University of California
Santa Barbara

A Theory of Collective Cell Migration and the Design of Stochastic Surveillance Strategies

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Mechanical Engineering

by

Mishel George

Committee in charge:

Professor Francesco Bullo (Advisor), Chair
Professor Otger Campàs (Advisor)
Professor Jeff Moehlis (Committee Member)
Professor Paul Atzberger (Committee Member)

June 2018

The Dissertation of Mishel George is approved.

Professor Otger Campàs (Advisor)

Professor Jeff Moehlis (Committee Member)

Professor Paul Atzberger (Committee Member)

Professor Francesco Bullo (Advisor), Committee Chair

March 2018

A Theory of Collective Cell Migration and
the Design of Stochastic Surveillance Strategies

by

Mishel George

To Pappa and Mummy.

Acknowledgements

First and foremost, I would like to thank both my advisors: Francesco Bullo and Otger Campàs for guiding me through my research at UCSB. Francesco's sharp wit and intuition has helped me find my way out of multiple dead ends. He has always been supportive of my choices and has given me the freedom to pursue exciting fundamental research. His efficiency, organization and unerring ability to catch typos is something I seek to emulate. He has been a wonderful mentor, both personally and professionally. The lessons learnt under his tutelage shall come in handy long after my time at UCSB.

Otger's infectious enthusiasm and energy for science spurred me to follow in his footsteps and aim for quality fundamental research. His emphasis on big-picture thinking as well as his guidance and motivation gave me inspiration to continue pursuing my problem even when I thought I had exhausted all possible avenues. Personally, he has been a very good friend and mentor. I am grateful for all the advice that he has given me over the years.

I would like to thank members of both the motion lab and the Campàs group. The exciting scientific conversations within the lab and the laid-back social outings have added tremendous value to my stay here. I have had a lot of fun collaborating with Rushabh, Saber and Xiaoming on various projects.

I would like to thank the several friends I have had the pleasure of meeting during my time at UCSB. I thank Weikang and Silvia for helping me acclimate during my early days at UCSB and for being nothing short of family ever since. I thank Eunhee and Brie for being excellent company and for being available as camping buddies at short-notice. I thank Ben and Anna for all the fun and the emotional support, and for introducing me to several essential American sports and activities. I thank Naveen and Nikita for the roadtrips and adventures that we have been on together. More generally, I would like

to thank all the friends I made at UCSB over the years: I believe that these friendships have been essential in helping me complete this Ph.D. In addition to people at UCSB, special thanks is in order to my friends from IIT, Varun and Abhishek. Their insightful advice helped me tackle several challenges along the way.

Finally, I would like to thank my Mom and my Dad for being the most wonderful parents that any one could ask for. I am eternally grateful for their persistence and meticulous effort which enabled me to pursue good education. Being given an opportunity to pursue a Ph.D. at UCSB has been a privilege and my parents deserve the majority of the credit for giving me this chance. This Ph.D. is a culmination of our combined efforts and I am glad to dedicate this thesis to my parents.

Curriculum Vitæ

Mishel George

Education

2018	**Ph.D. in Mechanical Engineering** University of California, Santa Barbara, Santa Barbara, CA, USA
2011	**B. Tech in Engineering Physics** Indian Institute of Technology Bombay Bombay, India.

Research Experience

2012-2018	**Graduate Student Researcher**, University of California, Santa Barbara. Investigated two independent research fronts: (1) the emergence of collective behavior in groups of cells and understanding of this behavior based on purely mechanical and contact-based interactions and (2) the design of unpredictable surveillance strategies for teams of robots which are robust to infiltration by intelligent intruders.
2011-2012	**Research Associate**, Indian Institute of Science, Bangalore. Designed and studied bio-inspired algorithms for formation control in groups of agents.

Teaching Experience

Winter 2018	**Teaching Assistant** Course: ME 6 *Electronic circuits* Mechanical Engineering, University of California, Santa Barbara.
Fall 2017	**Teaching Assistant** Course: ME 152A *Fluid Mechanics* Mechanical Engineering, University of California, Santa Barbara.
Spring 2017	**Teaching Assistant** Course: ME 179P *Robotics: Planning* Mechanical Engineering, University of California, Santa Barbara.
Summer 2015	**Teaching Assistant** Course: ENGR 3 *Introduction to Programming* College of Engineering, University of California, Santa Barbara.
Spring 2014	**Teaching Assistant** Course: ME 16 *Dynamics* Mechanical Engineering, University of California, Santa Barbara.

Summer 2013	**Teaching Assistant** Course: ME 179L *Robotics Design Lab* Mechanical Engineering, University of California, Santa Barbara.
Winter 2013	**Teaching Assistant** Course: ME 163 *Electronic circuits* Mechanical Engineering, University of California, Santa Barbara.
Fall 2012	**Teaching Assistant** Course: ME 106 *Mechatronics Lab* Mechanical Engineering, University of California, Santa Barbara.

Mentoring

Spring 2016	Jesus Castro Kevin Ton Asitha Kaduwela
Summer 2014	Dillon Azzam Nico Cristoffersen

Professional service

Reviewer	IEEE Transactions on Control of Networked Systems IEEE Transactions of Automatic Control IEEE Robotics and Automation Letters

Publications

- M. George and J. Saber and F. Bullo. Markov chains with Maximum Entropy for Robotic Surveillance. *IEEE Transactions on Automatic Control*. Note: Accepted.
- X. Duan and M. George and F. Bullo. Markov chains with Maximum Return Time Entropy for Robotic Surveillance. *IEEE Transactions on Automatic Control*. Note: Submitted.
- M. George and R. Patel and F. Bullo. The meeting time of Multiple Random Walks. *ESAIM: Probability & Statistics*. Note: Submitted.
- M. George and F. Bullo and O. Campàs. Connecting Individual to Collective Cell Migration. *Scientific Reports*. 7(9720), 2017.
- M. George and D. Ghose. Reducing Convergence Times of Self-Propelled Particles via Modified Nearest Neighbor Rules. *Physica A*. 391.16:4121-27, 2012.

Abstract

A Theory of Collective Cell Migration and
the Design of Stochastic Surveillance Strategies

by

Mishel George

In nature, complex emergent behavior arises in groups of biological entities often as a result of simple local interactions between neighbors in space or on a network. In such cases, scientific inquiry is typically aimed at inferring these local rules. Conversely, in teams of robots, the goal is to create decentralized control laws which results in efficient global behavior. These behaviors are designed for tasks such as maintaining formation control, performing effective coverage control or persistently monitoring an environment. With this in mind, we consider the following: 1> the emergence of collective cell migration from local contact and mechanical feedback and 2> the design of unpredictable surveillance strategies for teams of robots.

Collective cell migration is an essential part of tissue and organ morphogenesis during embryonic development, as well as of various disease processes, such as cancer. The vast majority of theoretical descriptions of collective cell behavior focus on large numbers of cells, but fail to accurately capture the dynamics of small groups of cells. Here we introduce a low-dimensional theoretical description that successfully describes single cell migration, cell collisions, collective dynamics in small groups of cells, and force propagation during sheet expansion, all within a common theoretical framework. We also explain the counter-intuitive observation that pairs of cells repel each other upon collision while they coordinate their motion in larger clusters.

Conventional monitoring strategies used by teams of robots are deterministic in nature

making it possible for intelligent intruders who study the motion of the patrolling agent to compromise the patrol route. This problem can be solved by designing random walkers on graphs which naturally incorporate unpredictability. Within this framework, we study and provide the first analytic expression for the first meeting time of multiple random walkers, in terms of their transition matrices. We also study two problems related to maximizing unpredictability: given graph and visit frequency constraints, 1> maximize the entropy rate generated by a Markov chain, and 2> maximize the return time entropy associated with the Markov chain, where the return time entropy is the weighted average over all graph nodes of the entropy of the first return times of the Markov chain.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Achieving the tight coordination observed in schools of fish [3] and flocks of birds [4] is possible in teams of robots due to advancements made in the field of multi-agent control [5, 6]. Groups of biological entities possess impressive dynamic capabilities based on purely localized interactions, and hence provide inspiration for decentralized decision-making schemes in robotic teams [7]. Even in the absence of knowledge concerning the global states of the system, biological networks have evolved local behaviors which translate robustly to global behaviors [8]. Theoretical descriptions of biological networks often try to demystify this emergence of global synchronized behavior by suggesting relatively simple local laws [9]. The objective in robotic teams when studied from this point of view is no different: design local control laws for each agent, wherein each agent possesses information limited to its own states and that of neighboring agents, that translate to a desired global behavior [10]. Thus, the study of emergent behavior in biological networks, while being of independent interest, also benefits the robotics community [11] and vice-versa the design of local control laws for robotic agents can shed light on how coordination in colonies of insects or groups of migrating animals might arise [12].

Collective motion in nature is observed across a breadth of length scales from clusters of macromolecules [13] to colonies of cells [14] and groups of humans [15]. Inspired by the

coordinated behavior observed in swarms of insects [16, 17], algorithms and heuristics designed for large teams of robots are called swarming algorithms or swarming behaviors, especially when the emphasis is on emergent behavior. While formation control is one of the many complex tasks that a group of robotic agents might need to achieve, robotic teams are designed with tasks in mind such as environmental monitoring [18], search and rescue [19] and intruder detection [20].

In this thesis, we study problems in coordinated biological and robotic behaviors: on one hand, we seek to unify collective behaviors observed across multiple cell types and on the other hand, we seek to design persistent surveillance strategies for groups of robots.

In the remainder of this chapter, a broad literature survey of each one of these fronts of research is described.

1.1 Literature relevant to collective cell migration

In this section, we present a literature survey of collective migration in groups of animal cells.

Collective migration in vivo: From embryonic development to tissue regeneration and wound healing, many processes of tissue (re)organization involve the coordinated migration of cells [21]. While some large scale migration processes involve the movements of hundreds of cells (e.g., neural crest cell migration [22]), many migratory events in developmental and disease processes involve small groups ($\sim$ 5-50) of cells [21, 23], including border cell migration [14] or lateral line formation [24]. Importantly, there is increasing evidence that cancer invasion and metastases rely on the migration of small clusters of cells rather than individual cells [25]. Despite the existing amount of information regarding the different migratory processes and their molecular control [26, 27, 28],

it is unclear how these different collective behaviors arise from the physical interactions among migrating cells, and how to connect the known individual behaviors of cells to their collective behavior in groups of different cell numbers.

Self-propelled particle models: When it comes to modeling collective behaviors in nature, two seminal pieces of work come to mind: the first widely-known flocking simulation published by Reynolds [29] and the statistical physics type approach by Vicsek et al. [9] which is now widely referred to as the "Vicsek Model". The standard Vicsek model consists of self-propelled particles (SPPs) with discrete-time evolution of their headings governed by the law that each agent computes an average of the headings of its neighbors with additive noise. Though the original objective of this model was to study second-order transition from order to disorder in the flocking behavior observed in the agents, this simple notion of agents trying to align velocities through local averaging laws spawned several models of collective motion in both biological [30, 31, 32] and physical systems [33, 34, 35].

Review of existing theoretical descriptions: Most experimental studies concerning the physical aspects of collective cellular movements have focused on the migration of thousands of cells, such as in *in vitro* wound healing assays [36, 37, 38, 39]. Accordingly, theoretical descriptions of these phenomena have been centered in the limit of very large numbers of cells, using both discrete approaches based on self-propelled particles (SPP) [40, 41, 42, 43, 44] and continuum theories [45, 43]. SPP and continuum phenomenological descriptions have provided important insights into the generic behaviors of collective cellular movements at length scales much larger than cell size and have been effective at providing scaling laws which describe such collective behaviors [45, 43].

1.2 Literature relevant to stochastic surveillance strategies

Here we present a broad literature review of surveillance strategies and Markov chains on graphs, which are essential tools in the design of unpredictable surveillance strategies. Topic-specific literature can be found in the chapters concerning stochastic surveillance strategies.

1.2.1 Multi-robot patrolling strategies

Some of the first work in multi-robot patrolling was done by Machado et al. [46], where several architectures for multi-agent patrolling are proposed and performance criteria specified. In particular, the architectures were split between centralized and decentralized strategies, and the delay between consecutive visits to nodes, termed the 'idleness', plays a key role in the design of both types of architectures. This notion of minimizing idleness was also studied using reinforcement learning methods in [47]. Interestingly, idleness has also been used to design a pheromone based swarm algorithm, where the evaporation of pheromones creates an oriented gradient following the chronology of cell visits [48]. Work by Elmaliach et al. [49] computes minimal cost cycles that visits all the points in the graph. Further, the proposed solution guarantees that each point is covered at the same optimal uniform frequency.

Stochastic vehicle routing strategies have the desirable property that an intruder can not predictably plan a path to avoid surveillance agents. The authors in [20, 50] use Markov chain Monte Carlo methods to design such surveillance strategies. A thorough study examining a range of strategies on the spectrum between purely random and purely deterministic is conducted in [51]. It was found that while deterministic strategies per-

form better against random attackers, introducing unpredictability increased the capture rate of intelligent attackers. Minimum hitting time Markov chains have been used in the design of stochastic surveillance strategies in [52] where a novel convex program formulation of the problem is considered. The notion of group hitting time for multiple random walkers is used in optimizing transition matrices for multiple agents in [53]. Further, Markov chains have been used in conjunction with specific notions of intelligent intruders to design stochastic strategies [54]. In [55] the mean hitting time in conjunction with multiple parallel instances of the CUSUM algorithm is used to devise a policy which ensures quickest average time to detection of anomalies. Finally, the work in [56] formulates an efficient algorithm based on Markov chains named PATROLGRAPH* which allows for effective extension to the multi-agent case.

1.2.2 Random walks on graphs

In the setup we consider for robotic surveillance, we model the environment as a graph and design random walks on this graph. Hence, we briefly review metrics related to random walks on graphs that are relevant to the design of robotic strategies, namely, speed of traversal and unpredictability in path. A relevant notion of speed of traversal of random walks on graphs is the mean hitting time which is the average time taken by a single random walker to travel between nodes of a graph. The hitting times of a finite irreducible Markov chain first appeared in [57], however, it was rediscovered for finite reversible Markov chains in [58]. Several bounds have been obtained for the hitting time for various graph topologies [59, 60]. Many closed-form formulations exist to compute the hitting time for a single random walker [59, 61, 52]. Another related notion that measures speed of traversal is that of the cover time which is the expected time taken by random walkers to hit every node on the graph [58, 62]. Markov chains are excellent tools for

introducing unpredictability into the motion of surveillance agents. A standard notion of unpredictability for discrete-time Markov chains is the entropy generated at each timestep called the entropy rate of a Markov chain [63]. Prior work has studied maximizing the entropy of discrete-time Markovian evolutions in various contexts [64, 65, 66]. Ekroot *et al.* studied the entropy of Markov trajectories in [67], i.e., the entropy of paths with specified initial and final states.

While the design of surveillance algorithms studied in this thesis are not directly inspired by biological behavior, the tools utilized in the design of these strategies, i.e., Markov chains on graphs, are used in the design of emergent behaviors in robotic swarms. More specifically, Markov chains on graphs have been utilized in the design of guidance schemes for large swarms [68, 69, 70] and finding solutions to task assignment problems in heterogeneous swarms of robots [71]. More broadly, random walks on networks appear in many areas of research: they are used to describe effective resistance in electrical networks [72, 73], for link-prediction and information propagation in social networks [74, 75], and in designing search algorithms on networks [76, 77].

1.3 Contribution and organization

There are several contributions of this thesis. In what follows, we detail the contributions of each chapter.

Chapter 2: We introduce a theoretical description that successfully describes the motion of groups of cells of arbitrary numbers, from single cell motion to the collective migration of small groups of cells and sheet migrations. First, we derive a novel theoretical description of the collective dynamics of groups of cells by balancing the forces in the system and specifying the dynamics of traction forces (or cell polarization) for individual cells, accounting for both contact inhibition of locomotion and force-induced

repolarization. Second, we show that small groups of cells (3 or more cells) display coherent collective behavior despite their effective repulsion during the collision of cells pairs, with persistence times that depend on the group size. Third, we find an optimal size for small groups of cells that maximizes the persistence of their coherent motion; this quantity depends on cellular adhesion and strength of traction forces generated. Fourth, we show that groups of identical cells can display coherent collective behavior or dispersal behavior by changing their confinement.

The theoretical description presented in the chapter, albeit 1-D in nature, is the first of its kind to successfully capture the motion of cells from an individual entity through to small groups and colonies. This description is obtained from understanding:(1) the essential dynamics of single cell motion, i.e., random diffusive motion due to the generation of sporadic ruffling protrusions which decay with a well-defined time-scale, and, (2) the essential interactions between cells, i.e, the fact that protrusions collapse on contact with other cells and protrusions are generated in response to pulling forces with the same characteristic time-scale and an additional force-scale. With this in mind, the chapter also details how experimental inquiry could be structured to test other predictions borne of this theoretical description. The most exciting contribution of this chapter is insight into a long-standing paradox with the cell migration community: how do groups of cells migrate collectively when experiments on pairs of cells indicate that cells repel each other upon contact?

In addition, we make a set of predictions: for single cells, the diffusion constant of cellular movements quadriatically depend on the cell's traction force; for pairs of cells, the dynamics of contact-based repolarization are strongly based on the ratio of adhesion strength to traction force and the ratio of traction repolarization time to mechanical relaxation time.

Chapter 3: There are several key contributions in this chapter. First, we provide

a set of necessary and sufficient conditions which characterize when the meeting times between a single pursuer and a single evader is finite for two arbitrary random walks modeled as Markov chains. To the best of our knowledge the bounds in the literature were obtained for meeting times between ergodic Markov chains where the meeting times are guaranteed to be finite. However we extend the notion to generic transition matrices as opposed to equal neighbor models which are studied in many works, and we discuss at length when the meeting times are finite based on the existence of walks of equal length to common nodes. Second, we provide a closed-form solution to the meeting time of two independent Markov chains by utilizing the Kronecker product of the transition matrices. Both these results are obtained using a technical approach which takes advantage of the properties of Kronecker products of graphs. We further use this closed-form expression to perform comparisons with existing bounds in the literature. Indeed we see that the bounds are very conservative for most graphs. Third, we provide a set of sufficient conditions in terms of the absorbing classes of the pursuer and evader chains which guarantee finite meeting times. Fourth, we extend the treatment to multiple pursuers and multiple evaders. Finally, we obtain conditions for the meeting times between two continuous-time Markov chains to be finite and provide closed-form results for this case, and further extend it to multiple pursuer and evader groups when dictated by multiple transition rate matrices.

To the best of our knowledge, this chapter provides the first closed-form solutions for the computation of the meeting time between two Markov chains for both discrete-time and continuous-time time indices. Two closely related references are as follows: first, a system of equations for computing meeting times for independent identical random walks on graphs with irreducible transition matrices, where the transition matrices are limited to equal-neighbor weights, were obtained using Laplace transform techniques in [78]. Second, Kronecker products and vectorization techniques have been used to compute the

Simrank of information networks which has interpretations in terms of meeting times [79]. Our work is different in several ways. First, we consider absolutely generic transition matrices which need not be identical. Second, we present expressions here which are valid for reducible transition matrices. Third, we present meeting time expressions for the case of multiple pursuers which would correspond to multiple infecting particles in [78]. Finally, we provide insight into when meeting times are finite by connecting this notion to the existence of walks on the Kronecker graph.

Chapter 4: This chapter makes contribution to Markov chain theory as well as to robotic surveillance. First, we show that the novel problem of entropy rate maximization subject to the fact that it is a random walk on a graph with n nodes and constraints on the visit frequency to each node is well-defined and is strictly convex. We show that the unique global solution is indeed an irreducible Markov chain. The irreducibility property implies that the solution has a well-defined stationary distribution identical to that posed in the stationary distribution constraint.

Second, as the main contribution of the chapter, we provide an iterative algorithm with rigorous convergence guarantees to compute an n-dimensional vector, called the so-called *maxentropic vector*. In turn, as a function of this maxentropic vector, we provide a closed-form formula for the maximum entropy rate Markov chain, referred to as the *maxentropic Markov chain with visit frequency constraints*. In other words, we compute maxentropic chains with arbitrary stationary distributions on a graph with n nodes using an n-dimensional vector instead of optimizing transition matrices in $\mathbb{R}^{n\times n}$.

Third, we establish various additional results, including (i) the reversibility of maxentropic Markov chains with prescribed stationary distributions, (ii) a formula for the maximum entropy rate subject to the constraints, and (iii) an equiprobable path property, which, prior to this work, was only known to hold for the maximal entropy random walk [65]. Additionally, for a few special choices of the constraints, we are able to char-

acterize interesting special cases. For example, we show that the equal neighbor random walk on a graph is equal to the maxentropic Markov chain with visit frequency at each node proportional to the degree of the node.

Fourth, we conduct a careful comparison between our proposed procedure and standard SDP methods across a range of graph topologies. Specifically, we conduct a worst-case complexity analysis of our procedure and compare it with interior point methods used to solve semidefinite programming formulations of the entropy rate maximization problem. Empirically and analytically, we show that our proposed procedure has significantly lower runtime than an SDP method to solve the optimization problem.

Finally, we demonstrate some example realizations of these maxentropic chains in robotic scenarios. A key simulation-based result is that maxentropic Markov chains perform better than minimum hitting time Markov chains for the important case of so-called intelligent waiting intruders with short attack durations. We also conduct simulations on a partitioned graph with multiple surveillance agents and find that this result appears to hold for the multi-agent case as well.

Chapter 5: In this chapter, we propose a new metric that measures the unpredictability of the Markov chains over a directed graph with travel times. This novel formulation is of interest in the general study of Markov chains as well as for its applications to robotic surveillance. The main contributions of this chapter are sixfold.

First, we introduce and analyze a discrete-time delayed linear system for the return time probabilities of the Markov chains. This system incorporates integer-valued travel times on the directed graph. Second, we propose to characterize the unpredictability of a Markov chain by the return time entropy and formulate an entropy maximization problem. Third, we prove the well-posedness of the return time entropy maximization problem, i.e., the objective function is continuous over a compact set and thus admits a global maximum. For the case of unitary travel times, we derive an upper bound for the

return time entropy and solve the problem analytically for the complete graph. Fourth, we compare the return time entropy with the entropy rate of Markov chains; specifically, we prove that the return time entropy is lower bounded by the entropy rate and upper bounded by the number of nodes times the entropy rate. Fifth, in order to compute Markov chains with maximum return time entropy numerically, we truncate the return time entropy and show that the truncated entropy is asymptotically equivalent to both the original objective and the practically useful conditional return time entropy. We also characterize the gradient of the truncated return time entropy and use it to implement a gradient projection method. Sixth, we apply our solution to different prototypical robotic surveillance scenarios and test cases and show that, for a model of rational intruder, the Markov chain with maximum return time entropy outperforms several existing Markov chains.

Each chapter includes an organization subsection which details the structure. As one might expect the mathematical notation used in describing the theory of cell migration is different from that of designing surveillance strategies. In the chapters on stochastic surveillance strategies (chapters 3, 4, 5), we use consistent mathematical notation which shall be built up over the course of those chapters.

Chapter 2

Connecting Single Cell to Collective Cell Migration

2.1 Introduction

2.1.1 Problem motivation

We seek a minimal theoretical description accounting for key phenomenological observations regarding cell-cell interactions. To this end, we first describe literature which specifically deals with cell-cell contact and single out two crucial phenomenon which are observed in multiple experiments. We then describe cells as particles and introduce pairwise physical interactions accounting for these phenomenon in a one-dimensional setting (Fig. 2.1A). While minimal, the 1D geometry has proven very useful to study collective cell migration at the experimental level [80, 1, 81], as it simplifies the system considerably while preserving the essential features of collective cell migration.

2.1.2 Relevant cell-cell interactions

During cell-cell contact, individual cells show very characteristic behaviors. Studies on the kinematics and physical interactions between two colliding cells have revealed that

cells retract their lamellipodium upon frontal contact with another cell, a phenomenon known as *Contact Inhibition of Locomotion* (CIL) [82, 83, 22, 84]. Studies of CIL have shown that cell pairs display an effective repulsion upon collision [84, 1, 80, 83] that is at odds with known coherent collective behavior of groups of cells both *in vitro* and *in vivo* [21, 23]. Recent experiments have shown that cells can display both CIL in collisions between cell pairs and the formation of coherently moving cells when in larger groups [1], suggesting that the same underlying mechanism of physical interaction between cells can give rise to both behaviors. In addition to CIL, recent *in vitro* studies indicate that cells repolarize away from pulling forces transmitted through cadherin-mediated cell adhesion and stabilize a lamellipodium in the opposite direction to the externally applied force [85, 86]. This *Force-Induced Repolarization* (FIR) establishes a mechanical feedback of cadherin-dependent adhesion forces from neighboring cells on the dynamics of cell polarization and traction forces. Both CIL and FIR play a major role in collective cell migration [87, 88, 83], as they couple cellular spatial configurations to the dynamics of cell traction forces via cell-cell contacts.

Discrete SPP models inspired by flocking or schooling behavior of animal groups can reproduce coherent collective cell behavior through local velocity alignment rules [42, 43]. These models have been shown to successfully reproduce important features of large scale collective cell behavior, but do not explain important features of the dynamics of small groups of cells in which the specific characteristics of cellular interactions, including behaviors such as CIL or FIR, may play an important role. In general, SPP models can be used to describe the dynamics of small groups of cells and study the effects of important cell behaviors and parameters. Indeed, models of SPP have started to explore the role of CIL in the collective dynamics of cells in 2D, but either focus on large 2D monolayers or do not account for FIR [89, 90, 91]. It remains unclear how cell behaviors such as CIL and FIR contribute to collective cell migration, especially for small groups

of cells, such as those observed in developing embryos or during cancer metastasis.

2.1.3 Organization

The rest of this chapter is organized as follows. In section 2.2, we derive a set of stochastic differential equations which model the behavior of cells in 1D. In section 2.3, we simulate the system for groups of varying sizes and describe the observed results. In section 2.4, we describe numerical methods and statistical quantities used to characterize the behavior of groups of cells. In section 2.5, we compare the results obtained from simulations to existing experimental observations. Finally, in section 2.6 we summarize the findings of this chapter.

2.2 Theoretical Description

We introduce one of the main contributions of the chapter which is a set of stochastic differential equations describing the dynamics of motion of individual cells as well as their contact-based interactions.

2.2.1 Particle-based description of single cell movements

In order to control their movements, cells regulate the forces they apply on their surroundings. A given cell generates a traction force $\vec{T}$ that causes its movement. Both dissipative processes inside the cell and friction with the substrate lead to a friction force opposing the cell movement which, in its most basic form, reads $-\xi\vec{v}$, with ξ being an effective friction coefficient and $\vec{v}$ the cell velocity. For the specific case of a single cell, it is instructive to consider also the effect of an external force $\vec{F}_{ext}$, as previously done experimentally by applying a controlled force with optical or magnetic tweezers [85, 86].

Neglecting inertial terms, force balance on the cell reads

$$\xi \vec{v} = \vec{T} + \vec{F}_{ext}\,, \tag{2.1}$$

and specifies the cell velocity $\vec{v}$ that results from the forces in the system. While the external force in Eq. 2.1 is given and fixed, traction forces are generated by lamellipodial protrusions and therefore controlled by their dynamics. The direction of lamellipodial extension, and consequently the direction of the traction force $\vec{T}$, depends on the direction of cell polarization as dictated by the intracellular localization of polarization factors such as RhoA, Cdc42 and Rac [26, 92]. In the absence of instructive external cues (biochemical or mechanical), cells constantly produce lamellipodial ruffles in random directions [92, 93, 94] that decay over a timescale τ_T (protrusion lifetime), which characterizes the persistence of traction along a specific spatial direction (Fig. 2.1B). The timescale τ_T accounts here for the time necessary to repolarize the cell at a molecular level (i.e., changing the molecular polarity of the cell) and physically (rebuilding the lamellipodium), and is therefore associated with the cell (traction) repolarization time. Accounting for FIR due to an externally applied force [85, 86], the dynamics of traction forces can be written as

$$\tau_T \frac{d\vec{T}}{dt} = -\vec{T} - T_M \hat{F}_{ext} + T_R\,\hat{\eta}\,, \tag{2.2}$$

where $\hat{F}_{ext} \equiv \vec{F}_{ext}/|\vec{F}_{ext}|$ is the direction of the applied external force, T_R is the characteristic force scale of individual lamellipodial ruffling and $\hat{\eta}$ is a random unit vector, denoting a delta-correlated white noise with unit variance, namely $< \eta_i(t)\eta_j(t') >= \delta_{ij}\delta(t-t')$. The force scale T_M represents the maximal force that a lamellipodium stabilized by the presence of an external cue can generate (Fig. 2.1B). While we do not consider the effect of external biochemical cues in this study, including them is straightforward.

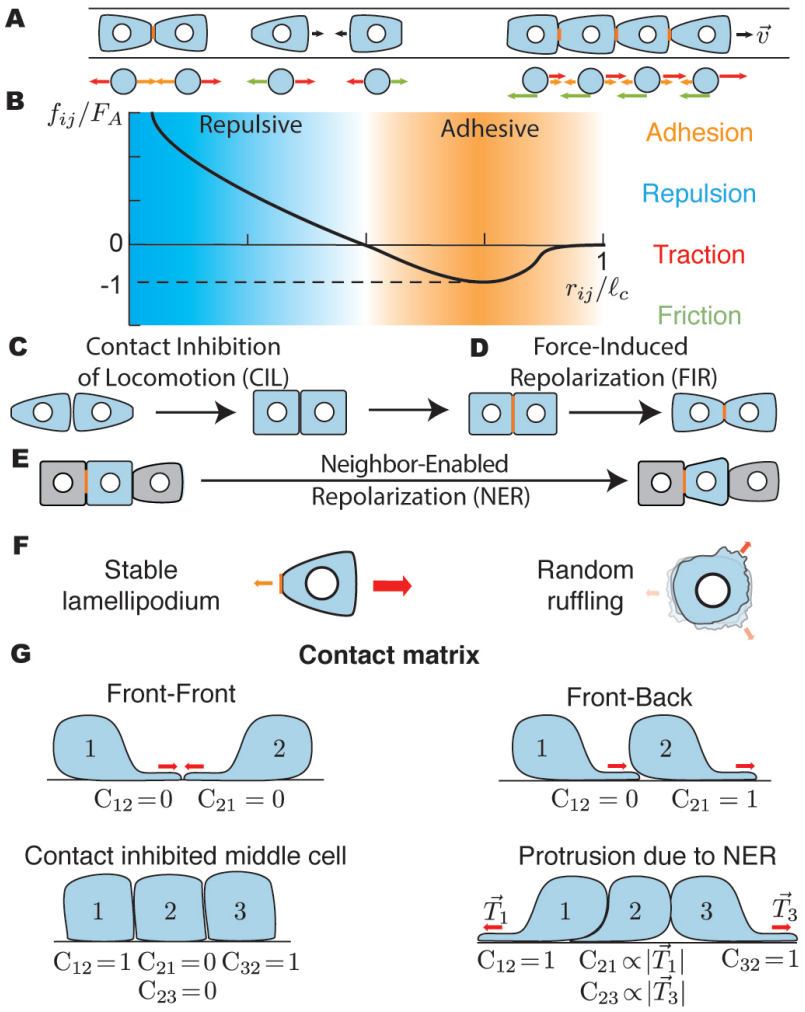

Figure 2.1: **Description of the system, interaction forces and phenomenological cell behaviors.** (A) Schematic representation of cells moving along a 1D strip (top) and particle-based representation of the system (bottom). Cells can be subject to adhesion forces (orange), excluded volume repulsion forces (blue) and friction forces (green), as well as generate traction forces (red). (B) Pairwise interaction forces f_{ij} between cells as a function of their relative distance. Schematic representation of CIL (C) and FIR (D), leading to an effective repulsion between cells. (E) Schematic representation of neighbor-enabled repolarization (NER). (F) Schematic representation of lamellipodial ruffling (right) and a stable lamellipodium (left). (G) Schematic representation of cellular configurations during collisions and the associated values of the contact matrix C_{ij} for each configuration and cell.

2.2.2 Systems with multiple cells

In a system with N cells (from $N = 2$ to $N \to \infty$), cells apply forces on each other that affect their dynamics at different levels. Considering the forces that cells apply on each other, force balance on cell i reads

$$\xi \vec{v}_i = \vec{T}_i + \sum_{j \neq i} \vec{f}_{ji} \,, \tag{2.3}$$

where $\vec{f}_{ji} = F_A \, f(r_{ji}) \hat{r}_{ji}$ is the force that cell j applies on cell i. From the perspective of cell i, $\vec{f}_{ji}$ is thus an external force along the direction $\hat{r}_{ji} = (\vec{r}_j - \vec{r}_i)/(|\vec{r}_j - \vec{r}_i|)$, where $\vec{r}_i$ and $\vec{r}_j$ are the cells' positions. In contrast to the constant external force considered above for the one cell case, the magnitude of the intracellular forces changes with the cells' configuration and we assume it depends only on the distance $r_{ji} = |\vec{r}_j - \vec{r}_i|$ between the cells. More specifically, it is characterized by a repulsive region, accounting for volume exclusion, and an attractive part, accounting for cell adhesion, with a attractive force F_A specifying the adhesion strength of the contact between two cells (Fig. 2.1C and Methods). We account for the finite size of the cell ℓ_c by setting a cutoff in the pairwise interaction force between cells at $r_{ji} = \ell_c$ ($f(r) = 0$ if $r > \ell_c$) that prevents cell interactions if separated by more than the cell size (Methods).

