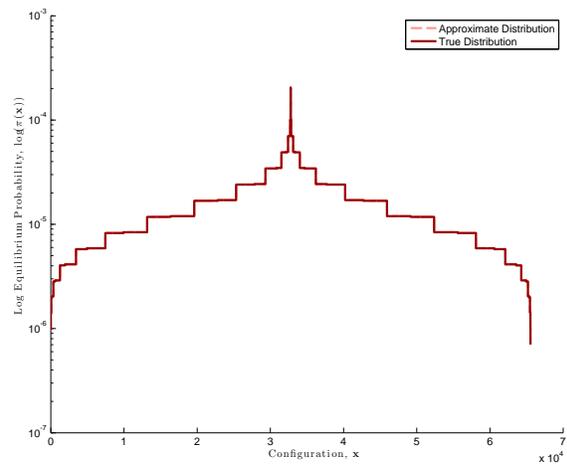
(c) Network C (TVD = 3.2392×10^{-6})



(d) Network D (TVD = 5.0208×10^{-6})



(e) Network E (TVD = 2.7487×10^{-5})



(f) Network F (TVD = 2.2729×10^{-5})

Figure 8.3: $\pi_e(\mathbf{x})$ and $\pi_{\text{approx}}(\mathbf{x})$ when $\frac{\lambda}{\mu} = 0.7$, $\Delta = 0.0023$

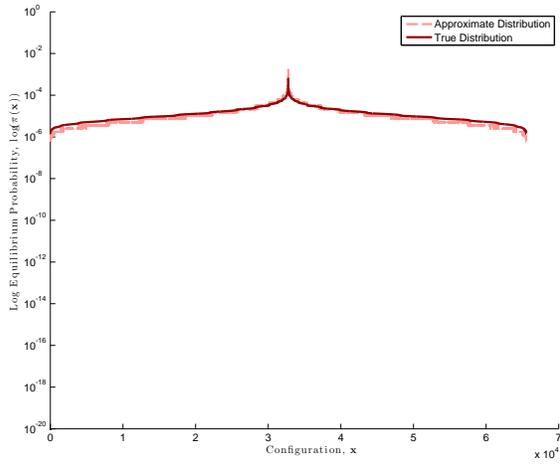
that many configurations have similar equilibrium probability for both distributions. For Networks D and E, TVDs are 0.4798 and 0.7898, respectively. These figures show that the approximate distribution tends to overestimate the probability of highly probable configurations but underestimates the low probability configurations. However, there is good correlation between the relative ordering of configurations in both distributions; configurations that are highly probable in $\pi_e(\mathbf{x})$ are also highly probable in $\pi_{\text{approx}}(\mathbf{x})$.

8.4.3 Results: TVD vs. Δ and $\frac{\lambda}{\mu}$

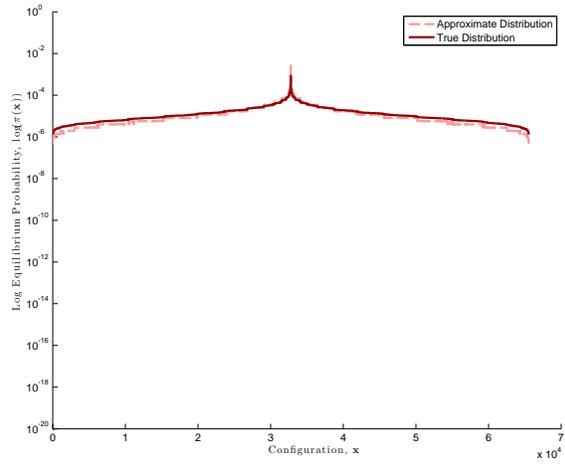
We consider here the approximation of $\pi_e(\mathbf{x})$ by $\pi_{\text{approx}}(\mathbf{x})$ as the infection and healing rates change. Like the scaled SIS process, we can interpret the extended contact process as consisting of a *topology-independent* process parameterized by $\frac{\lambda}{\mu}$ —it is topology independent because the exogenous infection rate λ and the healing rate μ are identical for all the agents in the network—and a *topology-dependent* process parameterized by the endogenous infection rate γ . When $\frac{\lambda}{\mu}$ is large, the topology-independent process exerts a larger effect on the equilibrium behavior of the network processes.

Figure 8.5 shows the TVD between the equilibrium distribution, $\pi_e(\mathbf{x})$, of the extended contact process and the approximate distribution, $\pi_{\text{approx}}(\mathbf{x})$, for different $\frac{\lambda}{\mu}$ and Δ values both below and well above the threshold Δ_u . Figure 8.5 considers the six different network topologies in Figure 8.2. We plot Δ along the X-axis and the TVD between $\pi_e(\mathbf{x})$ and $\pi_{\text{approx}}(\mathbf{x})$ along the Y-axis. Different curves in each figure correspond to equilibrium distributions with different $\frac{\lambda}{\mu}$ values.

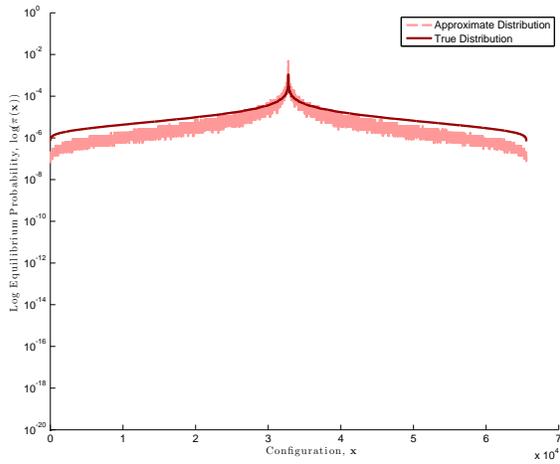
For the same Δ value, we observe that larger $\frac{\lambda}{\mu}$ correspond to smaller TVD. This holds for all the networks. Also, as we expect, for $\Delta \ll \Delta_u$, TVD is negligible for all the networks. The deviation between the true equilibrium distribution and the approximation increases as Δ moves toward Δ_u ; the rate of this increase differs for different topologies. Surprisingly, this increase is *not* monotonic for *all* network topologies. As Δ increases to values larger than Δ_u , TVD may actually decrease. We observe this decrease in TVD for both Network E in Figure 8.5(e) and Network F in Figure 8.5(f). In particular, Network F, which has the largest maximum degree of all the six networks, has relatively small deviation between $\pi_e(\mathbf{x})$ and $\pi_{\text{approx}}(\mathbf{x})$ compared to the other network topologies.



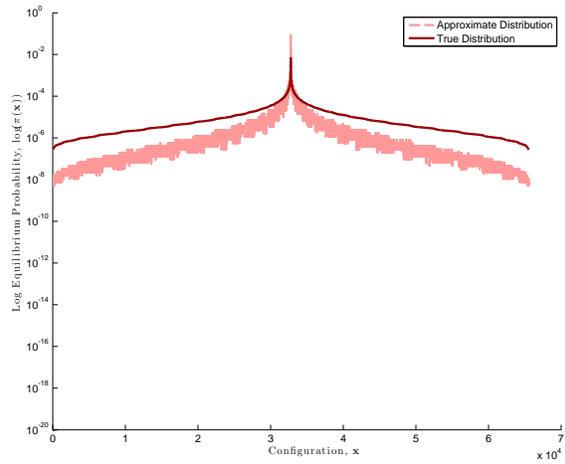
(a) Network A (TVD = 0.1073)



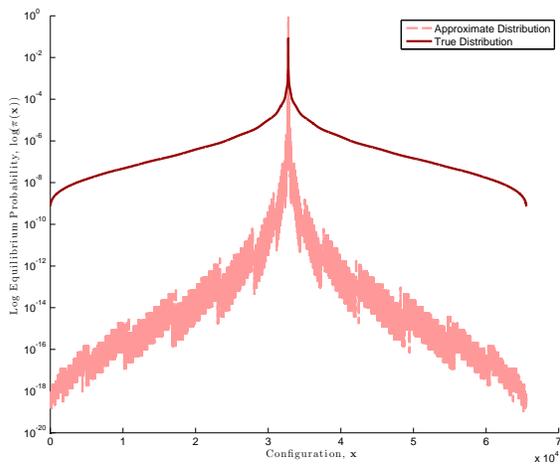
(b) Network B (TVD = 0.1186)



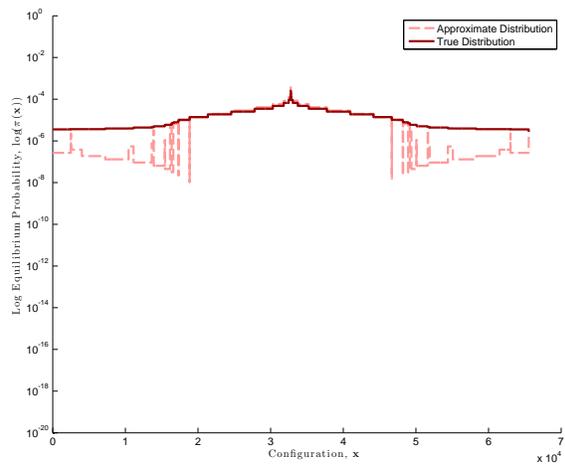
(c) Network C (TVD = 0.294)



(d) Network D (TVD = 0.4798)

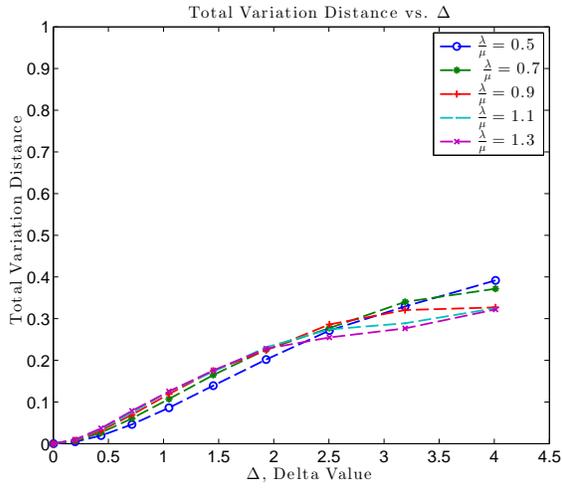


(e) Network E (TVD = 0.7898)

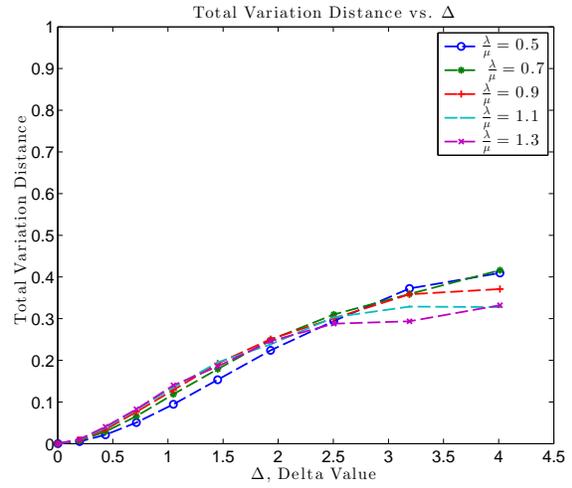


(f) Network F (TVD = 0.1330)

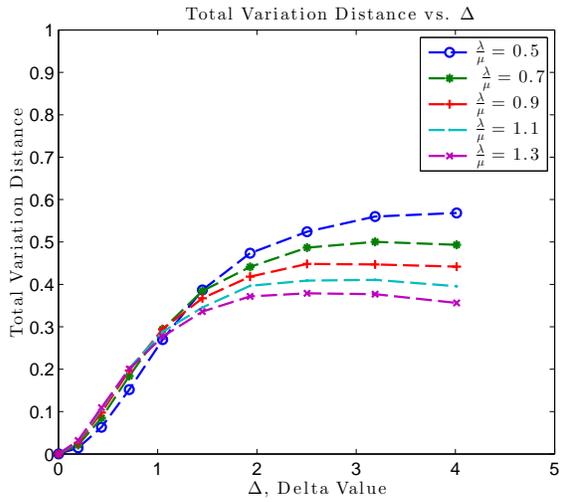
Figure 8.4: $\pi_e(\mathbf{x})$ and $\pi_{\text{approx}}(\mathbf{x})$ when $\frac{\lambda}{\mu} = 0.7$, $\Delta = 1.0496$



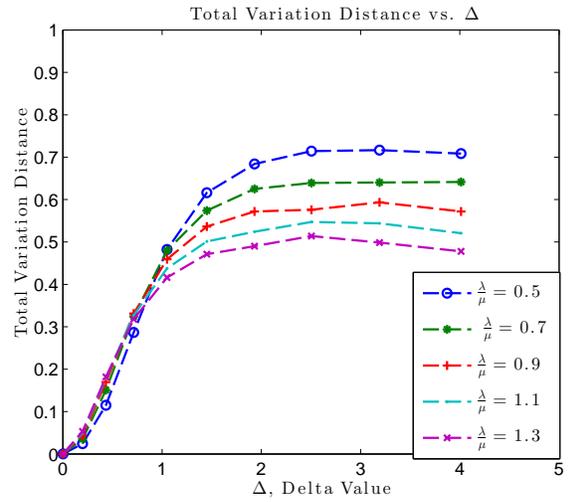
(a) Network A ($\Delta_u = 1$)



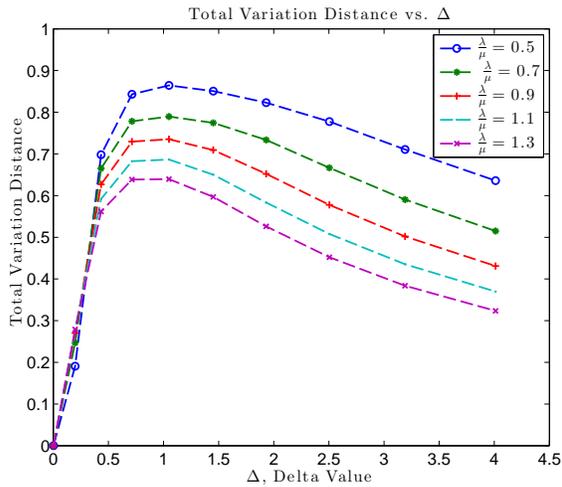
(b) Network B ($\Delta_u = 1$)



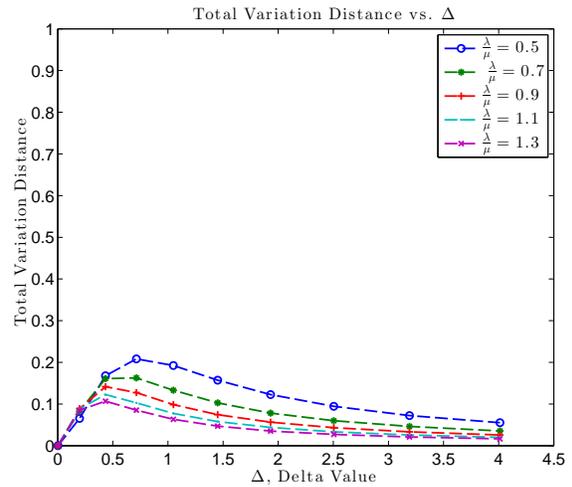
(c) Network C ($\Delta_u = 0.32$)



(d) Network D ($\Delta_u = 0.32$)



(e) Network E ($\Delta_u = 0.135$)



(f) Network F ($\Delta_u = 0.098$)

Figure 8.5: Dependence of $\text{TVD}(\pi_e, \pi_{\text{approx}})$ on Δ

8.5 Most-Probable Configuration

We showed in Figures 8.3 and 8.4 that, for a range of the dynamic parameters, the equilibrium distribution $\pi_e(\mathbf{x})$ of the extended contact process is well approximated by the equilibrium distribution $\pi_{\text{approx}}(\mathbf{x})$ of the scaled SIS process. In this section, we consider the problem of finding the most-probable configuration (i.e., configuration with the maximum equilibrium probability) of the extended contact process.