Beyond the direct effect that forces from neighboring cells have on the motion of a given cell (Eq. 2.3), these forces also act as cues for cell repolarization and, as a consequence, affect the dynamics of the traction force exerted by the cell. As described above for a single cell, we account for FIR in the dynamics of traction forces, namely

$$\tau_T \frac{d\vec{T}_i}{dt} = -\vec{T}_i - T_M \sum_{j \neq i} \mathrm{C}_{ij} \, \Theta \left[f(r_{ji}) \right] \hat{r}_{ji} + T_R \, \hat{\eta}_i \,, \tag{2.4}$$

where $\Theta(\cdot)$ is the Heaviside function and allows only pulling forces to cause FIR. In

addition to FIR, it is necessary to account for the effect of CIL (and other contact or exclusion effects) on traction forces when cells come into contact (Fig. 2.1D-F). We phenomenologically account for these processes using a contact matrix C_{ij} which we describe in details below.

When two cells collide, the observed lamellipodial retraction characteristic of CIL can be mathematically accounted for by expressing the contact matrix C_{ij} as $C_{ij} = (1 - \hat{r}_{ji} \cdot \hat{T}_i)/2$ (Fig. 2.1G). In 1D, C_{ij} is simply a Boolean matrix with zero values for configurations in which the lamellipodium frontally contacts the other cell, leading to lamellipodial retraction, and a value of one otherwise, allowing the formation of the lamellipodium (Fig. 2.1D,G and Eq. 2.4). Importantly, while we phenomenologically account for the observed retraction of the lamellipodium upon collision, we do not impose a repolarization of the lamellipodium away from the contact. We find that this repolarization, commonly associated with CIL [83, 22], occurs naturally within our description as a consequence of FIR (Fig. 2.1D,E), which causes tractions to repolarize away from pulling forces established between cells upon collision, as suggested in recent experiments [87, 88].

When $N > 2$, some cells may be contacted on all sides by other cells (Fig. 2.1A,E) and, according to CIL, these cells would not be able to generate any stable lamellipodium. However, experimental data from 1D cell clusters and 2D wound healing experiments suggests that cells contacted on all sides can generate cryptic (stable) lamellipodia [1, 95, 96]. In wound healing experiments, cells just behind the wound edge (second cell layer) generate stable lamellipodia in the same direction as that of already polarized cells at the leading edge. 3D imaging of cells in such expanding monolayers suggests that upon polarization, cells undergo shape changes that open spaces at their rear end, enabling neighboring trailing cells to protrude lamellipodia [96]. This effect, which we call *Neighbor-Enabled Repolarization* (NER), does not specify the direction of cell repolarization. It instead permits a cell i to protrude a cryptic lamellipodium if the neighboring cell j is polarized

away from cell i (Fig. 2.1E,G). While NER can be simply due to shape changes upon cell polarization, other biochemical mechanisms can effectively generate the same effect, as recently proposed [95, 88, 97]. We mathematically account for NER and CIL in the contact matrix C_{ij} (Fig. 2.1G) which, for 1D systems with arbitrary number of cells, can be written as

$$\mathrm{C}_{ij} = \frac{1 - \hat{r}_{ji} \cdot \hat{T}_i}{2} \left[1 + \sum_{k \neq j, k \neq i} \left[\left(\frac{1 - \hat{r}_{ji} \cdot \hat{T}_k}{2} \right) \frac{|\vec{T}_k|}{T_M} - 1 \right] \right] . \tag{2.5}$$

The movement of each cell in a system with N cells is governed by Eqs. 2.3, 2.4 and 2.5. Combining these equations and normalizing lengths with the cell size ℓ_c, forces with the adhesion force scale F_A, and time with the timescale $\tau_M = \xi \ell_c / \sigma$ associated with mechanical relaxation, we obtain three dimensionless parameters that control the dynamical regimes of the system, namely T_M/F_A, T_R/F_A and τ_T/τ_M. The parameters T_M/F_A and T_R/F_A compare the relative strengths of traction forces generated by stable lamellipodia and lamellipodial ruffling to adhesion forces, respectively. Finally, the ratio τ_T/τ_M compares the repolarization timescale τ_T to the time scale τ_M that a cell requires to reach mechanical equilibrium.

2.3 Simulation Results

We study the cellular movements in systems of N cells by numerically solving Eqs. 2.3, 2.4 and 2.5 (Methods), as analytical solutions are difficult to obtain due to the highly non-linear nature of the dynamics.

2.3.1 Single cell movements

In the absence of any external cues ($\vec{F}_{ext} = 0$), the cellular movements resulting from integrating Eqs. 2.1 and 2.2 are ballistic at short time scales ($t < \tau_T$), with average velocity T_R/ξ, and diffusive at timescales longer than the traction persistence time scale τ_T, with diffusion constant D given by $D = T_R^2\tau_T/\xi^2$. The timescale of velocity autocorrelation decay is τ_T, indicating that τ_T is indeed the persistence timescale of cellular motion. In the presence of an external pulling force (mechanical cue; $\vec{F}_{ext} \neq 0$), the cell polarizes away from the pulling force and, at time scales longer than τ_T, it generates a traction force $-T_M\,\hat{F}_{ext}$ opposing the external force (Eq. 2.2). Force balance (Eq. 2.1) shows that the average velocity of the cell is $\vec{v} = -(T_M/\xi)(1 - |\vec{F}_{ext}|/T_M)\hat{F}_{ext}$, indicating that the cell moves away from the pulling force at a speed that decreases linearly with the applied pulling force, with a maximal velocity T_M/ξ and a stall force T_M, analogous to molecular motors.

2.3.2 Collisions between two cells (N = 2)

Most cell-cell collision experiments measure the repolarization probabilities of two colliding cells at a fixed time after collision and for all possible initial cell-cell configurations before collision, namely front-front (F-F) and front-back (F-B) collisions [1, 80] (Fig. 2.2). Simulations of cell collisions indicate that cell repolarization is always faster in F-F collisions for any value of the different parameters in the system (Fig. 2.2A,B). In F-B collisions, the trailing cell (**F**-B) engages in a frontal collision with the leading cell (F-**B**), which is contacted at its back end, and always repolarizes faster than the leading cell, as observed experimentally [1]. When the force of adhesion is larger than the forces produced by stable lamellipodia ($F_A > T_M$), cells remain attached to each other after collision (Fig. 2.2E), with traction forces oriented away from each other (Fig. 2.2A,C). In

contrast, when traction forces are larger than adhesion forces ($T_M > F_A$), cells separate shortly after collision and move away from each other (Fig. 2.2B,D), with separation times being similar to the repolarization time scale τ_T (Fig. 2.2E). These configuration-dependent behaviors arise from the combined action of CIL and FIR, which depend on the mechanical state of each cell configuration.

Comparing published experimental data of repolarization probabilities in 1D collisions between NRK-52E cell pairs [1] to our theoretical predictions (Fig. 2.2F and Methods), we find that the minimal discrepancy is obtained for $T_M/T_R \simeq 2-3$ and $T_A/T_R \simeq 0.8-1.1$, indicating that NRK-52E cells generate stable traction forces T_M two to three times larger than adhesion forces F_A and that ruffling forces (T_R) alone are strong enough to separate the cells ($T_A/T_R \simeq 0.8-1.1$).

2.3.3 Small groups of cells ($2 < N \sim 10$)

To characterize the collective behavior of small groups of cells (or cell trains), we first simulate compact groups of identically polarized cells and study their persistence. When traction forces are larger than adhesive forces ($T_M > F_A$), the initially coherent train starts losing its persistence over a timescale τ_T, with cells at the trailing end repolarizing and detaching from the train (Fig. 2.3A). In contrast, when cell-cell adhesion is larger than traction ($F_A > T_M$), coherent cell trains with persistent average cell polarization exist (Fig. 2.3B) over timescales that depend on the number of cells in the train (Fig. 2.3C). We observe an optimal train size for each ratio T_M/F_A that maximizes the persistence time τ_p of the train, which can become orders of magnitude larger than τ_T (Fig. 2.3C). This optimal train size increases for increasing adhesion strength relative to the cell traction forces (Fig. 2.3D). Despite the existence of CIL, persistent trains with coherent polarization can exist because of NER. Importantly, the trailing cell in the train

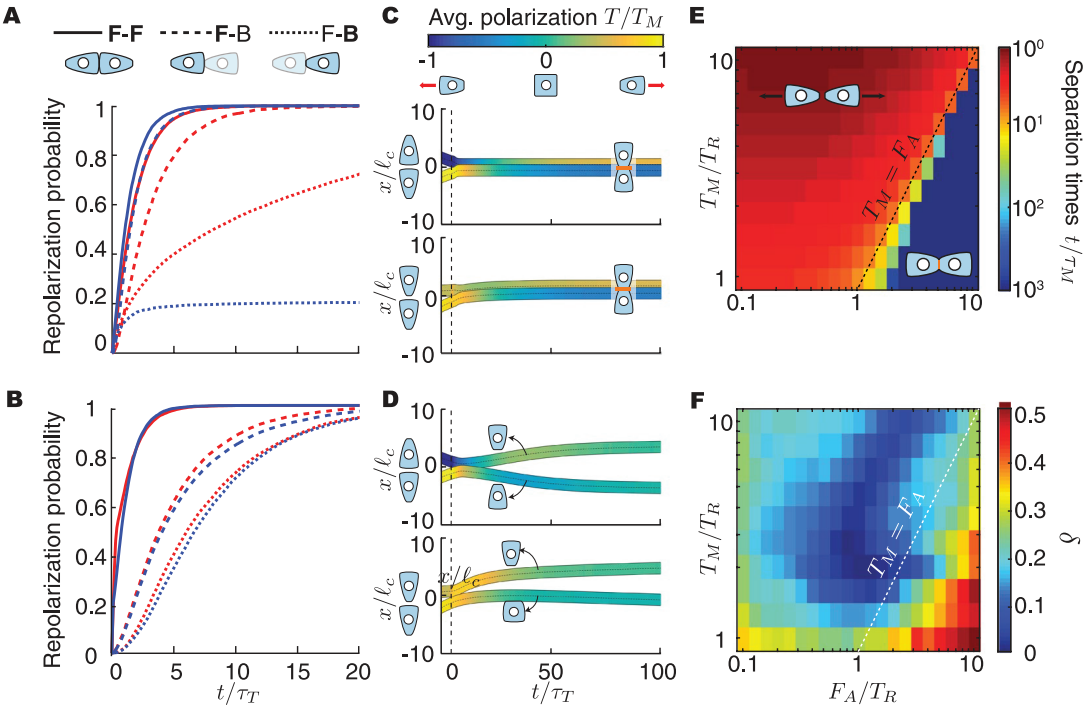

Figure 2.2: **Collision dynamics between two cells.** Cell repolarization probabilities (A,B) and average trajectories (C,D) of colliding cells after F-F (triangles) and F-B collisions for both trailing (**F**-B, squares) and leading (F-**B**, circles) cells, and high (A,C; $T_M/T_R = 1, F_A/T_R = 10$) and low (B,D; $T_M/T_R = 10, F_A/T_R = 1$) adhesion levels. Red and blue in (A,C) lines correspond to $\tau_T/\tau_M = 1$ and $\tau_T/\tau_M = 10$ respectively. Color code in (C,D) shows ensemble average of cell polarization during collision. Width of trajectory represents cell size ℓ_c. The value of $\tau_T/\tau_M = 1$ in panels (C,D). (E) Cell separation time (normalized to τ_M) for the different parameters in the problem. Cell separation times increase sharply, indicating that cells essentially remain attached, when $F_A > T_M$. (F) Comparison of theoretical predictions to experimental data in [1]. The measured discrepancy δ (Methods) between the experimental data and the theoretical predictions is shown (color coded) for varying values of T_M/T_R and F_A/T_R. Minimal values of discrepancy were found for $T_M/T_R \simeq 2-3$ and $F_A/T_R \simeq 0.8-1.1$. The value of $\tau_T/\tau_M = 1$ in panels (E,F).

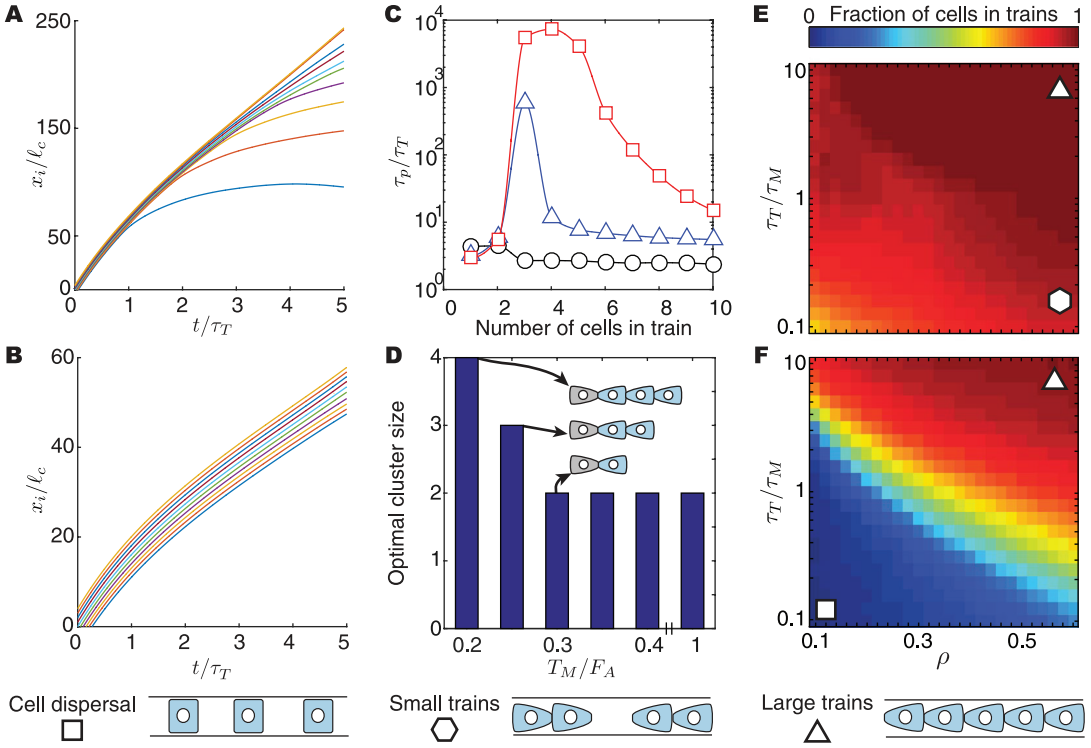

Figure 2.3: **Persistence and dynamics of cell trains.** (A, B) Position kymograph of 10-cell trains for low (A; $T_M/F_A = 0.2$) and high (B; $T_M/F_A = 0.8$) adhesion levels. (C) Average persistence time of trains for varying adhesion levels: $T_M/F_A = 0.2, 0.25, 0.8$ (red, blue and black, respectively). (D) Optimal train (cluster) size as a function of T_M/F_A. The value of $\tau_T/\tau_M = 10$ and $F_A/T_R = 10$ in panels A,B,C,D. (E, F) Dynamics of train formation as a function of density (confinement) and τ_T/τ_M for high (E; $T_M/F_A = 0.5, F_A/T_R = 10$) and low (F; $T_M/F_A = 2, F_A/T_R = 10$) adhesion levels.

always repolarizes away from the average train polarization because of CIL and is dragged forward by the collective train motion (Fig. 2.3D), as observed experimentally [1].

Beyond small groups of fixed number of cells, we study the collective behavior of cells moving along a 1D strip with periodic boundary conditions (ring geometry). In this case, the behavior of the system depends on the average cell density $\rho \equiv N\ell_c/L$ (with L being the perimeter of the ring), which parameterizes cellular confinement. If the adhesion strength between cells is much larger than the traction forces exerted by stably polarized cells ($F_A > T_M$), then either one or several groups of cells that move coherently

dominate the system for almost any value of initial density (Fig. 2.3E). In contrast, for small values of cell adhesion strength ($F_A < T_M$), we find cell dispersal behavior at low densities, with cells covering the entire length of the track and maximizing their average distance from each other (Fig. 2.3F), a result that could explain cell dispersal behaviors observed *in vivo* [98, 99]. Even in these low adhesion conditions, cells can form coherent trains at large densities. These trains are dynamic structures, with cells being added and removed from the train, but keeping a finite size. This is because at large densities the typical time scale of adding a new cell to a train can be shorter than the time scale τ_T for cells to repolarize and separate from the train. The transition between dispersal and coherent train formation occurs by solely changing the cell density, even if no cell parameters (traction, adhesion, polarization time, etc.) are changed. This indicates that a given cell type can display both dispersal behavior and coherent train formation at different densities (confinement conditions), as suggested in recent experiments [100].

2.3.4 Large cell colonies ($N \gg 10$)

We study large colonies of strongly adhesive cells ($F_A > T_M$) in 1D, as this situation mirrors sheet migration in 2D wound healing experiments. Cells are initialized in a configuration where they are attached to each other and have random polarizations. In all cases, cells at the edge develop polarizations away from the colony and start pulling on it. A polarization wave that propagates from the edge to the interior of the cell colony transfers the forces generated at the edge to cells deep in the colony (Fig. 2.4A). If the buildup of intercellular forces within the colony exceeds the maximal adhesion force between cells, the colony breaks, with the highest probability of breakage occurring where the intercellular forces are maximal on average (Fig. 2.4B). The possibility of colony breakage occurs because cells inside the colony can develop cryptic lamellipodia, contributing to a

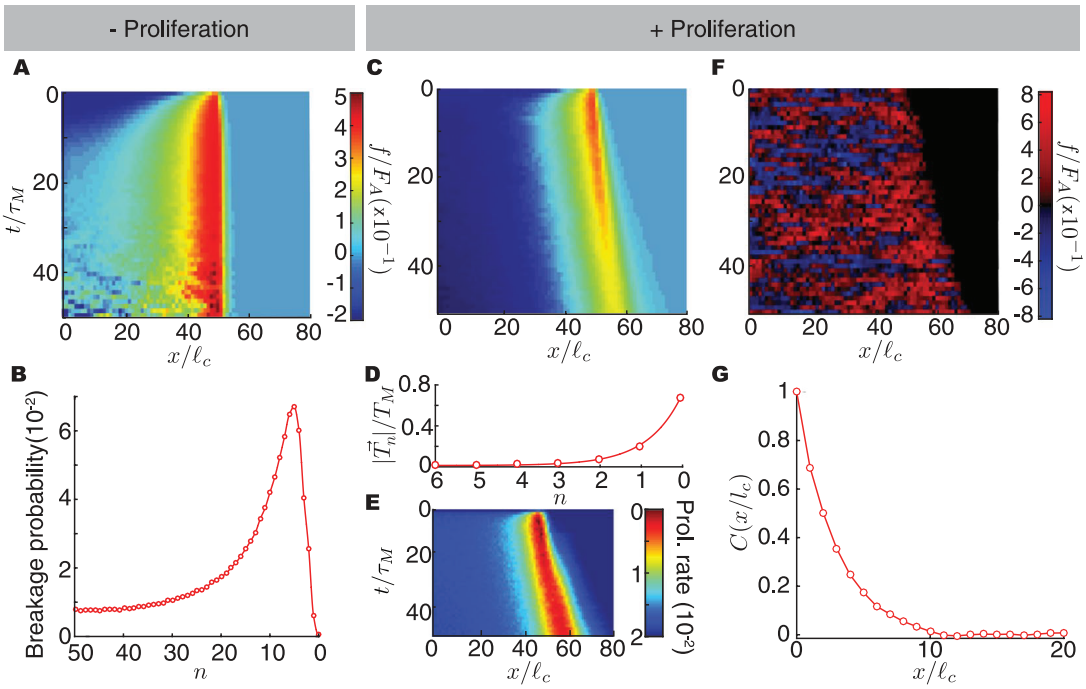

Figure 2.4: **Expansion of large cell colonies.** (A, C) Ensemble average intercellular force kymographs of expanding cell colonies in the absence (A) and presence (C) of cell proliferation. (B) Probability of colony breakage in the absence of cell proliferation as a function of the distance from the edge of the colony. (D) Spatial profile of traction forces from the edge of the colony in the presence of cell proliferation. (E) Ensemble average cell proliferation kymograph showing spatiotemporal variations during colony expansion. (F) Intercellular force kymograph from a single simulation run. (G) Ensemble spatial autocorrelation of intercellular force which can be fitted to an exponential with a characteristic length of $2.8l_c$. In all cases the value of $T_M/F_A = 0.25, F_A/T_R = 10$ and $\tau_T/\tau_M = 0.1$.

collective buildup of forces that must be sustained by adhesion at cell-cell junctions. If cryptic lamellipodia did not exist, only cells at the edge would generate traction forces and this would not lead to sufficient forces at cell-cell junctions to cause colony breakage (for strongly adhering cells, i.e., $F_A > T_M$). While both cryptic lamellipodia and the generation of traction forces inside the cell colony have been experimentally observed, colony breakage has not been reported. This can be for a number of reasons that we discuss in the *Discussion* section below.

All results above were obtained in the absence of cell proliferation. Since cell pro-

liferation is present in most experiments on colony expansion [38, 101], we study the role of cell proliferation in the propagation of intracellular forces within the colony. To this end, we simulate the dynamics of the colony as described above, but allowing cells to divide if the separation between them becomes larger than a critical length ℓ_d (the results described below do not qualitatively depend on the choice of ℓ_d). We find that proliferation prevents the buildup of large intercellular forces deep in the colony, enabling it to continuously grow (Fig. 2.4C). This effect is equivalent to a fluidization of the cell colony at time scales larger than proliferation times [102]. Both the traction (Fig. 2.4D) and proliferation (Fig. 2.4E) spatial profiles decay over just a few cell sizes from the edge of the colony, as observed experimentally [38]. The penetration (decay) length scale of intercellular forces and proliferation are considerably larger than that of traction forces, in agreement with experimental observations [101]. Importantly, as previously observed in the expansion of 2D cell monolayers [38], intercellular forces display large spatial heterogeneities (Fig. 2.4F). While these inhomogeneities are averaged out when performing ensemble averages over many simulations (Fig. 2.4A,C), they become apparent for single simulation runs (Fig. 2.4F). Calculation of the spatial autocorrelation function (Fig. 2.4G) and its associated autocorrelation length (Methods) indicate that the heterogeneities in intercellular forces span a few cell sizes, with the specific size of the inhomogeneities depending on the parameters of the system.

2.4 Methods

We describe the methods used to conduct simulations of cell movements described by the theoretical description derived in section 2.2. We also describe statistical quantities used to describe various features of the simulated collective behaviors.

2.4.1 Particle-based simulations

Cells are simulated with an intercellular force consisting of a repulsive core up to $\ell_c/2$ and an attractive region between $\ell_c/2$ and ℓ_c (Fig. 2.1B). The exact functional form used is

$$f(r_{ji}) = \begin{cases} c\left(\frac{l_c}{2r_{ij}} - \frac{l_c^3}{8r_{ij}^3}\right) & \text{when } 0 < r_{ij} < 3l_c/4, \\ c\left(\frac{l_c}{2r_{ij}} - \frac{l_c^3}{8r_{ij}^3}\right)\Psi(r_{ij}/l_c - 3/4, 1/4) & \text{when } 3l_c/4 \leq r_{ij} < l_c, \\ 0, & \text{when } l_c \leq r_{ij}, \end{cases}$$

where

$$\Psi(x,a) = \begin{cases} e^{-\left(\frac{1}{(1-(x/a-1)^2)}\right)} & \text{when } |x| < a, \\ 0 & \text{otherwise} \end{cases}$$

is a bump function which is smooth at a, hence guaranteeing that $f(r_{ji})$ is smooth when $r_{ji} = l_c$. The constant c is chosen such that $\max_{0<r_{ji}\leq l_c} f(r_{ji}) = 1$. Eq. 2.3 is solved using an explicit Euler scheme, and the stochastic differential equation associated with the generation of traction, Eq. 2.4, is solved using the Euler-Maruyama method [103]. The timesteps of simulation were chosen adaptively based on the parameters of the system. All simulations were performed with custom computer codes.

2.4.2 Simulations of collisions between two cells

Repolarization probabilities (Fig. 2.2A,B) were obtained by computing repolarization times in 25000 instances of collisions in a 1D box of length $100\ell_c$. The repolarization time is computed as the difference between the instant at which a cell is within ℓ_c distance of the other cell, and the instant at which it switches direction, as defined by a change in sign of its traction. To guarantee that repolarization was not transient, the

direction of traction was tracked for a time τ_T after switching direction and only events in which the sign of traction did not revert were taken into account. Trajectories are mean displacements over 25000 instances of collisions, starting with tractions $T_M/2$ for head-on collisions and tractions $3T_M/4, T_M/4$ for rear-end collisions. Since we average many simulations for each set of parameters (ensemble average), error bars associated to simulation results are very small and not shown; simulation results are plotted as continuous or dashed lines (Fig. 2.2A,B,C).

The average separation times for each combination of T_M/T_R and F_A/T_R (Fig. 2.2E) were computed as the ensemble average ($N = 10^4$) of the time required for two cells to separate. Cells were randomly initialized at a distance between $0.4l_c$ and $0.6l_c$ and the simulations were terminated after $10^4\tau_M$ timesteps. If cells were still attached at that point, their separation time was set to $10^4\tau_M$.

To calculate the difference between experimental data from Ref. [1] and theoretical predictions (Fig. 2.2F), we first simulated 2500 instances of F-F, F-**B**, **F**-B collisions between two cells for different values of the parameters T_M/T_R and F_A/T_R (and fixed $\tau_T/\tau_M = 0.1$) and obtain the cumulative repolarization probabilities $P^{sim}_{\text{F-F}}(t)$, $P^{sim}_{\mathbf{F}\text{-B}}(t)$ and $P^{sim}_{\text{F-}\mathbf{B}}(t)$ as a function of these parameters (using the same procedure as described above; Fig. 2.2A,B). The superscript *sim* refers to the fact that values were obtained from simulation. While in our simulations the cumulative repolarization probabilities depend on time, the experimental data from Ref. [1] reports the repolarization probability 2 hours after cells collided. Since neither of the timescales τ_T or τ_M are explicitly known for NRK-52E cells [1], the comparison of our simulation results to the experimental data requires comparing simultaneously at least two collision types, as for a single collision type it is always possible to find a time in the predicted cumulative repolarization probability that matches the experimental value of the repolarization probability. When comparing simulations and experimental data for two or more collision types simultaneously, there

are enough constraints to make the comparison meaningful. For this reason, we compare simultaneously the three collision types (F-F, F-**B**, **F**-B) and define the measure of discrepancy between experimental and theoretical values, δ, as

$$\delta(F_A/T_R, T_M/T_R) = \min_t \sqrt{\frac{1}{3}\left\{\frac{(P^{exp}_{\text{F-F}} - P^{sim}_{\text{F-F}}(t))^2}{(P^{exp}_{\text{F-F}} + P^{sim}_{\text{F-F}}(t))^2} + \frac{(P^{exp}_{\text{F-}\mathbf{B}} - P^{sim}_{\text{F-}\mathbf{B}}(t))^2}{(P^{exp}_{\text{F-}\mathbf{B}} + P^{sim}_{\text{F-}\mathbf{B}}(t))^2} + \frac{(P^{exp}_{\mathbf{F}\text{-B}} - P^{sim}_{\mathbf{F}\text{-B}}(t))^2}{(P^{exp}_{\mathbf{F}\text{-B}} + P^{sim}_{\mathbf{F}\text{-B}}(t))^2}\right\}},$$

where the values $P^{exp}_{\text{F-F}}$, $P^{exp}_{\text{F-}\mathbf{B}}$ and $P^{exp}_{\mathbf{F}\text{-B}}$ are the probabilities of repolarization 2 hours after collision, for each type of collision, reported in Desai et al. [1], namely: $P^{exp}_{\text{F-F}} \simeq 0.87$, $P^{exp}_{\text{F-}\mathbf{B}} \simeq 0.18$ and $P^{exp}_{\mathbf{F}\text{-B}} \simeq 0.59$. The measure δ finds the time for which the discrepancy between theory and experiments is minimal and reports, for each value of the parameters T_M/T_R and F_A/T_R, such discrepancy δ, which is shown in Fig. 2.2F.

2.4.3 Simulations of groups of cells

To obtain position kymographs and train persistence times, the cells were initialized with identical tractions T_M spaced at a distance of $0.5\ell_c$. The persistence time τ_p corresponds to the time until either breakage or repolarization of the train occurs, i.e., when the distance between adjacent cells in the group becomes larger than ℓ_c, or the time at which the cluster reverses direction, namely $\sum_i \vec{T}_i < 0$. The position kymographs and persistence times correspond to ensemble averages obtained from 10^4 runs. Dynamic train formation is computed by simulating N cells in a 1-D box of length $10^2\ell_c$ with periodic boundary conditions (equivalent to a ring geometry). The fraction of cells in trains were determined as the cells existing in clusters of length greater than 2 cells as compared to the total number of cells. The system was simulated for times $10^3 \max\{\tau_T, \tau_M\}$ and was repeated for 10000 instances. Train fractions were obtained as ensemble averages of

the ratio of the mean number of cells in trains to the total number of cells.

2.4.4 Simulations of cell colonies

In all cases, 10^2 cells were initialized in close proximity, with the distance between neighboring cells randomly chosen between $0.4\ell_c$ and $0.6\ell_c$. To quantify colony breakage, the cell position where the distance to neighboring cells exceeds ℓ_c is noted as the point of breakage. Breakage probabilities are breakage frequencies from 10^4 runs. Cell division is modeled as the inclusion of a new cell at the midpoint of the segment joining the centers of two adjacent cells whose distance has exceeded $0.75\ell_c$. The newly formed cell starts with no traction. The parameter T_M/F_A is chosen as 0.25 to mimic colony expansion in cell types with high adhesion (e.g., MDCK cells). Intercellular forces (Fig. 2.4A,C), breakage probabilities (Fig. 2.4B), traction profiles (Fig. 2.4D) and proliferation rates (Fig. 2.4E) correspond to ensemble averages over 10000 runs. The intercellular forces in the kymograph of Fig. 2.4F were obtained using a single simulation run of the same system.

The spatial autocorrelation function was calculated as

$$C(|x_0 - x|/l_c) = \frac{1}{N}\sum_{i=1}^{N} \frac{\sum_{t=1}^{t_{\max}}(f^{(i)}(x_0/l_c, t/\tau_M) - < f^{(i)}(x_0/l_c) >)(f^{(i)}(x/l_c, t/\tau_M) - < f^{(i)}(x/l_c) >)}{\sum_{t=1}^{t_{\max}}(f^{(i)}(x_0/l_c, t/\tau_M) - < f^{(i)}(x_0/l_c) >)^2},$$

where $f^{(i)}(x,t)$ refers to the intercellular force at position x at time t from the i^{th} simulation run and $< f^{(i)}(x) >$ refers to the time average of the intercellular force at position x for the i^{th} simulation run; x_0 was chosen to be $25l_c$ from the middle of the colony; $t_{\max}$, which is the maximum simulation time, was chosen to be $500\tau_M$; and $N = 100$ is the number of ensembles over which the autocorrelation function was computed. The

exact position of x_0 in the colony does not affect the value of the autocorrelation function obtained.

2.5 Discussion

At the single cell level, our predictions of diffusive movements at time scales longer than the traction persistence time, are in good agreement with experimental observations showing diffusive cell movements at long timescales [104, 105]. Our results predict that the diffusion constant of cellular movements depends quadratically with the cell's traction force. This prediction could be experimentally tested by measuring the magnitude of traction forces using traction force microscopy while monitoring cellular movements. In addition, the predicted dependence of the cell velocity on an applied external force can potentially be measured using magnetic tweezers in a similar way as in previous experiments [86].

Beyond single cell movements, the observed behaviors in collision experiments on CIL [87, 80, 1] arise naturally in our description if both CIL and FIR are taken into account. Importantly, in the theoretical description presented above, CIL involves only lamellipodial retraction but does not impose repolarization away from contact; repolarization is a consequence of the pulling forces acting on the cell via FIR. While it is typically assumed that CIL involves repolarization, our description highlights the importance of considering the separate effects of lamellipodial retraction and force-dependent repolarization, as suggested by recent experimental results [82, 87, 88]. Indeed, some cell types show lamellipodial retraction upon contact, but no repolarization [82, 106, 84]. Our predictions indicate that the dynamics of repolarization, characterized by the cumulative probability of repolarization (Fig. 2.2), are very different for distinct collision types and depend strongly on parameters such as adhesion strength or traction force (T_M/F_A)

as well as the traction repolarization time and the mechanical relaxation time (τ_T/τ_M). These parameters can be experimentally varied using drugs targeting force generation or cell adhesion, and the dynamics of cell polarization could be monitored with polarization markers. Having a quantitative understanding of the behavior of cells during collisions would considerably help understand their behavior in larger groups.