The most-probable configuration depends on the infection and healing rates and on the underlying network topology. It identifies the set of agents that are most likely to be infected in the long run. These are the more vulnerable agents in the network. If the most probable configuration is $\mathbf{x}^0 = [0, 0, \dots, 0]^T$, all agents are healthy, whereas if the most-probable configuration is $\mathbf{x}^N = [1, 1, \dots, 1]^T$, then all agents are at risk regardless of their location in the network.

The most-probable configuration of the extended contact process is

$$\mathbf{x}_e^* = \arg \max_{\mathbf{x} \in \mathcal{X}} \pi_e(\mathbf{x}),$$

where \mathcal{X} is the set of all 2^N possible network configurations. For the extended contact process, there is no closed-form description of the equilibrium distribution. Therefore, the most-probable configuration can only be computed by iterating through all the possible configurations; an approach that is intractable for large-scale networks.

As stated in Theorem 8.3.1, when $\Delta \ll \Delta_u$, the equilibrium distribution, $\pi_e(\mathbf{x})$, of the extended contact process is well approximated by the equilibrium distribution, $\pi_{\text{approx}}(\mathbf{x})$, of a scaled SIS process with endogenous infection rate $\gamma_s = 1 + \Delta$. We proved in [36] (see chapter 5) that, in this case, the most-probable configuration of the scaled SIS process can be solved in *polynomial-time* because it corresponds to solving for the minimum of a submodular function. It is therefore possible to identify vulnerable network substructures for networks with hundreds and thousands of agents.

From the simulation results in the previous section, we compare the most-probable configuration of the extended contact process with the most-probable configuration of the approximating scaled SIS process. Table 8.1 lists for the six networks in Figure 8.2, the TVD between the distributions, the corresponding most-probable configurations, and the probabilities of the most-probable

configuration for $\frac{\lambda}{\mu} = 0.9744$ and $\Delta = 0.02$. We observe that when the condition of Theorem 8.3.1 is satisfied:

1. the most-probable configuration, \mathbf{x}_e^* , of the extended contact process is the approximately the same as the most-probable configuration, $\mathbf{x}_{\text{approx}}^*$, of the scaled SIS process;
2. the probability of the most-probable configuration, $\pi_e(\mathbf{x}_e^*)$, of the extended contact process is the same as the probability of the most-probable configuration, $\pi_{\text{approx}}(\mathbf{x}_{\text{approx}}^*)$, of the scaled SIS process.

	$\text{TVD}(\pi_e, \pi_{\text{approx}})$	\mathbf{x}_e^*	$\mathbf{x}_{\text{approx}}^*$	$\pi_e(\mathbf{x}_e^*)$	$\pi_{\text{approx}}(\mathbf{x}_{\text{approx}}^*)$
Network A	1.0236×10^{-4}	$\mathbf{x}^0 = [0, 0, \dots, 0]^T$	$\mathbf{x}^0 = [0, 0, \dots, 0]^T$	1.7431×10^{-5}	1.7427×10^{-5}
Network B	1.1027×10^{-4}	$\mathbf{x}^0 = [0, 0, \dots, 0]^T$	$\mathbf{x}^0 = [0, 0, \dots, 0]^T$	1.7347×10^{-5}	1.7342×10^{-5}
Network C	3.3806×10^{-4}	see Figure 8.6(c)	see Figure 8.6(c)	1.7107×10^{-5}	1.7154×10^{-5}
Network D	5.2714×10^{-4}	$\mathbf{x}^N = [1, 1, \dots, 1]^T$	$\mathbf{x}^N = [1, 1, \dots, 1]^T$	1.8622×10^{-5}	1.8781×10^{-5}
Network E	0.0031	$\mathbf{x}^N = [1, 1, \dots, 1]^T$	$\mathbf{x}^N = [1, 1, \dots, 1]^T$	3.1073×10^{-5}	3.277×10^{-5}
Network F	0.0023	$\mathbf{x}^0 = [0, 0, \dots, 0]^T$	$\mathbf{x}^N = [0, 0, \dots, 0]^T$	1.7419×10^{-5}	1.7389×10^{-5}

Table 8.1: Most-Probable Configuration when $\frac{\lambda}{\mu} = 0.9744$ and $\Delta = 0.02$.

For Networks A, B, and F, the most-probable configuration for both the extended contact process and the approximate scaled SIS process is \mathbf{x}^0 , the configuration where all the agents are healthy. However, for the same infection and healing rate, the most-probable configuration for Networks D and E for both the extended contact process and the scaled SIS process is \mathbf{x}^N , the configuration where all the agents are infected. Figure 8.6(c) shows that the most-probable configuration for Network C is neither \mathbf{x}^0 nor \mathbf{x}^N , but a configuration where nine agents are infected while seven agents are healthy; we call most-probable configurations that are neither \mathbf{x}^0 nor \mathbf{x}^N *non-degenerate* most-probable configurations.

For an extended contact process with exogenous infection rate and healing rate, $\frac{\lambda}{\mu} = 0.9744$, and endogenous infection rate, $\gamma_c = \frac{\lambda}{\mu}\Delta = 0.019488$, the epidemic is minor in Networks A, B, and F, but should be of concern in Networks D and E. In Network C, subsets of agents are more at risk than others. Different networks have different risk levels because the propagation of contagious infection is dependent on the underlying network topology. The result in Figure 8.6(c) confirms for the extended contact process what we proved for the scaled SIS process in [36], namely, that in the most-probable configuration the infected agents belong to dense subgraphs in the network.

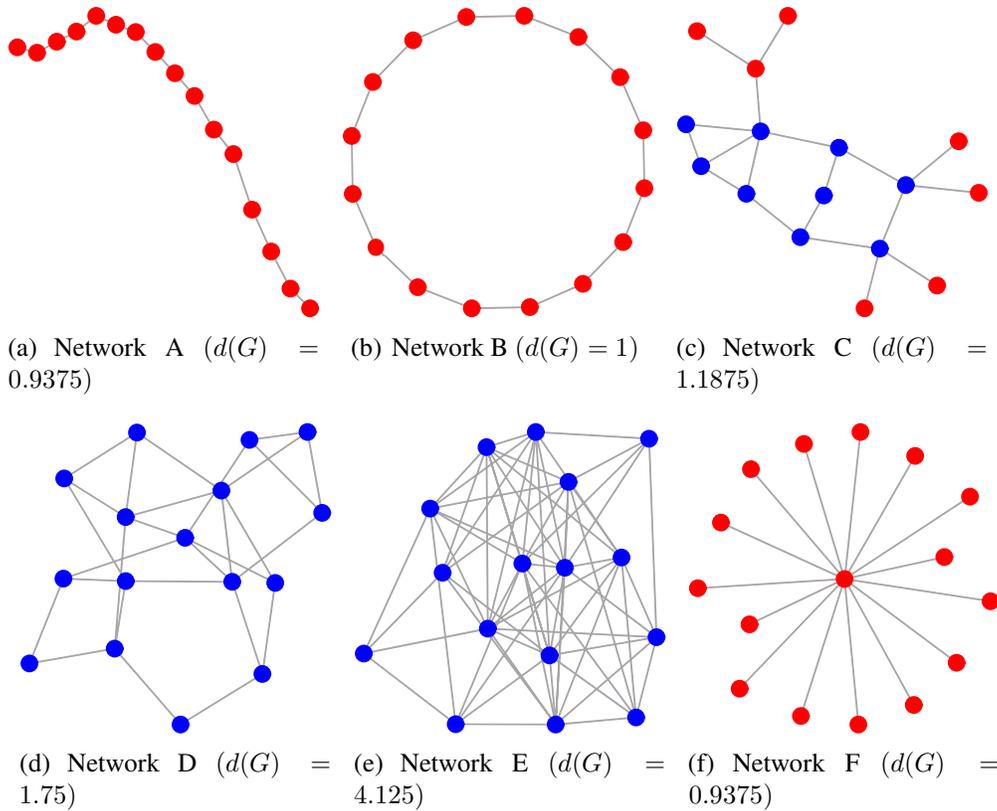


Figure 8.6: Most-Probable Configuration when $\frac{\lambda}{\mu} = 0.9744$ and $\Delta = 0.02$ (Blue = Infected, Red = Healthy)

Reference [44] defines density of a graph G by

$$d(G) = \frac{|E(G)|}{|V(G)|},$$

where $|E(G)|$ is the total number of edges and $|V(G)|$ is the total number of nodes. Networks that are more connected have higher densities than sparsely connected networks.

Networks with high density, such as Networks D and E are more at risk to contagion than networks with low density such as Networks A, B, and F. Network F, although it has the largest maximum degree, has the same density as Network A. It is difficult for infection to spread in Network F because the center agent is the only agent capable of transmitting the infection to its neighbors. We showed in [36] that the nine infected agents in Network C are more at risk of infection than the other agents because they form a subgraph that is denser than the overall network; these nine agents are especially well-connected in this network.

Table 8.2 lists the TVD between the distributions, the most-probable configurations for the

	$\text{TVD}(\pi_e, \pi_{\text{approx}})$	\mathbf{x}_e^*	$\mathbf{x}_{\text{approx}}^*$	$\pi_e(\mathbf{x}_e^*)$	$\pi_{\text{approx}}(\mathbf{x}_{\text{approx}}^*)$
Network A	0.0266	$\mathbf{x}^0 = [0, 0, \dots, 0]^T$	$\mathbf{x}^0 = [0, 0, \dots, 0]^T$	6.7989×10^{-5}	6.4085×10^{-5}
Network B	0.029	$\mathbf{x}^0 = [0, 0, \dots, 0]^T$	$\mathbf{x}^N = [1, 1, \dots, 1]^T$	6.2942×10^{-5}	6.1972×10^{-5}
Network C	0.0848	see Figure 8.6(c)	$\mathbf{x}^N = [1, 1, \dots, 1]^T$	7.0847×10^{-5}	1.214×10^{-4}
Network D	0.1505	$\mathbf{x}^N = [1, 1, \dots, 1]^T$	$\mathbf{x}^N = [1, 1, \dots, 1]^T$	2.5957×10^{-4}	0.0011
Network E	0.6652	$\mathbf{x}^N = [1, 1, \dots, 1]^T$	$\mathbf{x}^N = [1, 1, \dots, 1]^T$	0.0066	0.1849
Network F	0.1609	$\mathbf{x}^0 = [0, 0, \dots, 0]^T$	$\mathbf{x}^N = [0, 0, \dots, 0]^T$	5.988×10^{-5}	3.7915×10^{-5}

Table 8.2: Most-Probable Configuration when $\frac{\lambda}{\mu} = 0.7$ and $\Delta = 0.4333$.

extended contact process and the approximate scaled SIS process, and the probabilities of the most-probable configurations for $\frac{\lambda}{\mu} = 0.7$ and $\Delta = 0.4333$; with this value, the factor, Δ , no longer satisfies the condition of Theorem 8.3.1. As a result, the TVD between the distributions are larger and for Networks B and C, the most-probable configurations of the extended contact process and the approximate scaled SIS process are no longer the same. One intuition as to why for Network C, $\mathbf{x}_{\text{approx}}^* = \mathbf{x}^N$ (the configuration where all the agents are infected), while $\mathbf{x}_e^* = \mathbf{x}^0$ (the configuration where all the agents are healthy) is that the infection rate of the scaled SIS is exponentially dependent on the number of infected neighbors, while the infection rate of the extended contact process is linearly dependent on the number of infected neighbors; contagion is more virulent because the infection rate is higher in the scaled SIS process.

Note that the configuration in Figure 8.6(c), where nine agents are more at risk of infection than others, remains the most-probable configuration for Network C. Even though this configuration no longer has the highest equilibrium probability in the approximate distribution, it remains a highly probable configuration. This reinforces our observation from Figure 8.4 that configurations with high probabilities in the approximate scaled SIS distribution are also highly probable in the equilibrium distribution of the extended contact process. The substructures that are vulnerable for the scaled SIS process, the *non-degenerate* most-probable configurations, are also vulnerable substructures of the extended contact process.

8.6 Real-World Network

In chapter 5, we showed that the most-probable configuration of the scaled SIS process can be solved exactly in polynomial-time. In chapter 6 and chapter 7, we showed that the marginal probability of infection, $\widehat{P}(x_i = 1)$, and the expected fraction of infected agents, $\widehat{E}[Y]$, could also

be approximated effectively. Using Theorem 8.3.1, we can also find these values for the extended contact process for large real-world networks.

The US Western powergrid (see Figure 2.5) has 4941 nodes, making it impossible to solve for the equilibrium distribution of the extended contact process using the eigenvector-eigenvalue method. The maximum degree of the network is 19; therefore, the threshold condition is

$$\Delta \ll \sqrt{\frac{2}{19 * 18}} = 0.0765.$$

Theorem 8.3.1 states that the equilibrium distribution of the extended contact process with $\frac{\lambda}{\mu} = 0.9$ and endogenous infection rate $\gamma_e = 0.0387$ is approximated by the distribution

$$\pi_{\text{approx}}(\mathbf{x}) = \frac{1}{Z} (0.9)^{1^T \mathbf{x}} (1.043)^{\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2}}, \quad \mathbf{x} \in \mathcal{X}. \quad (8.8)$$

The expected fraction of infected agents at equilibrium for the extended contact process is

$$\widehat{E[Y]} = 0.488.$$

Figure 8.7(a) shows the most-probable configuration of (8.8), and Figure 8.7(b) shows the marginal probability of infection of the extended contact process. The marginal probability of infection ranges from 0.462 to 0.572.

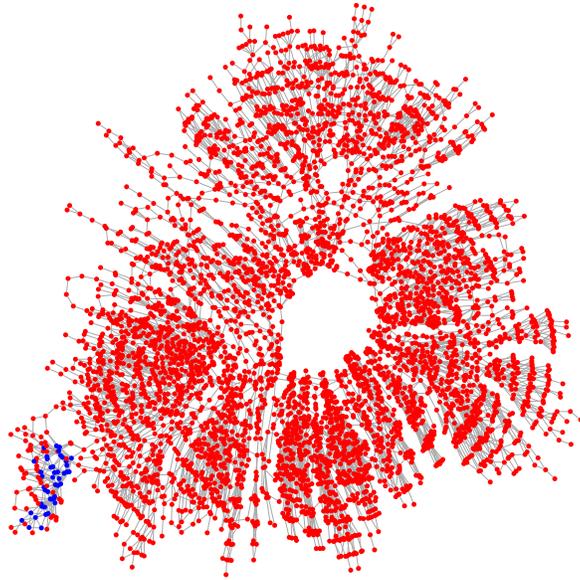
Alternatively, the equilibrium distribution of an extended contact process with $\frac{\lambda}{\mu} = 0.9$ and endogenous infection rate $\gamma_e = 0.0585$ is

$$\pi_{\text{approx}}(\mathbf{x}) = \frac{1}{Z} (0.9)^{1^T \mathbf{x}} (1.065)^{\frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{2}}, \quad \mathbf{x} \in \mathcal{X}. \quad (8.9)$$

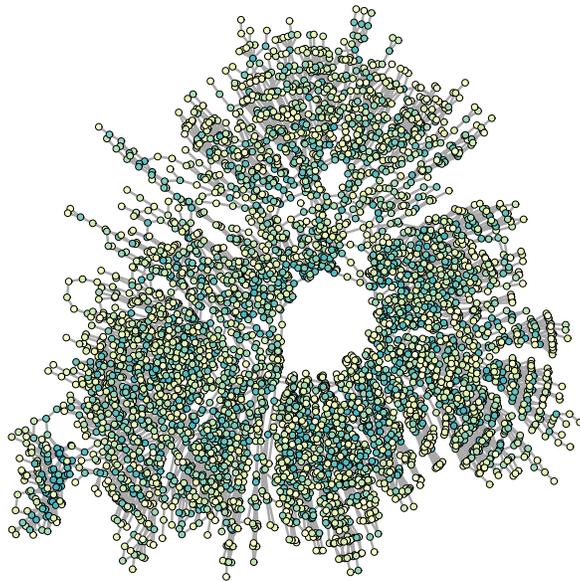
The expected fraction of infected agents at equilibrium of the extended contact process is

$$\widehat{E[Y]} = 0.495.$$

Figure 8.8(a) shows the most-probable configuration of (8.9), and Figure 8.8(b) shows the marginal probability of infection of the extended contact process. The marginal probability of infection ranges from 0.466 to 0.613. Since the endogenous infection rate is larger than the previous example, we expect both the number of infected agents in \mathbf{x} and the marginal probability of infection to be higher.