Several experimental works have shown that coherently moving cell groups emerge even for cells types that display repulsion upon collision [22, 1]. Our theoretical predictions indicate that this phenomenon can be explained by a stabilization of lamellipodia enabled by neighboring polarized cells (NER) through either physical or biochemical mechanisms. In the absence of NER, our analysis predicts that no coherently moving cell groups can exist, as CIL prevents their formation. Since NER is directly related to the existence of cryptic lamellipodia, experiments exploring the physical and biochemical cues enabling cells to generate lamellipodia when contacted on all sides may help understand their collective behavior. In particular, experiments to characterize how polarization of cells affects the ability of their neighbors to polarize and generate cryptic lamellipodia may help understand the role of NER.

We find an optimal group number that maximizes migration persistence of small groups of cells (Fig. 2.3C,D), which could explain why collective migration of small cell groups is often observed in developing embryos [14, 21, 23] and cancer metastasis [25]. This prediction can directly be tested in 1D systems by measuring either switches in the direction of group motion or group breakage for groups of cells of different numbers (no cell proliferation) and for different cell adhesion strength. Our results also indicate that by varying the cell density alone (or confinement), with no changes in cell specific parameters (for a given cell type), both coherently moving cell trains or cell dispersal behavior can be observed (Fig. 2.3C,F). These predictions suggest that several experimentally observed behaviors [100, 98, 99], such as cell dispersal and coordinated group migration, can be

achieved by varying cellular confinement. We also find that, in addition to cell density, cell specific parameters can control the ability of cells to form coherently moving groups or disperse (Fig. 2.3C,F). Experiments to test these results could be realized in 1D systems by controlling the cell seeding density and monitoring cellular movements in the absence of cell proliferation.

Beyond small groups, we find that in the absence of cell proliferation, large cell colonies break up into smaller groups as a consequence of large intercellular forces that build up within the colony. However, colony breakage has not been observed in 2D cell monolayers in the absence of cell proliferation. Since colony breakage occurs in our simulations when intercellular pulling forces become larger than the cell-cell adhesion strength, it is conceivable that the cells used in many of these experiments (such as MDCK cells) adhere so strongly to each other that breakage is never observed. Our predictions indicate that lowering mildly the adhesion strength between cells should enable portions of the colony located close to the migrating edge (where intercellular forces are predicted to be largest) to break off. Another possibility is that colony breakage is an effect observed only in 1D geometries, as in 2D cell monolayers, the larger number of neighbors per cell may be able to sustain the forces that build up within the monolayer and help prevent breakage. Our results show that, at least in 1D, the presence of cell proliferation can help avoid colony breakup by preventing the buildup of large intercellular forces. While the 1D system studied here is not equivalent to a 2D cell monolayer, the predicted profile of cell proliferation is in agreement with previous experimental observations [101].

2.6 Summary

We presented a theoretical description of cell migration that accounts for known individual cell behaviors, such as CIL and FIR, and is able to reproduce the motion of a single

cell, two cell collisions, small groups of cells and large colonies. This description provides a unified framework to connect a large number of experiments in different conditions and with different cell types. Moreover, it allows a direct connection between specific molecular perturbations in cell adhesion, cell polarization, the generation of traction forces and mechanical feedback, and predicts their effect on collective cell migration.

Chapter 3

The Meeting Time of Random Walks

3.1 Introduction

3.1.1 Problem description and motivation.

In this chapter, we examine the meeting time between two groups of random walkers. This problem is motivated by a group of pursuers trying to intercept a group of evaders. The meeting time, in the context of this chapter, describes the average time till a first encounter occurs between one of the pursuers and one of the evaders given initial positions of the pursuers and the evaders. This notion of two adversarial mobile groups wherein one of the groups is trying to intercept members of the other group appears under several names: pursuit-evasion games [107], predator-prey interactions [108], cops and robbers games [109] and princess-monster games [110]. Our primary motivation is the design of stochastic surveillance strategies for quickest detection of mobile intruders. Single and multi-agent surveillance strategies appear in environmental monitoring [111, 112], minimizing emergency vehicle response times [113], traffic routing and border patrol [114, 115]. Aside from our proposed application to stochastic surveillance, the meeting time has direct applications to information flow in distributed networks [116], self-stabilization

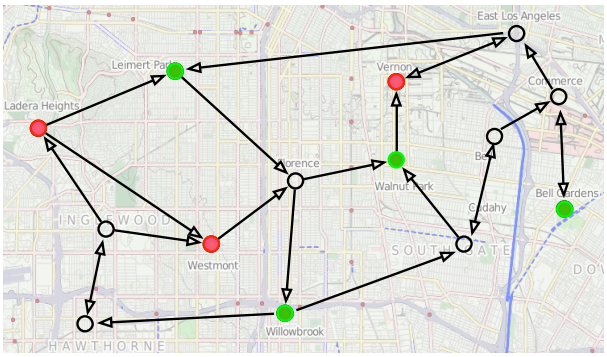

Figure 3.1: Multiple pursuers (green) and multiple evaders (red) performing random walks on a digraph.

of tokens [117] and measuring similarity of objects [118].

3.1.2 Applications of meeting times in various contexts

Early interest in meeting times was motivated by applications to self-stabilizing token management schemes [119]. In a token management scheme, only one of the many processors on a distributed network is enabled to change state or perform a particular task, and this processor is said to possess the token. If two tokens meet then they collapse into a single token. Israeli and Jalfon suggest a scheme in which the token is passed randomly to a neighbor [117]. In a general connected, undirected, n-vertex graph they were able to obtain an exponential bound for the meeting time of two tokens in terms of the maximum degree and the diameter of the graph. Coppersmith et al [120] improved the bound to be polynomial in the number of nodes by bounding the meeting time in terms of the pairwise hitting time from the starting nodes of the tokens to hidden vertices. In [120, 117, 119] the notion of the meeting time involves the tokens being moved asynchronously by an adversary whose objective is to maximize meeting time by playing only one of the two tokens. Bshouty et al [121] obtain a bound on the meeting time of several such tokens in terms of the meeting time of two tokens. Bounds for meeting times of two identical independent continuous-time reversible Markov chains in terms of the

pairwise hitting times of the chain are mentioned in the work by Aldous [122]. Several variations of "cat-mouse" games are discussed in [123] wherein bounds are obtained in terms of the pairwise hitting time or the variation-threshold time (a measure of rate of convergence to stationary distributions) depending on the Markov chains being discrete-time or continuous-time.

3.1.3 Organization.

This chapter is organized as follows. In Section 3.2 we introduce notation that is used throughout the chapter and review useful concepts. In Section 3.3 we introduce our formulation for the meeting times of pairs of Markov chains, and also define sets of pairs of matrices for which finite meeting times exist. In Section 3.4 we extend the notion of the meeting time to multiple pursuers and evaders. In Section 3.5 we obtain closed-form expressions for continuous-time Markov chains. Finally, in Section 3.6 we present conclusions.

3.2 Notation and review of known results

In this section we define various useful concepts and notation. We provide an overview of some facts and known results on Markov chains and the Kronecker product of Markov chains, while also introducing notation that will be used throughout the chapter to deal with vectors and matrices, and random walks on graphs.

3.2.1 Markov chains.

A *Markov chain* is a sequence of random variables taking value in the *finite* set $\{1, \ldots, n\}$ with the Markov property, namely that the future state depends only on the

present state.

Let $X_k \in \{1, \ldots, n\}$ denote the location of a random walker at time $k \in \{0, 1, 2, \ldots\}$, then a discrete-time Markov chain is *time-homogeneous* if $\mathbb{P}[X_{n+1} = j \,|\, X_n = i] = \mathbb{P}[X_n = j \,|\, X_{n-1} = i] = p_{i,j}$, where $P = [p_{i,j}] \in \mathbb{R}^{n \times n}$ is the *transition matrix* of the Markov chain. By definition, each transition matrix P is row-stochastic, i.e., $P\mathbb{1}_n = \mathbb{1}_n$. The period of a state is defined as the greatest common divisor of all t such that $\{t \geq 1 \,|\, \mathbb{P}[X_t = i \,|\, X_0 = i] \neq 0\}$. A state whose period is one is referred to as an *aperiodic* state. It can be shown that in a *communicating class*(defined below) all states share the same period. For more details on discrete-time Markov chains refer [124, Chapter 8].

Let $X_t \in \{1, \ldots, n\}$ denote the location of a random walker at time $t \in \mathbb{R}^+$, then a continuous-time Markov chain is *time-homogeneous* if $\mathbb{P}[X_{t'+t} = j \,|\, X_{t'} = i] = p^t_{i,j}$ for all $t \geq 0, t' \geq 0$, where $P(t) = p^t_{i,j} \in \mathbb{R}^{n \times n}$ is the transition matrix of the Markov chain. The evolution of the continuous-time Markov chain is determined by the solution to the first-order differential equation $P'(t) = P(t)Q$, where $P(t) = p^t_{i,j}$ and Q is a *transition rate matrix* which satisfies $Q\mathbb{1}_n = \mathbb{0}_n$. For more details on continuous-time Markov chains refer [125, Chapters 2 & 3]. A continuous-time Markov chain is said to be *ergodic* if it is irreducible.

Consider two states i and j belonging to a Markov chain. We say i *communicates* with j if $p^t_{i,j} \neq 0$ for some $t > 0$. For a subset of states $X \subset \{1, \ldots, n\}$, we say that X forms a *communicating class* if for every state $i, j \in X$ the states communicate with each other, i.e $\mathbb{P}[X_t = j \,|\, X_0 = i] \neq 0$ and $\mathbb{P}[X_{t'} = i \,|\, X_0 = j] \neq 0$ for some $t, t' \geq 0$. An *absorbing class* A of a Markov chain is a communicating class such that the probability of escaping the set is zero, i.e $p^t_{i,j} = 0$ for all $t > 0$ for all $i \in A, j \notin A$. If a communicating class is not absorbing, then it is called a *transient class*. In general, a Markov chain will have multiple absorbing and transient classes. If a Markov chain has only a single

absorbing class then it is referred to as a *single absorbing Markov chain.*

For any two start nodes i, j, the *first hitting time from node i to node j*, denoted by $T_{i,j}$, is the first time that a random walker reaches node j starting from node i. More formally,

$$T_{i,j} = \min\{t \geq 1 \mid X_t = j \text{ given that } X_0 = i\}.$$

Note that the first hitting time can be infinite when the Markov chain is reducible but is always finite when the Markov chain is irreducible. The *hitting time between nodes i and j* is given by $h_{i,j} = \mathbb{E}[T_{i,j}]$. If a Markov chain is single absorbing, then a unique stationary distribution π exists. The vector $\pi \in \mathbb{R}^{n\times 1}$ is a *stationary distribution* of a discrete-time Markov chain with transition matrix P if $\sum_{i=1}^{n} \pi_i = 1$ and $\pi^\top P = \pi^\top$ and of a continuous-time Markov chain with transition rate matrix Q if $\sum_{i=1}^{n} \pi_i = 1$ and $\pi^\top Q = 0$. A Markov chain is *irreducible* if the absorbing class is the entire set of states $\{1, \ldots, n\}$. A discrete-time Markov chain is said to be *ergodic* if it is irreducible and aperiodic.

3.2.2 Matrix notation.

We use the notation $A = [a_{i_1 \ldots i_l, j_1 \ldots j_m}]$ to denote the matrix generated by elements $a_{i_1 \ldots i_l, j_1 \ldots j_m}$, where the rows of A are determined by cycling through indices i_l followed by i_{l-1} and so on until i_1, and the columns of A are determined by cycling through indices j_m followed by j_{m-1} and so on until j_1. For example, consider $i_1, i_2, j_1, j_2 \in \{1, \ldots, n\}$, then

$$A = [a_{i_1 i_2, j_1 j_2}] = \begin{bmatrix} a_{11,11} & a_{11,12} & \dots & a_{11,1n} & a_{11,21} & \dots & a_{11,nn} \\ a_{12,11} & a_{12,12} & \dots & a_{12,1n} & a_{12,21} & \dots & a_{12,nn} \\ \vdots & \vdots & \dots & \dots & \dots & \vdots & \vdots \\ a_{1n,11} & a_{1n,12} & \dots & a_{1n,1n} & a_{1n,21} & \dots & a_{1n,nn} \\ a_{21,11} & a_{21,12} & \dots & a_{21,1n} & a_{21,21} & \dots & a_{21,nn} \\ \vdots & \vdots & \dots & \dots & \dots & \vdots & \vdots \\ a_{nn,11} & a_{nn,12} & \dots & a_{nn,1n} & a_{nn,21} & \dots & a_{nn,nn} \end{bmatrix}.$$

For the case where $A = [a_{i,j}]$ this corresponds to the classic interpretation with element $a_{i,j}$ in the i-th row and j-th column of A. We use the notation $\text{diag}[a]$ to indicate the diagonal matrix generated by vector a and $\text{vec}(A)$ to indicate the vectorization of a matrix $A \in \mathbb{R}^{n \times m}$ where $\text{vec}(A) = [A(1,1), \dots, A(n,1), A(1,2), \dots, A(n,2), \dots, A(m,1), \dots, A(n,m)]^\top$. In other words, even if we define A as $A = [a_{i_1 i_2, j_1 j_2}]$, the vector $\text{vec}(A) = \text{vec}([a_{i_1 i_2, j_1 j_2}])$ is simply a stacking of the columns of A.

Let $I_n \in \mathbb{R}^{n \times n}$ denote the identity matrix of size n, $\mathbb{1}_n \in \mathbb{R}^{n \times 1}$ denote the vector of ones of size n, and $\mathbb{e}_1, \mathbb{e}_2, \dots, \mathbb{e}_n \in \mathbb{R}^{n \times 1}$ denote vectors with unity in the row indicated by the subscript. We define a generalized Kronecker delta function $\delta_{i_1 i_2 \dots i_l, j_1 j_2 \dots j_m}$, by

$$\delta_{i_1 \dots i_l, j_1 j_2 \dots j_m} = \begin{cases} 1, & \text{if } \exists\ l', m' \text{ such that } i_{l'} = j_{m'} \text{ for any } 1 \le l' \le l, 1 \le m' \le m, \\ 0, & \text{otherwise.} \end{cases}$$

We use the subscript p, e or superscript (p), (e) to delineate between quantities associated with pursuers and evaders.

We are now ready to review some useful facts about Kronecker products. The Kronecker product, represented by the symbol $\otimes$, of two matrices $A \in \mathbb{R}^{n \times m}$ and $B \in \mathbb{R}^{q \times r}$

is an $nq \times mr$ matrix given by

$$A \otimes B = \begin{bmatrix} a_{1,1}B & \dots & a_{1,m}B \\ \vdots & \ddots & \vdots \\ a_{n,1}B & \ddots & a_{n,m}B \end{bmatrix}.$$

The Kronecker product is bilinear and has many useful properties, two of which are summarized in the following Lemma; see [126, Chapter 4] for more information.

Lemma 1 (Properties of the Kronecker product) *Given the matrices A, B, C and D, the following relations hold for the Kronecker product.*

(i) $(A \otimes B)(C \otimes D) = (AC) \otimes (BD)$,

(ii) $(B^\top \otimes A)\,\mathrm{vec}(C) = \mathrm{vec}(ACB)$,

where it is assumed that the matrices are of appropriate dimension when matrix multiplication or addition occurs.

3.2.3 Markov chains on graphs.

In this chapter, for discrete-time Markov chains we consider weighted digraphs $G = (V, E, P)$ with node sets $V := \{1, \dots, n\}$, edge set $E \subset V \times V$, and associated transition matrix $P = [p_{i,j}]$ with the property that $p_{i,j} \geq 0$ if $(i, j) \in E$ and $p_{i,j} = 0$ otherwise. The weight of the edge (i, j) is interpreted as the weight associated with the probability of transition from node i to node j. The nodes of the graph are equivalent to the states of the Markov chain. We say there exists a *walk* of *length* ℓ from node i_1 to node i_l if there exists a sequence of nodes $i_2, \dots, i_{\ell-2}$ such that $p_{i_k, i_{k+1}} > 0$ for $1 \leq k \leq \ell - 1$.

In this chapter, for continuous-time Markov chains we consider weighted digraphs $G = (V, E, Q)$ with node sets $V := \{1, \dots, n\}$, edge set $E \subset V \times V$, and associated

transition rate matrix $Q = [Q_{i,j}]$ with the property that $q_{i,j} \geq 0$ if $(i,j) \in E$, $q_{i,j} = 0$ otherwise and $q_{i,i} = -\sum_{j \in V} q_{i,j}$. The weight of the edge (i,j) is interpreted as the rate of transition from node i to node j. One could also look at the entry $-1/q_{i,i}$ as the average time at which the walker leaves node i and $1/q_{i,j}$ as the average time for a jump from i to j. We say there exists a walk from node i_1 to node i_l if there exists a sequence of nodes $i_2, \ldots, i_{\ell-1}$ such that $q_{i_l, i_{l+1}} > 0$ for $1 \leq l \leq \ell - 1$.

The following lemmas are used in the proofs of the main results introduced in this chapter.

Lemma 2 (Convergence of substochastic matrices) *Let $P \in \mathbb{R}^{n \times n}$ be a substochastic matrix with at least one row-sum $\sum_{j=1}^{n} P_{i,j} < 1$. If for every node there exists a walk to a node with row-sum less than 1, then P is convergent.*

Lemma 3 (Existence of walks on Kronecker products) *Let $P_1, P_2, \ldots, P_N \in \mathbb{R}^{n \times n}$ be stochastic matrices. If there exists a walk from $i_1 \to j_1$ in P_1, $i_2 \to j_2$ in P_2, ..., and $i_N \to j_N$ in P_N of equal length, then there exists a walk from $(i_1, i_2, \ldots, i_N)$ to $(j_1, j_2, \ldots, j_N)$ in $P_1 \otimes P_2 \otimes \ldots \otimes P_N$.*

Note that the Kronecker product of two Markov chains on a graph can be interpreted as a single Markov chain on the Kronecker product graph [127]. This property of random walks on graphs enables a direct translation of the notion of meeting time of multiple random walkers on a single graph to that of the hitting time of a single random walker on the Kronecker product graph.

3.3 The meeting time of two Markov chains.

In this section, we formulate and study the meeting time between two discrete-time Markov chains.

3.3.1 Main result

Consider the pursuer and evader performing random walks on a set of nodes $V := \{1, \ldots, n\}$ with digraphs $G_{\mathrm{p}} = (V, E_{\mathrm{p}}, P_{\mathrm{p}})$, $G_{\mathrm{e}} = (V, E_{\mathrm{e}}, P_{\mathrm{e}})$, edge sets $E_{\mathrm{p}}, E_{\mathrm{e}} \subset V \times V$, and transition matrices P_{p}, P_{e}. The matrix P_{p} satisfies $p^{(\mathrm{p})}_{i,j} \geq 0$ if $(i,j) \in E_{\mathrm{p}}$ and $p^{(\mathrm{p})}_{i,j} = 0$ if $(i,j) \notin E_{\mathrm{p}}$. Similarly P_{e} satisfies similar properties to be a well-defined transition matrix on G_{e}.

Let $X^{(\mathrm{p})}_t, X^{(\mathrm{e})}_t \in \{1, \ldots, n\}$ be the location of the two agents at time $t \in \{0, 1, 2, \ldots\}$. For any two start nodes i, j, the *first meeting time from i and j*, denoted by $T_{i,j}$, is the first time that both random walkers meet when starting from nodes i and j, respectively. More formally,

$$T_{i,j} = \min\{t \geq 1 \mid X^{(\mathrm{p})}_t = X^{(\mathrm{e})}_t \text{ given that } X^{(\mathrm{p})}_0 = i \text{ and } X^{(\mathrm{e})}_0 = j\}.$$

Note that the first meeting time can be infinite. It is easy to construct examples in which the two agents never meet. Let $m_{i,j} = \mathbb{E}[T_{i,j}]$ be the expected first meeting time starting from nodes i and j. For the sake of brevity, we shall refer to the expected first meeting time as just the meeting time.

Theorem 1 (The meeting time of two Markov chains) *Consider two Markov chains with transition matrices P_{p} and P_{e} defined on a digraph G with nodeset $V = \{1, \ldots, n\}$. The following statements are equivalent:*

(i) for each pair of nodes i, j, the meeting time $m_{i,j}$ from nodes i and j is finite,

(ii) for each pair of nodes i, j, there exists a node j and a length ℓ such that a walk of length ℓ exists from i to k and a walk of length ℓ exists from j to k,

(iii) for each pair of nodes i, j, there exists a walk in the digraph associated with the stochastic matrix $P_{\mathrm{p}} \otimes P_{\mathrm{e}}$ from (i,j) to a node (k,k), for some $k \in \{1, \ldots, n\}$, and

(iv) the sub-stochastic matrix $(P_{\mathrm{p}} \otimes P_{\mathrm{e}})E$ *is convergent and the vector of meeting times is given by*

$$\mathrm{vec}(M) = (I_{n^2} - (P_{\mathrm{p}} \otimes P_{\mathrm{e}})E)^{-1}\mathbb{1}_{n^2}, \tag{3.1}$$

where $M \in \mathbb{R}^{n\times n}$ *and* $E \in \mathbb{R}^{n^2\times n^2}$ *is a binary diagonal matrix with diagonal entries* $\mathbb{1}_{n^2} - \mathrm{vec}(I_n)$.

Proof: For the nodes i and j, the first meeting time satisfies the recursive formula

$$T_{i,j} = \begin{cases} 1, & \text{w.p. } \sum_k p^{(\mathrm{p})}_{i,k} p^{(\mathrm{e})}_{j,k}, \\ T_{k_1,h_1} + 1, & \text{w.p. } p^{(\mathrm{p})}_{i,k_1} p^{(\mathrm{e})}_{j,h_1}, k_1 \neq h_1. \end{cases}$$

Taking the expectation we have

$$\begin{aligned} \mathbb{E}[T_{i,j}] &= \sum_k p^{(\mathrm{p})}_{i,k} p^{(\mathrm{e})}_{j,k} + \sum_{k_1\neq h_1} p^{(\mathrm{p})}_{i,k_1} p^{(\mathrm{e})}_{j,h_1} (\mathbb{E}[T_{k_1,h_1}] + 1), \\ &= \sum_{k_1}\sum_{h_1} p^{(\mathrm{p})}_{i,k_1} p^{(\mathrm{e})}_{j,h_1} + \sum_{k_1\neq h_1} p^{(\mathrm{p})}_{i,k_1} p^{(\mathrm{e})}_{j,h_1} \mathbb{E}[T_{k_1,h_1}], \\ &= 1 + \sum_{k_1\neq h_1} p^{(\mathrm{p})}_{i,k_1} p^{(\mathrm{e})}_{j,h_1} \mathbb{E}[T_{k_1,h_1}]. \end{aligned}$$

Let $m_{i,j} = \mathbb{E}[T_{i,j}]$ for every $i, j \in \{1, \ldots, n\}$ and let $M = [m_{i,j}]$. Note that the entries of M can be written as

$$\begin{aligned} m_{i,j} &= 1 + \sum_{k_1\neq h_1} p^{(\mathrm{p})}_{i,k_1} p^{(\mathrm{e})}_{j,h_1} m_{k_1,h_1}, \\ \implies m_{i,j} &= 1 + \sum_{k_1=1}^{n} p^{(\mathrm{p})}_{i,k_1} \sum_{h_1=1}^{n} m_{k_1,h_1} p^{(\mathrm{e})}_{j,h_1} - \sum_{k=1}^{n} p^{(\mathrm{p})}_{i,k} p^{(\mathrm{e})}_{j,k} m_{k,k}, \\ \implies M &= \mathbb{1}_n \mathbb{1}_n^\top + P_{\mathrm{p}}(M - M_d) P_{\mathrm{e}}^\top, \end{aligned}$$

where $M_d \in \mathbb{R}^{n\times n}$ is a diagonal matrix with only the diagonal elements of M . We have used the property that $(ABC)_{i,j} = \sum_k A_{i,k} \sum_l B_{k,l}C_{l,j}$ to obtain the equation in matrix form. Rewriting the equation in vector form and using Lemma 1 gives

$$\begin{aligned}
\operatorname{vec}(M) &= \mathbb{1}_{n^2} + (P_{\mathrm{p}} \otimes P_{\mathrm{e}})(\operatorname{vec}(M) - \operatorname{vec}(M_d)),\\
\operatorname{vec}(M) &= \mathbb{1}_{n^2} + (P_{\mathrm{p}} \otimes P_{\mathrm{e}})(I_{n^2} - \operatorname{vec}(I_n))\operatorname{vec}(M)\\
\operatorname{vec}(M) &= \mathbb{1}_{n^2} + (P_{\mathrm{p}} \otimes P_{\mathrm{e}})E\operatorname{vec}(M)
\end{aligned}$$

If the matrix $I_{n^2} - (P_{\mathrm{p}} \otimes P_{\mathrm{e}})E$ is invertible then we have a unique solution to the meeting times. We shall now show that the finiteness of meeting times as in (i) is equivalent to the existence of walks of equal length to common nodes as mentioned in (ii) and in (iii), which guarantees invertibility of $I_{n^2} - (P_{\mathrm{p}} \otimes P_{\mathrm{e}})E$ in (iv).

We start by proving that $(i) \implies (ii)$. If we assume that $(i) \not\implies (ii)$, then there exists a pair of nodes i and j such that the meeting time is finite and there exists no walk of equal length to any node in V. However if there exists no walk of equal length to a common node, then the agents never meet and the meeting time is always infinite. Hence by contradiction $(i) \implies (ii)$.

Next, we show that $(ii) \iff (iii)$. The Kronecker product of the transition matrices gives a joint transition matrix for the agents over the set of nodes $V \times V$. The (i, j) entry of the matrix $P_{\mathrm{p}} \otimes P_{\mathrm{e}}$ corresponds to the states $X^{(\mathrm{p})} = i$ and $X^{(\mathrm{e})} = j$ [128]. The statement (ii) ensures the existence of a node k for every pair (i, j) which is reachable by a walk of equal length from i in P_{p} and j in P_{e}. This condition is equivalent to the node (k, k) being reachable from the pair (i, j) on the Kronecker product of the two Markov chains [129, Proposition 1].

Next, we show $(iii) \implies (iv)$. The stochastic matrix $P_{\mathrm{p}} \otimes P_{\mathrm{e}}$ has a walk from any node (i, j) to some node (k_1, h_1) where $\mathbb{P}[X^{(\mathrm{p})} = k, X^{(\mathrm{e})} = k \,|\, X^{(\mathrm{p})} = k_1, X^{(\mathrm{e})} = h_1] \neq 0$ as

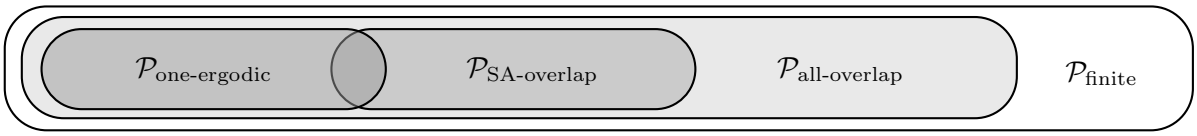

Figure 3.2: Sets of pairs of transition matrices with finite meeting times.

there exists a walk $(i,j) \to (k,k)$. Note that post-multiplying the square matrix $P_{\mathrm{p}} \otimes P_{\mathrm{e}}$ by E corresponds to setting the columns associated with nodes of the form (k,k) to $\mathbb{0}_{n^2}$. Thus the row associated with (k,h) has row-sum less than 1. Therefore every node (i,j) has a walk to a node whose row-sum is less than 1 which implies that the matrix $(P_{\mathrm{p}} \otimes P_{\mathrm{e}})E$ is convergent by virtue of Lemma 2.

From this we obtain equation (3.1). Since (iii) guarantees the existence of $(I_{n^2} - (P_{\mathrm{p}} \otimes P_{\mathrm{e}}))^{-1}$, we prove that $(iii) \implies (iv)$.

Note that the existence of $\mathrm{vec}(M)$ in (iv) gives $(iv) \implies (i)$. Thus we have shown that $(i) \implies (ii) \iff (iii) \implies (iv) \implies (i)$. Hence the four conditions are equivalent. ■

The above necessary and sufficient conditions give the most general set of pairs of matrices for which finite meeting times exist. These conditions are in practice difficult to use for designing transition matrices. Hence, we introduce a few sets of pairs of matrices for which the meeting times are guaranteed to be finite.

3.3.2 Sufficient conditions for finiteness.

Consider the following sets of pairs of matrices:

$\mathcal{P}_{\text{finite}}$: finite meeting times. Let $\mathcal{P}_{\text{finite}}$ be the set of pairs of transition matrices $P_{\mathrm{p}}, P_{\mathrm{e}}$ satisfying the conditions stated in Theorem 1 and therefore having finite meeting times.

$\mathcal{P}_{\text{all-overlap}}$: Markov chains with all-to-all overlapping absorbing classes.

Let $\mathcal{P}_{\text{all-overlap}}$ be the set of pairs of transition matrices $P_{\text{p}}, P_{\text{e}}$ with the following property: P_{p} has multiple absorbing classes $A_1^{(\text{p})}, A_2^{(\text{p})}, \ldots, A_q^{(\text{p})}$ with associated periods $d_1^{(\text{p})}, d_2^{(\text{p})}, \ldots, d_q^{(\text{p})}$, and P_{e} has multiple absorbing classes $A_1^{(\text{e})}, A_2^{(\text{e})}, \ldots, A_r^{(\text{e})}$ with associated periods $d_1^{(\text{e})}, d_2^{(\text{e})}, \ldots, d_r^{(\text{e})}$, and for each $q' \in \{1, \ldots, q\}$ and $r' \in \{1, \ldots, r\}$, $A_{q'}^{(\text{p})} \cap A_{r'}^{(\text{e})} \neq \phi$ and $\gcd(d_{q'}^{(\text{p})}, d_{r'}^{(\text{e})}) = 1$.

$\mathcal{P}_{\text{SA-overlap}}$: single absorbing Markov chains with overlapping absorbing classes. Let $\mathcal{P}_{\text{SA-overlap}}$ be the set of pairs of transition matrices $P_{\text{p}}, P_{\text{e}}$ with the following property: P_{p} has a single absorbing class $A^{(\text{p})}$ with period $d^{(\text{p})}$, and P_{e} has a single absorbing class $A^{(\text{e})}$ with period $d^{(\text{e})}$, and $A^{(\text{p})} \cap A^{(\text{e})} \neq \phi$ and $\gcd(d^{(\text{p})}, d^{(\text{e})}) = 1$.

$\mathcal{P}_{\text{one-ergodic}}$: one ergodic Markov chain. Let $\mathcal{P}_{\text{one-ergodic}}$ be the set of pairs of transition matrices $P_{\text{p}}, P_{\text{e}}$ such that one of the matrices P_{p} or P_{e} is ergodic.

Given the above definitions the following theorem holds.

Theorem 2 (Sufficient conditions for finite meeting times) *The sets of pairs of transition matrices $\mathcal{P}_{\text{finite}}$, $\mathcal{P}_{\text{all-overlap}}$, $\mathcal{P}_{\text{SA-overlap}}$, $\mathcal{P}_{\text{one-ergodic}}$ satisfy*

$$(\mathcal{P}_{\text{one-ergodic}} \cup \mathcal{P}_{\text{SA-overlap}}) \subset \mathcal{P}_{\text{all-overlap}} \subset \mathcal{P}_{\text{finite}}.$$

Proof: Before we prove the statement in the theorem we prove a minor result. Consider two Markov chains, each with transition matrices $P_{\text{p}}, P_{\text{e}} \in \mathbb{R}^{n \times n}$ defined on a digraph G with nodeset $V = \{1, \ldots, n\}$. Let the absorbing classes of P_{p} be $A_1^{(\text{p})}, A_2^{(\text{p})}, \ldots, A_q^{(\text{p})}$ with periods $d_1^{(\text{p})}, d_2^{(\text{p})}, \ldots, d_q^{(\text{p})}$ respectively, and let the absorbing classes of P_{e} be $A_1^{(\text{e})}, A_2^{(\text{e})}, \ldots, A_r^{(\text{e})}$ with periods $d_1^{(\text{e})}, d_2^{(\text{e})}, \ldots, d_r^{(\text{e})}$ respectively. If there exists an absorbing class $A_{q'}^{(\text{p})}$ in P_{p} and $A_{r'}^{(\text{e})}$ in P_{e} such that $A_{q'}^{(\text{p})} \cap A_{r'}^{(\text{e})} \neq \phi$ and $\gcd(d_{q'}^{(\text{p})}, d_{r'}^{(\text{e})}) = 1$, then there exists a walk

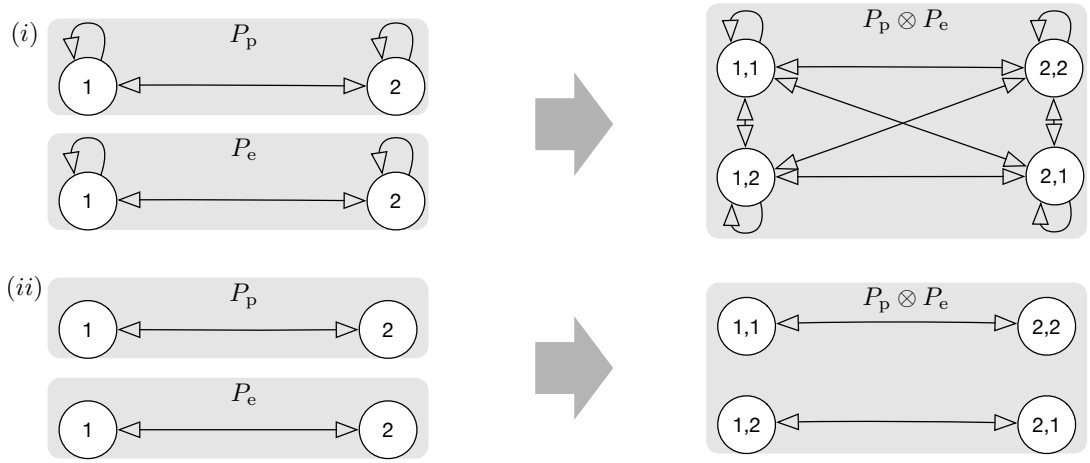

Figure 3.3: The pursuer-evader pair in (i) has finite meeting times as every node has a walk to the common nodes $(1,1)$ and $(2,2)$ in the Kronecker graph. However, in (ii) there exists no walks to common nodes from $(1,2)$ and $(2,1)$.

from all pairs (i,j), where i is any node from which there exists a walk to $A^{(\mathrm{p})}_{q'}$ and j is any node from which there exists a walk to $A^{(\mathrm{e})}_{r'}$, to a node (k,k) in the digraph associated with the transition matrix $P_{\mathrm{p}} \otimes P_{\mathrm{e}}$.