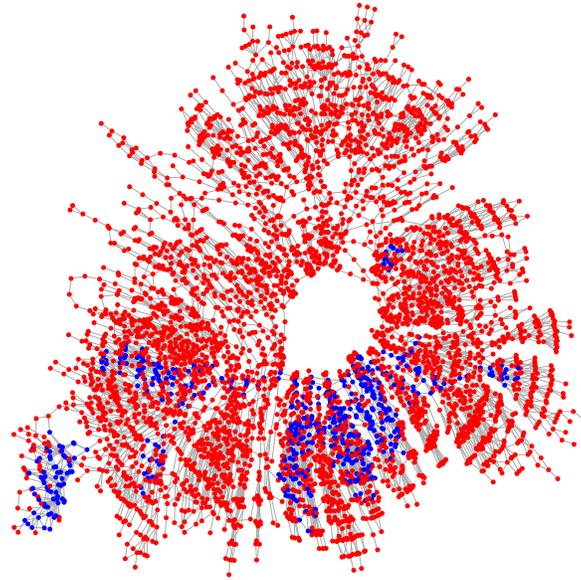


(a) Most-Probable Configuration, \mathbf{x}^* (Blue = Infected, Red = Healthy)

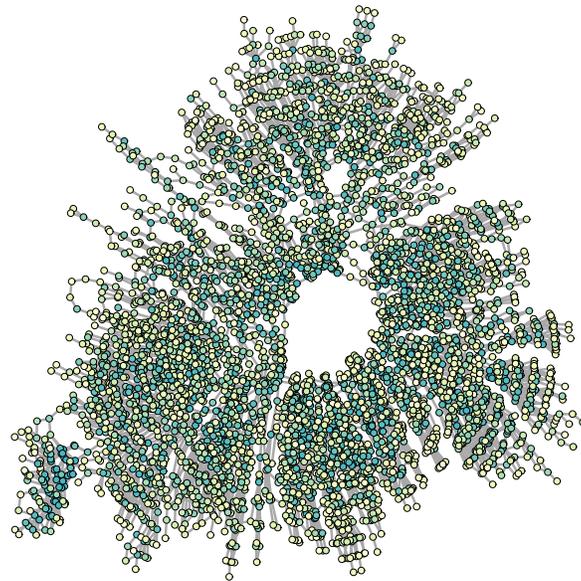


(b) Marginal Probability of Infection, $\hat{P}(x_i = 1)$, ranging from 0.462 to 0.572

Figure 8.7: Extended Contact Process $\frac{\lambda}{\mu} = 0.9$ and $\gamma_e = 0.0387$



(a) Most-Probable Configuration, \mathbf{x}^* (Blue = Infected, Red = Healthy)



(b) Marginal Probability of Infection, $\hat{P}(x_i = 1)$, ranging from 0.466 to 0.613

Figure 8.8: Extended Contact Process $\frac{\lambda}{\mu} = 0.9$ and $\gamma_e = 0.0585$

8.7 Conclusion

The scaled SIS process has a closed-form equilibrium distribution whereas the extended contact process does not. For a range of endogenous infection rates of the extended contact process, its equilibrium distribution is well approximated by that of a scaled SIS process with endogenous infection rate. We prove that this range depends on the maximum degree of the underlying network. With numerical simulations, we confirm that, within this parameter range, the total variational distance between the true equilibrium distribution of the extended contact process and its approximation, using the scaled SIS process, is very small (on the order of 10^{-5}). Further, we observe from experiments that, for certain network topologies, the approximation remains good even as the infection rate deviates from the established range.

Chapter 5 showed that agents belonging to subgraphs denser than the overall networks are more vulnerable to infection when the process follows the underlying assumptions of the scaled SIS process. Assuming that the infection rate is linearly dependent on the number of infected neighbors, as in the extended contact process, this intuition that densely connected agents are more vulnerable to infection is also applicable. The analysis results regarding the marginal probability of infection in chapter 6 and the expected fraction of infected agents in chapter 7 are also applicable to the extended contact process. When the equilibrium distribution of the extended contact process is no longer well approximated by that of the scaled SIS process, the most-probable configuration may differ between the two processes. We show through small numerical examples that configurations with high probability in the scaled SIS process remain highly probable in the extended contact process.

8.8 Appendix

Proof of Lemma 8.3.1

Lemma. *For any nonnegative integer m from 0 to d_{\max} , if*

$$\Delta^2 \ll \frac{2}{d_{\max}(d_{\max} - 1)},$$

then, for $\gamma = 1 + \Delta$,

$$\frac{\lambda}{\mu} \gamma^m = \frac{\lambda}{\mu} (1 + \Delta)^m \approx \frac{\lambda}{\mu} + \frac{\lambda}{\mu} \Delta m.$$

Proof. From the binomial series, for integer $m \in \{0, 1, \dots, d_{\max}\}$,

$$\begin{aligned} \frac{\lambda}{\mu} \gamma^m &= \frac{\lambda}{\mu} (1 + \Delta)^m = \frac{\lambda}{\mu} \left(\sum_{k=0}^m \binom{m}{k} \Delta^k \right) \\ &= \frac{\lambda}{\mu} \left(\binom{m}{0} \Delta^0 + \binom{m}{1} \Delta + \binom{m}{2} \Delta^2 + \binom{m}{3} \Delta^3 \dots + \binom{m}{m} \Delta^m \right) \\ &= \frac{\lambda}{\mu} \left(1 + m\Delta + \frac{m(m-1)}{2} \Delta^2 + \dots + m\Delta^{m-1} + \Delta^m \right). \end{aligned}$$

If

$$\left\{ \binom{m}{2} \Delta^2, \binom{m}{3} \Delta^3, \dots, \binom{m}{\frac{m}{2}} \Delta^{\frac{m}{2}} \right\} \ll 1, \quad \forall m \in \{0, 1, \dots, d_{\max}\}, \quad (8.10)$$

then the quadratic and higher order terms in the summation are negligible and we obtain the linear approximation

$$\frac{\lambda}{\mu} \gamma^m \approx \frac{\lambda}{\mu} + \frac{\lambda}{\mu} \Delta m,$$

which holds for all m .

Recognize that for $m \in \{0, 1, \dots, d_{\max}\}$, $m > \frac{m-1}{2} > \frac{m-2}{3} > \dots > \frac{\frac{m}{2}-1}{\frac{m}{2}}$. This means that

$$\binom{m}{2} \Delta^2 > \binom{m}{3} \Delta^3 > \dots > \binom{m}{\frac{m}{2}} \Delta^{\frac{m}{2}}, \quad \forall m \in \{0, 1, \dots, d_{\max}\}.$$

The largest possible upperbound is when $m = d_{\max}$. Therefore, condition (8.10) is satisfied when $\frac{d_{\max}(d_{\max}-1)}{2} \Delta^2 \ll 1$.

□

Proof of Theorem 8.3.1

Theorem. Consider the extended contact process exogenous infection rate λ , healing rate μ , and endogenous infection rate γ_e , over a static, simple, connected, undirected network of arbitrary topology, G , with maximum degree d_{\max} . Let $\gamma_e = \frac{\lambda}{\mu} \Delta$. If

$$\Delta^2 \ll \frac{2}{d_{\max}(d_{\max} - 1)},$$

then the equilibrium distribution of the extended contact process is well approximated by

$$\pi_{\text{approx}}(\mathbf{x}) = \frac{1}{Z} \left(\frac{\lambda}{\mu} \right)^{1^T \mathbf{x}} (1 + \Delta)^{\frac{\mathbf{x}^T A \mathbf{x}}{2}}, \quad \mathbf{x} \in \mathcal{X},$$

where A is the adjacency matrix of the network G , and Z is the partition function. The approximate distribution, $\pi_{\text{approx}}(\mathbf{x})$, is the equilibrium distribution (8.4) of an equivalent scaled SIS process over the same network G with exogenous infection rate λ , healing rate μ , and endogenous infection rate $\gamma_s = 1 + \Delta$.

Proof. From the theory of continuous-time Markov processes [49], the equilibrium distribution of the extended contact process is the left eigenvector of the transition rate matrix, \mathbf{Q}_e , corresponding to the 0 eigenvalue:

$$\pi \mathbf{Q}_e = 0$$

Entries of the matrix \mathbf{Q}_e correspond to the transition rates from one configuration $\mathbf{x} \in \mathcal{X}$ to another configuration according to the rates (8.1) and (8.3).

Lemma 8.3.1 gave the condition for when the infection rates (normalized by the healing rate) of the extended contact process are approximately the same as those of the scaled SIS process. As a result, the transition rate matrix of both processes are approximately the same. Therefore, the left eigenvector corresponding to the 0 eigenvalue of \mathbf{Q}_e is also the left eigenvector corresponding to the 0 eigenvalue of \mathbf{Q}_s , the transition rate matrix of the scaled SIS process with entries generated according to (3.1) and (3.2). We know that the left eigenvector of interest for the rate matrix, \mathbf{Q}_s , of the scaled SIS process is given by the closed-form equation (3.4). \square

Properties of Functions of Exponentially Distributed Random Variables [70]

Two Independent Random Variables

Let $A \sim \exp(\alpha)$, $B \sim \exp(\beta)$ be two independent exponentially distributed random variables. Then,

$$\begin{aligned}P(A \leq B) &= P(A - B \leq 0) \\&= \int_0^\infty \int_0^b \alpha e^{-\alpha a} \beta e^{-\beta b} da db \\&= \int_0^\infty \left(\int_0^b \alpha e^{-\alpha a} da \right) \beta e^{-\beta b} db \\&= \int_0^\infty (1 - e^{-\alpha b}) \beta e^{-\beta b} db \\&= \int_0^\infty \beta e^{-\beta b} db - \int_0^\infty \beta e^{-(\alpha+\beta)b} db \\&= 1 - \left(\frac{\beta}{\alpha + \beta} \right) \\&= \frac{\alpha}{\alpha + \beta}\end{aligned}$$

Multiple Independent Random Variables

Let $A \sim \exp(\alpha)$, $B_1 \sim \exp(\beta_1)$, $B_2 \sim \exp(\beta_2)$, \dots , $B_m \sim \exp(\beta_m)$ be independent exponentially distributed random variables. Let $C = \min\{B_1, B_2, \dots, B_m\}$, from properties of the exponential distribution, C is also an exponentially distributed random variable with rate $\beta_1 + \beta_2 + \dots + \beta_m$.

Therefore,

$$\begin{aligned}P(A \leq C) &= P(A - C \leq 0) \\&= 1 - \left(\frac{\beta_1 + \beta_2 + \dots + \beta_m}{\alpha + \beta_1 + \beta_2 + \dots + \beta_m} \right) \\&= \frac{\alpha}{\alpha + \beta_1 + \beta_2 + \dots + \beta_m}.\end{aligned}$$

The proof follows by induction on the number of independent exponentially distributed random variables.

Dynamic Bond Percolation Process

9.1 Introduction

The scaled SIS process models dynamic processes on networks; nodal states change in time depending on dynamics parameters and the states of neighboring nodes. For example, nodes represent agents in a population and the process models how virus spread in a population while accounting for infection and healing. A different but related problem is dynamical processes *of* networks. Rather than assuming that the network topology is static, the network structure itself now evolves according to some dynamical process.

Section 9.2 proposes the **dynamic bond percolation process**, a binary-state, continuous-time Markov process, where the states of edges change rather than the states of nodes. This is useful for modeling cascading failures of edges, where failed neighboring (i.e., adjacent) edges can lead to additional failures. For example, failure of transmission lines in the power grid can lead to failures of other lines. Unlike the scaled SIS process, there can be multiple variations of the dynamic bond percolation process. We consider two: 1) failure rate (i.e., edge closure rate) is exponentially dependent on the sum of the number of adjacent failed edges, 2) failure rate is exponentially dependent on the product of the number of adjacent failed edges. In the former case, failure rate does not depend on if the number of failed, adjacent edges at one end node is much larger than the number of failed, adjacent edges at the other end node. Section 9.3.1 proves that the dynamic bond percolation process has a closed-form equilibrium distribution and is therefore amenable to analysis similar to the scaled SIS process.

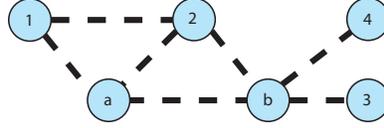


Figure 9.1: Maximal network represented by \mathbf{A}_{\max} . Dashed edges (i.e., bonds) are the only possible edges in the network.

9.2 Dynamic Bond Percolation Process

Consider a population of N agents represented as nodes (i.e., sites) in a network. The adjacency matrix, \mathbf{A}_{\max} , represents the restriction on the maximum number of edges (i.e., bonds) that can be supported by each node and to which node a given node is connected to. We assume that \mathbf{A}_{\max} represents a simple, connected, undirected graph, $G_{\max}(V_{\max}, E_{\max})$, where $V_{\max} = V(\mathbf{A}_{\max})$ is the set of nodes and $E_{\max} = E(\mathbf{A}_{\max})$ is the set of edges. We call the network represented by \mathbf{A}_{\max} the **maximal network**. Figure 9.1 shows an example of a maximal network.

The dynamic bond percolation process, $\{\mathbf{A}(t), t \geq 0\}$, is a stochastic process where the undirected edge $(a, b) \in E_{\max}$ can be in one of two states at any given time: open and closed. Depending on the application, we can interpret edge closure as edge removal, failure, or formation and edge opening as edge recovery or dissipation. At each time t , the state of the dynamic bond percolation process can be represented by the $N \times N$ adjacency matrix \mathbf{A} , where

$$\begin{aligned} \mathbf{A}_{i,j} &= 1 \text{ if edge } (i, j) \text{ is closed} \\ &= 0 \text{ if edge } (i, j) \text{ is open.} \end{aligned}$$

We will call \mathbf{A} the *configuration* or the *network state*. Recognize also then that $V(\mathbf{A}) = V_{\max}$ but $E(\mathbf{A}) \subseteq E_{\max}$.

Assumption 1: Network states that contain edges not in E_{\max} are invalid. For example, with respect to the maximal network in Figure 9.1, the network state in Figure 9.2 is not valid since edge $(2, 4)$ is not in the maximal network: edge closure and opening.

Figure 9.3 shows a possible path of the evolution of $\mathbf{A}(t)$ as edges close and open. Per convention in graph theory, we denote the rates related to edges with an apostrophe symbol. We consider two types of events:

Assumption 2: The dynamic bond percolation process, $\{\mathbf{A}(t), t \geq 0\}$, is a continuous-time Markov process.

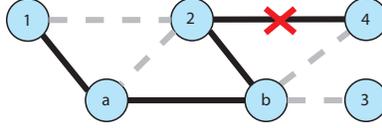


Figure 9.2: Invalid network state. Solid edges are *closed*. Dashed edges are *open*.

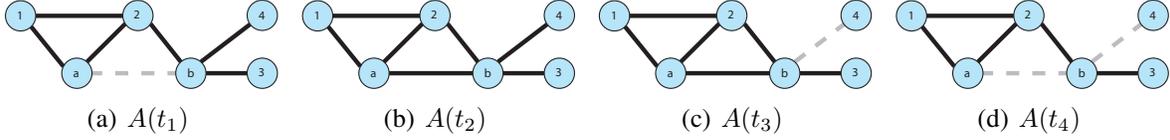


Figure 9.3: Example evolution of $\mathbf{A}(t)$. Solid edges are *closed*. Dashed edges are *open*.

The time between different events (e.g., $t_2 - t_1$ as in Figure 9.3) is exponentially distributed with a specific rate. By the properties of the exponential distribution, the average time between events is the inverse of the rate. See chapter 2 for a review of continuous-time Markov processes.

1. **Edge Opening.** Let the network state $T_{i,j}^- \mathbf{A}$ represent the state that is the same as \mathbf{A} except that edge (i, j) is open. As seen in Figure 9.5, the process goes from network state \mathbf{A} to $T_{i,j}^- \mathbf{A}$, where edge (i, j) goes from *closed* to *open*, with transition rate

$$q(\mathbf{A}, T_{i,j}^- \mathbf{A}) = \mu', \quad \mathbf{A} \neq T_{i,j}^- \mathbf{A}, \quad (9.1)$$

where $E(T_{i,j}^- \mathbf{A}) = E(\mathbf{A}) \setminus (i, j)$. We call $\mu' > 0$ the edge opening rate. We assume that μ' is the same for all edges. For applications where edge opening (i.e., recovery or dissipation) is rare, μ' can be arbitrarily small as long as it is not 0.

2. **Edge Closure.** An edge goes from *open* to *closed* in a length of time that is exponentially distributed. We separate the edge closure process into two types: *spontaneous* and *cascading*. Spontaneous edge closure means that edges close independently of one another; this is the scenario considered by standard bond percolation models [19]. On the other hand, cascading edge closures are dependent on the state of neighboring edges due to the underlying network structure. Spontaneous edge closure is akin to exogenous infection whereas cascading edge closure is similar to endogenous infection in the node-centric scaled SIS process (see chapter 3).

- a) *Spontaneous Edge Closure.* Edge (i, j) spontaneously goes from the *open* to the *closed* state, independently of all other edges.

- b) *Cascading Edge Closure*. Closure rate of edge (i, j) depends on the state of other edges in the network due to the cascading (i.e., infectious) effect.

Combining both spontaneous and cascading edge closures, the closure rates for the edges are no longer independent. Let the network state $T_{i,j}^+ \mathbf{A}$ represent the network state that is the same as \mathbf{A} except that (i, j) is closed. As seen in Figure 9.6, the process transitions from state \mathbf{A} to $T_{i,j}^+ \mathbf{A}$, where edge (i, j) goes from *open* to *closed* with transition rate:

$$q(\mathbf{A}, T_{i,j}^+ \mathbf{A}) = \lambda' \gamma'^{f(\mathcal{N}_a, \mathcal{N}_b)}, \quad \mathbf{A} \neq T_{i,j}^+ \mathbf{A}, \quad (9.2)$$

where $E(T_{i,j}^+ \mathbf{A}) = E(\mathbf{A}) \cup (i, j)$. The function $f(\mathcal{N}_i, \mathcal{N}_j)$ captures how the edge closure rate of (i, j) depends on the local neighborhood of edge (i, j) . We only consider closure rates to be dependent on the number of neighboring *closed* edges and of these specific forms:

- a) SUD (Sum-Dependent Dynamic Bond Percolation):

$$f(\mathcal{N}_i, \mathcal{N}_j) = |\mathcal{N}_i| + |\mathcal{N}_j| \quad (9.3)$$

- b) POD (Product-Dependent Dynamic Bond Percolation):

$$f(\mathcal{N}_i, \mathcal{N}_j) = |\mathcal{N}_i| |\mathcal{N}_j|. \quad (9.4)$$

Note that when $f(\mathcal{N}_i, \mathcal{N}_j) = 0$, the transition rate $q(\mathbf{A}, T_{i,j}^+ \mathbf{A}) = \lambda'$. Therefore, we consider λ' to be the spontaneous edge closure rate and γ' to be the cascading edge closure rate.

The SUD and POD models show different structure dependency behaviors as shown by the two scenarios in Figure 9.4. Under the SUD model, the closure rate of edge (i, j) is $\lambda' \gamma'^6$ for both Scenarios A and B. Under the POD model, the closure rate of edge (i, j) is $\lambda' \gamma'^9$ for Scenario A and $\lambda' \gamma'^5$ for Scenario B; the POD model discriminates imbalance between the closed edges on node a and node b whereas the SUD model does not.

Assumption 3: Multiple edge opening or closure can not occur simultaneously.

Since edges can spontaneously open and close, there are no absorbing states in the Markov process. Under these assumptions, the dynamic bond percolation process, $\{\mathbf{A}(t), t \geq 0\}$, is an

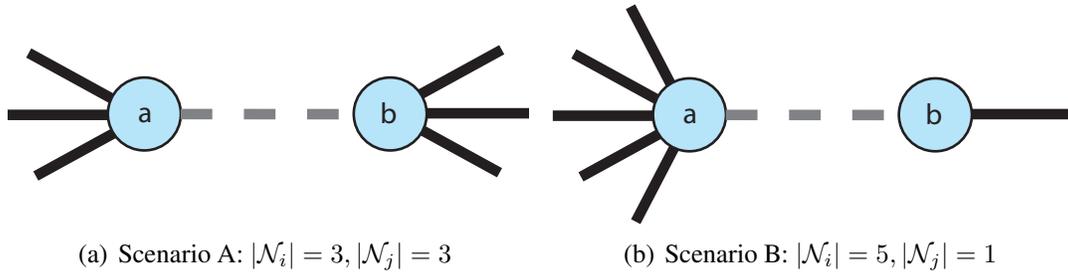


Figure 9.4: Different edge removal scenarios. Solid edges are *closed*. Dashed edges are *open*.

irreducible, continuous-time Markov process with finite state space, $\mathcal{A} = \{\mathbf{A}\}$; each state in the Markov process corresponds to a potential network state \mathbf{A} . Therefore,

$$|\mathcal{A}| = 2^{|E_{\max}|}.$$

For such a Markov process, the equilibrium distribution exists and is unique [50].

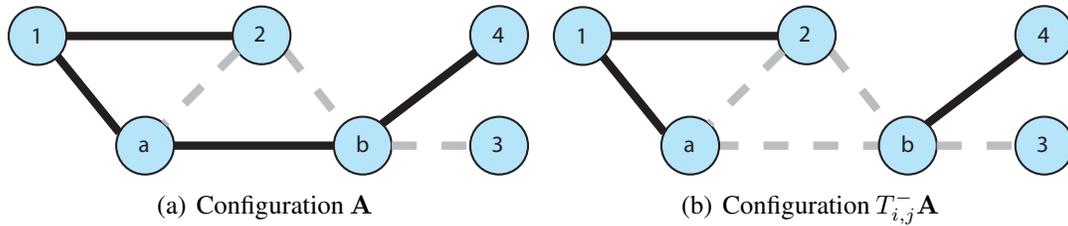


Figure 9.5: Example configurations \mathbf{A} and $T_{i,j}^- \mathbf{A}$. Solid edges are *closed*. Dashed edges are *open*.

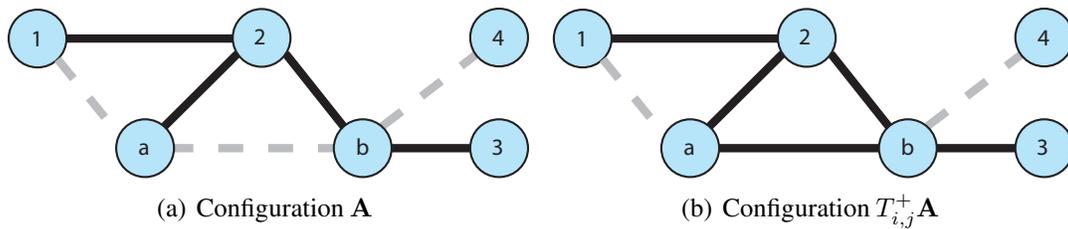


Figure 9.6: Example configurations \mathbf{A} and $T_{i,j}^+ \mathbf{A}$. Solid edges are *closed*. Dashed edges are *open*.

9.3 Equilibrium Distribution

Like the scaled SIS process, numerically solving for the equilibrium distribution of the dynamic bond percolation process, $\{\mathbf{A}(t), t \geq 0\}$, is intractable for large networks; the size of the state space $|\mathcal{A}| = 2^{|E_{\max}|}$. Under the rate assumptions (9.1) and (9.2), the dynamic bond percolation process

is also a reversible, continuous-time Markov process for which we will derive the equilibrium distribution in closed-form. Unlike the scaled SIS process, the sufficient statistics of $\mathbf{A}(t)$ relate to edges rather than nodes. First, we review some graph theory concepts [18, 42].

Definition 9.3.1. A walk is a list $v_0, e_1, v_1, e_2, \dots, e_k, v_k$ of vertices and edges. The length of the walk is the number of edges in the list. The number of walks in an undirected graph from node i to node j of length k is

$$(\mathbf{A}^k)_{i,j}, \tag{9.5}$$

where \mathbf{A} is the adjacency matrix of the undirected graph.

Definition 9.3.2. A path is a walk where all the vertices are distinct (although some literatures do not make this distinction between paths and walks). A graph that is a path is called a path graph and written as P_n , where n is the number of vertices (not edges) in the path. By convention, the path graph P_n is equivalent to a path of length $n - 1$. Figure 9.7 shows the P_3 subgraph and Figure 9.9 shows the P_4 subgraph.

Definition 9.3.3. A cycle is a path and where $v_0 = v_k$. A graph that is a cycle is called a cycle graph and written as C_n , where n is the number of vertices (not edges) in the cycle. By convention, the cycle graph C_n is equivalent to a cycle of length n . Figure 9.8 shows the C_3 subgraph.



Figure 9.7: P_3 graph

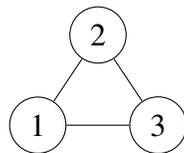


Figure 9.8: C_3 graph

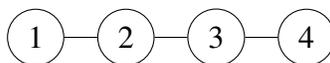


Figure 9.9: P_4 graph

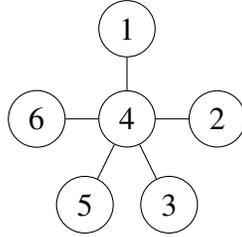


Figure 9.10: S_6 star graph

Definition 9.3.4. A star graph, $S_n(i)$, has n vertices that are only connected to a center vertex i . Figure 9.10 shows the $S_6(4)$ subgraph.

Definition 9.3.5. [42] A matching, \mathcal{M} of the graph $G(V, E)$, also called the Independent Edge Set, is a subset of edges E such that no vertex in V is incident to more than one edge in \mathcal{M} (see Figure 9.11(a)). Maximum Matching is a matching with the maximum number of edges (Figure 9.11(b)).

The number of edges in the maximum matching is known as the matching number, $\nu(G)$.

In this chapter, we want to introduce star matching, a generalization of the matching set.

Definition 9.3.6. A star matching, \mathcal{S} of the graph $G(V, E)$, is a subset of edges E such that these edges form a collection of disconnected star graphs (see Figure 9.11(c)). Maximum star matching is a star matching with the maximum number of edges (see Figure 9.11(d)). Note that $\mathcal{M} \subset \mathcal{S}$.

9.3.1 Reversibility and Equilibrium Distribution

Some Markov process possess the property that the process forward in time is statistically the same as the process backward in time. These Markov processes are called *reversible* Markov processes. There exists a necessary and sufficient theorem that makes solving for the equilibrium distribution of reversible Markov processes easier than for general Markov processes:

Theorem 9.3.7 (From [50]). A stationary Markov process is reversible if and only if there exists a collection of positive number $\pi(j), j \in \mathcal{L}$, summing to unity that satisfy the detailed balance conditions

$$\pi(j)q(j, k) = \pi(k)q(k, j), \quad j, k, \in \mathcal{L},$$

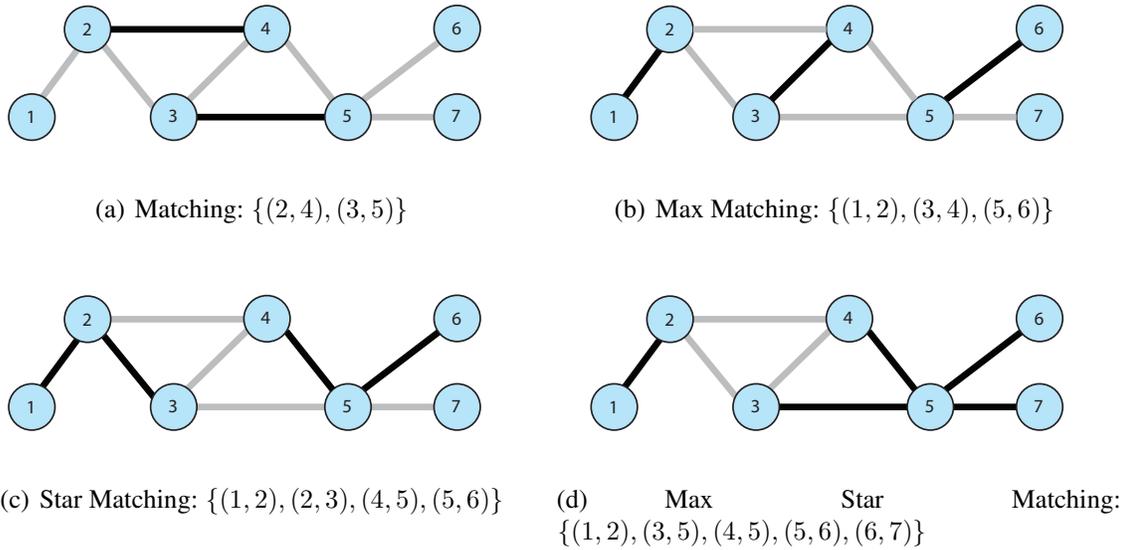


Figure 9.11: Matching and Star Matching

where $q(\cdot, \cdot)$ is a transition rate of the Markov process. When there exists such a collection $\pi(j), j \in \mathcal{L}$, it is the equilibrium distribution of the process.

Using this theorem, we will show that the equilibrium distribution for the SUD and POD dynamic bond percolation process has the same general form:

$$\pi(\mathbf{A}) = \frac{1}{Z} \left(\frac{\lambda'}{\mu'} \right)^{|E(\mathbf{A})|} \gamma'^{g(E(\mathbf{A}))}, \quad \mathbf{A} \in \mathcal{A},$$

where the partition function, Z , is

$$Z = \sum_{\mathbf{A} \in \mathcal{A}} \left(\frac{\lambda'}{\mu'} \right)^{|E(\mathbf{A})|} \gamma'^{g(E(\mathbf{A}))}. \quad (9.6)$$

Let $E(\mathbf{A})$ denote the set of closed edges in network state \mathbf{A} , then $g(E(\mathbf{A}))$ is the number of network structures induced by the set of closed edges $E(\mathbf{A})$.

As we will prove later, the function $g(E(\mathbf{A}))$ will change depending on how we account for nearest neighbor dependencies. However, the structure of the equilibrium distribution remains the same; it is also the product of three terms: the partition function, a *topology-free* term, and a *topology-dependent* term that accounts for the network structure of interest. The term $\left(\frac{\lambda'}{\mu'} \right)^{|E(\mathbf{A})|}$ is topology-free because it cares only about the number of edges in network state \mathbf{A} instead of what network structures are formed by these edges. The function $g(E(\mathbf{A}))$ on the other hand,

considers what network structures are induced by the edges in $E(\mathbf{A})$. Like the scaled SIS process, the equilibrium distribution of the dynamic bond percolation process is also a Gibbs distribution.

9.3.2 (SUD) Sum-Dependent Dynamic Bond Percolation Process

Theorem 9.3.8. *The Sum-Dependent dynamic bond percolation process, $\{\mathbf{A}(t), t \geq 0\}$, is a reversible Markov process and the equilibrium distribution is*

$$\pi(\mathbf{A}) = \frac{1}{Z} \left(\frac{\lambda'}{\mu'} \right)^{|E(\mathbf{A})|} \gamma'^{g(E(\mathbf{A}))}, \quad \mathbf{A} \in \mathcal{A}, \quad (9.7)$$

where $g(E(\mathbf{A}))$ is the number of paths of length 2 formed by the set of edges, $E(\mathbf{A})$, of the network represented by the adjacency matrix \mathbf{A} ; this is equivalent to the number of P_3 subgraphs induced by the closed edges in \mathbf{A} .