The proof of this result is as follows. Since $A^{(\mathrm{p})}_{q'} \cap A^{(\mathrm{e})}_{r'} \neq \phi$ there exists at least one node k which is accessible from both i and j. Since k belongs to the absorbing class $A^{(\mathrm{p})}_{q'}$, starting from the node i there exists all walks of length $u_1 d^{(\mathrm{p})}_{q'} + v_1$ to the node k for all $u_1 \geq U_1$, for some $U_1 \in \mathbb{N}$ sufficiently large and some $v_1 \in \mathbb{N}$ such that $0 \leq v_1 \leq d^{(\mathrm{p})}_{q'}$. Similarly, since k also belongs to the absorbing class $A^{(\mathrm{e})}_{r'}$, starting from the node j there exists all walks of length $u_2 d^{(\mathrm{e})}_{r'} + v_2$ to the node k for all $u_2 \geq U_2$, for some $U_2 \in \mathbb{N}$ sufficiently large and some $v_2 \in \mathbb{N}$ such that $0 \leq v_2 \leq d^{(\mathrm{e})}_{r'}$. Since $\gcd(d^{(\mathrm{p})}_{q'}, d^{(\mathrm{e})}_{r'}) = 1$ we can always find u_1 and u_2 such that $u_1 d^{(\mathrm{p})}_{q'} + v_1 = u_2 d^{(\mathrm{p})}_{r'} + v_2$. Thus there exists a walk of equal length to the node k from both i and j which ensures that (k,k) is accessible from (i,j).

To prove $\mathcal{P}_{\text{all-overlap}} \subset \mathcal{P}_{\text{finite}}$ we utilize statement (iii) in Theorem 1 to show that for every pair of nodes (i,j), where i,j are nodes in the Markov chain associated with $P_{\mathrm{p}}, P_{\mathrm{e}}$, there must exist a walk to a common node of the form (k,k). Consider a pair

of Markov chains $(P_{\mathrm{p}}, P_{\mathrm{e}}) \in \mathcal{P}_{\text{all-overlap}}$. The states of the Markov chain associated with the transition matrix P_{p} can be split into a set of absorbing classes $A_1^{(\mathrm{p})}, A_2^{(\mathrm{p})}, \ldots, A_q^{(\mathrm{p})}$ and transient classes $T_1^{(\mathrm{p})}, T_2^{(\mathrm{p})}, \ldots, T_s^{(\mathrm{p})}$. Similarly for the Markov chain associated with P_{e}, the states can be split into a set of absorbing classes $A_1^{(\mathrm{e})}, A_2^{(\mathrm{e})}, \ldots, A_r^{(\mathrm{e})}$ and transient classes $T_1^{(\mathrm{e})}, T_2^{(\mathrm{e})}, \ldots, T_t^{(\mathrm{e})}$. We begin by first proving the case for pairs of states belonging to (1) the absorbing classes of both chains, (2) the transient classes of both chains, and finally, (3) transient states of one chain paired with absorbing classes from the other chain.

Now we will use this result to prove $\mathcal{P}_{\text{all-overlap}} \subset \mathcal{P}_{\text{finite}}$. We shall show that for pairs of matrices belonging to $\mathcal{P}_{\text{all-overlap}}$ the meeting times are finite by concluding that statement (iii) of Theorem 1 is satisfied. Consider a pair of Markov chains $(P_{\mathrm{p}}, P_{\mathrm{e}}) \in \mathcal{P}_{\text{all-overlap}}$. The states of the Markov chain associated with the transition matrix P_{p} can be split into a set of absorbing classes $A_1^{(\mathrm{p})}, A_2^{(\mathrm{p})}, \ldots, A_q^{(\mathrm{p})}$ and transient classes $T_1^{(\mathrm{p})}, T_2^{(\mathrm{p})}, \ldots, T_s^{(\mathrm{p})}$. Similarly for the Markov chain associated with P_{e}, the states can be split into a set of absorbing classes $A_1^{(\mathrm{e})}, A_2^{(\mathrm{e})}, \ldots, A_r^{(\mathrm{e})}$ and transient classes $T_1^{(\mathrm{e})}, T_2^{(\mathrm{e})}, \ldots, T_t^{(\mathrm{e})}$. For statement (iii) in Theorem 1 to be satisfied, for every pair of nodes (i, j), where i, j are nodes in the Markov chain associated with $P_{\mathrm{p}}, P_{\mathrm{e}}$, there must exist a walk to a node of the form (k, k). To do so we shall initially consider pairs of states belonging to the absorbing classes of both chains, followed by the transient classes of both chains, and finally, transient states of one chain paired with absorbing classes from the other chain, and show that in each case we show a common node exists to which there is a walk of equal length using the proven result.

First, consider nodes (i, j) such that i belongs to an absorbing class $A_{q'}^{(\mathrm{p})}$ where $q' \in \{1, \ldots, q\}$ and j belongs to an absorbing class $A_{r'}^{(\mathrm{e})}$ where $r' \in \{1, \ldots, r\}$. By definition, every node in an absorbing class has walks to every other node in its class. $\mathcal{P}_{\text{all-overlap}}$ gives that $A_{q'}^{(\mathrm{p})} \cap A_{r'}^{(\mathrm{e})} \neq \phi$ and $\gcd(d_{q'}^{(\mathrm{p})}, d_{r'}^{(\mathrm{e})}) = 1$. Hence the provisions for the result are

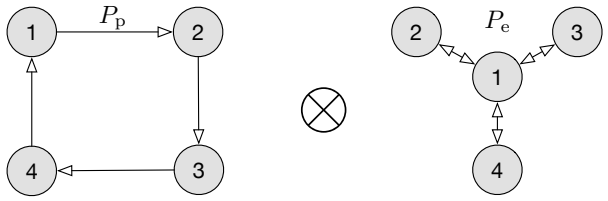

Figure 3.4: The periods associated with the pair of Markov Chains shown here are not co-prime; P_{p} is a period 4 chain and P_{e} is a period 2 chain. However, the meeting times are finite as they satisfy conditions in Theorem 1.

satisfied for nodes $(i,j) \in A_{q'}^{(\mathrm{p})} \times A_{r'}^{(\mathrm{e})}$ as $A_{q'}^{(\mathrm{p})} \cap A_{r'}^{(\mathrm{e})} \neq \phi$ and $\gcd(d_{q'}^{(\mathrm{p})}, d_{r'}^{(\mathrm{e})}) = 1$ for every $q' \in \{1,\ldots,q\}$ and every $r' \in \{1,\ldots,r\}$.

Second, consider nodes (i,j) such that i belongs to a transient class $T_{s'}^{(\mathrm{p})}$ where $s' \in \{1,\ldots,s\}$ and j belongs to a transient class $T_{t'}^{(\mathrm{e})}$ where $t' \in \{1,\ldots,t\}$. Since i belongs to a transient class, there must exist a walk to one of the absorbing classes, say $A_{q'}^{(\mathrm{p})}$. Similarly, since j belongs to a transient class, there must exist a walk to one of the absorbing classes, say $A_{r'}^{(\mathrm{e})}$. Hence by the proven result, for each node $(i,j) \in T_{s'}^{(\mathrm{p})} \times T_{t'}^{(\mathrm{e})}$ for every $s' \in \{1,\ldots,s\}$ and $t' \in \{1,\ldots,t\}$ there exists walks to a node of the form (k,k).

Finally, consider nodes (i,j) such that i belongs to a transient class $T_{s'}^{(\mathrm{p})}$ and j belongs to an absorbing class $A_{r'}^{(\mathrm{e})}$. Since i belongs to a transient class, there must exist a walk starting from i to an absorbing class, say $A_{q'}^{(\mathrm{p})}$. Thus once again we can apply the earlier stated result for nodes $(i,j) \in T_{s'}^{(\mathrm{p})} \times A_{r'}^{(\mathrm{e})}$ for every $s' \in \{1,\ldots,s\}$ and $r' \in \{1,\ldots,r\}$. Similarly, the case of nodes belong to absorbing classes in P_{p} and transient classes in P_{e} also follows.

Thus we have exhausted all pairs (i,j) in $P_{\mathrm{p}} \otimes P_{\mathrm{e}}$ and for each pair found a node of the form (k,k). Therefore P_{p} and P_{e} satisfy the conditions stated in statement (iii) of Theorem 1, hence guaranteeing finite meeting times and proving that $\mathcal{P}_{\text{all-overlap}} \subseteq S$. To show that $\mathcal{P}_{\text{all-overlap}} \neq S$, we present a counter-example in Figure 3.3.2. This concludes the proof for $\mathcal{P}_{\text{all-overlap}} \subset \mathcal{P}_{\text{finite}}$.

Now, we prove $\mathcal{P}_{\text{SA-overlap}} \subset \mathcal{P}_{\text{all-overlap}}$. The pairs of matrices $(P_{\mathrm{p}}, P_{\mathrm{e}}) \in$

$\mathcal{P}_{\text{SA-overlap}}$ is obtained by considering the subset of matrices which only have a single absorbing class. Thus $\mathcal{P}_{\text{SA-overlap}} \subset \mathcal{P}_{\text{all-overlap}}$.

Finally, to prove $\mathcal{P}_{\text{one-ergodic}} \subset \mathcal{P}_{\text{all-overlap}}$ let us assume without loss of generality that P_{p} is irreducible and aperiodic. This would imply that the entire nodeset V is an absorbing state and $d^{(\text{p})} = 1$. One can see that P_{p} paired with any other matrix P_{e} belongs to $\mathcal{P}_{\text{all-overlap}}$. Thus $\mathcal{P}_{\text{one-ergodic}} \subset \mathcal{P}_{\text{all-overlap}}$. ■

3.3.3 Mean meeting time and relation to hitting times.

Before we define the mean meeting time for two random walkers, we introduce a minor result.

Remark 1 *Consider two random walkers moving with transition matrices $P_{\text{p}}, P_{\text{e}}$ starting from nodes i, j respectively, then the meeting time*

$$m_{i,j} = (\mathbb{e}_1 \otimes \mathbb{e}_2)(I_{n^2} - (P_{\text{p}} \otimes P_{\text{e}})E)^{-1}\mathbb{1}_{n^2}. \tag{3.2}$$

Note that the expression above is a direct result of equation (3.1). entries m_{ii} as the value represents the meeting time assuming the agents do not meet at time $t = 0$.

We are now in a position to define the *mean meeting time* of two random walkers. Stationary distributions are well-defined for both P_{p} and P_{e} when each transition matrix has a single absorbing class. Further the meeting times for matrices with this property are finite only if the absorbing classes overlap and the periods are co-prime as is the case for pairs of transition matrices in $\mathcal{P}_{\text{SA-overlap}}$. Hence we have the following result.

Corollary 1 (Mean meeting time) *Consider two transition matrices $P_{\text{p}}, P_{\text{e}}$ with stationary distributions $\pi_{\text{p}}, \pi_{\text{e}}$. The mean meeting time*

$$\mathcal{M}(P_{\text{p}}, P_{\text{e}}) = \pi_{\text{p}}^{\top} M \pi_{\text{e}} = (\pi_{\text{p}} \otimes \pi_{\text{e}})^{\top} \operatorname{vec}(M), \tag{3.3}$$

where M is the matrix of meeting times, is finite if the pair of transition matrices $(P_{\mathrm{p}}, P_{\mathrm{e}}) \in \mathcal{P}_{\text{SA-overlap}}$.

Proof: The mean meeting time can be obtained from the meeting times as

$$\begin{aligned}\mathcal{M}(P_{\mathrm{p}}, P_{\mathrm{e}}) &= \sum_i \sum_j \pi_{\mathrm{p}}^{(i)} \pi_{\mathrm{e}}^{(j)} m_{i,j} \\ &= \sum_i \sum_j (\pi_{\mathrm{p}}^{(i)} \mathbb{e}_i \otimes \pi_{\mathrm{e}}^{(j)} \mathbb{e}_j)(I_{n^2} - (I_n \otimes P)E)^{-1} \mathbb{1}_{n^2} \\ &= (\pi_{\mathrm{p}} \otimes \pi_{\mathrm{e}})(I_{n^2} - (P_{\mathrm{p}} \otimes P_{\mathrm{e}})E)^{-1} \mathbb{1}_{n^2}.\end{aligned}$$

■

Further as the following result shows, the hitting times of a Markov chain are equal to the meeting times for the case of a mobile pursuer and stationary evader.

Corollary 2 (Connection to hitting times and meeting times with stationary evader) *Consider a stationary evader with distribution π_e and a pursuer with an irreducible transition matrix P_p and stationary distribution π_p, then the following properties hold:*

(i) the meeting times between the stationary evader and the pursuer are equal to the pairwise hitting times of P_p and are given by

$$h_{i,j} = m_{i,j} = (\mathbb{e}_1 \otimes \mathbb{e}_2)^\top (I_{n^2} - (I_n \otimes P_p)E)^{-1} \mathbb{1}_{n^2}, \tag{3.4}$$

where $h_{i,j}$ is the expected time to travel from node i to node j and

(ii) the mean meeting time between the stationary evader and the pursuer is given by

$$\mathcal{M}_{\text{stationary}}(\pi_e, P_p) = (\pi_e \otimes \pi_p)^\top (I_{n^2} - (I_n \otimes P_p)E)^{-1} \mathbb{1}_{n^2}. \tag{3.5}$$

Proof: A stationary evader can be described by the transition matrix I_n. However, note that the identity matrix has non-unique stationary distribution hence the evader stationary distribution can be arbitrarily defined given that $\sum_{i=1}^{n} \pi_e^{(i)} = 1$. Since P is irreducible, the pair of matrices (I_n, P) belongs to $\mathcal{P}_{\text{one-ergodic}}$ and hence meeting times are finite. Further the expression for meeting times in this context is identical to that of pairwise hitting times [53, Theorem 2.3(i)]. The mean first meeting time in such a case is

$$\begin{aligned}
\mathcal{M}_{\text{stationary}}(\pi_e, P_p) &= \sum_i \sum_j \pi_e^{(i)} \pi_p^{(j)} m_{i,j} \\
&= \sum_i \sum_j (\pi_e^{(i)} \mathbb{e}_i \otimes \pi_p^{(j)} \mathbb{e}_j)(I_{n^2} - (I_n \otimes P_p)E)^{-1} \mathbb{1}_n \\
&= (\pi_e \otimes \pi_p)(I_{n^2} - (I_n \otimes P_p)E)^{-1} \mathbb{1}_n.
\end{aligned}$$

■

When the stationary distribution of the evader is equal to the stationary distribution of the pursuer the expression for the meeting time is identical to the mean first passage time of the Markov chain P_p [53, Theorem 2.3(i)].

3.3.4 Comparison to existing bounds.

In this section we provide comparisons with existing bounds from literature, a summary of which is presented in Table 3.3.4. We present numerics for a variety of graphs and compare the exact value of the worst meeting time, denoted as $\mathcal{M}_{\max}$, with bounds on the same quantity from refs. [130, 120] and with the worst hitting time computed using the formula in [53] and also a bound on the worst hitting time as described in [131].

Most of the bounds discussed here are for random walks i.e., equal probability of transition from a node to every neighbor. The bounds by Aldous [122] also hold for

all reversible Markov chains. Hence in this section we consider transition matrices only corresponding to random walks. We include self-loops in all transition matrices to ensure aperiodicity. In general meeting times for transition matrices can be significantly smaller than the values discussed here. For example, using transition matrices which are permutation matrices one could obtain $O(n)$ meeting times on all graphs.

Bounding the worst meeting time as discussed in [122] in terms of the worst pairwise hitting time gives estimates which are of the same order. The computational complexity of exactly obtaining the worst hitting time is $O(n^3)$ [53, Theorem 2.3(i)] as compared to $O(n^6)$ for computing worst meeting times. Thus for small to medium graphs the worst hitting time can be a useful proxy.

The polynomial bound from Coppersmith et al [120, Theorem 3] is for sequential motion of the tokens i.e, one of the two tokens moves followed by the other. In order to compare this bound with the expression in equation (3.1) which is for simultaneous motion of the two random walkers, we divide the bound by two. This bound while easy to compute only provides a maximal estimate of the worst case meeting times.

The bound from Lovász [131, Corollary 3.3] is a bound for the worst hitting time. The bounds from Cooper et al [130, Theorem 1] and Lovász, both of which involve the spectral gap of the transition matrix, behave similarly in most cases. In general the estimates tend to be one or two orders of magnitude off. The complexity of computing the spectral gap can be cost-effective as this operation can be performed in worst-case $O(n^3)$ and for certain types of matrices in $O(n^2)$.

3.4 The meeting time of multiple random walkers

Here we extend the results from the previous section to the case of a group of pursuers trying to intercept a group of evaders.

	Meeting time			Hitting time	
Quantity	$\mathcal{M}_{\max}$	Bound by Cooper et al [130]	Bound by Coppersmith et al [120]	$\mathcal{H}_{\max}$ [53]	Bound by Lovász [131]
Complexity	$O(n^6)$	$O(n^3)$	$O(1)$	$O(n^3)$	$O(n^3)$
Ring	83.7	2488.8	856.0	150.0	2451.8
Path	174.8	9249.0	856.0	551.0	17308.6
Star	8.0	161.6	856.0	58.0	304.0
Lollipop	224.0	1376.3	856.0	483.8	2107.1
Lattice	35.9	805.6	856.0	83.7	1233.0
Random geometric graph (dense)	22.7	342.2	856.0	92.6	1098.8
Random geometric graph (sparse)	77.0	3587.1	856.0	319.6	10138.9

Table 3.1: Comparison of exact value of worst meeting time with bounds from literature and worst hitting times for random walks on various graphs of size 20 nodes. Values shown for random geometric graphs are averages over 100 instances.

3.4.1 Main result

Now consider L pursuers and M evaders. Let $X_t^{(\mathrm{p},1)}, X_t^{(\mathrm{p},2)}, \dots X_t^{(\mathrm{p},L)} \in \{1,\dots,n\}$ denote the locations of the L pursuers at time $t \in \{0,1,2,\dots\}$. Let $X_t^{(\mathrm{e},1)}, X_t^{(\mathrm{e},2)}, \dots X_t^{(\mathrm{e},M)} \in \{1,\dots,n\}$ denote the locations of the M evaders at time $t \in \{0,1,2,\dots\}$. For an L-tuple of nodes associated with the pursuers $(i_1, i_2, \dots i_L)$ and an M-tuple of nodes associated with the evaders $(j_1, j_2, \dots, j_M)$, *the first meeting time among L pursuers and M evaders*, denoted by $T_{i_1 i_2 \dots i_L, j_1 j_2 \dots j_M}$, is the first time that one of the pursuers meets one of the evaders. More formally, $T_{i_1 i_2 \dots i_L, j_1 j_2 \dots j_M}$ is

$$\min\{t \geq 1 | X_t^{(\mathrm{p},a)} = X_t^{(\mathrm{e},b)} \text{ for some } a \in \{1,\dots,L\} \text{ and } b \in \{1,\dots,M\}$$
$$\text{given that } X_0^{(\mathrm{p},l)} = i_l \,\forall\, l \in \{1,\dots,L\} \text{ and } X_0^{(\mathrm{e},m)} = j_m \,\forall\, m \in \{1,\dots,M\}\}.$$

Let the transition matrices associated with the L pursuers be $P_{\mathrm{p}}^{(1)}, P_{\mathrm{p}}^{(2)}, \dots, P_{\mathrm{p}}^{(L)}$ and the transition matrices associated with the M evaders be $P_{\mathrm{e}}^{(1)}, P_{\mathrm{e}}^{(2)}, \dots, P_{\mathrm{e}}^{(M)}$. The following theorem gives necessary and sufficient conditions for the the first expected meeting time between the L pursuers and M evaders $m_{i_1 i_2 \dots i_L, j_1 j_2 \dots j_M} = \mathbb{E}[T_{i_1 i_2 \dots i_L, j_1 j_2 \dots j_M}]$. For the sake

of brevity, we shall refer to the first expected meeting time in this context as the *group meeting time.*

Theorem 3 (The group meeting time of multiple Markov chains) *Consider Markov chains with transition matrices* $P_{\mathrm{p}}^{(1)}, P_{\mathrm{p}}^{(2)} \ldots, P_{\mathrm{p}}^{(L)}, P_{\mathrm{e}}^{(1)}, P_{\mathrm{e}}^{(2)}, \ldots, P_{\mathrm{e}}^{(M)}$ *defined on a digraph* G *with nodeset* $V = \{1, \ldots, n\}$*. The following statements are equivalent:*

(i) for every $i_1, i_2, \ldots i_L, j_1, j_2, \ldots, j_M \in \{1, \ldots, n\}$*, the group meeting time* $m_{i_1 i_2 \ldots i_L, j_1 j_2 \ldots j_M}$ *is finite,*

(ii) for every $i_1, i_2, \ldots, i_L, j_1, j_2, \ldots, j_M \in \{1, \ldots, n\}$*, there exists a node* k *and a length* ℓ *such that a walk of length* ℓ *exists from one of the nodes* $i_1, i_2, \ldots, i_L$ *to* k *in one of the transition matrices* $P_{\mathrm{p}}^{(1)}, P_{\mathrm{p}}^{(2)}, \ldots, P_{\mathrm{p}}^{(L)}$ *and a walk of length* ℓ *exists from one of the nodes* $j_1, j_2, \ldots, j_M$ *to* k *in one of the transition matrices in* $P_{\mathrm{e}}^{(1)}, P_{\mathrm{e}}^{(2)}, \ldots, P_{\mathrm{e}}^{(M)}$*,*

(iii) for every $i_1, i_2, \ldots, i_L, j_1, j_2, \ldots, j_M \in \{1, \ldots, n\}$*, there exists a walk in the digraph associated with the stochastic matrix* $P_{\mathrm{p}}^{(1)} \otimes P_{\mathrm{p}}^{(2)} \ldots \otimes P_{\mathrm{p}}^{(L)} \otimes P_{\mathrm{e}}^{(1)} \otimes P_{\mathrm{e}}^{(2)} \ldots \otimes P_{\mathrm{e}}^{(M)}$ *from a node* $(i_1, i_2, \ldots, i_L, j_1, j_2, \ldots, j_M)$ *to a node of the form* $(i_1^{'}, i_2^{'}, \ldots, k, \ldots, i_L^{'}, j_1^{'}, j_2^{'}, \ldots, k, \ldots, j_M^{'})$*, for some* $k \in \{1, \ldots, n\}$*, and*

(iv) the substochastic matrix PE *is convergent and the vector of group meeting times is given by*

$$\mathrm{vec}(M) = (I_{n^{L+M}} - PE)^{-1} \mathbb{1}_{n^{L+M}}, \tag{3.6}$$

where $M \in \mathbb{R}^{n^L \times n^M}$*,* $P = P_{\mathrm{p}}^{(1)} \otimes P_{\mathrm{p}}^{(2)} \ldots \otimes P_{\mathrm{p}}^{(L)} \otimes P_{\mathrm{e}}^{(1)} \otimes P_{\mathrm{e}}^{(2)} \ldots \otimes P_{\mathrm{e}}^{(M)}$ *and* E *is a binary diagonal matrix with entries* $\mathbb{1}_{n^{L+M}} - \mathrm{vec}([\delta_{i_1 i_2 \ldots i_L, j_1 j_2 \ldots j_M}])$*.*

Proof: For the nodes $i_1, i_2, \ldots, i_L, j_1, j_2, \ldots, j_M$, the group meeting time satisfies the recursive formula

$$T_{i_1 i_2 \ldots i_L, j_1 j_2 \ldots j_M} = \begin{cases} 1, & \text{w.p.} \sum_{k=1}^{n} \big(1 - \prod_{a=1}^{L}(1 - p_{i_a,k}^{(\mathrm{p},a)})\big)\big(1 - \prod_{b=1}^{M}(1 - p_{j_b,k}^{(\mathrm{e},b)})\big), \\ T_{k_1 k_2 \ldots k_L, h_1 h_2 \ldots h_M} + 1, & \text{w.p.} \sum_{k_a \neq h_b} \prod_{a=1}^{L} p_{i_a,k_a}^{(\mathrm{p},a)} \prod_{b=1}^{M} p_{j_b,h_b}^{(\mathrm{e},b)}. \end{cases}$$

Note that the symbol $\sum_{k_a \neq h_b}$ is a summation over the indices $k_1, k_2, \ldots, k_L, h_1, h_2, \ldots, h_M$ such that $k_a \neq h_b$ for every $a \in \{1, \ldots, L\}$ and $b \in \{1, \ldots, M\}$. The quantity $(1 - \prod_{a=1}^{L}(1 - p_{i_a,k}^{(\mathrm{p},a)}))$ indicates the probability that one of the pursuers will move to node k and $(1 - \prod_{b=1}^{M}(1 - p_{i_b,k}^{(\mathrm{e},b)}))$ indicates the probability that one of the evaders will move to node k. Therefore $\sum_{k=1}^{n}(1 - \prod_{a=1}^{L}(1 - p_{i_a,k}^{(\mathrm{p},a)}))(1 - \prod_{b=1}^{M}(1 - p_{j_b,k}^{(\mathrm{e},b)}))$ is the probability that one of the pursuers encounters one of the evaders at a common node.

Taking the expectation we have

$$\begin{aligned} \mathbb{E}[T_{i_1 i_2 \ldots i_L, j_1 j_2 \ldots j_M}] &= \sum_{k=1}^{n}(1 - \prod_{a=1}^{L}(1 - p_{i_a,k}^{(\mathrm{p},a)}))(1 - \prod_{b=1}^{M}(1 - p_{j_b,k}^{(\mathrm{e},b)})) \\ &\quad + \sum_{k_a \neq h_b} \prod_{a=1}^{L} p_{i_a,k_a}^{(\mathrm{p},a)} \prod_{b=1}^{M} p_{j_b,h_b}^{(\mathrm{e},b)} (\mathbb{E}[T_{k_1 k_2 \ldots k_L, h_1 h_2 \ldots h_M}] + 1), \\ \implies \mathbb{E}[T_{i_1 i_2, \ldots i_L, j_1 j_2 \ldots j_M}] &= 1 + \sum_{k_a \neq h_b} \prod_{a=1}^{L} p_{i_a,k_a}^{(\mathrm{p},a)} \prod_{b=1}^{M} p_{j_b,h_b}^{(\mathrm{e},b)} (\mathbb{E}[T_{k_1 k_2 \ldots k_L, h_1 h_2 \ldots h_M}]). \end{aligned}$$

Let $m_{i_1 i_2 \ldots i_L, j_1 j_2 \ldots j_M} = \mathbb{E}[T_{i_1 i_2 \ldots i_L, j_1 j_2 \ldots j_M}]$ for every $i_1, i_2, \ldots, i_L, j_1, j_2, \ldots, j_M \in \{1, \ldots, n\}$

and let $M = [m_{i_1 i_2 \dots i_l, j_1 j_2 \dots j_M}]$. Note that the entries of M can be written as

$$\begin{aligned} m_{i_1 i_2 \dots i_l, j_1 j_2 \dots j_M} &= 1 + \sum_{k_a \neq h_b} \prod_{a=1}^{L} p^{(\mathrm{p},a)}_{j_a,k_a} \prod_{b=1}^{M} p^{(\mathrm{e},b)}_{j_b,h_b} m_{k_1 k_2 \dots k_l, h_1 h_2 \dots h_M} \\ &= 1 + \sum_{\substack{k_1,k_2,\dots,k_L \\ h_1,h_2,\dots,h_M}}^{n} \prod_{a=1}^{L} p^{(\mathrm{p},a)}_{i_a,k_a} \prod_{b=1}^{M} p^{(\mathrm{e},b)}_{j_b,h_b} m_{k_1 k_2 \dots k_l, h_1 h_2 \dots h_M} - \\ &\qquad \sum_{\substack{k_1,k_2,\dots,k_L \\ h_1,h_2,\dots,h_M}} \delta_{i_1 i_2 \dots i_L, j_1 j_2 \dots j_M} \prod_{a=1}^{L} p^{(\mathrm{p},a)}_{i_a,k_a} \prod_{b=1}^{M} p^{(\mathrm{e},b)}_{j_b,h_b} m_{k_1 k_2 \dots k_l, h_1 h_2 \dots h_M}, \end{aligned}$$

where we have rewritten the summation $\sum_{k_a \neq h_b}$ in terms of the generalized kronecker delta function. This equation can be written in vector form as

$$\begin{aligned} \operatorname{vec}(M) =& \mathbb{1}_{n^{L+M}} + (P^{(1)}_{\mathrm{p}} \otimes P^{(2)}_{\mathrm{p}} \otimes \cdots \otimes P^{(L)}_{\mathrm{p}} \otimes P^{(1)}_{\mathrm{e}} \otimes P^{(2)}_{\mathrm{e}} \otimes \cdots \otimes P^{(M)}_{\mathrm{e}}) \operatorname{vec}[M] \\ & - (P^{(1)}_{\mathrm{p}} \otimes P^{(2)}_{\mathrm{p}} \otimes \cdots \otimes P^{(L)}_{\mathrm{p}} \otimes P^{(1)}_{\mathrm{e}} \otimes P^{(2)}_{\mathrm{e}} \otimes \cdots \otimes P^{(M)}_{\mathrm{e}}) \\ & \operatorname{vec}([\delta_{i_1 i_2 \dots i_L, j_1 j_2 \dots j_M}]) \operatorname{vec}[M], \end{aligned}$$

$$\implies \operatorname{vec}(M) = \mathbb{1}_{n^{L+M}} + P(I_{n^{L+M}} - \operatorname{vec}([\delta_{i_1 i_2 \dots i_L, j_1 j_2 \dots j_M}])) \operatorname{vec}[M],$$

$$\implies \operatorname{vec}(M) = \mathbb{1}_{n^{L+M}} + PE \operatorname{vec}(M).$$

If the matrix $I_{n^{L+M}} - PE$ is invertible then we have a unique solution to the meeting times. We shall now show that the finiteness of group meeting times as in (i) is equivalent to the existence of walks of equal length to common nodes as mentioned in (ii) and in (iii), which leads to invertibility of $I_{n^{L+M}} - PE$ in (iv).

We start by proving $(i) \implies (ii)$. If we assume that $(i) \not\implies (ii)$ then there exists an L-tuple of nodes $i_1, i_2, \dots, i_L$ and an M-tuple of nodes $j_1, j_2, \dots, j_M$ such that the group meeting time between groups of agents starting from these positions is finite and there exists no walk of equal length to a common node for any possible pursuer-evader pairs. However if there exists no walk of equal length from one of the nodes in $i_1, i_2, \dots, i_L$ and

one of the nodes in $j_1, j_2, \ldots, j_M$ to a common node, then none of the agents ever meet and the group meeting time is infinite. Hence by contradiction $(i) \implies (ii)$.