We can find the number of edges (i.e., the number of open edges), $|E(\mathbf{A})|$, using the adjacency matrix by solving

$$|E(\mathbf{A})| = \frac{\mathbf{1}^T \mathbf{A} \mathbf{1}}{2},$$

where $\mathbf{1} = [1, 1, \dots, 1]^T$, and the number of P_3 subgraphs induced by the edges $E(\mathbf{A})$ (see section 9.9 for derivation)

$$g(E(\mathbf{A})) = \sum_{i=1}^N \sum_{j>i} (\mathbf{A}^2)_{i,j} = \sum_{i=1}^N \binom{k_i}{2},$$

where k_i is the number of incident closed edges at node i .

In the SUD model, the sufficient statistics are the total number of closed edges, $|E(\mathbf{A})|$ and the total number of paths of length 2 (i.e., number of P_3 subgraphs) induced by the closed edges. The proof for Theorem 9.3.8 was presented in [37]. The number of P_3 subgraphs induced by the closed edges is related to the degree of the network \mathbf{A} . This reveals the complicated nature of networks: a structure that alternates between edges and nodes. A network process where failures cascade from edge to edge has a sufficient statistics that is a nodal characteristic.

9.3.3 (POD) Product-Dependent Dynamic Bond Percolation Process

Theorem 9.3.9. *The Product-Dependent dynamic bond percolation process, $\{\mathbf{A}(t), t \geq 0\}$, is a reversible Markov process and the equilibrium distribution is*

$$\pi(\mathbf{A}) = \frac{1}{Z} \left(\frac{\lambda'}{\mu'} \right)^{|E(\mathbf{A})|} \gamma'^{g(E(\mathbf{A}))}, \quad \mathbf{A} \in \mathcal{A}, \quad (9.8)$$

where $g(E(\mathbf{A}))$ is the number triangles and paths of length 3 formed by the set of closed edges, $E(\mathbf{A})$, of the network represented by the adjacency matrix \mathbf{A} ; this is equivalent to the number of C_3 and P_4 subgraphs induced by the closed edges in \mathbf{A} .

We can find the number of edges (i.e., the number of closed edges), $|E(\mathbf{A})|$, using the adjacency matrix by solving

$$|E(\mathbf{A})| = \frac{\mathbf{1}^T \mathbf{A} \mathbf{1}}{2},$$

where $\mathbf{1} = [1, 1, \dots, 1]^T$, and the number of C_3 and P_4 subgraphs induced by the edges $E(\mathbf{A})$ (see section 9.9 for detail)

$$g(E(\mathbf{A})) = \sum_{i=1}^N \frac{(\mathbf{A}^3)_{i,i}}{6} + \sum_{i=1}^N \sum_{j>i} \left[(\mathbf{A}^3)_{i,j} - (\mathbf{A}_{i,j}) \left((\mathbf{A}^2)_{i,i} + \sum_{k=1, k \neq i, j}^N \mathbf{A}_{k,j} \right) \right].$$

In the POD model, the sufficient statistics are the total number of closed edges and the total number of C_3 and P_4 subgraphs induced by the closed edges. The POD model and SUD model do not have the same sufficient statistics.

Proof. We want to prove that the equilibrium distribution of the Product-Dependent (POD) Dynamic Bond Percolation process is (9.8).

Using Theorem 3.4.2, the equilibrium distribution $\pi(\mathbf{A})$ must satisfy the detailed balance equations:

$$\pi(\mathbf{A})q(\mathbf{A}, T_{a,b}^+ \mathbf{A}) = \pi(T_{a,b}^+ \mathbf{A})q(T_{a,b}^+ \mathbf{A}, \mathbf{A}) \quad (9.9)$$

and

$$\pi(\mathbf{A})q(\mathbf{A}, T_{a,b}^- \mathbf{A}) = \pi(T_{a,b}^- \mathbf{A})q(T_{a,b}^- \mathbf{A}, \mathbf{A}). \quad (9.10)$$

We will prove condition (9.9) first. The LHS of (9.9) is

$$\pi(\mathbf{A})q(\mathbf{A}, T_{a,b}^+ \mathbf{A}) = \frac{1}{Z} \left(\frac{\lambda'}{\mu'} \right)^{|E(\mathbf{A})|} \gamma'^{g(E(\mathbf{A}))} \left(\lambda' \gamma'^{k_a k_b} \right) \quad (9.11)$$

$$= \frac{1}{Z} \left(\frac{\lambda'^{|E(\mathbf{A})|+1}}{\mu'^{|E(\mathbf{A})|}} \right) \gamma'^{g(E(\mathbf{A})) + k_a k_b}. \quad (9.12)$$

Note that $|E(T_{a,b}^+ \mathbf{A})| = 1 + |E(\mathbf{A})|$ since edge $\{a, b\}$ is closed. Second $g(E(T_{a,b}^+ \mathbf{A})) = k_a k_b + g(E(\mathbf{A}))$, as illustrated by Figure 9.6.

As shown in Figure 9.6(a), there are 0 C_3 and 3 P_4 (i.e., $\{1, a, 2, b\}$, $\{a, 2, b, 3\}$, $\{a, 2, b, 4\}$) subgraphs; therefore $g(E(\mathbf{A})) = 3$. Then $g(E(T_{a,b}^+\mathbf{A})) = 9$, since we can confirm that, after turning edge $\{a, b\}$ on, there are now 1 C_3 (i.e., $\{a, 2, b\}$) and 8 P_4 (i.e., $\{1, a, 2, b\}$, $\{a, 2, b, 3\}$, $\{a, 2, b, 4\}$, $\{1, a, b, 3\}$, $\{1, a, b, 4\}$, $\{2, a, b, 3\}$, $\{2, a, b, 4\}$, $\{1, a, b, 2\}$) subgraphs as shown in Figure 9.6(b).

Therefore, the RHS of (9.9) is

$$\pi(T_{a,b}^+\mathbf{A})q(T_{a,b}^+\mathbf{A}, \mathbf{A}) = \frac{1}{Z} \left(\frac{\lambda'}{\mu'} \right)^{|E(T_{a,b}^+\mathbf{A})|} \gamma'^{g(E(T_{a,b}^+\mathbf{A}))} \mu' \quad (9.13)$$

$$= \frac{1}{Z} \left(\frac{\lambda'^{|E(\mathbf{A})|+1}}{\mu'^{|E(\mathbf{A})|}} \right) \gamma'^{g(E(\mathbf{A}))+k_a k_b}. \quad (9.14)$$

The LHS and RHS of (9.9) are equivalent. Similar reasoning holds for (9.10).

□

9.4 At-Risk Edges and the Most-Probable Network

Problem

The advantage of the Dynamic Bond Percolation Mode is that it is an analyzable, microscopic model that couples the dynamic process of edge closing (i.e., failure) and opening (i.e., recovery) with the underlying network topology with no approximation of the network structure. Although the marginal probability of failure of each edge is difficult to compute for large maximal networks, we can solve for the network state (i.e., joint configuration of *all* the edges) with the maximum equilibrium probability and use this to find at-risk edges. We define the set of at-risk edges as the set edges that are closed in the network state with the highest equilibrium probability, \mathbf{A}^* :

$$\mathbf{A}^* = \arg \max_{\mathbf{A} \in \mathcal{A}} \pi(\mathbf{A}) = \arg \max_{\mathbf{A} \in \mathcal{A}} \left(\frac{\lambda'}{\mu'} \right)^{|E(\mathbf{A})|} \gamma'^{g(E(\mathbf{A}))}. \quad (9.15)$$

We call (9.15) the Most-Probable Network Problem and \mathbf{A}^* , the *most-probable network*. We can see from the formulation that the set of at-risk edges will be different for the SUD and POD model; this is because the function $g(E(\mathbf{A}))$. Furthermore, the set of at-risk edges depends on the dynamics of the process through the closure and opening rates. This is intuitive as in systems with high closure (i.e., failure) rate, more edges will be at-risk compared to systems with low closure rate.

Like in the scaled SIS process, the solution space of the Most-Probable Network Problem depends on the dynamics parameters. There are 4 parameter regimes:

I) **Recovery Dominant:** $0 < \frac{\lambda'}{\mu'} \leq 1, 0 < \gamma' \leq 1$

II) **Cascading Failure:** $0 < \frac{\lambda'}{\mu'} \leq 1, \gamma' > 1$

III) **Cascading Prevention:** $\frac{\lambda'}{\mu'} > 1, 0 < \gamma' \leq 1$

IV) **Removal Dominant:** $\frac{\lambda'}{\mu'} > 1, \gamma' > 1$

We are especially interested in comparing the difference between the solutions of the Most-Probable Network Problem, therefore the set of at-risk edges, for the SUD and POD models.

9.5 Regime I) Recovery Dominant and Regime IV)

Removal Dominant

In these regimes, the most-probable network of the SUD and POD model are the same. When $0 < \frac{\lambda'}{\mu'} \leq 1$, the average time an edge is open is longer than the average time the edge is closed; edges prefer to be open. Therefore, the most-probable network is the network state where the set of closed edges is empty: $E(\mathbf{A}) = \emptyset$. In other words, all the edges are open and $\mathbf{A}^* = \mathbf{A}_0$. When $0 < \gamma' \leq 1$, the term $\gamma'^{E(\mathbf{A})}$ is also a decreasing exponential; therefore, the most-probable network is also \mathbf{A}_0 . In regime I) **Recovery Dominant**, the topology-free and the topology-dependent terms are both decreasing exponential function of the number of closed edges, therefore, the most-probable network for both SUD and POD model is \mathbf{A}_0 ; for both regimes, the opening rate is high enough that none of the edges in the maximal network are considered at-risk.

When $\frac{\lambda'}{\mu'} > 1$, the average time an edge is open is shorter than the average time the edge is closed; edges prefer to be closed. Therefore, the most-probable network is the network where all the edges are closed: $E(\mathbf{A}) = E_{\max}$. Since the underlying network is restricted by the maximal network, the most-probable network is \mathbf{A}_{\max} . When $\gamma' > 1$, then term $\gamma'^{g(\mathbf{A})}$ is also an increasing exponential with increasing number of closed edges; therefore, the most-probable network is also \mathbf{A}_{\max} . In regime IV), the topology-dependent and the topology-free term are both increasing exponential functions, therefore, the most-probable network for both SUD and POD model is \mathbf{A}_{\max} ;

for these regimes, the closure rate is high enough that all of the edges in the maximal network are considered at-risk.

On the other hand, the most-probable networks for SUD and POD model are different for regime II) and regime III). We will focus on these more interesting regimes.

9.6 Regime II) of the Most-Probable Network Problem

In regime II) **Cascading Failure**, since $0 < \frac{\lambda'}{\mu'} \leq 1$ and $\gamma' > 1$, there is competition between edge recovery, which is driven by the topology-free process, and edge removal, which is driven by the topology-dependent process. Therefore, the solution space of the most-probable network exhibits phase transition behavior depending on if edge removal or edge recovery dominates. When the edge removal process dominates, the most-probable network of both SUD and POD model will be driven toward \mathbf{A}_{\max} ; if the edge recovery dominates, the most-probable network for both models will be driven toward \mathbf{A}_0 . This is what we expect from analysis of Regime I) and Regime IV). However, there are also solutions to the Most-Probable Network Problem that are neither \mathbf{A}_0 nor \mathbf{A}_{\max} . We call these solutions the *non-degenerate* most-probable networks; these solutions help us identify edges that are more at-risk to removal (i.e., failure) during cascading failures.

To find these non-degenerate solutions, we have to solve the Most-Probable Network Problem, a combinatorial optimization problem. In general, such computation is NP-hard [54]. But it was proved in [55], that the minimization of submodular functions could be solved in polynomial time. For both SUD and POD model, the Most-Probable Network Problem can be transformed to an equivalent submodular minimization problem in Regime II). The proof follows similar steps as in the proof for the Most-Probable Configuration Problem of the scaled SIS process (see chapter 5). Therefore, we only provide an outline of the proof.

Definition 9.6.1 ([53]). *A set function, $g : \mathcal{P}(E) \rightarrow \mathcal{R}$, is submodular if and only if for any $E(\mathbf{A}_1), E(\mathbf{A}_2) \subseteq E$ with $E(\mathbf{A}_2) \subseteq E(\mathbf{A}_1), i \notin E \setminus E(\mathbf{A}_1)$:*

$$g(E(\mathbf{A}_1) \cup \{i\}) - g(E(\mathbf{A}_1)) \leq g(E(\mathbf{A}_2) \cup \{i\}) - g(E(\mathbf{A}_2)).$$

The Most-Probable Network Problem is a set maximization problem where we want to find the subset of closed edges, $E(\mathbf{A}) \subseteq E_{\max}$, which maximizes the function

$$h(E(\mathbf{A})) = \left(\frac{\lambda'}{\mu'} \right)^{|E(\mathbf{A})|} \gamma'^{g(E(\mathbf{A}))}.$$

Section 5.2 presents the detailed proof for a variant model. Here, we outline the proof to argue that when $\gamma' \geq 1$, the function $-\log(h(E(\mathbf{A})))$ for both SUD and POD model is submodular. The Most-Probable Network Configuration Problem in Regime II) Cascading Failure is the maximization of a supermodular function and can therefore be solved in polynomial-time for both models.

Lemma 9.6.1. *Consider two sets of closed edges, $E(\mathbf{A}_1), E(\mathbf{A}_2) \subseteq E_{\max}$ and an additional closed edge $i \in E_{\max} \setminus \{E(\mathbf{A}_1), E(\mathbf{A}_2)\}$. We can see that*

$$|E(\mathbf{A}_1) \cup \{i\}| = |E(\mathbf{A}_1)| + 1,$$

and

$$|E(\mathbf{A}_2) \cup \{i\}| = |E(\mathbf{A}_2)| + 1.$$

The numbers of network structures (P_3 subgraphs in the SUD model or C_3, P_4 subgraphs in the POD model) induced by the closed edges in $E(\mathbf{A}_1)$ and $E(\mathbf{A}_2)$ are $g(E(\mathbf{A}_1))$ and $g(E(\mathbf{A}_2))$, respectively. Let the number of network structures induced by

$$E(\mathbf{A}_1) \cup \{i\} = g(E(\mathbf{A}_1)) + m_1$$

and

$$E(\mathbf{A}_2) \cup \{i\} = g(E(\mathbf{A}_2)) + m_2.$$

Therefore m_1 is the number of network structures created with the inclusion of edge i in $E(\mathbf{A}_1)$ and m_2 is the number of additional network structures created with the inclusion of edge i in $E(\mathbf{A}_2)$. If $E(\mathbf{A}_2) \subseteq E(\mathbf{A}_1)$, then:

1. $|E(\mathbf{A}_1)| \geq |E(\mathbf{A}_2)|$.
2. $g(E(\mathbf{A}_1)) \geq g(E(\mathbf{A}_2))$.
3. $m_1 \geq m_2$.

Proof. 1. When $E(\mathbf{A}_2) \subset E(\mathbf{A}_1)$, $E(\mathbf{A}_2)$ must have strictly fewer number of edges than $E(\mathbf{A}_1)$. When $E(\mathbf{A}_2) = E(\mathbf{A}_1)$, then they contain the same number of edges. Hence, $e_1 \geq e_2$.

2. When $E(\mathbf{A}_2) \subset E(\mathbf{A}_1)$, edges in $E(\mathbf{A}_2)$ can not induce more network structures (P_3 subgraphs in the SUD model or C_4, P_4 subgraphs in the POD model) than edges in $E(\mathbf{A}_1)$. When $E(\mathbf{A}_2) = E(\mathbf{A}_1)$, then the edge in $E(\mathbf{A}_1)$ and $E(\mathbf{A}_2)$ will induce the same number of network structures. Hence, $g(E(\mathbf{A}_1)) \geq g(E(\mathbf{A}_2))$.
3. Every edge in $E(\mathbf{A}_2)$ is also an edge in $E(\mathbf{A}_1)$. Every new closed edge connecting the closed edges $j \in E(\mathbf{A}_2)$ with i is also a new infected edge in $E(\mathbf{A}_1) \cup \{i\}$. However, some edge may also have $j \in E(\mathbf{A}_1)$. Hence, $m_1 \geq m_2$.

Every edge in $E(\mathbf{A}_2)$ is also an edge in $E(\mathbf{A}_1)$. Therefore, adding edge i to $E(\mathbf{A}_2)$ will generate the same or less number of network structures as adding edge i to $E(\mathbf{A}_1)$. Hence, $m_1 \geq m_2$.

□

Theorem 9.6.2. *If $\frac{\lambda'}{\mu'}$ and γ' are in Regime II) $0 < \frac{\lambda'}{\mu'} \leq 1, \gamma' > 1$, then $-\log(h(E(\mathbf{A})))$ is a submodular function, where*

$$-\log(h(E(\mathbf{A}))) = -|E(\mathbf{A})| \log\left(\frac{\lambda'}{\mu'}\right) - g(E(\mathbf{A})) \log(\gamma').$$

We will outline the proof for Theorem 9.6.2 because it is similar to the proof we presented in section 5.2. To prove that $-\log(h(E(\mathbf{A})))$ is submodular, we need to prove that it satisfies Definition 9.6.1. We show that

$$-\log(h(E(\mathbf{A}_1 \cup \{i\}))) + \log(h(E(\mathbf{A}_1))) \leq -\log(h(E(\mathbf{A}_2 \cup \{i\}))) + \log(h(E(\mathbf{A}_2))),$$

for any

$$E(\mathbf{A}_1) \subset E_{\max}, E(\mathbf{A}_2) \subseteq E(\mathbf{A}_1), i \notin E(\mathbf{A}_1).$$

Through algebra, we can show that this expression reduces to

$$-\log\left(\frac{\lambda'}{\mu'}\right) - m_1 \log(\gamma') \leq -\log\left(\frac{\lambda'}{\mu'}\right) - m_2 \log(\gamma').$$

This is a true statement since, from Lemma 9.6.1, we know that if $E(\mathbf{A}_2) \subseteq E(\mathbf{A}_1)$, then $m_1 \geq m_2$ and $\gamma' > 1$ implies that $\log(\gamma') > 0$. Therefore, $-\log(h(E(\mathbf{A})))$ is submodular. Since this is true regardless of the network structure in question. We can solve for the most-probable network for both the SUD and POD model in regime II) in polynomial time.

9.6.1 SUD vs. POD Most-Probable Network

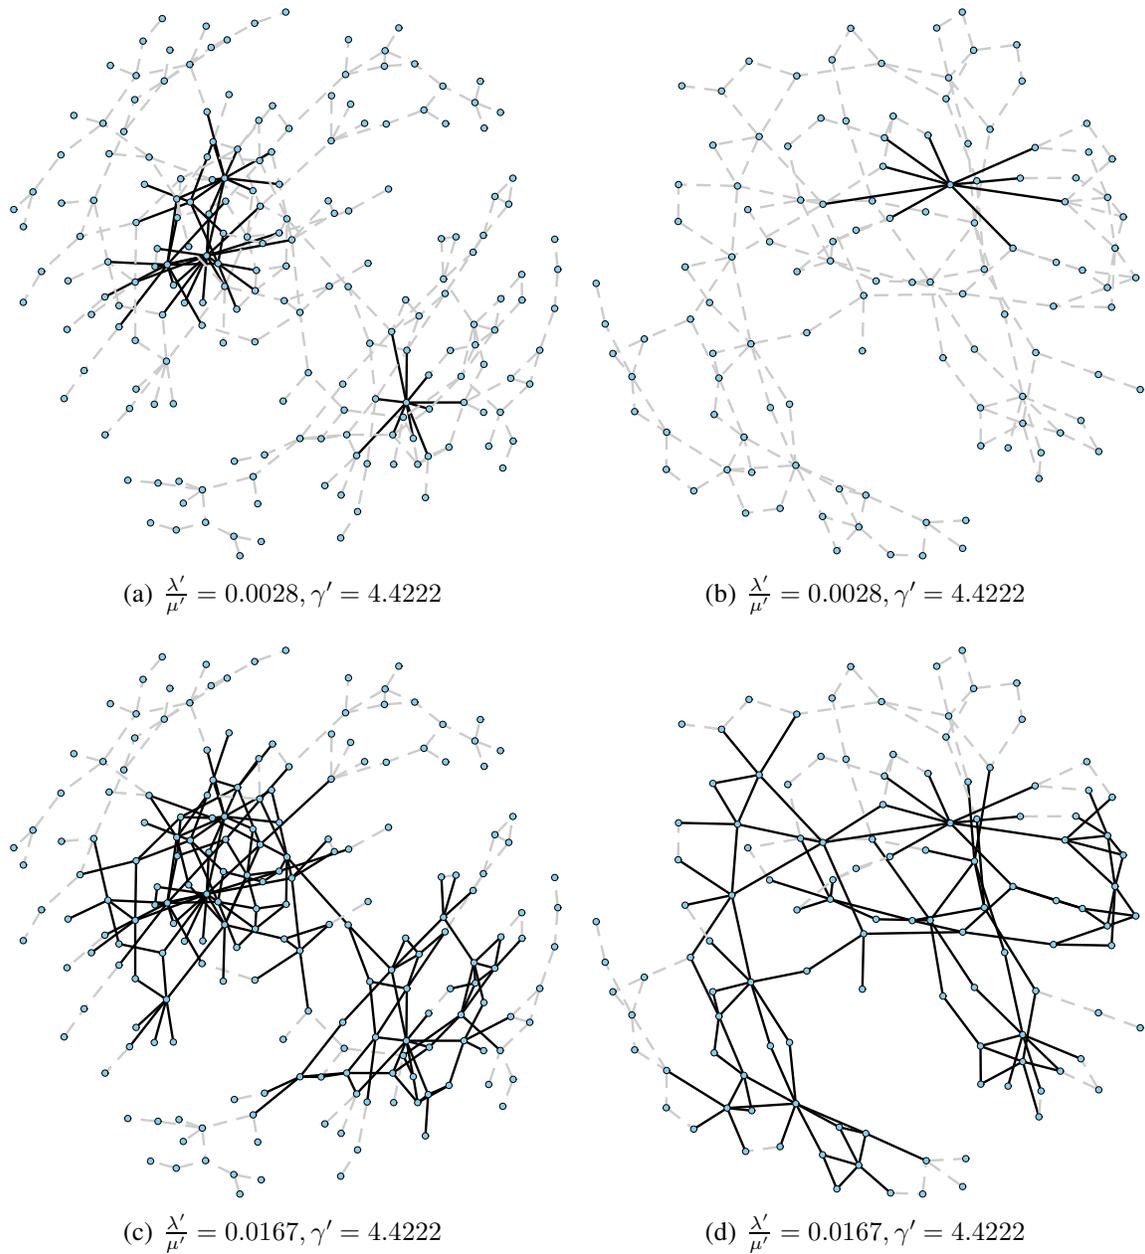


Figure 9.12: Most-Probable Network SUD (solid edges = closed, dashed edges = open)

Figure 9.12 and Figure 9.13 show the most-probable networks for the SUD and POD models under different edge closure and opening rates for 2 real-world networks: a 198-node social network [60] and the 118-bus IEEE power test system [71]. Unlike regime I) and regime IV), the solutions of the non-degenerate most-probable network are not the same for the SUD and POD

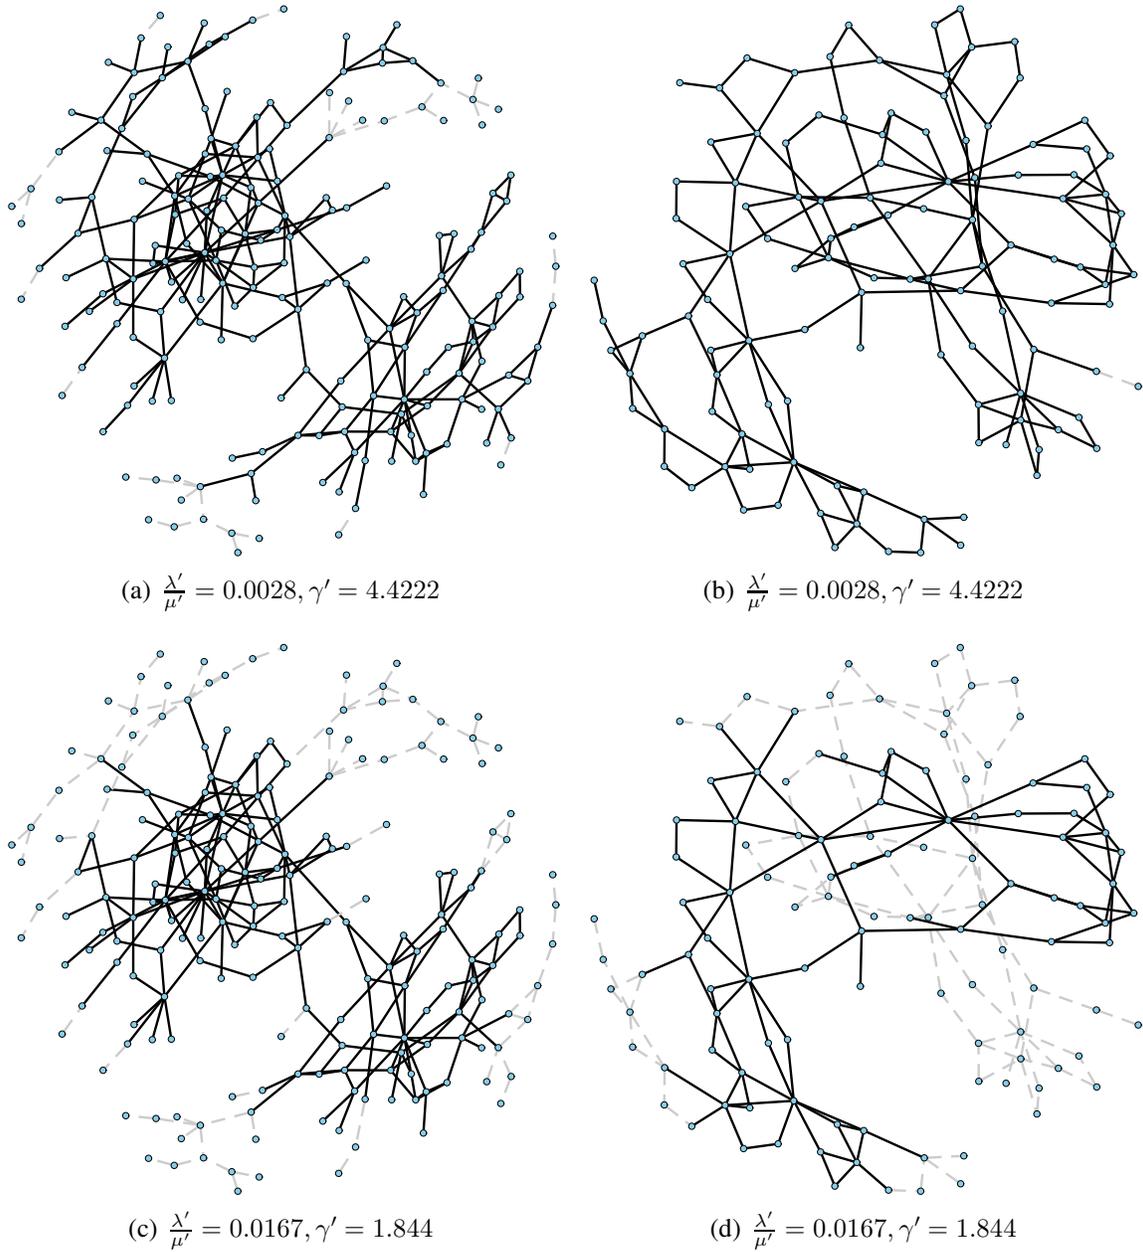


Figure 9.13: Most-Probable Network POD (solid edges = closed, dashed edges = open)

model. In regime II), the SUD model wants to minimize the number of closed edges, $|E(\mathbf{A})|$, while maximizing the number of P_3 subgraphs induced by the closed edges. The most-probable network of the SUD model consists of many star subgraphs (i.e., many P_3 s), few number of triangles, and tend to be disconnected.

On the other hand, the POD model wants to minimize the number of closed edges, $|E(\mathbf{A})|$, while maximizing the number of induced C_3 and P_4 subgraphs; we see that the most-probable networks of the POD model have more triangles (i.e., C_3 s) than the most-probable network of the SUD model. Additionally, the most-probable network of POD models are more likely to be connected.

For the same value of $\frac{\lambda'}{\gamma'}$ and γ' , the most-probable network of the SUD model have fewer number of closed edges than the most-probable network of the POD model. This is due to the fact that the cascading effect of the POD model is stronger because the edge closure rate (9.4) is dependent on the product rather than the sum of the number of neighboring closed edges.

9.7 Regime III) of the Most-Probable Network Problem

In regime III) **Cascading Prevention**, since $\frac{\lambda'}{\mu'} > 1$ and $0 < \gamma \leq 1$, like regime II) **Cascading Failure**, there is also competition between the topology-free process and the topology-dependent process; therefore, the solution space of the most-probable network will exhibit phase transition behavior depending on if edge recovery or edge removal dominates. Unlike in regime II), since the cascading edge closure rate γ is less than 1, this means that the average time an edge is open is longer with increasing number of closed edges on its end nodes; contagion, instead of driving cascading failures, prevents edge removal. Therefore, this regime is called regime III) **Cascading Prevention**.

Unfortunately, since $0 < \gamma' \leq 1$ in regime III), the Most-Probable Network Problem can not be transformed to an equivalent submodular problem. However, we will show that we can still solve for the most-probable network in polynomial time for a range of parameter values in regime III) for the SUD variation of the Dynamic Bond Percolation process. For the POD model, we do not know if the most-probable network can be solved in polynomial-time.

9.7.1 SUD Model and Maximum Matching

Recall that the equilibrium probability of network state \mathbf{A} is

$$\pi(\mathbf{A}) = \frac{1}{Z} \left(\frac{\lambda'}{\mu'} \right)^{|E(\mathbf{A})|} \gamma'^{g(E(\mathbf{A}))},$$

where $E(\mathbf{A})$, is the set of edges in \mathbf{A} (also the set of closed edges in \mathbf{A}_{\max} . Since in Regime III), $0 < \gamma' \leq 1$, we want $g(E(\mathbf{A}))$, the number of induced network structures to be as small as possible while maximizing $|E(\mathbf{A})|$.

For the SUD model, this means that we want to maximize the number of closed edges, $|E(\mathbf{A})|$, while minimizing the number of P_3 subgraphs; this means that we want to avoid paths of length 2 and only allow for paths of length 1. As a result, for a range of closing and opening rates, the most-probable network is a maximum matching (see Definition 9.3.5). It is known that the maximum matching can be found in polynomial time for arbitrary, undirected graphs [72, 73]. However the maximum matching may not be unique.

Theorem 9.7.1. *If $\frac{\lambda'}{\mu'}$ and γ' are in Regime III) $\frac{\lambda'}{\mu'} > 1, 0 < \gamma' \leq 1$ and $\lambda'\gamma' < \mu'$, then the most-probable network, \mathbf{A}^* , is the configuration(s) where $E(\mathbf{A}^*)$ is a maximum matching (see Definition 9.3.5).*

Proof. If $\frac{\lambda'}{\mu'}$ and γ' are in Regime III) $\frac{\lambda'}{\mu'} > 1, 0 < \gamma' \leq 1$ and $\lambda'\gamma' < \mu'$, then $E(\mathbf{A}^*)$ is a maximum matching.

Let \mathcal{A}' be the set of network states whose set of closed edges are maximum matching:

$$\mathcal{A}^* = \{\mathbf{A} \in \mathcal{A} : g(E(\mathbf{A})) = 0, |E(\mathbf{A})| \text{ is maximum}\}.$$

Proof by contradiction. Suppose that the most-probable network is $\mathbf{A}' \in \mathcal{A} \setminus \mathcal{A}^*$. Then there are two possibilities for \mathbf{A}' :

1. \mathbf{A}' is the network state such that $E(\mathbf{A}')$ is a matching but it is not the maximum matching.
2. \mathbf{A}' is the network state such that $E(\mathbf{A}')$ is not a matching.