Next we show that $(ii) \implies (iii)$. The Kronecker product of the L pursuer transition matrices and the M evader transition matrices gives a joint transition matrix for the agents. The node $(i_1, i_2, \ldots, i_L, j_1, j_2, \ldots, j_M)$ corresponds to the states $X_{\mathrm{p}}^{(1)} = i_1, X_{\mathrm{p}}^{(2)} = i_2, \ldots, X_{\mathrm{p}}^{(L)} = i_L$ and $X_{\mathrm{e}}^{(1)} = j_1, X_{\mathrm{e}}^{(2)} = j_2, \ldots, X_{\mathrm{e}}^{(M)} = j_M$. Statement (ii) ensures that there exists a node k to which there is a walk of length ℓ from one of the nodes $i_1, i_2, \ldots, i_L$ in a pursuer transition matrix $P_{\mathrm{p}}^{L'}$ and from one of the nodes $j_1, j_2, \ldots, j_M$ in an evader transition matrix $P_{\mathrm{e}}^{M'}$. Starting from any node i_l for any $l \in \{1, \ldots, L\}$ there exists a node $i_l^{'}$ such that there exists a walk of length ℓ from i_l to $i_l^{'}$ in the transition matrix $P_l^{(\mathrm{p})}$. Similarly starting from j_m there exists some node $j_m^{'}$ to which there exists a walk of length ℓ for some $m \in \{1, \ldots, M\}$ in the transition matrix $P_m^{(\mathrm{e})}$. Thus there exists walks of length ℓ : $i_1 \to i_1^{'}$, $i_2 \to i_2^{'}$, $\ldots$, $i_{L'} \to k$, $\ldots$, $i_L \to i_L^{'}$, $j_1 \to j_1^{'}$, $j_2 \to j_2^{'}$, $\ldots$, $j_{M'} \to k$, $\ldots$, and $j_M \to j_M^{'}$. Using Lemma 3 there exists a walk from $(i_1, i_2, \ldots, i_{L'}, \ldots, i_L, j_1, j_2, \ldots, j_{M'}, \ldots, j_M)$ to a node of the form $(i_1^{'}, i_2^{'}, \ldots, k, \ldots, i_L^{'}, j_1^{'}, j_2^{'}, \ldots, k, \ldots, j_M^{'})$, thus proving $(ii) \implies (iii)$.

Next, we show $(iii) \implies (iv)$. Note that post-multiplying the Kronecker product of all transition matrices $P_{\mathrm{p}}^{(1)} \otimes P_{\mathrm{p}}^{(2)} \otimes \ldots \otimes P_{\mathrm{p}}^{(L)} \otimes P_{\mathrm{e}}^{(1)} \otimes P_{\mathrm{e}}^{(2)} \otimes \ldots \otimes P_{\mathrm{e}}^{(M)}$ by E sets columns associated with nodes of the form $(i_1^{'}, i_2^{'}, \ldots, k, \ldots, i_L^{'}, j_1^{'}, j_2^{'}, \ldots, k, \ldots, j_M^{'})$ to $\mathbb{0}_{n^{L+M}}$. Therefore every node has a walk to a node whose row-sum is less than 1 which implies that the matrix PE is convergent. From this we obtain equation (3.6). Since (iii) guarantees the existence of $(I_{n^{L+M}} - PE)^{-1}$, we prove that $(iii) \implies (iv)$.

Note that the existence of $\mathrm{vec}(M)$ in (iv) gives $(iv) \implies (i)$. Thus we have shown that $(i) \implies (ii) \implies (iii) \implies (iv) \implies (i)$. Hence the four conditions are equivalent. ■

The above necessary and sufficient conditions give the most general set of tuples of

matrices for which finite meeting times exist. Similar to the single pursuer and single evader case, we present sufficient conditions on the transition matrices which ensure that the meeting times between two groups is finite.

3.4.2 Sufficient conditions for finiteness and mean group meeting time.

Consider the following sets of $L+M$-tuples of matrices:

$\mathcal{P}^{\mathbf{L,M}}_{\mathbf{finite}}$**: finite group meeting times.** Let $\mathcal{P}^{\mathrm{L,M}}_{\mathrm{finite}}$ be the set of $L+M$-tuples of transition matrices $P_{\mathrm{p}}^{(1)}, P_{\mathrm{p}}^{(2)}, \ldots, P_{\mathrm{p}}^{(L)}, P_{\mathrm{e}}^{(1)}, P_{\mathrm{e}}^{(2)}, \ldots, P_{\mathrm{e}}^{(M)}$ satisfying the conditions stated in Theorem 3 and therefore having finite group meeting times.

$\mathcal{P}^{\mathbf{L,M}}_{\mathbf{all\text{-}overlap}}$**: Markov chains with all-to-all overlapping absorbing classes.** Let $\mathcal{P}^{\mathrm{L,M}}_{\mathrm{all\text{-}overlap}}$ be the set of $L+M$-tuples of transition matrices $P_{\mathrm{p}}^{(1)}, P_{\mathrm{p}}^{(2)}, \ldots, P_{\mathrm{p}}^{(L)}, P_{\mathrm{e}}^{(1)}, P_{\mathrm{e}}^{(2)}, \ldots, P_{\mathrm{e}}^{(M)}$ with the following property: for each transition matrix $P_{\mathrm{p}}^{(l)}, l \in \{1, \ldots, L\}$ there exists a transition matrix $P_{\mathrm{e}}^{(m)}$ for some $m \in \{1, \ldots, M\}$ such that $(P_{\mathrm{e}}^{(l)}, P_{\mathrm{p}}^{(m)}) \in \mathcal{P}_{\mathrm{all\text{-}overlap}}$.

$\mathcal{P}^{\mathbf{L,M}}_{\mathbf{SA\text{-}overlap}}$**: single absorbing Markov chains with overlapping absorbing classes.** Let $\mathcal{P}^{\mathrm{L,M}}_{\mathrm{SA\text{-}overlap}}$ be the set of $L+M$-tuples of transition matrices $P_{\mathrm{p}}^{(1)}, P_{\mathrm{p}}^{(2)}, \ldots, P_{\mathrm{p}}^{(L)}, P_{\mathrm{e}}^{(1)}, P_{\mathrm{e}}^{(2)}, \ldots, P_{\mathrm{e}}^{(M)}$ with the following property: for each transition matrix $P_{\mathrm{p}}^{(l)}, l \in \{1, \ldots, L\}$ there exists a transition matrix $P_{\mathrm{e}}^{(m)}$ for some $m \in \{1, \ldots, M\}$ such that $(P_{\mathrm{e}}^{(l)}, P_{\mathrm{p}}^{(m)}) \in \mathcal{P}_{\mathrm{SA\text{-}overlap}}$.

$\mathcal{P}^{\mathbf{L,M}}_{\mathbf{one\text{-}ergodic}}$**: one ergodic Markov chain.** Let $\mathcal{P}^{\mathrm{L,M}}_{\mathrm{one\text{-}ergodic}}$ be the set of $L+M$-tuples of transition matrices $P_{\mathrm{p}}^{(1)}, P_{\mathrm{p}}^{(2)}, \ldots, P_{\mathrm{p}}^{(L)}, P_{\mathrm{e}}^{(1)}, P_{\mathrm{e}}^{(2)}, \ldots, P_{\mathrm{e}}^{(M)}$ such that one of the transition matrices is ergodic.

Given the above, we are now in a position to define the *mean group meeting time* of two sets of random walkers, L pursuers and M evaders. The group meeting times for matrices with single absorbing classes are finite when the $L+M$-tuple $(P_{\mathrm{p}}^{(1)}, P_{\mathrm{p}}^{(2)}, \ldots, P_{\mathrm{p}}^{(L)}, P_{\mathrm{e}}^{(1)}, P_{\mathrm{e}}^{(2)}, \ldots, P_{\mathrm{e}}^{(M)}) \in \mathcal{P}^{\mathrm{L,M}}_{\text{SA-overlap}}$. Hence, we have the following result.

Corollary 3 (Mean group meeting time) *Consider $L+M$ transition matrices $P_{\mathrm{p}}^{(1)}, P_{\mathrm{p}}^{(2)}, \ldots, P_{\mathrm{p}}^{(L)}, P_{\mathrm{e}}^{(1)}, P_{\mathrm{e}}^{(2)}, \ldots, P_{\mathrm{e}}^{(M)}$ with stationary distributions $\pi_{\mathrm{p}}^{(1)}, \pi_{\mathrm{p}}^{(2)}, \ldots, \pi_{\mathrm{p}}^{(L)}, \pi_{\mathrm{e}}^{(1)}, \pi_{\mathrm{e}}^{(2)}, \ldots, \pi_{\mathrm{e}}^{(M)}$. The mean group meeting time*

$$\mathcal{M}_{L,M} = (\pi_{\mathrm{p}}^{(1)} \otimes \pi_{\mathrm{p}}^{(2)} \otimes \ldots \otimes \pi_{\mathrm{p}}^{(L)} \otimes \pi_{\mathrm{e}}^{(1)} \otimes \pi_{\mathrm{e}}^{(2)} \otimes \ldots \otimes \pi_{\mathrm{e}}^{(M)})^{\top} \operatorname{vec}(M), \tag{3.7}$$

where M is the matrix of group meeting times, is finite if the $L+M$-tuple of transition matrices $(P_{\mathrm{p}}^{(1)}, P_{\mathrm{p}}^{(2)}, \ldots, P_{\mathrm{p}}^{(L)}, P_{\mathrm{e}}^{(1)}, P_{\mathrm{e}}^{(2)}, \ldots, P_{\mathrm{e}}^{(M)}) \in \mathcal{P}^{L,M}_{\text{SA-overlap}}$.

A word on the computational complexity for the multiple pursuer-evader case: since the general expression for the group meeting time among groups of pursuers and evaders involve extensive use of the Kronecker product, the memory and computational resources necessary are significantly affected by the curse of dimensionality. The matrix $(P_{\mathrm{p}}^{(1)} \otimes P_{\mathrm{p}}^{(2)} \otimes \ldots \otimes P_{\mathrm{p}}^{(L)} \otimes P_{\mathrm{e}}^{(1)} \otimes P_{\mathrm{e}}^{(2)} \otimes \ldots \otimes P_{\mathrm{e}}^{(M)})$ contains $n^{(L+M)}$ elements. Inversion of a full matrix would require $O(k^3)$ operations lending an undesirable complexity of $O(n^{3(L+M)})$ [132]. Most practical solutions to the transition matrices benefit from the sparse nature of the graphs. A sparse system of equations can be solved with complexity $O(nnz)$ where nnz is the number of non-zero elements. For the Kronecker product of $L+M$ transition matrices defined on the same graph the number of non-zero elements is $|E|^{(L+M)}$, and with a sparse solver the group meeting time can be computed in $O(|E|^{(L+M)})$ operations. For further details see [133].

3.5 Meeting times for continous-time Markov chains

In this section we formulate the meeting time between two continuous-time Markov chains. The essence of the proof is to use the fact that the meeting time of two continuous-time Markov chains is equivalent to the hitting time on the joint transition matrix generated by the Kronecker product of the digraphs associated with the two chains.

3.5.1 The meeting time of two continuous-time Markov chains

Consider the pursuer and evader performing random walks on a set of nodes $V := \{1, \ldots, n\}$ with digraphs $G_{\mathrm{p}} = (V, E_{\mathrm{p}}, Q_{\mathrm{p}})$, $G_{\mathrm{e}} = (V, E_{\mathrm{e}}, Q_{\mathrm{e}})$, edge sets $E_{\mathrm{p}}, E_{\mathrm{e}} \subset V \times V$, and transition rate matrices Q_{p}, Q_{e}. Let $P_{\mathrm{p}}(t), P_{\mathrm{e}}(t)$ denote the transition matrices of the purser and evader at time t. Let $m_{i,j}$ denote the expected first meeting time for a pursuer starting from node i and an evader starting from node j, which shall be referred to simply as the meeting time. Then the following theorem holds.

Theorem 4 (The meeting time of two continuous-time Markov chains) *Consider two Markov chains with transition rate matrices Q_{p} and Q_{e} defined on a digraph G with nodeset $V = \{1, \ldots, n\}$. The following statements are equivalent:*

(i) for each pair of nodes i, j, the expected first meeting time $m_{i,j}$ from nodes i and j is finite,

(ii) for each pair of nodes i, j, there exists a node k such that a walk exists from i to k and a walk exists from j to k,

(iii) for each pair of nodes i, j, there exists a walk in the digraph associated with the transition rate matrix $Q_{\mathrm{p}} \otimes I_n + I_n \otimes Q_{\mathrm{e}}$ from (i, j) to a node (k, k), for some $k \in \{1, \ldots, n\}$, and

(iv) the matrix $E(I_{n^2} - (Q_{\mathrm{p}} \otimes I_n + I_n \otimes Q_{\mathrm{e}})) - I_{n^2}$ *is invertible and the first meeting times are given by the unique solution to*

$$\mathrm{vec}(M) = (E(I_{n^2} - (Q_{\mathrm{p}} \otimes I_n + I_n \otimes Q_{\mathrm{e}})) - I_{n^2})^{-1} E \mathbb{1}_{n^2} \tag{3.8}$$

where $M \in \mathbb{R}^{n \times n}$ *and* $E \in \mathbb{R}^{n^2 \times n^2}$ *is a binary diagonal matrix with diagonal entries* $\mathbb{1}_{n^2} - \mathrm{vec}(I_n)$.

Proof: Consider the joint evolution of the two continuous-time Markov chains on the Kronecker product graph given by the Kronecker product $P_{\mathrm{p}}(t) \otimes P_{\mathrm{e}}(t)$. The transition rate matrix for this Markov chain is easy to derive. Consider

$$\begin{aligned}\frac{d}{dt}(P_{\mathrm{p}}(t) \otimes P_{\mathrm{e}}(t)) =& P_{\mathrm{p}}(t)Q_{\mathrm{p}} \otimes P_{\mathrm{e}}(t) + P_{\mathrm{p}}(t) \otimes P_{\mathrm{e}}(t)Q_{\mathrm{e}} \\ =& (P_{\mathrm{p}}(t) \otimes P_{\mathrm{e}}(t))(Q_{\mathrm{p}} \otimes I_n + I_n \otimes Q_{\mathrm{e}}),\end{aligned}$$

where we have used the product rule for derivatives and Lemma 1 to obtain the joint transition rate matrix as $Q_{\mathrm{p}} \otimes I_n + I_n \otimes Q_{\mathrm{e}}$.

The $n \times n$ block entries of the joint transition rate matrix

$$Q_{\mathrm{p}} \otimes I_n + I_n \otimes Q_{\mathrm{e}} = \begin{bmatrix} q_{1,1}^{(\mathrm{p})} I_n + Q_{\mathrm{e}} & q_{1,2}^{(\mathrm{p})} I_n & \dots & q_{1,n}^{(\mathrm{p})} I_n \\ q_{2,1}^{(\mathrm{p})} I_n & q_{2,2}^{(\mathrm{p})} I_n + Q_{\mathrm{e}} & \dots & q_{2,n}^{(\mathrm{p})} I_n \\ \vdots & \vdots & \ddots & \vdots \\ q_{n,1}^{(\mathrm{p})} I_n & q_{n,2}^{(\mathrm{p})} I_n & \dots & q_{n,n}^{(\mathrm{p})} I_n + Q_{\mathrm{e}} \end{bmatrix}.$$

The meeting times for the two transition rate matrices correspond to hitting times from nodes on the joint transition rate matrix $Q_{\mathrm{p}} \otimes I_n + I_n \otimes Q_{\mathrm{e}}$ to the set of common nodes of the form (k, k). The solution to hitting times for continuous-time Markov chains is given in [125, Theorem 3.3.3]. We restate the result here for the sake of completeness.

Given a transition rate matrix $Q = [q_{a,b}]$ defined on a set of nodes A and a subset $S \subset A$, the expected meeting times starting from a node $a \in A$ to the set S denoted by h_a^S is given by the solution to the system of equations

$$\begin{cases} h_a^S = 0 & \text{for } a \in S \\ -\sum_{b \in A} q_{a,b} h_b^S = 1 & \text{for } a \notin S. \end{cases} \tag{3.9}$$

The meeting times can be obtained as the solution to the system of equations above with transition rate matrix given by $Q_{\mathrm{p}} \otimes I_n + I_n \otimes Q_{\mathrm{e}}$ and $S = \{(k,k) \,|\, k \in V\}$. Denoting $Q_{\mathrm{p}} \otimes I_n + I_n \otimes Q_{\mathrm{e}}$ and $S = \{(k,k) \,|\, k \in V\}$ by Q^{eff} and S_{common}, respectively, the system of equations in (3.9) can be written as

$$\begin{cases} m_{i,j} = 0 & \text{for } (i,j) \in S_{\mathrm{common}} \\ -\sum_{k \in V} \sum_{h \in V} Q^{\mathrm{eff}}_{(i,j),(k,h)} m_{k,h} = 1 & \text{for } (i,j) \notin S_{\mathrm{common}}. \end{cases}$$

These equations can be re-written in vector form as

$$-EQ^{\mathrm{eff}} \operatorname{vec}(M) = E\mathbb{1}_{n^2}, \qquad (E - I_{n^2}) \operatorname{vec}(M) = \mathbb{0}_{n^2}.$$

Adding the above two equations we obtain equation (3.8). If the matrix $E(I_{n^2} - Q^{\mathrm{eff}}) - I_{n^2}$ is invertible then we have a unique solution to the meeting times. We shall now show that the finiteness of meeting times as in (i) is equivalent to the existence of walks to common nodes as mentioned in (ii) and (iii), which leads to invertibility of $E(I_{n^2} - Q^{\mathrm{eff}}) - I_{n^2}$ in (iv).

We start by proving that $(i) \implies (ii)$. If we assume that $(i) \not\Longrightarrow (ii)$, then there exists a pair of nodes i and j such that the expected first meeting time is finite and there exists no walk to a common node in V. However if there exists no walk to a common

node, then the agents never meet and the first meeting time is always infinite. Hence by contradiction $(i) \implies (ii)$.

Now we shall prove that $(ii) \implies (iii)$. Since the matrix Q_{e} is added to every diagonal block of the joint transition rate matrix if there exists a walk in the transition rate matrix Q_{e} from j to k then there exists a walk in $Q_{\mathrm{p}} \otimes I_n + I_n \otimes Q_{\mathrm{e}}$ from (i,j) to (i,k) for every $i \in \{1,\ldots,n\}$. Also note that the off diagonal block elements are of the form $q^{(p)}_{i',j'} I_n$. One can verify that because of this structure if there is a walk from $j \to k$ in Q_{p} then in $Q_{\mathrm{p}} \otimes I_n + I_n \otimes Q_{\mathrm{e}}$ there exists a walk from $(i,j) \to (k,j)$ for every $j \in \{1,\ldots,n\}$. Hence if there is a walk $i \to k$ in Q_{p} and $j \to k$ in Q_{e} then there exists walks $(i,j) \to (i,k) \to (k,k)$, thus proving $(ii) \implies (iii)$.

Finally we shall prove $(iii) \implies (iv)$. First consider the modified transition rate matrix $E(Q_{\mathrm{p}} \otimes I_n + I_n \otimes Q_{\mathrm{e}})$. The matrix E sets the rows corresponding to nodes of the form (j,j) to $\mathbb{0}^{\top}_{n^2}$. The rank of a transition rate matrix is $n-d$ where d is the number of sinks in the transition rate matrix [134]. The matrix $E(Q_{\mathrm{p}} \otimes I_n + I_n \otimes Q_{\mathrm{e}})$ has at least n sinks corresponding to the elements (k,k) for every $j \in \{1,\ldots,n\}$. If every node has a path to a node of the form (k,k) as in (iii), then there are only exactly n sinks. This is because there are exactly n nodes of the form (k,k). Thus the rank of $E(Q_{\mathrm{p}} \otimes I_n + I_n \otimes Q_{\mathrm{e}})$ is $n^2 - n$ implying that this matrix has n null eigenvectors. One can verify that the null eigenvectors (and basis vectors for the kernel) are given by $\mathbb{e}_1, \mathbb{e}_{n+1}, \mathbb{e}_{2n+2}, \ldots, \mathbb{e}_{n^2}$. Let the other eigenvectors be $v_1, v_2, \ldots, v_{n^2-n}$. Since the kernel of $E(Q_{\mathrm{p}} \otimes I_n + I_n \otimes Q_{\mathrm{e}})$ is spanned by $\mathbb{e}_1, \mathbb{e}_{n+1}, \mathbb{e}_{2n+2}, \ldots, \mathbb{e}_{n^2}$, the eigenvectors of the same matrix can be uniquely constructed by ensuring they are orthogonal to the kernel, i.e. $v_p^{\top} \mathbb{e}_q = 0$ for every $p \in \{1,\ldots,n^2-n\}$ and $q \in \{1, n+1, \ldots, n^2\}$. Let us denote the eigenvalues associated with these eigenvectors as $\lambda_1, \lambda_2, \ldots, \lambda_{n^2-n}$. Consider the matrix $E(Q_{\mathrm{p}} \otimes I_n + I_n \otimes Q_{\mathrm{e}}) + (I_{n^2} - E)$. We shall show that $E(Q_{\mathrm{p}} \otimes I_n + I_n \otimes Q_{\mathrm{e}}) + (I_{n^2} - E)$ has the same eigenvectors as $E(Q_{\mathrm{p}} \otimes I_n + I_n \otimes Q_{\mathrm{e}})$. It is easy to see that now the

eigenvectors $\mathbb{e}_1, \mathbb{e}_{n+1}, \ldots, \mathbb{e}_{n^2}$ have eigenvalues 1. One can verify v_q is still an eigenvector but with eigenvalue $\lambda_q + 1$. Note that $E(Q_\mathrm{p} \otimes I_n + I_n \otimes Q_\mathrm{e})$ is positive semi-definite from Gershgorin's disk theorem [124]. Since $E(Q_\mathrm{p} \otimes I_n + I_n \otimes Q_\mathrm{e})$ has all non-negative eigenvalues we are assured that $E(Q_\mathrm{p} \otimes I_n + I_n \otimes Q_\mathrm{e}) + (I_{n^2} - E)$ has all positive eigenvalues and is invertible. Thus if (iii) holds $-(E(Q_\mathrm{p} \otimes I_n + I_n \otimes Q_\mathrm{e}) + (I_{n^2} - E))$ has full rank and is invertible. Thus equation (3.8) gives the unique solution to the meeting times. Therefore $(iii) \implies (iv)$.

Note that the existence of $\mathrm{vec}(M)$ in (iv) gives $(iv) \implies (i)$. Thus we have shown that $(i) \implies (ii) \implies (iii) \implies (iv) \implies (i)$. Hence the four conditions are equivalent. ■

One can derive sets of pairs of transition rate matrices for which meeting times are guaranteed to be finite akin to the discrete-time case: $\mathcal{Q}_{\text{all-overlap}}$, $\mathcal{Q}_{\text{SA-overlap}}$ and $\mathcal{Q}_{\text{one-ergodic}}$. The sets are almost identical in description except for the fact that periodicity conditions are no longer necessary. A notion of mean meeting time is applicable to the set of transition rate matrices belonging to $\mathcal{Q}_{\text{SA-overlap}}$.

3.5.2 The group meeting times of multiple continuous-time Markov chains.

The setup for multiple pursuers and evaders following continuous-time Markov chains on a common graph is identical to the multiple pursuers and multiple evaders in the discrete time case. Consider pursuer transition rate matrices $Q_\mathrm{p}^{(1)}, Q_\mathrm{p}^{(2)}, \ldots, Q_\mathrm{p}^{(L)}$ and evader transition rate matrices $Q_\mathrm{e}^{(1)}, Q_\mathrm{e}^{(2)}, \ldots, Q_\mathrm{e}^{(M)}$.

Theorem 5 (The group meeting time of multiple continuous-time Markov chains) *Consider Markov chains with transition rate matrices $Q_\mathrm{p}^{(1)}, Q_\mathrm{p}^{(2)}, \ldots, Q_\mathrm{p}^{(L)}$, $Q_\mathrm{e}^{(1)}$, $Q_\mathrm{e}^{(2)}$, $\ldots, Q_\mathrm{e}^{(M)}$ defined on a digraph G with nodeset $V = \{1, \ldots, n\}$. The following statements*

are equivalent:

(i) for every $i_1, i_2, \ldots i_L, j_1, j_2, \ldots, j_M \in \{1, \ldots, n\}$*, the expected first meeting time* $m_{i_1 i_2 \ldots i_L, j_1 j_2 \ldots j_M}$ *is finite,*

(ii) for every $i_1, i_2, \ldots, i_L, j_1, j_2, \ldots, j_M \in \{1, \ldots, n\}$*, there exists a node* k *such that there exists a walk from one of the nodes* $i_1, i_2, \ldots, i_L$ *to* k *in one of the transition rate matrices* $Q_{\mathrm{p}}^{(1)}, Q_{\mathrm{p}}^{(2)}, \ldots, Q_{\mathrm{p}}^{(L)}$ *and a walk exists from one of the nodes* $j_1, j_2, \ldots, j_M$ *to* k *in one of the transition matrices in* $Q_{\mathrm{e}}^{(1)}$, $Q_{\mathrm{e}}^{(2)}$, $\ldots$, $Q_{\mathrm{e}}^{(L)}$,

(iii) for every $i_1, i_2, \ldots, i_L, j_1, j_2, \ldots, j_M \in \{1, \ldots, n\}$*, there exists a walk in the digraph associated with the transition rate matrix* $\sum_{l=1}^{L} I_{n^{l-1}} \otimes Q_{\mathrm{p}}^{(l)} \otimes I_{n^{L+M-l}} + \sum_{m=1}^{M} I_{n^{L+m-1}} \otimes Q_{\mathrm{e}}^{(m)} \otimes I_{n^{M-m}}$ *from a node* $(i_1, i_2, \ldots, i_L, j_1, j_2, \ldots, j_M)$ *to a node of the form* $(i_1^{'}, i_2^{'}, \ldots, k, \ldots, i_L^{'}, j_1^{'}, j_2^{'}, \ldots, k, \ldots, j_M^{'})$*, for some* $k \in \{1, \ldots, n\}$*, and*

(iv) the matrix $E(I_{n^{L+M}} - Q) - I_{n^{L+M}}$ *is invertible and the expected first meeting time is given by*

$$\mathrm{vec}(M) = (E(I_{n^{L+M}} - Q) - I_{n^{L+M}})^{-1} \mathbb{1}_{n^{L+M}}, \tag{3.10}$$

where $M \in \mathbb{R}^{n^L \times n^M}$, $\sum_{l=1}^{L} I_{n^{l-1}} \otimes Q_{\mathrm{p}}^{(l)} \otimes I_{n^{L+M-l}} + \sum_{m=1}^{M} I_{n^{L+m-1}} \otimes Q_{\mathrm{e}}^{(m)} \otimes I_{n^{M-m}}$ *and* E *is a binary diagonal matrix with entries* $\mathbb{1}_{n^{L+M}} - \mathrm{vec}([\delta_{i_1 i_2 \ldots i_L, j_1 j_2 \ldots j_M}])$.

We state this result without proof as it utilizes the same technique as in the proof of Theorem 4. The proof of this result involves constructing the joint transition rate matrix of all agents on the Kronecker digraph and then computing the hitting time to the set of $(L + M)$ tuples of nodes such that one of the first L entries is the same as one of the next M entries. The complexity of computing meeting times using equation (3.10) for

continuous-time Markov chains is $O(n^{3(L+M)})$ as it involves inversion of a matrix which has n^{L+M} elements, which is identical to the discrete-time case.

3.6 Summary

We have studied the meeting time of multiple random walkers on a graph and have presented necessary and sufficient conditions for finiteness and novel closed-form expressions for the expected time to meeting between a single pursuer and a single evader, multiple pursuers and multiple evaders, and extended the treatment to continuous-time chains. We also provide sufficient conditions for certain pairs (or tuples) of Markov chains that satisfy conditions on their absorbing classes to have finite meeting times. Finally, we discuss connections to other metrics relevant to Markov chains such as the hitting time.

Chapter 4

Markov chains with Maximum Entropy Rate

4.1 Introduction

4.1.1 Problem description

The entropy rate of a Markov chain is a measure of information and unpredictability generated with each time-step [63]. In this chapter, we study Markov chains with maximal entropy generation subject to two constraints: (i) allowable transitions are specified by a given irreducible adjacency matrix and (ii) the stationary distribution of the Markov chain is given. It is customary to refer to Markov chains with maximum entropy rate as *maxentropic*. Maxentropic Markov chains with stationary distribution constraints are of interest in surveillance strategies as they maximize the uncertainty in the path of the surveillance agent. Aside from applications to stochastic surveillance, the notion of maxentropic Markov chains is useful for example in link-prediction [135], community detection [136] and image processing [137].

4.1.2 Prior applications of maxentropic Markov chains

To the best of our knowledge, maxentropic Markov chains first appeared in [64] as the solution to the optimization problem of maximizing the entropy rate given the first and second moments of the Markov chain. More recently, Burda et al. [65] provide a closed form solution for maxentropic Markov chains subject solely to graph constraints. This Markov chain, referred to as the *maximal entropy random walk* (MERW), possesses the property that all walks of equal length with given start and end node are equiprobable. The solution we provide is for Markov chains subject to stationary distribution constraints in addition to graph constraints. In what follows, we review prior applications of maxentropic Markov chains: (i) detection of features in images and (ii) design of metrics on large graphs and complex networks.

Image analysis. Based of the notion of maximal entropy random walks in [65], several applications have been proposed in image analysis. The MERW is utilized instead of the equal neighbor random walk to detect visually salient features in [137]. The MERW has also been utilized to implement a probabilistic object localization scheme in [138]. Korus and Huang [139] successfully adopt the MERW for localizing forgeries in digital images.

Metrics on large networks. The MERW is used to design unsupervised methods for link prediction in [135]. Ochab and Burda study the feasibility of using the MERW in algorithms for community detection [136]. Furthermore, the MERW is used to study the trapping problem in dendrimers, i.e., artificial macromolecules with treelike structures [140]. More recently, a relation between entropy rate and congestion in complex networks was established and a method was proposed to mitigate congestion using MERW in [141].

4.1.3 Relevance to stochastic surveillance

The setup we consider is one in which the area to be surveilled has been sampled to obtain a robotic roadmap represented by a graph. The nodes of the graph designate points of high priority and the edges indicate whether it is possible to move between different nodes. (Restrictions might be imposed by obstacles, no-fly zones, etc.) The graph structure is captured by a binary adjacency matrix and the relative importance of each node is given by a normalized vector which indicates a desired visit frequency to each node. Markov chains modeled by transition matrices are well suited to designing random walks on graphs with visit frequency constraints. The left-dominant eigenvector of the transition matrix, referred to as the stationary distribution, gives the visit frequency of a random walker who moves according to the Markov chain. Graph and stationary distribution constraints are linear and hence can be enforced quite effectively in optimization problems involving cost functions with various robotic motivations such as maximizing speed of traversal, minimizing the expected reward for an intruder or convergence to a desired swarm formation [52, 54, 142].

While prior work with the same framework has emphasized the speed of the Markov chain or optimizing the probability of capture given an intruder model, the transition matrices obtained as solutions to such formulations need not necessarily be unpredictable (e.g., permutation matrices, which have zero entropy rate, are the fastest Markov chains when a Hamiltonian tour exists). The notion of maximum entropy rate Markov chains is valuable as it translates directly to maximum unpredictability in the path of the surveillance agent. The specification of a stationary distribution, which serves as a prior for where the intruder might be located, makes our approach more suited than the MERW which has a fixed stationary distribution.

4.1.4 Applications in other areas

The methods described in this chapter are potentially useful for developing novel methods of conducting image analysis. The maxentropic Markov chain with visit frequency specification provides a natural way of incorporating prior knowledge of where an object or an anomaly is likely to be located within an image and hence can be used in place of the MERW in [138, 137] when such knowledge is available. Further the fact that the method described in this chapter scales well with the graph dimension enables its use for analysis of large images.

Finally, the methods developed in this chapter could also aid in the design of novel metrics for complex networks. In refs. [139, 135, 136], novel metrics are designed by using the MERW to evaluate properties of the network. With the ability to specify visit frequencies in the random walk it becomes possible to evaluate some of these metrics in a weighted sense. For example, if one specifies that visit frequencies of the random walk be a function of the degree of each node, the metric thus obtained will incorporate such a weighting.

4.1.5 Organization

This chapter is organized as follows. In section 4.2 we introduce notation and review known results. In section 4.3 we derive some preliminary results. In section 4.4 we introduce the main result of this chapter which is the maxentropic chain with prescribed visit frequencies. In section 4.5 we show realizations of the maxentropic Markov chain over prototypical roadmaps. Finally, in section 4.6 we present conclusions.

4.2 Notation and review of known results

4.2.1 Notation

For $x \in \mathbb{R}^n$, let $\|x\|_1 = \sum_{i=1}^n |x_i|$, let $[x]$ denote the diagonal matrix with diagonal entries x, that is,

$$[x] = \begin{bmatrix} x_1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & x_n \end{bmatrix}.$$

A matrix $A \in \mathbb{R}^{n\times n}$ is irreducible if, for all partitions $\{I, J\}$ of the index set $\{1, \dots, n\}$, there exists $i \in I$ and $j \in J$ such that $a_{ij} \neq 0$. Here $\{I, J\}$ is a partition of the index set if $I \cup J = \{1, \dots, n\}$ and $I \cap J = \phi$.

Given $x, y \in \mathbb{R}^n$, we define the component-wise vector product $x \circ y \in \mathbb{R}^n$ by $(x \circ y)_i = x_i y_i$ for $i \in \{1, \dots, n\}$. We note the simple equalities:

$$[x]y = x \circ y = [y]x \quad \text{and} \quad [[x]y] = [x][y] = [x \circ y]. \tag{4.1}$$

Define the set of positive n-tuples by $\mathbb{R}^n_{>0} = \{x \in \mathbb{R}^n \mid x_i > 0, i \in \{1, \dots, n\}\}$ and the probability simplex of order n by $\Delta_n = \{v \in \mathbb{R}^n \mid \sum_{i=1}^n v_i = 1, v_i \geq 0 \text{ for } i \in \{1, \dots, n\}\}$.

Consider a graph G with nodeset $V = \{1, \dots, n\}$ then a walk from node i_1 to i_2 and so on until node i_k for $\{i_1, i_2, \dots, i_k\} \in V$ is denoted as $i_1 \to i_2 \to \dots \to i_k$.

For notation and definitions related to Markov chains refer section 3.2. We modify the notation slightly to suit the contents of this chapter in the following manner: transition probabilities from node i to node j shall be denoted as p_{ij} instead of $p_{i,j}$. A similar change in notation is applied to all other quantities including hitting times. This is appropriate considering we are dealing with a single random walker throughout this chapter.

The following lemma and its proof are included here for completeness.

Lemma 4 *Let $S \subseteq \mathbb{R}^n$ be a compact convex set, $\|.\|$ be a matrix norm on $\mathbb{R}^n$, and $h : S \to S$ be a continuously differentiable map. If $\|\partial h/\partial x(x)\| < 1$ for all $x \in S$, then h is a contraction mapping with respect to the norm $\|.\|$ and has a unique fixed point in S.*

Proof: Because S is compact and h is C^1, there exists $c \in (0,1)$ such that

$$\|\partial h/\partial x(x)\| \leq c, \qquad \text{for all } x \in S.$$

By the Mean Value Inequality [143, Proposition 2.4.8], for every $x, y \in S$, there exists $\eta \in S$ such that

$$\|h(y) - h(x)\| \leq \|\partial h/\partial x(\eta)\|\|y - x\|.$$

Therefore, for every $x, y \in S$, we know

$$\|h(y) - h(x)\| \leq c\|y - x\|.$$

Since $0 < c < 1$, this inequality shows that $h : S \to S$ is a contraction with respect to the norm $\|\cdot\|$. By the Banach Contraction Theorem [144, Theorem 3.4.1], h has a unique fixed point in S. ■

4.2.2 Review of maxentropic Markov chains

Throughout the chapter we model the transition matrix of a Markov chain as a row-stochastic matrix. Given a Markov chain with an irreducible transition matrix $P \in \mathbb{R}^{n\times n}$ (i.e., an irreducible row-stochastic matrix), the entropy rate of the Markov chain is given by

$$\mathcal{H}(P) = -\sum_{i,j=1}^{n} \pi_i(P) p_{ij} \log p_{ij}, \tag{4.2}$$

where $\pi(P) \in \text{interior}(\Delta_n)$ is the stationary distribution of P (whose existence, uniqueness, and positivity are established by the Perron-Frobenius Theorem for irreducible matrices).

Problem 1 (Maximizing entropy rate) *Given a connected undirected unweighted graph G, compute the matrix $P \in \mathbb{R}^{n\times n}$ satisfying*

$$\begin{aligned} \max \quad & \mathcal{H}(P) \\ \textit{subj. to} \quad & P \textit{ is row stochastic, i.e., } P \geq 0 \textit{ and } P\mathbb{1}_n = \mathbb{1}_n \\ & p_{ij} = 0, \textit{ if } \{i,j\} \textit{ is not an edge of } G. \end{aligned}$$

Theorem 6 (The maxentropic Markov chain [65, 64]) *Given a symmetric, irreducible $A \in \{0,1\}^{n\times n}$ with associated undirected graph G, let $\lambda > 0$ and $v \in \mathbb{R}^n_{>0}$ be the dominant eigenvalue and eigenvector of A (whose existence and uniqueness are established by the Perron-Frobenius Theorem).*

Then the solution to Problem 1 is unique, is called maxentropic Markov chain over *G, and is given by*

$$P^* = \frac{1}{\lambda}[v]^{-1}A[v], \tag{4.3}$$

or, in components, by

$$P^*_{ij} = \frac{a_{ij}}{\lambda}\frac{v_j}{v_i}.$$

Moreover, P^* has the following properties:

(i) its stationary distribution is $v \circ v / \|v \circ v\|_1$,

(ii) its paths are equiprobable *in the following sense: pick a start node i and a path*

length k. The probability of traversal for a path from i of length $k \geq 1$ is

$$\frac{1}{\lambda^k}\frac{v_j}{v_i}, \tag{4.4}$$

where j is the final node in the path. Note that all paths from i to j of length k have the same probability.

The following example illustrates a maxentropic chain.

Example 1 *Consider the adjacency matrix associated with a 4-node ring and the maxentropic Markov chain associated with this graph,*

$$A = \begin{bmatrix} 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \end{bmatrix}, P^* = \begin{bmatrix} 1/3 & 1/3 & 0 & 1/3 \\ 1/3 & 1/3 & 1/3 & 0 \\ 0 & 1/3 & 1/3 & 1/3 \\ 1/3 & 0 & 1/3 & 1/3 \end{bmatrix}.$$

In general one can show that the Markov chain that maximizes entropy on the ring is the transition matrix that randomizes the position of the random walker at the subsequent timestep between its current location and the two adjacent nodes on the ring.

4.3 Maxentropic maps and their properties

In this section we introduce and characterize two maps: the *maxentropic matrix map* and *maxentropic vector map*. These maps shall be used in the construction of Markov chains with maximum entropy subject to graph and stationary distribution constraints.

4.3.1 The maxentropic matrix map and its properties

Given a symmetric, irreducible, binary matrix $A \in \{0,1\}^{n\times n}$, define the *maxentropic matrix map* $\Phi_A : \mathbb{R}^n_{>0} \to \mathbb{R}^{n\times n}_{\geq 0}$ by

$$\Phi_A(x) = [Ax]^{-1} A[x], \tag{4.5}$$

or, in components, by

$$\big(\Phi_A(x)\big)_{ij} = a_{ij} \frac{x_j}{\sum_{k=1}^n a_{ik} x_k}.$$

The maxentropic Markov chain subject to graph and stationary distribution constraints can be generated from the maxentropic matrix map for a suitable choice of x. In the remainder of this section, we only characterize the maxentropic matrix map. The connection to maxentropic Markov chains shall become clear in Section 4.4.

Theorem 7 (Properties of the maxentropic matrix map) *Given a symmetric, irreducible, binary matrix $A \in \{0,1\}^{n\times n}$ and a vector $x \in \mathbb{R}^n_{>0}$, the maxentropic matrix map has the following properties:*

(i) $\Phi_A(x)$ *is well defined, nonnegative, and row-stochastic,*

(ii) $\Phi_A(x)$ *has the same irreducible zero/positive pattern as* A*,*

(iii) the left dominant eigenvector of $\Phi_A(x)$ *is*

$$\pi(x) = \frac{1}{\|[x]Ax\|_1}[x]Ax, \tag{4.6}$$

(iv) $\Phi_A(x)$ *is reversible, i.e.,* $[\pi(x)]\Phi_A(x) = \Phi_A(x)^\top[\pi(x)]$.

Proof: First, we know $Ax > 0$ because $x > 0$ and because A being irreducible implies each row of A has at least one positive entry. Hence, the diagonal matrix $[Ax]$ is

invertible and $\Phi_A(x)$ is well defined and nonnegative. Finally, $[x]\mathbb{1}_n = x$ implies

$$\Phi_A(x)\mathbb{1}_n = [Ax]^{-1}A[x]\mathbb{1}_n = [Ax]^{-1}Ax = \mathbb{1}_n.$$

This concludes the proof of statement (i).

Next, note that $\Phi_A(x)$ is equal to the matrix A pre- and post-multiplied by two diagonal matrices with positive diagonal; hence $\Phi_A(x)$ has the same zero/positive pattern as A and is irreducible. This concludes the proof of statement (ii).

Regarding statement (iii), by the Perron-Frobenius Theorem for irreducible nonnegative matrices we know that $\Phi_A(x)$ has a unique left dominant eigenvector, i.e., a vector $\pi(x)$ satisfying $\pi(x)^\top\Phi_A(x) = \pi(x)^\top$ and $\mathbb{1}^\top\pi(x) = 1$.

It suffices to show $\pi(x)^\top\Phi_A(x) = \pi(x)^\top$. Recalling the equalities (4.1), we compute

$$\begin{aligned}
\pi(x)^\top\Phi_A(x) &= \frac{1}{\|[x]Ax\|_1}\Big([x]Ax\Big)^\top\Big([Ax]^{-1}A[x]\Big) \\
&= \frac{1}{\|[x]Ax\|_1}\Big([Ax]x\Big)^\top\Big([Ax]^{-1}A[x]\Big) \\
&= \frac{1}{\|[x]Ax\|_1}x^\top[Ax][Ax]^{-1}A[x] \\
&= \frac{1}{\|[x]Ax\|_1}x^\top A[x] = \frac{1}{\|[x]Ax\|_1}\Big([x]Ax\Big)^\top.
\end{aligned}$$

This concludes the proof of statement (iii).

Finally, again recalling the equalities (4.1) and assuming $\|[x]Ax\|_1 = 1$ without loss of generality, we compute

$$\begin{aligned}
[\pi(x)]\Phi_A(x) &= [[x]Ax][Ax]^{-1}A[x] \\
&= ([x][Ax])[Ax]^{-1}A[x] = [x]A[x], \\
\Phi_A(x)^\top[\pi(x)] &= [x]A[Ax]^{-1}[[x]Ax] = [x]A[x].
\end{aligned}$$

This concludes the proof of statement (iv). ■

4.3.2 The maxentropic vector map and its properties

Next, we study the left dominant eigenvector of the row-stochastic matrix $\Phi_A(x)$. Given a binary, symmetric, irreducible matrix A with unit diagonal entries, define the *maxentropic vector map* $\phi_A : \mathbb{R}^n_{>0} \to \mathbb{R}^n_{>0}$ by

$$\phi_A(x) = [x]Ax,$$

or, in components, by

$$\big(\phi_A(x)\big)_i = x_i \sum_{k=1}^{n} a_{ik}x_k.$$

In what follows, we use the notion of proper maps to establish that the maxentropic vector map is a global diffeomorphism. A map $h : X \to Y$ is proper if for every compact set $C \subset Y$, the preimage $h^{-1}(C) \subset X$ is compact.

Theorem 8 (Properties of the maxentropic vector map) *Given a symmetric, irreducible, binary matrix $A \in \{0,1\}^{n\times n}$ with unit diagonal entries, the maxentropic vector map ϕ_A has the following properties:*

(i) the Jacobian of ϕ_A satisfies $\partial\phi_A/\partial x(x) = [x]A + [Ax]$ and is full rank at all $x \in \mathbb{R}^n_{>0}$,

(ii) ϕ_A is a proper map, and

(iii) ϕ_A is a global diffeomorphism, in particular, for every $\pi \in \mathbb{R}^n_{>0}$, there exists a unique $x^* \in \mathbb{R}^n_{>0}$ such that $\phi_A(x^*) = \pi$.

Proof: Regarding property (i), clearly ϕ_A is analytic. Elementary calculations based

also on the equalities (4.1) show that $\partial\phi_A/\partial x(x) = [x]A + [Ax]$. One can show

$$\frac{\partial(\phi_A)_i}{\partial x_i}(x) = a_{ii}x_i + \sum_{j=1}^{n} a_{ij}x_j$$
$$> \sum_{j=1, j\neq i}^{n} a_{ij}x_j = \sum_{j=1, j\neq i}^{n} \frac{\partial(\phi_A)_i}{\partial x_j}(x),$$

for all $x \in \mathbb{R}^n_{>0}$, because $a_{ii} = 1 > 0$ for all $i \in \{1, \ldots, n\}$. Hence, the Jacobian matrix $\partial\phi_A/\partial x(x)$ is strictly row diagonally dominant and, therefore, invertible for all $x \in \mathbb{R}^n_{>0}$.

Before continuing, it is convenient to define the map $\widehat{\phi_A} : \mathbb{R}^n_{\geq 0} \to \mathbb{R}^n_{\geq 0}$ by $\widehat{\phi_A}(x) = [x]Ax$, so that ϕ_A is the restriction of the map $\widehat{\phi_A}$ to $\mathbb{R}^n_{>0}$. We claim that, for every $S \subseteq \mathbb{R}^n_{>0}$, we have $\phi_A^{-1}(S) = \widehat{\phi_A}^{-1}(S)$. We establish this claim as follows. By the property of the restriction map, we can easily show that $\phi_A^{-1}(S) \subseteq \widehat{\phi_A}^{-1}(S)$. Now suppose that there exists a vector $\mathbf{v} = (v_1, v_2, \ldots, v_n)^\top$ such that $\mathbf{v} \in \widehat{\phi_A}^{-1}(S)$ and $\mathbf{v} \notin \phi_A^{-1}(S)$. This implies that $\widehat{\phi_A}(\mathbf{v}) \in S$. Since $\mathbf{v} \notin \phi_A^{-1}(S)$, there exists some $i \in \{1, \ldots, n\}$, such that $v_i = 0$. This implies that $([\mathbf{v}]A\mathbf{v})_i = 0$ and therefore we have $\left(\widehat{\phi_A}(\mathbf{v})\right)_i = 0$. However, this means that $\widehat{\phi_A}(\mathbf{v}) \notin S$. Which is a contradiction. Therefore, we have $\phi_A^{-1}(S) = \widehat{\phi_A}^{-1}(S)$.

Regarding property (ii), let C be a compact set in $\mathbb{R}^n_{>0}$. Then it is a compact set in $\mathbb{R}^n_{\geq 0}$. Therefore, C is closed in $\mathbb{R}^n_{\geq 0}$. Since $\widehat{\phi_A}$ is a continuous map, $\widehat{\phi_A}^{-1}(C)$ is closed in $\mathbb{R}^n_{\geq 0}$. We show that $\widehat{\phi_A}^{-1}(C)$ is bounded in $\mathbb{R}^n_{\geq 0}$. Since all diagonal elements of A are one, we have the following inequality

$$\|x\|_\infty^2 \leq \|[x]Ax\|_\infty, \qquad \text{for all } x \in \mathbb{R}^n_{\geq 0}.$$

Since C is compact, there exists $M \in \mathbb{R}_{>0}$ such that, for every $y \in C$, we have $\|y\|_\infty < M$. Thus, for every $x \in \widehat{\phi_A}^{-1}(C)$, we have

$$\|x\|_\infty^2 \leq \|[x]Ax\|_\infty = \|\widehat{\phi_A}(x)\|_\infty < M.$$

Therefore, $\widehat{\phi_A}^{-1}(C)$ is bounded in $\mathbb{R}^n_{\geq 0}$. This implies that $\widehat{\phi_A}^{-1}(C)$ is compact in $\mathbb{R}^n_{\geq 0}$. Recall that we established $\phi_A^{-1}(C) = \widehat{\phi_A}^{-1}(C)$. Therefore $\phi_A^{-1}(C)$ is a compact set in $\mathbb{R}^n_{>0}$.

Finally, regarding property (iii), we start by noting that property (i) implies, by the Inverse Function Theorem, that ϕ_A is a local diffeomorphism. Therefore, using property (ii) the map ϕ_A is a proper local diffeomorphism and [143, Theorem 2.5.17] implies that ϕ_A is a global diffeomorphism. ■

Remark 2 (The maxentropic vector map is ill-posed without self-loops) *The following example shows that the statements (ii) and (iii) in Theorem 8 do not generally hold for graphs without self-loops. Consider the adjacency matrix*

$$A = \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix}.$$

Define the vectors, $x = [x_1 \;\; x_2]^\top$ *and* $\pi = [\pi_1 \;\; \pi_2]^\top$*. The maxentropic vector map is given by* $\phi_A(x) = [x_1 x_2 \;\; x_1 x_2 + x_2^2]^\top$*. One can solve for the inverse of the map* ϕ_A *explicitly in this case obtaining*

$$\phi_A^{-1}(\pi) = \begin{bmatrix} \dfrac{\pi_1}{\sqrt{\pi_2 - \pi_1}} & \sqrt{\pi_2 - \pi_1} \end{bmatrix}.$$

Consider the compact set $\Pi = \{[\pi_1 \;\; \pi_2]^\top \mid \pi_1 + \pi_2 = 1,\; 0.25 \leq \pi_1 \leq 0.5\}$*. The preimage of the set* Π *under the maxentropic vector map* ϕ_A^{-1} *is not bounded, and hence this set is not compact in* $\mathbb{R}^n$*. Also note that the* $\phi_A^{-1}(\pi)$ *is empty when* $\pi_1 > \pi_2$ *and hence the map is not a diffeomorphism.* ◇

In what follows, we characterize the inverse function of ϕ_A at π. In other words, given a point $\pi \in \text{interior}(\Delta_n)$, we compute $x = \phi_A^{-1}(\pi)$ as the solution to the algebraic

equation

$$\phi_A(x) = [x]Ax = \pi. \tag{4.7}$$

Theorem 9 (Inverse of the maxentropic vector map) *Given a symmetric, irreducible, binary matrix $A \in \{0,1\}^{n\times n}$ with unit diagonal entries, pick $\pi \in \text{interior}(\Delta_n)$.*

(i) For $A = \mathbb{1}_n\mathbb{1}_n^\top$, the algebraic equation (4.7) admits the unique solution π.

(ii) Define the constants $\eta = \max_i\{\sum_{j=1}^n a_{ij}\sqrt{\pi}_j\}$, $\xi = \max_i\{\sum_{j=1}^n a_{ij}\pi_j\}$, and the vector $x^0 = \frac{\pi}{\sqrt{\xi}}$. Then the sequence $\{x^k\}_{k\in\mathbb{N}}$ defined by linear iteration

$$x^{k+1} = x^k - \frac{1}{2\eta}\left([x^k]Ax^k - \pi\right), \qquad \textit{for all } k \in \mathbb{N}, \tag{4.8}$$

converges to the unique solution of equation (4.7).

Proof: Regarding statement (i), note that $A = \mathbb{1}_n\mathbb{1}_n^\top$ implies $A\pi = \mathbb{1}_n$. Therefore, if we set x equal to π into equation (4.7), we get

$$[\pi]A\pi = [\pi]\mathbb{1}_n = \pi.$$

Regarding statement (ii), we first define the nonempty compact convex domain

$$\Omega_\pi = \Big\{\mathbf{y} \in \mathbb{R}^n_{>0} \;\Big|\; \frac{\pi}{\eta} \le \mathbf{y} \le \sqrt{\pi}\Big\}.$$

We first show that $x^0 \in \Omega_\pi$. Since A is a binary matrix, for every $i \in \{1,\dots,n\}$, we have

$$\sqrt{\sum_{j=1}^n a_{ij}\pi_j} \le \sum_{j=1}^n a_{ij}\sqrt{\pi}_j.$$

Therefore, one can deduce that

$$\sqrt{\max_i \Big\{ \sum_{j=1}^{n} a_{ij}\pi_j \Big\}} \leq \max_i \Big\{ \sum_{j=1}^{n} a_{ij}\sqrt{\pi}_j \Big\}.$$

Moreover, matrix A has unit diagonal entries so that

$$\sqrt{\pi_i} \leq \sqrt{\sum_{j=1}^{n} a_{ij}\pi_j} \leq \sqrt{\max_i \Big\{ \sum_{j=1}^{n} a_{ij}\pi_j \Big\}}.$$

Therefore, we have

$$\frac{\pi}{\eta} \leq \frac{\pi}{\sqrt{\xi}} \leq \sqrt{\pi},$$

so that $x^0 \in \Omega_\pi$. Next, define the map $f_\pi : \Omega_\pi \to \mathbb{R}^n$ by

$$f_\pi(x) = x - \frac{1}{2\eta}\left([x]Ax - \pi\right).$$

We aim to show that Ω_π is invariant under the map f_π, i.e., $f_\pi(\Omega_\pi) \subseteq \Omega_\pi$. Consider a point $x \in \Omega_\pi$. We have

$$\frac{\pi_i}{\eta} \leq x_i \leq \sqrt{\pi_i}, \qquad \text{for all } i \in \{1, \ldots, n\}.$$

Therefore, for every $i \in \{1, \ldots, n\}$, we compute

$$\begin{aligned}
\left(f_\pi(x)\right)_i &= x_i - \frac{1}{2\eta}\left(([x]Ax)_i - \pi_i\right) \\
&= x_i - \frac{1}{2\eta}x_i^2 - \frac{1}{2\eta}\left(\sum_{j=1, i\neq j}^{n} a_{ij}x_j \right) + \frac{1}{2\eta}\pi_i \\
&\leq x_i - \frac{1}{2\eta}x_i^2 + \frac{1}{2\eta}\pi_i.
\end{aligned}$$

Note that, for every $i \in \{1, \ldots, n\}$, we have $x_i \leq \sqrt{\pi_i} < \eta$. This implies that the maximum of the function $x_i - \frac{1}{2\eta} x_i^2$ is $\sqrt{\pi_i} - \frac{1}{2\eta}\pi_i$. Hence, we have

$$(f_\pi(x))_i \leq x_i - \frac{1}{2\eta} x_i^2 + \frac{1}{2\eta}\pi_i \leq \sqrt{\pi_i} - \frac{1}{2\eta}\pi_i + \frac{1}{2\eta}\pi_i = \sqrt{\pi_i}.$$

On the other hand, for every $i \in \{1, \ldots, n\}$, we have

$$\sum_{i=1}^{n} a_{ij} x_j \leq \sum_{i=1}^{n} a_{ij}\sqrt{\pi_j} \leq \eta$$

Therefore, for every $i \in \{1, \ldots, n\}$, we have

$$\begin{aligned}
(f_\pi(x))_i &= x_i - \frac{1}{2\eta}([x]Ax)_i + \frac{1}{2\eta}\pi_i \\
&= x_i \left(1 - \frac{1}{2\eta}\sum_{j=1}^{n} a_{ij} x_j\right) + \frac{1}{2\eta}\pi_i \\
&\geq \frac{\pi_i}{\eta}\left(1 - \frac{\eta}{2\eta}\right) + \frac{1}{2\eta}\pi_i \geq \frac{\pi_i}{\eta}.
\end{aligned}$$

This shows that $f_\pi(x) \in \Omega_\pi$ and therefore Ω_π is an invariant set for the map f_π. Next, we show that the map f_π is a contraction mapping on Ω_π. The derivative of f_π satisfies

$$\frac{\partial f_\pi}{\partial x}(x) = I_n - \frac{1}{2\eta}\left([x]A + [Ax]\right), \qquad \text{for all } x \in \Omega_\pi.$$

Also, we have

$$\begin{aligned}
\left\| \frac{\partial f_\pi}{\partial x}(x) \right\|_1 &= \left\| I_n - \frac{1}{2\eta}\left([x]A + [Ax]\right) \right\|_1 \\
&= \max_i \left\{ \left| 1 - \frac{a_{ii} x_i}{2\eta} - \sum_{j=1}^{n} \frac{a_{ij} x_j}{2\eta} \right| + \left| \sum_{j=1, j\neq i}^{n} \frac{a_{ij} x_j}{2\eta} \right| \right\}.
\end{aligned}$$

Since $x \in \Omega_\pi$ implies $x \leq \sqrt{\pi}$, we deduce that, for every $i \in \{1, \ldots, n\}$, we have

$a_{ii}x_i + \sum_{j=1}^{n} a_{ij}x_j \leq 2\eta$. This implies that, for every $i \in \{1, 2, \ldots, n\}$, we have

$$\left| 1 - \frac{a_{ii}x_i}{2\eta} - \sum_{j=1}^{n} \frac{a_{ij}x_j}{2\eta} \right| = 1 - \frac{a_{ii}x_i}{2\eta} - \sum_{j=1}^{n} \frac{a_{ij}x_j}{2\eta}.$$

Thus, for every $i \in \{1, 2, \ldots, n\}$, we get

$$\left| 1 - \frac{a_{ii}x_i}{2\eta} - \sum_{j=1}^{n} \frac{a_{ij}x_j}{2\eta} \right| + \left| \sum_{j=1, j\neq i}^{n} \frac{a_{ij}x_j}{2\eta} \right| = 1 - \frac{a_{ii}x_i}{\eta}.$$

Therefore, we obtain

$$\left\| \frac{\partial f_\pi}{\partial x}(x) \right\|_1 = \max_i \left\{ \left| 1 - \frac{a_{ii}x_i}{\eta} \right| \right\} < 1.$$

Now, using Lemma 4, the map f_π has a unique fixed point in the domain Ω_π and, for $x^0 = \frac{\pi}{\sqrt{\xi}} \in \Omega_\pi$, the sequence defined by the linear iteration (4.8) converges to this unique fixed-point. The proof of the theorem is complete if one notes that x^* is the unique fixed point of f_π if and only if x^* is the unique solution to the algebraic equation (4.7). ■

Remark 3 (Solution to the maxentropic vector map on the complete graph) *If $A = \mathbb{1}_n \mathbb{1}_n^\top$, then we have $\eta = 1$ and the initial condition $x^0 = \frac{\pi}{\sqrt{\eta}} = \pi$ in statement (ii) is the fixed-point of the linear iteration* (4.8) *and the unique solution to the algebraic equation* (4.7). ◇

Remark 4 (Newton-Raphson iteration) *The Newton-Raphson iteration for the non-linear equation $\phi_A(x) = \pi$ is*

$$x^{k+1} = x^k - ([x^k]A + [Ax^k])^{-1}([x^k]Ax^k - \pi). \tag{4.9}$$

In simulations, this iteration appears to always converge for a wide variety of graphs,

from random initial conditions, and for arbitrary choices of $\pi \in \text{interior}(\Delta_n)$ *— even if we are unable to provide a convergence proof. We postpone to Section 4.4.3 a runtime comparison between the linear iteration* (4.8) *and this Newton-Raphson iteration* (4.9).$\diamond$

4.4 Maxentropic Markov chains with prescribed stationary distributions

In this section, we define the optimization problem whose solution we characterize. We then prove uniqueness and existence of the solution before we introduce the main result of the chapter which is a closed-form expression for the maxentropic Markov chain at given stationary distribution, following which we perform computational comparisons with standard convex program solvers and provide proofs for the main result.

4.4.1 Problem statement

Recall that the solution to Problem 1, i.e., the maximum entropy problem subject to purely graph constraints, is the Markov chain given by equation (4.3) in Theorem 6. In what follows, we introduce a new optimization problem by imposing additional stationary distribution constraints on Problem 1. Before we state the problem definition, we remind the reader that given $\pi \in \text{interior}(\Delta_n)$ and given a Markov chain with an irreducible transition matrix $P \in \mathbb{R}^{n\times n}$ (i.e., an irreducible row-stochastic matrix), the entropy rate of the Markov chain P at fixed π is given by

$$\mathcal{H}_\pi(P) = -\sum_{i,j=1}^{n} \pi_i p_{ij} \log p_{ij}. \tag{4.10}$$

Problem 2 (Maximizing entropy rate with a stationary distribution constraint) *Given a symmetric, irreducible, binary matrix $A \in \{0,1\}^{n\times n}$ with unit diagonal entries and given a positive vector $\pi \in \text{interior}(\Delta_n)$, compute the transition matrix $P \in \mathbb{R}^{n\times n}$ satisfying*

$$\begin{aligned} \max \quad & \mathcal{H}_\pi(P) && (4.11)\\ \textit{subj. to} \quad & P \geq 0, && (4.12)\\ & p_{ij} = 0, \ \textit{if } a_{ij} = 0, && (4.13)\\ & P\mathbb{1}_n = \mathbb{1}_n, && (4.14)\\ & \pi^\top P = \pi^\top. && (4.15) \end{aligned}$$

Remark 5 *Problem 2 is a disciplined convex program and hence the numerical solution of this program can be computed in CVX [145].* ◇

Remark 6 (Problem 2 is ill-posed without self-loops) *For given graph topologies without self-loops and for many corresponding instances of stationary distributions, CVX returns that Problem 2 is infeasible. For all such cases, we find that the linear iteration in equation* (4.8) *diverges (recall Remark 2). For example, consider once again the adjacency matrix*

$$A = \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix}.$$

CVX returns that the program is infeasible for any stationary distribution constraint (4.15)*, $[\pi_1 \ \pi_2]^\top P = [\pi_1 \ \pi_2]^\top$, in which $\pi_1 > \pi_2$. Additionally, for this setting the linear iteration* (4.8) *diverges.*

Indeed, the graph topology embodied by A dictates that, whenever the surveillance agent visits node 1, then the agent visits node 2 in the subsequent timestep. Hence, the

visit frequency at node 2 is necessarily greater than or equal to the visit frequency at node 1. ◇

Several such preliminary results indicate that imposing graph constraints can restrict the set of stationary distribution achieved by Markov chains and inspired the conjecture in Section 4.6. We do not pursue this potentially interesting direction of research in this chapter as within the framework of designing surveillance strategies, self-loops can be naturally incorporated. Hence, we proceed under the assumption that all nodes have self-loops.

Remark 7 *In the absence of self-loops in G, the set of irreducible Markov chains over G with prescribed stationary distribution might be empty; see [146] and the conjecture in Section 4.6 for additional context.*

Before we introduce the main result, we prove that the optimizer is irreducible. The optimizer being irreducible ensures existence of the solution to Problem 2 as only irreducible stochastic matrices have well-defined stationary distributions. In addition, we also prove that the optimizer assigns a positive transition probability to every edge in the graph, which is a property that shall be utilized in the proof of the main result.

Theorem 10 (Maxentropic Markov chains are well defined) *Given a symmetric, irreducible, binary matrix $A \in \{0,1\}^{n\times n}$ with unit diagonal entries and given a positive vector $\pi \in \text{interior}(\Delta_n)$, Problem 2 satisfies the following properties:*

(i) the cost function is strictly concave and the constraint set is compact and convex. Hence, its global maximum solution P^* exists and is unique;

(ii) the optimizer P^* satisfies $p^*_{ij} > 0$ whenever $\{i,j\}$ is an edge of the graph G associated to A. Hence, P^* is irreducible and has a well-defined stationary distribution that must be equal to π.

Because of statement (ii), we refer to P^* as the maxentropic Markov chain over G with stationary distribution π.

Note that, for a symmetric, irreducible $A \in \{0,1\}^{n\times n}$ with unit diagonal entries, the graph associated to A is undirected, unweighted, and connected and has self-loops at each node.

Proof: [Proof of Theorem 10] Regarding statement (i), the function $-p\log(p)$ is strictly concave with a strictly positive second-derivative for $p > 0$. The entropy rate is a linear combination of strictly concave functions and hence $\mathcal{H}(P)$ is strictly concave.