Case 1) is the easiest to prove. It implies that $|E(\mathbf{A}')| < |E(\mathbf{A}^*)|$, for $\mathbf{A}^* \in \mathcal{A}^*$. Since $\frac{\lambda'}{\mu'} > 1$, then

$$\frac{\lambda'}{\mu'}^{|E(\mathbf{A}^*)|} > \frac{\lambda'}{\mu'}^{|E(\mathbf{A}')|}.$$

Therefore, \mathbf{A}' can not be the most-probable network.

Case 2) implies that $g(E(\mathbf{A}')) > 0$. There are then two possibilities: 1) $|E(\mathbf{A}')| \leq |E(\mathbf{A}^*)|$, or 2) $|E(\mathbf{A}')| > |E(\mathbf{A}^*)|$.

If $|E(\mathbf{A}')| \leq |E(\mathbf{A}^*)|$, then

$$\left(\frac{\lambda'}{\mu'}\right)^{|E(\mathbf{A}')|} \gamma'^{g(E(\mathbf{A}'))} < \left(\frac{\lambda'}{\mu'}\right)^{|E(\mathbf{A}^*)|}.$$

Therefore, \mathbf{A}' can not be the most-probable network.

The most complicated case is if $g(E(\mathbf{A}')) > 0$ and $|E(\mathbf{A}')| > |E(\mathbf{A}^*)|$. Let $|E(\mathbf{A}^*)| = k$ and

$$\mathcal{A}_1 = \{\mathbf{A} \in \mathcal{A} : g(E(\mathbf{A})) \geq 1, |E(\mathbf{A})| = k + 1\}$$

Similarly, we can define another set of network states, \mathcal{A}_2 , where $|E(\mathbf{A}_2)| = k + 2$ for all $\mathbf{A}_2 \in \mathcal{A}_2$.

Realize that

$$g(E(\mathbf{A}_2)) > g(E(\mathbf{A}_1)) \forall \mathbf{A}_1 \in \mathcal{A}_1$$

since k is the size of the maximum matching set. Therefore,

$$\mathcal{A}_2 = \{\mathbf{A} \in \mathcal{A} : g(E(\mathbf{A})) \geq 2, |E(\mathbf{A})| = |E(\mathbf{A}^*)| + 2\}.$$

Since $0 < \gamma' \leq 1$, the configuration with the maximum equilibrium probability in set \mathcal{A}_1 is

$$\pi(\mathbf{A}_1^*) = \frac{1}{Z} \left(\frac{\lambda'}{\mu'}\right)^{k+1} \gamma',$$

and the configuration with the maximum equilibrium probability in set \mathcal{A}_2 is

$$\pi(\mathbf{A}_2^*) = \frac{1}{Z} \left(\frac{\lambda'}{\mu'}\right)^{k+2} \gamma'^2.$$

Since $\lambda'\gamma' < \mu'$, this implies that $\pi(\mathbf{A}_1^*) > \pi(\mathbf{A}_2^*)$. Similar argument will show that $\pi(\mathbf{A}_2^*) > \pi(\mathbf{A}_3^*)$, etc.

However

$$\pi(\mathbf{A}^*) = \frac{1}{Z} \left(\frac{\lambda'}{\mu'}\right)^k,$$

which means that $\pi(\mathbf{A}^*) > \pi(\mathbf{A}_1^*)$.

□

9.7.2 POD Model and Maximum Star Matching

For the POD model, the network structure of interest, $g(E(\mathbf{A}))$, is the number of C_3 and P_4 subgraphs induced by the closed edges. To maximize the equilibrium distribution, we want to maximize $|E(\mathbf{A})|$ while minimizing the number of C_3 and P_4 subgraphs; this means that we want to avoid paths of length 3. Unlike the SUD model, the set of closed edges can form paths of length 1 or paths of length 2. As we want to maximize the number of closed edges, the most-probable configuration will be biased toward the set of closed edges that maximizes the number of paths of length 2.

Theorem 9.7.2. *In Regime III), if $\lambda'\gamma' < \mu'$, then the most-probable network, \mathbf{A}^* , is the configuration(s) where $E(\mathbf{A}^*)$ is a maximum star matching (see Definition 9.3.6).*

The proof follows the same argument as Theorem 9.7.1. Currently we do not know if the maximum star matching can be solved in polynomial-time.

9.8 Conclusion

The dynamic bond percolation process shows the framework of the scaled SIS process applied to edges instead of nodes. It is also a binary-state, reversible continuous-time Markov process for which we can find the closed-form equilibrium distribution instead of solving for an eigenvector of the $2^{|E|} \times 2^{|E|}$ transition rate matrix. Like for the scaled SIS process, the equilibrium distribution is a Gibbs distribution. Under the SUD (sum-dependent) assumption, the sufficient statistics of the dynamic bond percolation process are the total number of closed edges and the total number of P_3 (path graphs with 3 nodes) subgraphs. Under the POD (product-dependent) assumption, the sufficient statistics are the total number of closed edges and the total number of P_4 (path graphs with 4 node) and C_3 (cycle graphs with 3 nodes) subgraphs.

Finding the configuration with the maximum equilibrium probability, like solving the Most-Probable Configuration Problem scaled SIS process, is also supermodular in the regime II): individual edges prefer the open state but closure rate increases with additional adjacent edge closures. We show the existence of non-degenerate solutions, configurations where only some of the edges in the network are closed but others are open. *Unlike* the scaled SIS process, this combinatorial optimization problem can also be solved in polynomial-time for regime III): individual edges prefer

the closed state but closure rate decreases with additional adjacent edge closures. This is because the **maximum matching** –the edge-centric analogy to the maximum independent set problem –is solvable in polynomial-time solvable for general graph. The most-probable configuration in regime III) under the SUD assumption is related to maximum matching. The most-probable configuration under the POD assumption is related to a new concept that we call maximum star matching.

9.9 Appendix

Determining $g(E(\mathbf{A}))$ for SUD

For the Sum-Dependent dynamic bond percolation process, $g(E(\mathbf{A}))$ is the number paths of length 2 formed by the set of edges, $E(\mathbf{A})$, of the network represented by the adjacency matrix \mathbf{A} . The number of walks of length 2 from node i to node $j \neq i$ is

$$\sum_{i=1}^N \sum_{j>i} (\mathbf{A}^2)_{i,j}.$$

Realize that this is also equivalent to the number of paths of length 2 from node i to node $j \neq i$.

Determining $g(E(\mathbf{A}))$ for POD

For the Product-Dependent dynamic bond percolation process, $g(E(\mathbf{A}))$ is the number of triangles and paths of length 3 formed by the set of edges, $E(\mathbf{A})$, of the network represented by the adjacency matrix \mathbf{A} . The number walks of length 3 [18] is

$$\sum_{i=1}^N \frac{(\mathbf{A}^3)_{i,i}}{6}.$$

We need to find the number of paths of length 3. We know that the number of walks of length 3 from node i to node $j \neq i$ is

$$\sum_{i=1}^N \sum_{j>i} (\mathbf{A}^3)_{i,j}.$$

This number is *larger* than the number of paths of length 3 because there are walks from node i to node j that are not paths. Figure 9.14, Figure 9.15, and Figure 9.16 illustrate the three cases of walks of length 3 that are not paths of length 3 because the vertices repeat.

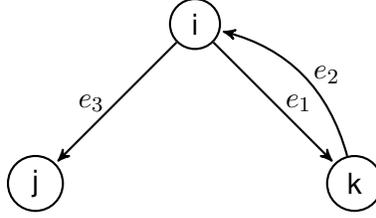


Figure 9.14: Walks of Length 3 that are not Paths of Length 3: $i, e_1, k, e_2, i, e_3, j$

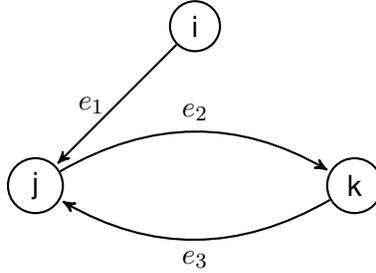


Figure 9.15: Walks of Length 3 that are not Paths of Length 3: $i, e_1, j, e_2, k, e_3, j$

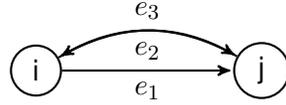


Figure 9.16: Walks of Length 3 that are not Paths of Length 3: $i, e_1, j, e_3, i, e_2, j$

Therefore, the number of paths of length 3 from node i to node $j \neq i$ is the total number of walks of length 3 from i to j minus the above scenarios:

$$\sum_{i=1}^N \sum_{j>i} (\mathbf{A}^3)_{i,j} - (\mathbf{A}_{i,j}) \sum_{k=1, k \neq i, j}^N \mathbf{A}_{i,k} - (\mathbf{A}_{i,j}) \sum_{k=1, k \neq i, j}^N \mathbf{A}_{j,k} - \mathbf{A}_{i,j} \quad (9.16)$$

$$= \sum_{i=1}^N \sum_{j>i} (\mathbf{A}^3)_{i,j} - (\mathbf{A}_{i,j}) \left(\sum_{k=1, k \neq i, j}^N \mathbf{A}_{i,k} + \sum_{k=1, k \neq i, j}^N \mathbf{A}_{j,k} + 1 \right) \quad (9.17)$$

$$= \sum_{i=1}^N \sum_{j>i} (\mathbf{A}^3)_{i,j} - (\mathbf{A}_{i,j}) \left((\mathbf{A}^2)_{i,i} + \sum_{k=1, k \neq i, j}^N \mathbf{A}_{k,j} \right). \quad (9.18)$$

Therefore,

$$g(E(\mathbf{A})) = \sum_{i=1}^N \frac{(\mathbf{A}^3)_{i,i}}{6} + \sum_{i=1}^N \sum_{j>i} (\mathbf{A}^3)_{i,j} - (\mathbf{A}_{i,j}) \left((\mathbf{A}^2)_{i,i} + \sum_{k=1, k \neq i, j}^N \mathbf{A}_{k,j} \right).$$

Thesis Summary and Future Work

The thesis studied network processes, which are dynamical processes on networks. The network represents interactions or dependencies between multiple agents, units, or components. Individual units do not act independently of one another. Network process models are more realistic than traditional models in the sense that they can account for interdependencies that may lead to phenomena such as epidemics, blackouts, or the sudden explosion of popularity of a new Internet meme.

In real-world application of interest, network processes usually involve thousands or millions of agents, making them difficult to study efficiently through experiments or simulations. The number of possible states of a network process grows exponentially with the size of the network. Many models either rely on Monte Carlo methods or approximate the underlying network with simpler structure.

Through the scaled SIS process, we developed a network process model, which is tractable for unweighted, undirected networks of arbitrary topology and size. The scaled SIS process, unlike previous network process models such as the contact process, assumes that the infection rate of a healthy agent is exponentially, instead of linearly, dependent on the number of infected neighbors. As a mathematical model, infection in the scaled SIS process can also stand for failure and healing for recovery of components. In chapter 3, we proved that the scaled SIS process is a reversible, continuous-time Markov process and derived the closed-form equilibrium distribution, $\pi(\mathbf{x})$.

The equilibrium distribution showed that the scaled SIS process can be decomposed into two separate processes: a topology-independent process, which models the preference of individual

agents and a topology-dependent process, which models the effect of the network structure on the process. The dynamics parameters of the scaled SIS process, $\frac{\lambda}{\mu}, \gamma$ can be divided into 4 regimes. Complex equilibrium behavior occurs when the topology-independent and the topology-dependent processes oppose each other as in regime II) **Endogenous Infection Dominant:** $0 < \frac{\lambda}{\mu} \leq 1, \gamma > 1$ and regime III) **Exogenous Infection Dominant:** $\frac{\lambda}{\mu} > 1, 0 < \gamma \leq 1$. From an application perspective, regime II) is the regime of interest since it best models traditional epidemics behaviors: individual agents prefer the healthy state, but additional number of infected neighbors increases an agent's susceptibility to infection.

Chapters 4 to 7 used the scaled SIS process to analyze vulnerability of the population to infection at 3 different scales: 1) microscopic by finding the marginal probability of infection of individual agents, $P(x_i = 1)$; 2) mesoscopic by solving for the most-probable configuration, \mathbf{x}^* ; and 3) macroscopic by finding the expected fraction of infected agents, $E[Y]$. Although these inference problems are NP-hard in general, they can be solved exactly or approximated in regime II) in polynomial-time.

We showed in chapter 4 that in regime III), the vulnerable substructures in the network are related to the maximum independent set whereas in regime II), they are related to dense subgraphs in the network. Many solutions of the Most-Probable Configuration Problem in regime II) are non-degenerate, where only some nodes in the network are infected. Chapter 5 proved that these infected nodes induce subgraphs whose density is higher than the overall network and that they are subgraph unique. This shows that the mesoscale structures of interest for epidemics are dense subgraphs. Furthermore, the structure can be found efficiently; solving for the exact most-probable configuration in regime II) of the scaled SIS process on a network with approximately 5000 nodes, corresponding to a network process with 2^{5000} possible configurations, takes less than 1 sec on a standard desktop using Max-Flow/Min-Cut algorithm.

Chapter 6 discussed the marginal probability of infection, $P(x_i = 1)$, which characterizes the susceptibility of individual agents to infection. We show using examples that when $0 < \frac{\lambda}{\mu} \ll 1$ and $\gamma \approx 1$, nodal degree is indicative of the marginal probability of infection. However, when $0 < \frac{\lambda}{\mu} \ll 1$ and $\gamma \gg 1$, membership in dense subgraphs is more indicative than nodal degree for $P(x_i = 1)$. As a result, depending on its location in the network, agent A may have higher probability of being infected at equilibrium than agent B for one set of dynamics parameters and

lower probability of being infected for another set of dynamics parameters. Finding $P(x_i = 1)$ is NP-hard since it requires marginalization and summing over 2^N configurations to find the partition function, Z . In regime II) however, $P(x_i = 1)$ can be approximated efficiently using Perturb-and-MAP sampling. This Monte Carlo sampling method relies on the condition that the most-probable configuration can be solved in polynomial-time in regime II). It is an effective sampling method and performs better than loopy belief propagation (LBP) in experiments.

Chapter 7 considered the expected fraction of infected agents, $E[Y]$, which characterizes the vulnerability of the entire network. We proved that the expected fraction of infected agents, $E[Y]$, of the scaled SIS process at equilibrium is the average of $P(x_i = 1)$; this shows that the macroscopic characteristics of the process can be derived from the microscopic characteristics but not the other way around. Further, we proved that the macroscopic characterization (i.e., the state of the entire population) is indicative of the microscopic characterization (i.e, the state of individuals) only for special cases such as when $\gamma = 1$, or when the underlying network is a vertex-transitive graph. In regime II), $E[Y]$ can also be approximated efficiently using Perturb-and-MAP sampling. In this regime, networks that are more densely connected have larger $E[Y]$ than sparser networks for the same dynamics parameters. The dependence of $E[Y]$ on the topology-independent parameter $\frac{\lambda}{\mu}$ grows more nonlinear with increasing γ until saturation at $E[Y] = 1$. The dependence of $E[Y]$ on the topology-dependent parameter γ is a sigmoid function.

Lastly, chapters 8 and 9 analyzed the relationships between the scaled SIS process and other network process models. We showed that the equilibrium distribution for a subclass of extended contact process can be approximated by that of an equivalent scaled SIS process in chapter 8; consequently, analysis of the scaled SIS process also extends to the extended contact process in these cases.

In chapter 9, we then presented the dynamic bond percolation process, which is also a reversible, continuous-time Markov process similar to the scaled SIS process. The dynamic bond percolation process models changing edge states instead of node states. While the equilibrium distribution can also be described in closed-form, the sufficient statistics of the dynamic bond percolation process differ from the scaled SIS process. The sufficient statistics are related to edges and the number of small subgraphs (i.e., motifs) in the network such as P_3 and C_3 subgraphs. Furthermore, the configuration with the highest equilibrium probability can be solved in polynomial-time

for both regime II) and regime III). The dynamic bond percolation process provides an alternative approach to network processes that is edge-centric. This is useful for applications where the network structure changes depending on some process, for example, cascading failure of transmission lines in the power grid.