Regarding statement (ii), we first show, using a contradiction, that the diagonal entries of P^* can not be zero. Assume that exactly one of the diagonal elements of P^* is zero, i.e., there exists a single $k \in \{1,\dots,n\}$ such that $p^*_{kk} = 0$. We try to find a contradiction. For every $0 < \epsilon < \frac{1}{2}$, define the matrix-valued function $\tilde{P}^*(\epsilon) = (1-\epsilon)P^* + \epsilon I_n$. Note that, for every $\epsilon \in (0,\frac{1}{2})$, we have $\tilde{P}^*(\epsilon) \geq 0$. Also, for every $\epsilon \in (0,\frac{1}{2})$, we have $\tilde{p}^*_{ij} \geq 0$ if $\{i,j\}$ is an edge of G and $\tilde{p}^*_{ij} = 0$ otherwise. One can check that $\tilde{P}^*(\epsilon)\mathbb{1}_n = \mathbb{1}_n$ and $\pi^\top \tilde{P}^*(\epsilon) = \pi^\top$. These facts imply that, for every $\epsilon \in (0,\frac{1}{2})$, the matrix $\tilde{P}^*(\epsilon)$ is in the feasible set of Problem 2. By the Mean Value Theorem [147, Theorem 5.10], for every $\epsilon \in (0,\frac{1}{2})$, there exists $c_\epsilon \in (0,\epsilon)$ such that

$$\mathcal{H}_\pi(\tilde{P}^*(\epsilon)) - \mathcal{H}_\pi(P^*) = \left.\frac{\partial \mathcal{H}_\pi(\tilde{P}^*)}{\partial \epsilon}\right|_{c_\epsilon} \epsilon.$$

Note that, for $i,j \in \{1,\dots,n\}$, we have

$$\frac{\partial \mathcal{H}_\pi(\tilde{P}^*)}{\partial \tilde{p}^*_{ij}} = \pi_i(\log((1-\epsilon)p^*_{ij}) + 1), \qquad \forall i \neq j,$$
$$\frac{\partial \mathcal{H}_\pi(\tilde{P}^*)}{\partial \tilde{p}^*_{ii}} = \pi_i(\log((1-\epsilon)p^*_{ii} + \epsilon) + 1).$$

Using the chain rule, we compute

$$\begin{aligned}\frac{\partial \mathcal{H}_\pi(\tilde{P}^*)}{\partial \epsilon} &= \sum_{i=1}^{n}\sum_{j=1}^{n} \frac{\partial \mathcal{H}_\pi(\tilde{P}^*)}{\partial \tilde{p}^*_{ij}} \frac{\partial \tilde{p}^*_{ij}}{\partial \epsilon} \\ &= \sum_{i=1}^{n} \frac{\partial \mathcal{H}_\pi(\tilde{P}^*)}{\partial \tilde{p}^*_{ii}} \frac{\partial \tilde{p}^*_{ii}}{\partial \epsilon} + \sum_{i=1}^{n}\sum_{j=1,j\neq i}^{n} \frac{\partial \mathcal{H}_\pi(\tilde{P}^*)}{\partial \tilde{p}^*_{ij}} \frac{\partial \tilde{p}^*_{ij}}{\partial \epsilon} \\ &= -\sum_{i=1}^{n} \pi_i(\log((1-\epsilon)p^*_{ii}+\epsilon)+1)(1-p^*_{ii}) \\ &\quad + \sum_{i=1}^{n}\sum_{j=1,j\neq i}^{n} \pi_i(\log((1-\epsilon)p^*_{ij})+1)(p^*_{ij})\end{aligned}$$

Using the fact that $p^*_{kk} = 0$, we get

$$\begin{aligned}\frac{\partial \mathcal{H}_\pi(\tilde{P}^*)}{\partial \epsilon} &= -\pi_k(\log(\epsilon)+1) \\ &\quad - \sum_{i=1,i\neq k}^{n} \pi_i(\log\left((1-\epsilon)p^*_{ii}+\epsilon\right)+1)(1-p^*_{ii}) \\ &\quad + \sum_{i=1}^{n}\sum_{j=1,j\neq i}^{n} \pi_i(\log((1-\epsilon)p^*_{ij})+1)(p^*_{ij}).\end{aligned}$$

Hence we obtain

$$\begin{aligned}\frac{1}{\epsilon}\Big(\mathcal{H}_\pi(\tilde{P}^*(\epsilon)) - \mathcal{H}_\pi(P^*)\Big) &= -\pi_k(\log(c_\epsilon)+1) \\ &\quad - \sum_{i=1,i\neq k}^{n} \pi_i\big(\log\left((1-c_\epsilon)p^*_{ii}+c_\epsilon\right)+1\big)(1-p^*_{ii}) \\ &\quad + \sum_{i=1}^{n}\sum_{j=1,j\neq i}^{n} \pi_i(\log((1-c_\epsilon)p^*_{ij})+1)p^*_{ij}.\end{aligned}$$

Since $c_\epsilon \in (0, \frac{1}{2})$ and $p^*_{ii} \neq 0$, for every $i \neq k$, the term $\sum_{k\neq i} \pi_i(\log((1-c_\epsilon)p^*_{ii}+1)+1)(1-p^*_{ii})$ is bounded. Similarly, since for every $i \neq j$, $p^*_{ij} \neq 0$, the term $\sum\sum_{i\neq j} \pi_i(\log((1-c_\epsilon)p^*_{ij})+1)p^*_{ij}$ is bounded. Thus, by choosing ϵ small enough, one can make c_ϵ small enough

and, therefore, the term $-\pi_k(\log(c_\epsilon)+1)$ large enough. Thus, there exists $\epsilon^* \in (0, \frac{1}{2})$ such that

$$\mathcal{H}_\pi(\tilde{P}^*(\epsilon)) - \mathcal{H}_\pi(P^*) > 0, \qquad \text{for all } \epsilon \in (0, \epsilon^*].$$

This is a contradiction, since we assumed that P^* is the solution to Problem 2. It is straightforward to generalize this argument to the case when we have more zeros on the diagonal of P^* Therefore, all the diagonal entries of P^* are strictly positive.

Next, assuming that all diagonal elements of P^* are positive, we show that, for every i, j with $i \neq j$, if $a_{ij} > 0$, then $p^*_{ij} > 0$. Assume that there exists exactly one pair (k, l) such that $a_{kl} > 0$ but $p^*_{kl} = 0$. We try to find a contradiction. Define the matrix $\Gamma \in \mathbb{R}^{n\times n}$ with all zero entries except for:

$$\Gamma_{kk} = -1, \quad \Gamma_{kl} = 1, \quad \Gamma_{lk} = \frac{\pi_k}{\pi_l}, \quad \text{and } \Gamma_{ll} = -\frac{\pi_k}{\pi_l}.$$

Define $\eta = \min\{\frac{1}{1+p^*_{kk}}, \frac{\pi_l}{\pi_l p^*_{ll}+\pi_k}\}$. For every $\epsilon \in [0, \eta)$, we define the matrix function $\tilde{P}^*(\epsilon) = (1-\epsilon)P^* + \epsilon\Gamma$. One can show that, for every $\epsilon \in (0, \eta)$, we have $\tilde{P}^*(\epsilon) \geq 0$. Moreover, for every $\epsilon \in (0, \eta)$, we have $\tilde{p}^*_{ij} \geq 0$ if $\{i, j\}$ is an edge of G and $\tilde{p}^*_{ij} = 0$ otherwise. One can check that $\tilde{P}^*(\epsilon)\mathbb{1}_n = \mathbb{1}_n$ and $\pi^\top \tilde{P}^*(\epsilon) = \pi^\top$. This implies that, for every $\epsilon \in (0, \eta)$, the matrix $\tilde{P}^*(\epsilon)$ is in the feasible set of Problem 2. By the Mean Value Theorem [147, Theorem 5.10], for every $\epsilon \in (0, \eta)$, there exists $c_\epsilon \in (0, \epsilon)$ such that

$$\mathcal{H}_\pi(\tilde{P}^*(\epsilon)) - \mathcal{H}_\pi(P^*) = \left.\frac{\partial \mathcal{H}_\pi(\tilde{P}^*)}{\partial \epsilon}\right|_{c_\epsilon} \epsilon.$$

Using the chain rule, we compute

$$\frac{1}{\epsilon}\left(\mathcal{H}_\pi(\tilde{P}^*(\epsilon)) - \mathcal{H}_\pi(P^*)\right) = \sum_{i=1}^{n}\sum_{j=1}^{n} \frac{\partial \mathcal{H}_\pi(\tilde{P}^*)}{\partial \tilde{p}^*_{ij}} \frac{\partial \tilde{p}^*_{ij}}{\partial \epsilon}\bigg|_{c_\epsilon}$$

$$\begin{aligned}
&= -\pi_k(\log(c_\epsilon)+1)+\pi_k(\log((1-c_\epsilon)p^*_{kk}-c_\epsilon)+1)(1+p^*_{kk}) \\
&\quad - \pi_l(\log((1-c_\epsilon)p^*_{lk}+\frac{\pi_k}{\pi_l}c_\epsilon)+1)\left(\frac{\pi_k}{\pi_l}-p^*_{lk}\right) \\
&\quad + \pi_l(\log((1-c_\epsilon)p^*_{ll}-\frac{\pi_k}{\pi_l}c_\epsilon)+1)\left(p^*_{ll}+\frac{\pi_k}{\pi_l}\right) \\
&\quad + \sum_{i\notin\{k,l\}}\sum_{j\notin\{k,l\}} \pi_i(\log((1-c_\epsilon)p^*_{ij})+1)p^*_{ij}. \qquad (4.16)
\end{aligned}$$

Note $p^*_{ij}=0$ if and only if $(i,j)=(k,l)$. Since $c_\epsilon \in (0,\eta)$, by choosing ϵ small enough, one can make $-\pi_k(\log(c_\epsilon)+1)$ large enough while the remaining terms in the right hand side of (4.16) are bounded. Therefore, there exists $\epsilon^* \in (0,\eta)$ such that

$$\mathcal{H}_\pi(\tilde{P}^*(\epsilon)) - \mathcal{H}_\pi(P^*) > 0, \qquad \text{for all } \epsilon \in (0,\epsilon^*].$$

This contradicts the fact that P^* is the solution to Problem 2. The generalization of this argument to the case where we have more zeros in P^* is straightforward. Hence, $p^*_{ij}=0$ if and only if $\{i,j\}$ is not an edge of the graph G. ■

4.4.2 Main result

Having motivated the problem of finding the maximum entropy Markov chain subject to graph and stationary distribution constraints and having obtained some preliminary results, we finally present the solution to Problem 2.

Theorem 11 (Maxentropic Markov chains with prescribed stationary distribution) *Consider a symmetric, irreducible, binary matrix $A \in \{0,1\}^{n\times n}$ with unit diagonal en-*

tries and a positive vector $\pi \in \text{interior}(\Delta_n)$. *Let* $x = \phi_A^{-1}(\pi)$ *denote the solution to* $[x]Ax = \pi$ *(whose existence, uniqueness, positivity, and computation algorithm are given in Theorems 8 and 9).*

Then the maxentropic Markov chain over G *with stationary distribution* π *is*

$$P^* = \Phi_A(\phi_A^{-1}(\pi)) = [Ax]^{-1}A[x]. \tag{4.17}$$

Moreover, P^* *is reversible and its entropy rate is*

$$\mathcal{H}(P^*) = -2x^\top A[x]\log(x) + \pi^\top \log(\pi). \tag{4.18}$$

We postpone the proof of this theorem to Section 4.4.4.

Remark 8 *Theorem 11 implies the following result: if* G *has self-loops at each node, then, for all* $\pi \in \text{interior}(\Delta_n)$*, there exists at least one Markov chain over* G *with stationary distribution* π. ◇

We provide a corollary describing notable choices of the maxentropic vector in Theorem 11.

Corollary 4 (Remarkable special cases) *Given a symmetric, irreducible, binary matrix* $A \in \{0,1\}^{n\times n}$ *with unit diagonal entries, let* $d = A\mathbb{1}_n$ *and* $D = [A\mathbb{1}_n]$ *denote its degree vector and matrix, and let* v *and* λ *denote its dominant eigenvector and eigenvalue. Then*

(i) the maxentropic Markov chain with stationary distribution $(\mathbb{1}_n^\top d)^{-1}d$ *is*

$$P^* = \Phi_A(\mathbb{1}_n) = [A\mathbb{1}_n]^{-1}A,$$

with entropy rate

$$\mathcal{H}(P^*) = (\mathbb{1}_n^\top d)^{-1}d^\top \log(d);$$

(This is the so-called equal neighbor random walk.)

(ii) the maxentropic Markov chain with stationary distribution $v \circ v / \|v \circ v\|_1$ *is*

$$P^* = \Phi_A(v) = \frac{1}{\lambda}[v]^{-1} A[v],$$

with entropy rate

$$\mathcal{H}(P^*) = \log \lambda;$$

(This is the maxentropic Markov chain characterized in Theorem 6 as the solution to Problem 1.)

(iii) if $A = \mathbb{1}_n \mathbb{1}_n^\top$ *and* π *is arbitrary, then the maxentropic Markov chain over the complete graph with stationary distribution* π *is*

$$P^* = \Phi_A(\pi) = \mathbb{1}_n \pi^\top,$$

with entropy rate

$$\mathcal{H}(P^*) = -\pi^\top \log(\pi).$$

(The maxentropic vector for the complete graph is shown to be π *in Theorem 9.)*

Finally, we present an interesting property associated with maxentropic Markov chains with prescribed stationary distribution, that is an extension of Theorem 6(ii).

Lemma 5 (All allowed permutations of a walk are equiprobable) *Under the same assumptions as in Theorem 11, consider a start node* i *and a final node* j *on the graph* G *for which there exists a path* $i, l_1, l_2, \ldots, l_k, j$ *and a path* $i, \sigma(l_1), \sigma(l_2), \ldots, \sigma(l_k), j$ *for a permutation* σ*. The following equiprobable path traversal property holds for maxentropic*

Markov chains:

$$\begin{aligned}&\mathbb{P}[i \to l_1 \to l_2 \to \ldots \to l_k \to j] \\ &= \mathbb{P}[i \to \sigma(l_1) \to \sigma(l_2) \to \ldots \to \sigma(l_k) \to j].\end{aligned}$$

We postpone the proof of this lemma to Section 4.4.4.

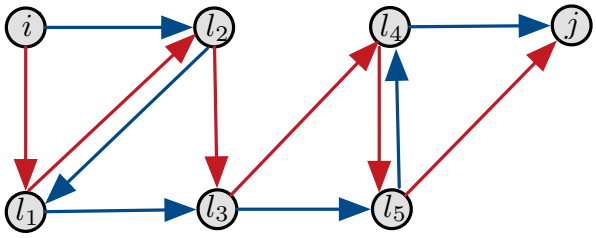

Figure 4.1: The maxentropic Markov chain satisfies a weak version of equiprobability: the probability of a walk through a set of nodes is equal, irrespective of the order of the nodes. The probability of the red and blue walks depicted here are equal for a maxentropic walk.

4.4.3 Computational complexity

In this subsection we show how our proposed procedure to compute maxentropic chains with prescribed stationary distributions is useful not only to reveal their structure and properties but also serves as a valuable method in terms of reducing computational complexity. In short our claim is that:

> *To compute maxentropic Markov chains (as in Problem 2), the linear iteration* (4.8) *and equality* (4.17) *(as stated in Theorem 11) are in general computationally faster than general-purpose convex program solvers.*

We establish this claim in two ways. First, we consider a variety of graphs, we fix a given tolerance, and we observe empirically that that our proposed method has significantly smaller runtimes than the standard CVX solver, see Table 4.1.

Graph	**Linear iteration** (4.8) **& equality** (4.17)	**Newton-Raphson iteration** (4.9)	**CVX**
Line	0.02s	0.01s	44.24s
Star	0.23s	0.01s	54.97s
Ring	0.02s	0.01s	37.53s
Lattice	0.01s	0.02s	40.42s
Complete*	0.01s	0.03s	575.58s

Table 4.1: Average runtimes of various methods over 100 runs on standard graph topologies with 100 nodes to compute maxentropic Markov chains with a randomly chosen stationary distribution for each run. Tolerance is fixed as 10^{-8} in all cases. Computations were performed on a 2.9GHz processor.
*We delete one edge from the complete graph as the iteration in Theorem 9 starts with the solution to the complete graph.

Second, we analyze the computational complexity of the competing algorithms in their two parts: the cost per iteration, and the number of iterations required to get to within a specific tolerance of the optimal solution. In what follows we analyze each algorithm and report the results in Table 4.2.

Method	**Cost per iteration**	**No. of iterations**
Linear iteration (4.8) & equality (4.17)	$O(n) - O(n^2)$	$O(1)$-$O(n)$
Newton-Raphson iteration (4.9)	$O(n^3)$	$O(1)^*$
CVX	$O(n^3)$-$O(n^6)$	$O(\sqrt{n})$

Table 4.2: Computational complexity of various method to compute maxentropic Markov chains with given stationary distribution.
* We can prove this bound for sparse graphs and in simulations the bound holds for complete graphs.

For the linear iteration (4.8) in Theorem 9, each iteration consists of only matrix multiplications with the adjacency matrix or a diagonal matrix, whose cost per iteration is $O(n)$ when the adjacency matrix A is sparse and $O(n^2)$ when A is dense. Also, a careful study of the Banach Fixed Point Theorem and the estimates in Theorem 9 shows that the number of iterations for a fixed tolerance depends on the maximum degree of nodes in the graph. In particular, it can be shown that for sparse graphs such as ring graphs and lattice graph, where the maximum degree does not change with the size of

the graph, the number of iterations is $O(1)$. However, for star graphs and dense graphs such as the complete graph, the number of iterations is of order $O(n)$. In short, the effective worst case complexity across graph topologies for a fixed tolerance is $O(n^3)$.

For the Newton–Raphson iteration (4.9), the factorization of the Jacobian at each step leads to $O(n^3)$ number of operations for each iteration. It can be shown that the number of iterations necessary to obtain a solution within a fixed tolerance is $O(1)$ for sparse graphs. In simulations it is observed that the number of iterations is only weakly dependent on the problem size even for dense graphs and is essentially a constant. We are unable to provide an effective worst-case analysis for the number of iterations necessary when the topology is dense, but across different graph topologies it appears safe to assume that the worst case complexity for fixed tolerance is $O(n^3)$.

In general using a convex program solver would be computationally more expensive as the search space for the convex program is $\mathbb{R}^m_{\geq 0}$, where $m = n^2 - (2n-1) - n_e$, where n_e is the number of edge constraints. Note that the stationary and stochastic constraints in equations (4.14) and (4.15) effectively sum up to $2n-1$ constraints (it can be shown that one of the constraints is redundant). When the graph is sparse $m = O(n)$, otherwise $m = O(n^2)$. Interior point methods used by convex program solvers would need to compute the factorization of an $O(m) \times O(m)$ matrix at every iterations resulting in a runtime complexity for each iteration of $O(n^3)$ for sparse graphs and $O(n^6)$ when the graph is dense (see refs. [148, 149]). The worst-case dependence on problem size is $O(\sqrt{n})$ [150]. Even assuming a constant dependence on problem size as is observed in practice in most semidefinite program interior point solvers, the effective worst-case runtime complexity for fixed tolerance is $O(n^3)$ for sparse graphs and $O(n^6)$ for dense graphs. Also, note that CVX uses a successive approximation scheme to approximate exponential and logarithmic functions [145, Section 11.3]. While this does not affect the computational complexity of the procedure, there are no theoretical guarantees for

convergence to the optimal solution for such an approximation.

Although this chapter presents numerical comparisons only with general purpose convex program solvers, we note that convex programs with linear constraints can be solved efficiently using first-order methods such as mirror descent [151]. Such methods have the same worst-case computational complexity as our proposed linear iteration (4.8). We expect our linear iteration to have lower constant factors than first-order methods for dense graphs for the following reason: as our linear iteration operates on an n-dimensional manifold whereas any first-order convex programming method operates on the space of transition probabilities which is $O(n^2)$ for the case of dense graphs.

4.4.4 Proofs

Consider a Markov chain with transition matrix P on a graph G with binary adjacency matrix A. Let the random variable Y_t denote the observed transition on the graph G at time t which can assume values on $\{1, \ldots, m\}$, where $m = \sum_i \sum_j a_{ij}$ is the total number of edges in the graph. If P is an irreducible Markov chain with stationary distribution π, then for very large times t the probability that a transition from node i to node j occurs is given by

$$\begin{aligned} \lim_{t\to\infty} \mathbb{P}[Y_t = \{i,j\}] &= \lim_{t\to\infty} \left(\mathbb{P}[X_{t+1} = j \,|\, X_t = i]\mathbb{P}[X_t = i]\right) \\ &= \mathbb{P}[X_{t+1} = j \,|\, X_t = i] \lim_{t\to\infty} \mathbb{P}[X_t = i] \\ &= \pi_i p_{ij}. \end{aligned} \tag{4.19}$$

This calculation motivates the following definition.

Definition 1 *For an irreducible transition matrix P with stationary distribution π, define the* ergodic flow matrix *by*

$$Q = [\pi]P. \tag{4.20}$$

Remark 9 *The ergodic flow matrix Q is symmetric if and only if the associated Markov chain P is reversible.* ◇

Let q_{ij} denote the entries of Q. Note that q_{ij} is the probability associated with observing a transition along an edge (i, j) at very large times t according to the calculation in equation (4.19) and the sum of this probability over all edges is 1. The entropy associated with this random variable is

$$\mathcal{H}(Q) = -\sum_{i,j=1}^{n} q_{ij} \log(q_{ij}). \tag{4.21}$$

Lemma 6 (Relation between entropy rate and entropy of ergodic flow matrix) *For an irreducible transition matrix P with stationary distribution π,*

$$\mathcal{H}_\pi(P) = \mathcal{H}(Q) - \mathcal{H}(\pi), \tag{4.22}$$

where $\mathcal{H}(\pi) = \sum_i^n \pi_i \log(\pi_i)$.

Proof: The entropy rate of an irreducible Markov chain P with a stationary distribution π is given by

$$\begin{aligned}
\mathcal{H}_\pi(P) &= -\sum_{i,j=1} \pi_i p_{ij} \log p_{ij} \\
&= -\sum_{i,j=1}^{n} \pi_i p_{ij} (\log(\pi_i p_{ij}) - \log(\pi_i)) \\
&= -\sum_{i,j=1}^{n} q_{ij} \log(q_{ij}) + \sum_{i}^{n} \pi_i \log(\pi_i) \\
&= \mathcal{H}(Q) - \mathcal{H}(\pi).
\end{aligned}$$

■

Consider the convex program which maximizes the entropy of the random variable associated with the ergodic flow matrix.

Problem 3 (Maximize entropy of ergodic flows) *Given a connected undirected unweighted graph G and a positive vector $\pi \in \text{interior}(\Delta_n)$, compute the ergodic flow matrix $Q \in \mathbb{R}^{n\times n}$ satisfying*

$$\begin{aligned} \max \quad & \mathcal{H}(Q) && (4.23)\\ \textit{subj. to} \quad & Q \geq 0, && (4.24)\\ & q_{ij} = 0, \ \textit{if } \{i,j\} \textit{ is not an edge of } G, && (4.25)\\ & Q\mathbb{1}_n = \pi, && (4.26)\\ & Q^\top \mathbb{1}_n = \pi. && (4.27) \end{aligned}$$

Note that the matrix Q is well-defined only when its associated transition matrix has a stationary distribution π. Hence an optimization algorithm might encounter instances where $q_{ij} = 0$ when $a_{ij} = 1$ and hence its associated transition matrix is possibly reducible. In such a case the matrix Q might not have the correct interpretation as the ergodic flow matrix associated with its transition matrix P. However, as a result of Theorem 10 we are guaranteed that the optimal solution Q^* will have the appropriate interpretation as the ergodic flow matrix associated with its transition matrix P^*. Further, since the ergodic flow matrix and its associated transition matrix are closely related we have the following result.

Lemma 7 (Equivalence of Problem 2 and Problem 3) *Given a stationary distribution π, Problem 2 is equivalent to Problem 3 in the following sense:*

(i) if p^*_{ij} is the optimal solution to Problem 2, then $q^*_{ij} = \pi_i p^*_{ij}$ is the optimal solution to Problem 3, and

(ii) if q^*_{ij} is the optimal solution to Problem 3, then $p^*_{ij} = q^*_{ij}/\pi_i$ is the optimal solution to Problem 2.

Proof: First, we shall show that the constraints (4.12)-(4.15) in Problem 2 are equivalent to constraints (4.24)-(4.27) in Problem 3. Note that equation (4.20) and the fact that $\pi \in \text{interior}(\Delta_n)$ implies that Q has the same zero/positive pattern as P. Hence constraints (4.12), (4.13) are equivalent to constraints (4.24), (4.25) respectively. Note that

$$P\mathbb{1}_n = \mathbb{1}_n \quad \Longrightarrow [\pi]P\mathbb{1}_n = [\pi]\mathbb{1}_n \quad \Longrightarrow Q\mathbb{1}_n = \pi.$$

Hence constraint (4.14) is equivalent to constraint (4.26). Also constraint (4.15) is equivalent to

$$P^\top\pi = P^\top\pi \quad \Longrightarrow P^\top[\pi]\mathbb{1}_n = \pi \quad \Longrightarrow Q^\top\mathbb{1}_n = \pi.$$

Hence constraint (4.15) is equivalent to constraint (4.27). This completes the proof of equivalency of constraints.

Second, we shall show that the maximization of the objective function in Problem 2 is equivalent to the maximization of the objective function in Problem 2 subject to the same constraints. For a given stationary distribution π, as a result of Lemma 6 $\mathcal{H}_\pi(P)$ and $\mathcal{H}(Q)$ differ by a constant quantity $\mathcal{H}(\pi)$. Hence the maximization of the objective functions in the two problems are equivalent. Given an optimal solution P^* to Problem 2 one can construct an ergodic flow matrix Q^* using equation (4.20) and vice-versa. Thus (i) and (ii) hold. ■

Problem 4 (Relaxed convex program to maximize entropy of ergodic flows) *Given a connected undirected unweighted graph G and a positive vector $\pi \in \text{interior}(\Delta_n)$, com-*

pute Q such that

$$\max \quad \mathcal{H}(Q) \tag{4.28}$$

$$\textit{subj. to} \quad Q \geq 0, \tag{4.29}$$

$$q_{ij} = 0, \ \textit{if } \{i,j\} \textit{ is not an edge of } G, \tag{4.30}$$

$$Q\mathbb{1}_n + Q^\top \mathbb{1}_n = 2\pi. \tag{4.31}$$

One can show that when the graph G has self-loops at each node, the optimal values of Problem 3 and Problem 4 are the same.

Theorem 12 (Equality of solutions to Problem 3 and Problem 4) *Let G be a connected undirected graph with self-loop at each node, A be the binary adjacency matrix associated to G, and $\pi \in \text{interior}(\Delta_n)$ be a positive vector. Denote the optimal value of Problem 4 by Q_r^* and the optimal value of Problem 3 by Q^*. Then the following statements hold:*

(i) $Q_r^* = Q^*$,

(ii) there exists a vector $x \in \mathbb{R}^n_{>0}$ such that $Q^* = [x]A[x]$.

Proof:

Note that graph and stationary constraints are identical in both problem formulations. Further constraint (4.31) in Problem 4 is obtained by adding (4.26) and (4.27). Therefore, the feasible set of Problem 4 is larger than the feasible set of Problem 3 and thus we have $\mathcal{H}(Q^*) \leq \mathcal{H}(Q_r^*)$.

Using similar arguments to the proof of Theorem 10, one can show that if $Q_r^* = [q_{ij}^*]$ is the solution for Problem 4, then we have $q_{ij}^* > 0$ if and only if $a_{ij} = 1$. This implies

that Q_r^* is the critical point of the Lagrange dual function $\mathcal{L} : \mathbb{R}^{|\mathcal{E}|} \times \mathbb{R}^n \to \mathbb{R}$ defined by

$$\mathcal{L}(Q, \lambda) = -\sum_{\{i,j\}\in\mathcal{E}} q_{ij} \log q_{ij} - \sum_{i=1}^{n} \sum_{\{i,j\}\in\mathcal{E}} \lambda_i (q_{ij} + q_{ji} - 2\pi_i),$$

where $\mathcal{E}$ is the edge set of the graph G. Setting the partial derivatives of $\mathcal{L}$ to zero, for every $\{i, j\} \in \mathcal{E}$, we obtain

$$\frac{\partial \mathcal{L}}{\partial q_{ij}} = 1 + \log q_{ij} - \lambda_i - \lambda_j.$$

Introducing new Lagrange multipliers $\tilde{\lambda}_i = \lambda_i + 1/2$, the solution $Q_r^* = [q_{ij}^*]$ satisfies $q_{ij}^* = a_{ij} \exp^{-\tilde{\lambda}_i} \exp^{-\tilde{\lambda}_j}$.

Let $x_i = \exp^{-\tilde{\lambda}_i}$ then $q_{ij}^* = a_{ij} x_i x_j$ or in matrix notation $Q_r^* = [x]A[x]$. Substituting this solution into the constraints in Problem 4 and using the fact that $a_{ij} = a_{ji}$,

$$\begin{aligned} &\sum_j a_{ij} x_i x_j + \sum_j a_{ji} x_j x_i = 2\pi_i \\ \implies\ & x_i \sum_j a_{ij} x_j = \pi_i \\ \implies\ & [x]Ax = \pi. \end{aligned}$$

Note that A is symmetric, binary matrix with unit diagonal entries. Thus, by Theorem 8, there exists a unique $x^* \in \mathbb{R}^n_{>0}$ such that $[x^*]Ax^* = \pi$. Therefore the global maximum of the concave function $\mathcal{H}$ is given by

$$Q_r^* = [x^*]A[x^*]. \tag{4.32}$$

One can verify that the solution Q_r^* also satisfies constraints (4.26) and (4.27) in Problem 3. This, together with the fact that the feasible set of Problem 4 is larger than the

feasible set of the Problem 3, implies that $Q_r^* = Q^*$. This completes the proof of the part (i). Part (ii) of the theorem follows from part (i) and equation 4.32. ■

Now we have the requisite results to prove Theorem 11.

Proof: [Proof of Theorem 11] Using Lemma 7 and Theorem 12, the solution $P^* = [p_{ij}^*]$ to Problem 2 is given by $p_{ij}^* = (Ax)_i^{-1} a_{ij} x_j$ or in matrix notation as $P^* = [Ax]^{-1} A[x] = \Phi_A(x)$. Also as a result of (iv) in Theorem 4.5 P^* is reversible.

The entropy of the ergodic flow matrix Q^* is given by

$$\begin{aligned}\mathcal{H}(Q^*) &= -\sum_{i,j=1}^{n} q_{ij}^* \log q_{ij}^* \\ &= -\sum_{i,j=1}^{n} a_{ij} x_i x_j (\log x_i + \log x_j) \\ &= -2\sum_{i}^{n} x_i \sum_{j}^{n} a_{ij} x_j \log(x_j) = -2x^\top A([x]\log(x)).\end{aligned}$$

The entropy rate of P^* is given by $\mathcal{H}_\pi(P^*) = \mathcal{H}(Q^*) - \mathcal{H}(\pi)$. The quantity $\mathcal{H}(\pi) = -\pi^\top \log \pi$ in vector notation and hence the result in equation (4.18). ■

Proof: [Proof of Lemma 5] Note that from equations (4.20) and (4.32) we can write the probability of transition from s to t as

$$p_{st} = \frac{a_{st} x_s x_t}{\pi_s}.$$

For the sake of brevity let $\sigma_m = \sigma(l_m)$ for every $1 \leq m \leq k$. Consider the probability of

any valid permutation of the path $i, l_1, l_2, \ldots, l_k, j$ being traversed. This is given by

$$\begin{aligned} p_{i\sigma_1} p_{\sigma_1\sigma_2} \ldots p_{\sigma_k j} &= a_{i\sigma_1} a_{\sigma_1\sigma_2} \ldots a_{\sigma_k j} \frac{x_i x_{\sigma_1} x_{\sigma_1} x_{\sigma_2} \ldots x_{\sigma_k} x_j}{\pi_i \pi_{\sigma_1} \pi_{\sigma_2} \ldots \pi_{\sigma_k}} \\ &= \frac{x_i x_j}{\pi_i} \frac{x_{\sigma_1}^2 x_{\sigma_2}^2 \ldots x_{\sigma_k}^2}{\pi_{\sigma_1} \pi_{\sigma_2} \ldots \pi_{\sigma_k}} \\ &= p_{ij} p_{\sigma_1\sigma_1} p_{\sigma_2\sigma_2} \ldots p_{\sigma_k\sigma_k}. \end{aligned}$$

The quantity $p_{\sigma_1\sigma_1} p_{\sigma_2\sigma_2} \ldots p_{\sigma_k\sigma_k}$ is invariant to permutations of the sequence $\{l_m\}_{1 \leq m \leq k}$. Hence all such paths are equiprobable. ■

4.5 Application to robotic surveillance

In this section, we apply maxentropic chains with non-uniform stationary distributions to the design of robotic surveillance strategies over graphs.

4.5.1 Setup

We consider scenarios in which (i) surveillance agents move on a roadmap (i.e., an undirected graph) according to a discrete-time random walk, (ii) intruders appear at random locations on the roadmap at random times, (iii) intruders can observe the local presence/absence of the surveillance agent(s) and decide when to attack and (iv) the intruder attack is detected precisely when a surveillance agent and the intruder are at the same location during the intruder attack. We consider the following settings: a single agent on a ring, a single agent on a lattice (see Fig. 4.2), and multiple agents on a partitioned map of a realistic environment (see Fig. 4.3).