The contributions of the thesis are:

1. Developed the scaled SIS process and derived the closed-form equilibrium distribution
2. Showed that subgraph density characterizes vulnerability at the mesoscale. Subgraphs that are denser than the overall network are more vulnerable to infection
3. Analyzed how network topology and dynamics parameters affect the vulnerability of individual agents through the marginal probability of infection
4. Analyzed how network topology and dynamics parameters affect the vulnerability of the entire network through the expected fraction of infected agents
5. Used scaled SIS process to analyze the extended contact process
6. Developed the dynamic bond percolation process and derived the closed-form equilibrium distribution

10.1 Future Works

A major open question regarding the equilibrium distribution of the scaled SIS process is the **mixing time**, which is the amount of time it takes for the continuous-time Markov process to reach equilibrium. It is intuitive that mixing time would increase with the size of the network. But is the dependence logarithmic, linear, or exponential? Since the equilibrium distribution does not depend on the initial configuration, knowing if a network process has or has not reached equilibrium can help to quantify if the initial configuration can be found. For example, solving for the initial configuration of the scaled SIS process is akin to solving for the initial source of the epidemics.

The analysis presented in this thesis focused on regime II). It is computationally feasible to study the scaled SIS process for large networks in this regime. Regime III) **Exogenous Infection Dominant:** $\frac{\lambda}{\mu} > 1, 0 < \gamma \leq 1$ on the other hand, relates to a known NP-hard problem (i.e.,

Maximum Independent Set). Since $0 < \gamma \leq 1$, this means that additional infected neighbors decreases an agent's susceptibility to infection. However, this is also realistic in the sense that it models defensive actions; when people observe that other people are getting sick, they often take steps to reduce their chance of catching the same illness. Analyzing the behavior of the scaled SIS process in regime III) would be an interesting extension of the thesis. The challenge being that we proved that solving for the most-probable configuration is related to solving the Maximum Independent Set Problem, which is a known NP-hard problem.

Using the scaled SIS process, we proved that very structured graphs, such as k -regular or complete multipartite graphs, regardless of size, do not contain subgraphs that are denser than the overall network. This indicates that analysis of network processes can yield results on the underlying graph structure itself, similar to how the behavior of systems can be studied with dynamical inputs (i.e., pulse sequences). An extension of this is to study network topology using dense subgraph structures. We showed with examples that real-world networks have many non-degenerate solutions. It would be interesting to extend this observation to classes of networks such as scale-free networks; perhaps the power-law associated complexity of scale-free networks is related to the possibility that they are more likely to contain denser subgraphs than Erdős-Rényi random graphs.

It would be of interest to explore the problems of network visualization and graph sampling as related to dense subgraphs. In many applications, the visualization of large networks is not informative and often misleading. How to design graph layout algorithms such that community structures like dense subgraphs, corresponding to the non-degenerate solutions of the Most-Probable Configuration Problem, are readily apparent? When networks have millions or even billions of nodes, even polynomial-time algorithms are too slow. It is an open question as to how to sample extremely large networks while retaining the fundamental structures of the network.

This thesis focuses on analysis problems. In many situations, the underlying network structure is not known and/or the dynamics parameters of the process are not known. A major challenge is to migrate from analyzing the behaviors of network processes to designing behaviors of network processes. We need methods to incorporate longitudinal data (i.e., data with both time and spatial information) with network process models such as the scaled SIS process or the dynamic bond percolation process. How can parameters such as the network structure or dynamic parameters be estimated using data? What is the computational complexity of such problem? What approxima-

tion and tools are needed for such problems to become tractable?

Bibliography

- [1] M. J. Keeling and K. T. Eames, “Networks and epidemic models,” *Journal of the Royal Society Interface*, vol. 2, no. 4, pp. 295–307, 2005.
- [2] D. J. Daley, J. Gani, and J. M. Gani, *Epidemic Modelling: an Introduction*. Cambridge University Press, 2001, vol. 15.
- [3] A. Barrat, M. Barthelemy, and A. Vespignani, *Dynamical Processes on Complex Networks*. Cambridge University Press, 2008, vol. 1.
- [4] R. Durrett, “Some features of the spread of epidemics and information on a random graph,” *Proceedings of the National Academy of Sciences*, vol. 107, no. 10, pp. 4491–4498, 2010. [Online]. Available: <http://www.pnas.org/content/107/10/4491.abstract>
- [5] T. E. Harris, “Contact interactions on a lattice,” *The Annals of Probability*, pp. 969–988, 1974.
- [6] T. M. Liggett, *Stochastic Interacting Systems: Contact, Voter and Exclusion Processes*. Springer, 1999, vol. 324.
- [7] P. Van Mieghem and E. Cator, “Epidemics in networks with nodal self-infection and the epidemic threshold,” *Phys. Rev. E*, vol. 86, p. 016116, Jul 2012. [Online]. Available: <http://link.aps.org/doi/10.1103/PhysRevE.86.016116>
- [8] M. E. Hochberg, “Non-linear transmission rates and the dynamics of infectious disease,” *Journal of theoretical biology*, vol. 153, no. 3, pp. 301–321, 1991.

- [9] U. Brandes, J. Pfeffer, and I. Mergel, *Studying Social Networks: A Guide to Empirical Research*. Campus Verlag, 2012.
- [10] M. Chiang, *Networked Life: 20 Questions and Answers*. Cambridge University Press, 2012.
- [11] I. Dobson, B. A. Carreras, V. E. Lynch, and D. E. Newman, “Complex systems analysis of series of blackouts: Cascading failure, critical points, and self-organization,” *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 17, no. 2, p. 026103, 2007.
- [12] D. J. Watts and S. H. Strogatz, “Collective dynamics of small-world networks,” *Nature*, vol. 393, no. 6684, pp. 440–442, 1998.
- [13] M. J. Keeling and K. T. Eames, “Networks and epidemic models,” *Journal of the Royal Society Interface*, vol. 2, no. 4, pp. 295–307, 2005.
- [14] F. Ball, “The threshold behaviour of epidemic models,” *Journal of Applied Probability*, pp. 227–241, 1983.
- [15] R. Pastor-Satorras and A. Vespignani, “Epidemic dynamics in finite size scale-free networks,” *Physical Review E*, vol. 65, no. 3, p. 035108, 2002.
- [16] J. Kunegis, “Konekt: the koblenz network collection,” in *Proceedings of the 22nd International Conference on World Wide Web Companion*. International World Wide Web Conferences Steering Committee, 2013, pp. 1343–1350.
- [17] J. Leskovec and A. Krevl, “SNAP Datasets: Stanford large network dataset collection,” <http://snap.stanford.edu/data>, Jun. 2014.
- [18] M. O. Jackson, *Social and Economic Networks*. Princeton University Press, 2008.
- [19] M. Newman, *Networks: an Introduction*. Oxford University Press, 2010.
- [20] V. E. Lee, N. Ruan, R. Jin, and C. Aggarwal, “A survey of algorithms for dense subgraph discovery,” in *Managing and Mining Graph Data*. Springer, 2010, pp. 303–336.
- [21] P. Csermely, A. London, L.-Y. Wu, and B. Uzzi, “Structure and dynamics of core/periphery networks,” *Journal of Complex Networks*, vol. 1, no. 2, pp. 93–123, 2013.

- [22] J. P. Gleeson, “Binary-state dynamics on complex networks: pair approximation and beyond,” *Physical Review X*, vol. 3, no. 2, p. 021004, 2013.
- [23] M. A. Porter and J. P. Gleeson, “Dynamical systems on networks: A tutorial,” *arXiv preprint arXiv:1403.7663*, 2014.
- [24] D. Chakrabarti, Y. Wang, C. Wang, J. Leskovec, and C. Faloutsos, “Epidemic thresholds in real networks,” *ACM Trans. Inf. Syst. Secur.*, vol. 10, no. 4, pp. 1:1–1:26, 2008.
- [25] A. Santos and J. M. Moura, “Emergent behavior in large scale networks,” in *50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*. IEEE, 2011, pp. 4485–4490.
- [26] Y. Moreno, M. Nekovee, and A. F. Pacheco, “Dynamics of rumor spreading in complex networks,” *Physical Review E*, vol. 69, no. 6, p. 066130, 2004.
- [27] C.-K. Chau, V. Pappas, K.-W. Lee, and A. Tantawi, “Self-organizing processes of collective behavior in computer networks,” 2009.
- [28] M. Draief, A. Ganesh, and L. Massoulié, “Thresholds for virus spread on networks,” in *Proceedings of the 1st International Conference on Performance Evaluation Methodologies and Tools*, ser. *valuetools '06*. New York, NY, USA: ACM, 2006. [Online]. Available: <http://doi.acm.org/10.1145/1190095.1190160>
- [29] A. Ganesh, L. Massoulié, and D. Towsley, “The effect of network topology on the spread of epidemics,” in *Proceedings of the Annual Joint Conference of the IEEE Computer and Communications Societies*, Miami, USA, Mar. 2005, pp. 1455–1466 vol. 2.
- [30] J. Abello, S. Butenko, P. M. Pardalos, and M. G. C. Resende, “Finding independent sets in a graph using continuous multivariable polynomial formulations,” *Journal of Global Optimization*, vol. 21, no. 2, pp. 111–137, 2001.
- [31] J. Zhang and J. M. F. Moura, “Accounting for topology in spreading contagion in non-complete networks,” in *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, March 2012, pp. 2681–2684.

- [32] N. Kashtan, S. Itzkovitz, R. Milo, and U. Alon, “Efficient sampling algorithm for estimating subgraph concentrations and detecting network motifs,” *Bioinformatics*, vol. 20, no. 11, pp. 1746–1758, 2004.
- [33] J. Zhang and J. M. F. Moura, “Threshold behavior of epidemics in regular networks,” in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2013, pp. 5411–5414.
- [34] —, “Diffusion in social networks as SIS epidemics: Beyond full mixing and complete graphs,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, no. 4, pp. 537–551, Aug 2014.
- [35] —, “Subgraph density and epidemics over networks,” in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2014, pp. 1125–1129.
- [36] —, “Role of subgraphs in epidemics over finite-size networks under the scaled SIS process,” *Journal of Complex Networks*, 2015.
- [37] —, “Dynamic bond percolation in networks,” in *Proceedings of IEEE Global Conference on Signal and Information Processing (GlobalSIP)*. IEEE, 2014.
- [38] —, “Contact process with exogenous infection and the scaled SIS process,” *submitted for publication*, 2015.
- [39] J. Mohammadi, J. Zhang, S. Kar, G. Hug, and J. M. F. Moura, “Multilevel distributed approach for DC optimal power flow,” in *Proceedings of IEEE Global Conference on Signal and Information Processing (GlobalSIP)*. IEEE, to be published 2015.
- [40] G. Papandreou and A. Yuille, “Perturb-and-map random fields: Using discrete optimization to learn and sample from energy models,” in *2011 IEEE International Conference on Computer Vision (ICCV)*, Nov 2011, pp. 193–200.
- [41] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*. MIT press, 2009.

- [42] D. B. West *et al.*, *Introduction to Graph Theory*. Prentice hall Upper Saddle River, 2001, vol. 2.
- [43] G. R. C. Godsil, *Algebraic Graph Theory*. Springer-Verlag, 2001.
- [44] S. Khuller and B. Saha, “On finding dense subgraphs,” in *Automata, Languages and Programming*. Springer, 2009, pp. 597–608.
- [45] S. Wasserman, *Social Network Analysis: Methods and Applications*. Cambridge University Press, 1994, vol. 8.
- [46] S. P. Borgatti, “The key player problem,” in *Dynamic Social Network Modeling and Analysis: Workshop Summary and Papers*. National Academies Press, 2003, p. 241.
- [47] H. Jeong, S. P. Mason, A.-L. Barabási, and Z. N. Oltvai, “Lethality and centrality in protein networks,” *Nature*, vol. 411, no. 6833, pp. 41–42, 2001.
- [48] S. Coulomb, M. Bauer, D. Bernard, and M.-C. Marsolier-Kergoat, “Gene essentiality and the topology of protein interaction networks,” *Proceedings of the Royal Society of London B: Biological Sciences*, vol. 272, no. 1573, pp. 1721–1725, 2005.
- [49] J. R. Norris, *Markov Chains*. Cambridge University Press, 1998.
- [50] F. P. Kelly, *Reversibility and Stochastic Networks*. Cambridge University Press, 2011.
- [51] A. K. Hartmann and M. Weigt, *Phase Transitions in Combinatorial Optimization Problems: Basics, Algorithms and Statistical Mechanics*. John Wiley & Sons, 2006.
- [52] Y. Okamoto, T. Uno, and R. Uehara, “Linear-time counting algorithms for independent sets in chordal graphs,” in *Graph-theoretic concepts in computer science*. Springer, 2005, pp. 433–444.
- [53] A. Billionnet and M. Minoux, “Maximizing a supermodular pseudoboolean function: A polynomial algorithm for supermodular cubic functions,” *Discrete Applied Mathematics*, vol. 12, no. 1, pp. 1–11, 1985.

- [54] E. Boros and P. L. Hammer, “Pseudo-boolean optimization,” *Discrete Applied Mathematics*, vol. 123, no. 1, pp. 155–225, 2002.
- [55] M. Grötschel, L. Lovász, and A. Schrijver, “The ellipsoid method and its consequences in combinatorial optimization,” *Combinatorica*, vol. 1, no. 2, pp. 169–197, 1981.
- [56] L. Lovász, “Submodular functions and convexity,” in *Mathematical Programming The State of the Art*. Springer, 1983, pp. 235–257.
- [57] V. Kolmogorov and R. Zabini, “What energy functions can be minimized via graph cuts?” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 2, pp. 147–159, 2004.
- [58] B. A. Cipra, “An introduction to the ising model,” *American Mathematical Monthly*, vol. 94, no. 10, pp. 937–959, 1987.
- [59] ———, “The Ising model is NP-complete,” *SIAM News*, vol. 33, no. 6, pp. 1–3, 2000.
- [60] M. R. Weeks, S. Clair, S. P. Borgatti, K. Radda, and J. J. Schensul, “Social networks of drug users in high-risk sites: finding the connections,” *AIDS and Behavior*, vol. 6, no. 2, pp. 193–206, 2002.
- [61] A. Krause, “SFO: A toolbox for submodular function optimization,” *The Journal of Machine Learning Research*, vol. 11, pp. 1141–1144, 2010.
- [62] G. Csardi and T. Nepusz, “The igraph software package for complex network research,” *InterJournal*, vol. Complex Systems, p. 1695, 2006. [Online]. Available: <http://igraph.sf.net>
- [63] S. P. Borgatti and M. G. Everett, “Models of core/periphery structures,” *Social Networks*, vol. 21, no. 4, pp. 375–395, 2000.
- [64] G. Strang, *Linear Algebra and its Applications*. Brooks/Cole, 1988.
- [65] T. Hazan and T. Jaakkola, “On the partition function and random maximum a-posteriori perturbations,” *arXiv preprint arXiv:1206.6410*, 2012.
- [66] M. Schmidt, “UGM: A matlab toolbox for probabilistic undirected graphical models,” 2010.

- [67] L. L. Scharf, *Statistical Signal Processing*. Addison-Wesley Reading, MA, 1991, vol. 98.
- [68] E. Estrada and J. A. Rodriguez-Velazquez, “Subgraph centrality in complex networks,” *Physical Review E*, vol. 71, no. 5, p. 056103, 2005.
- [69] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov Chains and Mixing Times*. American Mathematical Soc., 2009.
- [70] A. Papoulis and S. U. Pillai, *Probability, Random Variables, and Stochastic Processes*. Tata McGraw-Hill Education, 2002.
- [71] R. Christie. (1993, May) Power system test archive. [Online]. Available: https://www.ee.washington.edu/research/pstca/pf118/pg_tca118bus.htm
- [72] S. Micali and V. V. Vazirani, “An $O(v|v|c|e|)$ algorithm for finding maximum matching in general graphs,” in *Foundations of Computer Science, 1980., 21st Annual Symposium on*. IEEE, 1980, pp. 17–27.
- [73] N. J. Harvey, “Algebraic algorithms for matching and matroid problems,” *SIAM Journal on Computing*, vol. 39, no. 2, pp. 679–702, 2009.