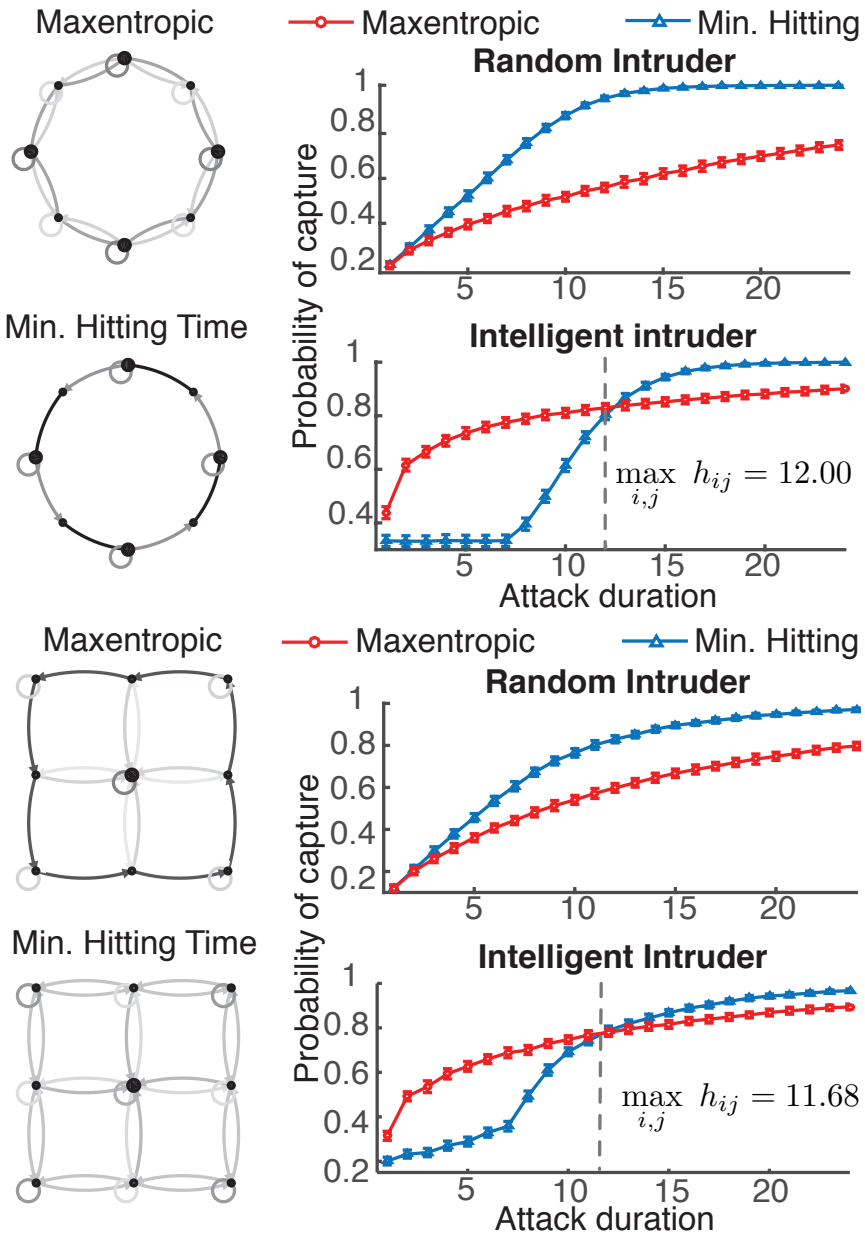

Figure 4.2: Comparison of maxentropic Markov chain strategy with minimum hitting time strategy. The size of the nodes indicate the stationary distribution value associated with the node and opacity of arrows indicate magnitude of transition probability. The worst hitting time for the minimum hitting time Markov chain is denoted by $\max_{i,j} h_{ij}$.

4.5.2 Intruder models

Given a probability vector $\pi \in \Delta_n$, we consider the following intruders models.

(i) *The Random Intruder*: The random intruder has no knowledge of the position of the surveillance agent(s). Such an intruder selects a node i with probability π_i. The attack takes an arbitrary duration which is quantified by the number of transitions performed by the agent(s).

(ii) *The Intelligent Intruder*: The intelligent intruder selects a node i with probability π_i, waits for a/the surveillance agent to arrive at the node, and commences an attack lasting for an arbitrary duration in the timestep immediately following the visit of the surveillance agent. The attack duration is quantified by the number of transitions performed by the agent(s). (Intelligent intruders have been previously studied for example by [51, 54].)

We visualize and design the probability vector π as follows. In Fig. 4.2 the size of the nodes depicts the importance of the node and hence the desired visit frequency. For the ring graph, the north, east, west and south nodes have been assigned twice the priority of the remaining nodes. For the lattice graph, the central node has twice the priority of the peripheral nodes. In Fig. 4.3, which depicts the multi-agent case, we pre-partition the graph for four agents and specify visit frequencies such that all nodes have the same priority. Equal priority with overlapping subgraphs is achieved by specifying a non-uniform visit frequency for each agent on their individual subgraphs. We do this by splitting the visit frequency load equally between the agents for shared nodes.

4.5.3 Surveillance strategies

First, we assume that the intruders and surveillance agents assign the same level of priority to nodes in the graph. In other words, visit frequencies by surveillance agents are biased in a manner so as to be proportional to intruder attacks. Second, we consider two policies for the surveillance agent:

(i) *The maxentropic agent*: The surveillance agent adopts a policy which is the maxentropic Markov chain with visit frequencies proportional to the importance of the node, i.e., the solution described in Theorem 11.

(ii) *The minimum hitting time agent*: Let $\{h_{ij}(P)\}_{ij}$ denote the matrix of mean hitting times for the Markov chain modeled by the transition matrix P. Consider the following optimization program.

Problem 5 (Nonlinear program to minimize mean hitting time) *Given a connected undirected unweighted graph G and a positive vector $\pi \in \text{interior}(\Delta_n)$, compute P such that*

$$\begin{aligned} \min \quad & \sum_i \sum_j \pi_i \pi_j h_{ij}(P) \\ \textit{subj. to} \quad & P \geq 0, \\ & p_{ij} = 0, \ \textit{if } a_{ij} = 0, \\ & P\mathbb{1}_n = \mathbb{1}_n, \\ & \pi^\top P = \pi^\top. \end{aligned}$$

The solution to this nonlinear program is the Markov chain adopted by the minimum hitting time agent. The numerical optimization is conducted using a sequential quadratic programming solver as implemented by the KNITRO/TOMLAB

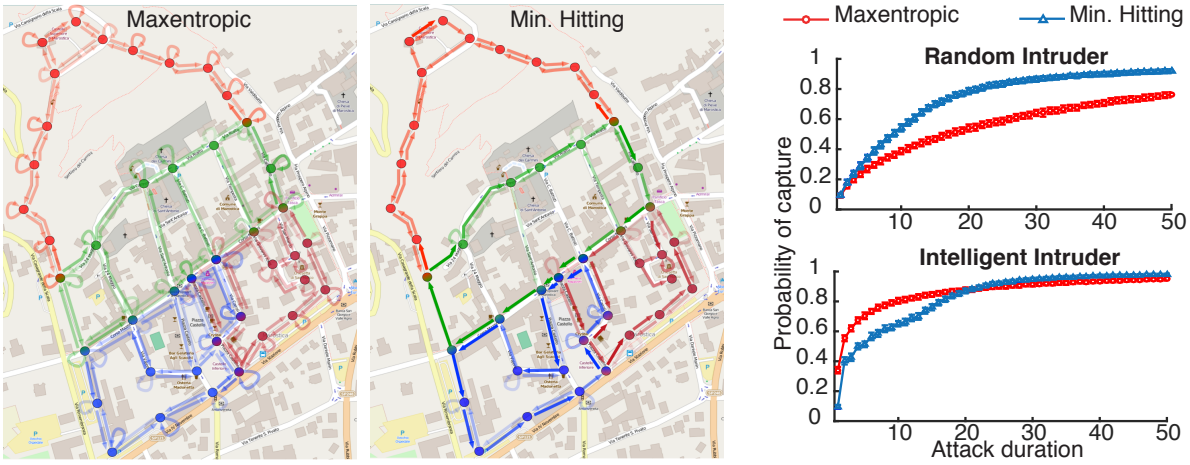

Figure 4.3: Comparison of maxentropic Markov chain strategy with minimum hitting time strategy for the multi-agent case on a partitioned graph. The color(s) of the node indicates which agent(s) are surveilling the node and the opacity of the arrows indicate transition probabilities.

package; for the graph sizes of interest here, this package reliably computes the global minimum solutions. This nonlinear program is identical to the formulation in [53, Problem 1] for a single agent.

For the multi-agent case, each agent performs either the maxentropic Markov chain strategy or the minimum hitting time strategy on their respective subgraphs. There is no actual coordination among the agents (except the joint specification of individual visit frequencies).

4.5.4 Simulation Results

Results for Random Intruders. For the random intruder, for all choices of visit frequencies on a variety of graph topologies, we find that the minimum hitting time agent outperforms the maxentropic agent for all attack durations. The minimum hitting time Markov chain results in faster travel times through the graph. In the absence of knowledge of attack durations, simulations indicate that a strategy with emphasis on fast travel times (small hitting times) performs better than one with emphasis on unpredictability

such as the maxentropic chain. Maxentropic Markov chains by their reversible nature have mixing times of $O(D^2)$ where D is the diameter of the graph [152]. For reversible Markov chains bounds exist on the mixing time and the mean hitting time showing that these notions are equivalent [123]. Thus the maxentropic agent has mean hitting time $O(D^2)$ whereas it is likely that minimum hitting time agent achieves hitting times of $O(D)$ in all cases, though such a result remains to be proved.

Results for Intelligent Intruders. For the intelligent intruder with relatively short attack durations, we note that the maxentropic strategy outperform the minimum hitting time strategy. It stands to reason that, for attack durations larger than the worst hitting time of the chain, the capture of the intelligent intruder is very probable for the minimum hitting time strategy (capture is certain for cases on the ring with uniform stationary distribution where the minimum hitting time strategy is a clockwise or a counterclockwise traversal). In simulations, it is observed that for attack durations which are larger than the worst hitting times of the minimum hitting time chain, the minimum hitting time strategy performs better (see Fig. 4.2). Analysis of the hitting times of the maxentropic chain might reveal an exact condition of the regime of attack durations wherein each strategy leads to higher capture rates.

We summarize these results in Table 4.3. In short, these results indicate that introducing unpredictability into surveillance strategies is appropriate in the important and realistic setting where (i) the intruder uses knowledge of the agents' locations to plan its attacks (e.g., attacking as soon as an agent leaves), and (ii) attack have sufficiently short duration so that they are not detectable by simple fast surveillance agents.

	Random Intruder	**Intelligent Intruder**
Maxentropic Agent	Low capture rate	High capture rate when attack duration is low
Min. Hitting Agent	High capture rate	High capture rate when attack duration is high

Table 4.3: Qualitative summary of results for intruder and agent models.

4.6 Summary

In this chapter we considered the optimization problem of maximizing the entropy rate of a Markov chain with prescribed stationary distribution. We showed this problem is strictly convex with a unique global optimizer. We provided a fast iterative algorithm with rigorous convergence guarantees to compute the so-called entropic vector; as a function of this entropic vector, we provide a closed-form formula for the maximum entropy Markov chain with prescribed stationary distribution. We then characterized several properties of maxentropic chains. The interest for Markov chains with maximum entropy and prescribed stationary distributions arises naturally in robotic surveillance; accordingly we showed some realizations of optimal chains for prototypical robotic roadmaps.

Chapter 5

Markov Chains with Maximum Return Time Entropy

5.1 Introduction

5.1.1 Problem description and motivation

Given a Markov chain, the first return time of a given node is the first time that the random walker returns to the starting node; this is a discrete random variable with infinite support and whose randomness is measured by its entropy. In this chapter, given a strongly connected directed graph with integer-valued travel times (weights) and a prescribed stationary distribution, we study Markov chains with maximum return time entropy. Here the return time entropy of a Markov chain is a weighted average of the entropy of different states' return times with weights equal to the stationary distribution.

We design stochastic surveillance strategies with an entropy maximization objective in order to thwart intruders who plan their attacks based on observations of the surveillance agent. The randomness in the first return time is desirable because an intelligent intruder observing the inter-visit times of the surveillance agent is confronted with a maximally unpredictable return pattern by the surveillance agent. The maximization of randomness in return times is expected to be more effective than maximization of unpredictability

in the sequence of locations as achieved by the maximum entropy rate Markov chain in chapter 4.

5.1.2 Organization

This chapter is organized as follows. We formulate the return time entropy maximization problem in Section 5.2. We establish the properties of the return time entropy in Section 5.3. The approximation analysis and the gradient formulas are provided in Section 5.4. We present the simulation results regarding the robotic surveillance problem in Section 5.5. Section 5.6 concludes the chapter.

5.1.3 Notation and useful lemmas

Let $\mathbb{R}$, $\mathbb{Z}_{\geq 0}$, and $\mathbb{Z}_{>0}$ denote the set of real numbers, nonnegative and positive integers, respectively. In addition to building on matrix notation from the previous two chapters, we utilize the notation $[S]$ to denote a diagonal matrix with diagonal elements being S if S is a vector, or being the diagonal of S if S is a square matrix. The following lemmas are useful.

Lemma 8 *(A uniform bound for stable matrices [153, Proposition D.3.1]) Assume the matrix subset $\mathcal{A} \subset \mathbb{R}^{n\times n}$ is compact and satisfies*

$$\rho_{\mathcal{A}} := \max_{A\in\mathcal{A}} \rho(A) < 1.$$

Then for any $\lambda \in (\rho_{\mathcal{A}}, 1)$ and for any induced matrix norm $\|\cdot\|$, there exists $c > 0$ such that

$$\|A^k\| \leq c\lambda^k, \quad \text{for all } A \in \mathcal{A} \text{ and } k \in \mathbb{Z}_{\geq 0}.$$

Lemma 9 *(Weierstrass M-test [147, Theorem 7.10]) Given a set $\mathcal{X}$, consider the sequence of functions $\{f_k : \mathcal{X} \to \mathbb{R}\}_{k\in\mathbb{Z}_{>0}}$. If there exists a sequence of scalars $\{M_k \in \mathbb{R}\}_{k\in\mathbb{Z}_{>0}}$ satisfying $\sum_{k=1}^{\infty} M_k < \infty$ and*

$$|f_k(x)| \leq M_k, \quad \text{for all } x \in \mathcal{X}, k \in \mathbb{Z}_{>0},$$

then $\sum_{k=1}^{\infty} f_k$ converges uniformly on $\mathcal{X}$.

Lemma 10 *(Geometric distribution generates maximum entropy [154]) Given a discrete random variable $Y \in \mathbb{Z}_{>0}$ and $\mathbb{E}[Y] = \mu \geq 1$, the probability distribution with maximum entropy is*

$$\mathbb{P}[Y = k] = (1 - \frac{1}{\mu})^{k-1}\frac{1}{\mu}, \quad k \in \mathbb{Z}_{>0},$$

with entropy

$$\mathbb{H}(Y) = \mu \log \mu - (\mu - 1)\log(\mu - 1). \tag{5.1}$$

5.2 Problem formulation

We introduce definitions and some preliminary results which lead to the formulation of the maximum return time entropy problem.

5.2.1 Return time of random walks

In this chapter, we consider a strongly connected directed weighted graph $\mathcal{G} = \{V, \mathcal{E}, W\}$, where V denotes the set of n nodes $\{1, \ldots, n\}$, $\mathcal{E} \subset V \times V$ denotes the set of edges, and $W \in \mathbb{Z}_{\geq 0}^{n\times n}$ is the integer-valued weight (travel time) matrix with w_{ij} being the one-hop travel time from node i to node j. If $(i, j) \notin \mathcal{E}$, then $w_{ij} = 0$; if $(i, j) \in \mathcal{E}$, then $w_{ij} \geq 1$. Let $w_{\max} = \max_{i,j}\{w_{ij}\}$ be the maximum travel time.

Given the graph $\mathcal{G} = \{V, \mathcal{E}, W\}$, let $X_k \in \{1, \ldots, n\}$ denote the location of a random walker on $\mathcal{G}$ following a transition matrix P at time $k \in \mathbb{Z}_{\geq 0}$. For any pair of nodes $i, j \in V$, we recall that the *first hitting time* from i to j, denoted by T_{ij}, is the first time the random walk reaches node j starting from node i, that is

$$T_{ij} = \min \Big\{ \sum_{k'=0}^{k-1} w_{X_{k'} X_{k'+1}} \,|\, X_0 = i, X_k = j, k \geq 1 \Big\}. \tag{5.2}$$

In particular, the *return time* T_{ii} of node i is the first time the random walk returns to node i starting from node i. Let the (i, j)-th element of the *first hitting time probability matrix* F_k denote the probability that the random walk reaches node j for the first time in exactly k time units starting from node i, i.e., $F_k(i, j) = \mathbb{P}[T_{ij} = k]$.

5.2.2 Return time entropy of random walks

For an irreducible Markov chain, the return time T_{ii} of each state i is a well-defined random variable over $\mathbb{Z}_{>0}$. We define the *return time entropy of state* i by

$$\begin{aligned} \mathbb{H}(T_{ii}) &= -\sum_{k=1}^{\infty} \mathbb{P}[T_{ii} = k] \log \mathbb{P}[T_{ii} = k] \\ &= -\sum_{k=1}^{\infty} F_k(i, i) \log F_k(i, i), \end{aligned} \tag{5.3}$$

where the logarithm is the natural logarithm and $0 \log 0 = 0$.

Remark 10 *(Coprime travel times) The return time entropy of states does not change when we scale the travel times on all edges simultaneously by the same factor. Therefore, we assume the weights on the graph are coprime.*

Definition 2 *(The set of Markov chains ϵ-conforming to a graph) Given a strongly connected directed weighted graph $\mathcal{G} = \{V, \mathcal{E}, W\}$ with n nodes and the stationary distribution*

$\pi > 0$, *pick a minimum edge weight* $\epsilon > 0$, *the set of Markov chains* ϵ-conforming *to* $\mathcal{G}$ *is defined by*

$$\begin{aligned}\mathcal{P}^{\epsilon}_{\mathcal{G},\pi} = \{P \in \mathbb{R}^{n\times n} \,|\, & p_{ij} \geq \epsilon \text{ if } (i,j) \in \mathcal{E}, \\ & p_{ij} = 0 \text{ if } (i,j) \notin \mathcal{E}, \\ & P\mathbb{1}_n = \mathbb{1}_n, \pi^\top P = \pi^\top\}.\end{aligned}$$

Definition 3 *(Return time entropy) Given a set* $\mathcal{P}^{\epsilon}_{\mathcal{G},\pi}$, *define the* return time entropy *function* $\mathcal{J} : \mathcal{P}^{\epsilon}_{\mathcal{G},\pi} \mapsto \mathbb{R}_{\geq 0}$ *by*

$$\mathcal{J}(P) = \sum_{i=1}^{n} \pi_i \mathbb{H}(T_{ii}). \tag{5.4}$$

Remark 11 *(The expectation of the first return time) For an irreducible Markov chain defined over a weighted graph with travel times, [52, Theorem 6] states*

$$\mathbb{E}[T_{ii}] = \frac{\pi^\top (P \circ W)\mathbb{1}_n}{\pi_i}, \tag{5.5}$$

where $\circ$ *is the Hadamard element-wise product. For unitary travel times, this formula reduces to the usual* $\mathbb{E}[T_{ii}] = 1/\pi_i$. *In both cases, the first return times expectations are inversely proportional to the entries of* π.

In general, it is difficult to obtain the closed-form expression for the return time entropy function.

Examples 13 *(Two special cases with unitary travel times) The elementary proofs of the following results are omitted in the interest of brevity.*

(i) (Two-node complete graph case) Given a two-node complete graph $\mathcal{G}$ *with unit*

weights, if the transition matrix $P \in \mathcal{P}^{\epsilon}_{\mathcal{G},\pi}$ *has the following form*

$$P = \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{bmatrix},$$

then the return time entropy function is

$$\begin{aligned} \mathcal{J}(P) = &-2\pi_1 p_{11}\log(p_{11}) - 2\pi_2 p_{22}\log(p_{22}) \\ &- 2\pi_1 p_{12}\log(p_{12}) - 2\pi_2 p_{21}\log(p_{21}). \end{aligned}$$

(ii) (Complete graph case with special structure) Given an $n \geq 2$*-node complete graph* $\mathcal{G}$ *with unit weights and the stationary distribution* $\pi = \frac{1}{n}\mathbb{1}_n$*, if the transition matrix* $P \in \mathcal{P}^{\epsilon}_{\mathcal{G},\pi}$ *has the form*

$$P = (a-b)I_n + b\mathbb{1}_n\mathbb{1}_n^{\top},$$

for any $a \geq 0$ *and* $b > 0$ *satisfying* $a + (n-1)b = 1$*, then the return time entropy function is*

$$\begin{aligned} \mathcal{J}(P) = &-a\log(a) - (n-1)b\log\left((n-1)b^2\right) \\ &- (n-1)(1-b)\log(1-b). \end{aligned}$$

In this chapter, we are interested in the following problem.

Problem 6 *(Maximization of the return time entropy) Given a strongly connected directed weighted graph* $\mathcal{G} = \{V, \mathcal{E}, W\}$ *and the stationary distribution* $\pi > 0$*, pick a mini-*

mum edge weight $\epsilon > 0$, the maximization of the return time entropy is as follows.

$$\begin{aligned} &\textit{maximize} && \mathcal{J}(P) \\ &\textit{subject to} && P \in \mathcal{P}^{\epsilon}_{\mathcal{G},\pi} \end{aligned}$$

5.3 Properties of the return time entropy

5.3.1 Dynamical model for hitting time probabilities

In this subsection, we characterize a dynamical model for the first hitting time probabilities and establish several important properties of the model.

Theorem 14 *(Linear dynamics for the first hitting time probabilities) Consider a transition matrix $P \in \mathbb{R}^{n\times n}$ that is nonnegative, row-stochastic and irreducible. Then*

(i) the hitting time probabilities F_k, $k \in \mathbb{Z}_{>0}$, satisfy the discrete-time delayed linear system with a finite number of impulse inputs:

$$\operatorname{vec}(F_k) = \operatorname{vec}(P \circ \mathbf{1}_{\{k\mathbb{1}_n\mathbb{1}_n^\top = W\}}) + \sum_{i=1}^{n}\sum_{j=1}^{n} p_{ij}(E_j \otimes \mathbb{e}_i\mathbb{e}_j^\top)\operatorname{vec}(F_{k-w_{ij}}), \tag{5.6}$$

where $E_i = [\mathbb{1}_n - \mathbb{e}_i] \in \mathbb{R}^{n\times n}$, and the initial conditions are $\operatorname{vec}(F_k) = \mathbb{0}_{n^2}$ for all $k \leq 0$;

(ii) if the weights are unitary, i.e., $w_{ij} \in \{0, 1\}$, then the hitting time probabilities satisfy

$$\operatorname{vec}(F_k) = (I_n \otimes P)(I_{n^2} - [\operatorname{vec}(I_n)])\operatorname{vec}(F_{k-1}), \tag{5.7}$$

where the initial condition is $F_1 = P$.

Proof: By definition in (5.2), $F_k(i,j)$ satisfies the following recursive formula for $k \in \mathbb{Z}_{>0}$

$$F_k(i,j) = p_{ij}\mathbf{1}_{\{k=w_{ij}\}} + \sum_{h=1,h\neq j}^{n} p_{ih}F_{k-w_{ih}}(h,j), \tag{5.8}$$

where $\mathbf{1}_{\{\cdot\}}$ is the indicator function and $F_k(i,j) = 0$ for all $k \leq 0$ and $i,j \in V$.

Let $D_k(i) \in \mathbb{R}^{n\times n}$ be a matrix associated with node i at time k that has the form

$$D_k(i) = \sum_{j\in\mathcal{N}_i} \mathbb{e}_j\mathbb{e}_j^\top F_{k-w_{ij}},$$

where $\mathcal{N}_i$ is the set of out-going neighbors of node i. Then, (5.8) can be written in the following matrix form

$$F_k = P \circ \mathbf{1}_{\{k\mathbb{1}_n\mathbb{1}_n^\top = W\}} + \sum_{i=1}^{n} \mathbb{e}_i\mathbb{e}_i^\top P(D_k(i) - [D_k(i)]). \tag{5.9}$$

Vectorizing both sides of (5.9), we have

$$\begin{aligned}\operatorname{vec}(F_k) =& \operatorname{vec}(P \circ \mathbf{1}_{\{k\mathbb{1}_n\mathbb{1}_n^\top = W\}}) \\ &+ \sum_{i=1}^{n} (I_n \otimes \mathbb{e}_i\mathbb{e}_i^\top P)(I_{n^2} - [\operatorname{vec}(I_n)])\operatorname{vec}(D_k(i)).\end{aligned}$$

Note that

$$\operatorname{vec}(D_k(i)) = \sum_{j\in\mathcal{N}_i} (I_n \otimes \mathbb{e}_j\mathbb{e}_j^\top)\operatorname{vec}(F_{k-w_{ij}}),$$

and

$$(I_{n^2} - [\operatorname{vec}(I_n)])(I_n \otimes \mathbb{e}_j\mathbb{e}_j^\top) = E_j \otimes \mathbb{e}_j\mathbb{e}_j^\top.$$

Therefore, we have (5.6).

Moreover, if the travel times are unitary, then $F_1 = P$ and

$$\sum_{i=1}^{n}\sum_{j=1}^{n} p_{ij}(E_j \otimes \mathbb{e}_i\mathbb{e}_j^\top) = (I_n \otimes P)(I_{n^2} - [\mathrm{vec}(I_n)]). \tag{5.10}$$

Thus, equation (5.7) follows. ■

The dynamical system (5.6) can be transformed to an equivalent homogeneous linear system by restarting the system at $k = w_M$ with same system matrices and appropriate initial conditions. Moreover, we can augment the system and obtain a discrete-time linear system without delays. This equivalent augmented system is useful for example in studying stability properties. For $k \geq 1$, we have

$$\begin{bmatrix} \mathrm{vec}(F_{k+w_{\max}}) \\ \mathrm{vec}(F_{k+w_{\max}-1}) \\ \vdots \\ \mathrm{vec}(F_{k+1}) \end{bmatrix} = \Psi \begin{bmatrix} \mathrm{vec}(F_{k+w_{\max}-1}) \\ \mathrm{vec}(F_{k+w_{\max}-2}) \\ \vdots \\ \mathrm{vec}(F_k) \end{bmatrix}, \tag{5.11}$$

where

$$\Psi = \begin{bmatrix} \Phi_1 & \Phi_2 & \cdots & \cdots & \Phi_{w_{\max}} \\ I_{n^2} & \mathbb{0}_{n^2\times n^2} & \cdots & \cdots & \mathbb{0}_{n^2\times n^2} \\ \mathbb{0}_{n^2\times n^2} & I_{n^2} & \cdots & \cdots & \mathbb{0}_{n^2\times n^2} \\ \vdots & \vdots & \ddots & \cdots & \mathbb{0}_{n^2\times n^2} \\ \mathbb{0}_{n^2\times n^2} & \cdots & \cdots & I_{n^2} & \mathbb{0}_{n^2\times n^2} \end{bmatrix}, \tag{5.12}$$

and for $h \in [1, w_{\max}]$,

$$\Phi_h = \sum_{i=1}^{n}\sum_{j=1}^{n} p_{ij}(E_j \otimes \mathbb{e}_i\mathbb{e}_j^\top)\mathbf{1}_{\{w_{ij}=h\}}. \tag{5.13}$$

The initial conditions for (5.11) can be computed using (5.6). For brevity, we denote

$\begin{bmatrix} \mathrm{vec}(F_{k+w_{\max}-1}) & \cdots & \mathrm{vec}(F_k) \end{bmatrix}^\top$ by $\mathrm{vec}(\tilde{F}_k)^\top$.

Lemma 11 *(Properties of the linear dynamics for the first hitting time probabilities) If* $P \in \mathbb{R}^{n\times n}$ *is nonnegative, row-stochastic and irreducible, then*

(i) the matrix $(I_n \otimes P)(I_{n^2} - [\mathrm{vec}(I_n)])$ *is row-substochastic with* $\rho\big((I_n \otimes P)(I_{n^2} - [\mathrm{vec}(I_n)])\big) < 1$.

(ii) the delayed discrete-time linear system with a finite number of impulse inputs (5.6) *is asymptotically stable;*

(iii) $\mathrm{vec}(F_k) \geq 0$ *for* $k \in \mathbb{Z}_{>0}$ *and* $\sum_{k=1}^{\infty} \mathrm{vec}(F_k) = \mathbb{1}_{n^2\times 1}$.

Proof: Regarding (i), note that the matrix $(I_n \otimes P)(I_{n^2} - [\mathrm{vec}(I_n)])$ is block diagonal with the i-th block being PE_i. Since P is irreducible, there is at least one positive entry in each column of P. Therefore PE_i's are row-substochastic and so is $(I_n \otimes P)(I_{n^2} - [\mathrm{vec}(I_n)])$. By [53, Lemma 2.2], $\rho(PE_i) < 1$ for all $i \in \{1, \ldots, n\}$ and $\rho((I_n \otimes P)(I_{n^2} - [\mathrm{vec}(I_n)])) = \max_i \rho(PE_i) < 1$.

Regarding (ii), since we can rewrite (5.6) as (5.11) with appropriate initial conditions and Φ_i's are nonnegative, by the stability criterion for delayed linear systems [155, Theorem 1], (5.6) is asymptotically stable if

$$\rho\Big(\sum_{i=1}^{w_{\max}} \Phi_i\Big) = \rho\big((I_n \otimes P)(I_{n^2} - [\mathrm{vec}(I_n)])\big) < 1,$$

which is true by (i).

Regarding (iii), first note that all the system matrices are nonnegative, thus $\mathrm{vec}(F_k) \geq 0$ for all $k \in \mathbb{Z}_{>0}$. Moreover, due to (ii), the delayed linear system (5.6) is asymptotically

stable. Summing both sides of (5.6) over k, we have

$$\begin{aligned}\sum_{k=1}^{\infty}\operatorname{vec}(F_k) &= \operatorname{vec}(P) + \sum_{i=1}^{n}\sum_{j=1}^{n} p_{ij}(E_j \otimes \mathbb{e}_i\mathbb{e}_j^\top)\sum_{k=1}^{\infty}\operatorname{vec}(F_k)\\ &= \operatorname{vec}(P) + (I_n \otimes P)(I_{n^2} - [\operatorname{vec}(I_n)])\sum_{k=1}^{\infty}\operatorname{vec}(F_k),\end{aligned}$$

which implies that $\sum_{k=1}^{\infty}\operatorname{vec}(F_k) = \mathbb{1}_{n^2\times 1}$. ■

5.3.2 Well-posedness of the optimization problem

We here show that the function $\mathcal{J}$ is continuous over the compact set $\mathcal{P}^{\epsilon}_{\mathcal{G},\pi}$. Then, by the extreme value theorem, $\mathcal{J}$ has a (possibly non-unique) maximum point in the set and thus Problem 6 is well-posed.

Lemma 12 *(Continuity of the return time entropy function) Given the compact set $\mathcal{P}^{\epsilon}_{\mathcal{G},\pi}$, the following statements hold:*

(i) there exist constants $\lambda_{\max} \in (0,1)$ and $c > 0$ such that

$$F_k(i,i) \le c\lambda_{\max}^k, \quad \text{for all } k \in \mathbb{Z}_{>0}, i \in \{1,\dots,n\};$$

(ii) the return time entropy functions $\mathbb{H}(T_{ii})$, $i \in \{1,\dots,n\}$, and $\mathcal{J}(P)$ are continuous on the compact set $\mathcal{P}^{\epsilon}_{\mathcal{G},\pi}$; and

(iii) Problem 6 is well-posed in the sense that a global optimum exists.

Proof: Regarding (i), for $k \ge w_M + 1$, since the spectral radius $\rho(\Psi)$ is a continuous function of Ψ [124, Example 7.1.3], where Ψ is given in (5.12), and Ψ is a continuous function of P, $\rho(\Psi)$ is a continuous function of P. Hence, by Lemma 11(ii) and the

extreme value theorem, there exists a $\rho_{\max} < 1$ such that

$$\rho_{\max} = \max_{P \in \mathcal{P}^{\epsilon}_{\mathcal{G},\pi}} \rho(\Psi) < 1.$$

Therefore, for $k \geq w_M + 1$ and $i \in \{1, \ldots, n\}$, by Lemma 8, there exist $c_1 > 0$ and $\rho_{\max} < \lambda_{\max} < 1$ such that

$$\begin{aligned} F_k(i,i) &\leq \| \operatorname{vec}(\tilde{F}_{k-w_{\max}+1})\|_\infty \\ &= \|(\Psi)^{k-w_{\max}} \operatorname{vec}(\tilde{F}_1)\|_\infty \\ &\leq \|(\Psi)^{k-w_{\max}}\|_\infty \| \operatorname{vec}(\tilde{F}_1)\|_\infty \\ &\leq c_1 \lambda_{\max}^{k-w_{\max}} = \frac{c_1}{\lambda_{\max}^{w_{\max}}} \lambda_{\max}^k. \end{aligned}$$

Let $c = \max\{\frac{c_1}{\lambda_{\max}^{w_{\max}}}, \frac{1}{\lambda_{\max}^{w_{\max}}}\}$, then we have for $k \geq w_M + 1$,

$$F_k(i,i) \leq \frac{c_1}{\lambda_{\max}^{w_{\max}}} \lambda_{\max}^k < c\lambda_{\max}^k.$$

For $k \leq w_M$,

$$c\lambda_{\max}^k \geq c\lambda_{\max}^{w_{\max}} \geq 1 \geq F_k(i,i).$$

Therefore, we have (i).

Regarding (ii), due to (i), there exists a positive integer K that does not depend on the elements of $\mathcal{P}^{\epsilon}_{\mathcal{G},\pi}$ such that when $k \geq K$, $c\lambda_{\max}^k \leq e^{-1}$. Since $x \mapsto -x \log x$ is an increasing function for $x \in [0, e^{-1}]$, when $k \geq K$,

$$-F_k(i,i) \log F_k(i,i) \leq -c\lambda_{\max}^k \log(c\lambda_{\max}^k) \coloneqq M_k.$$

For $k < K$, $-F_k(i,i)\log F_k(i,i) \le e^{-1} \coloneqq M_k$. Then

$$\sum_{k=1}^{K-1} M_k = \frac{K-1}{e},$$

and

$$\begin{aligned}\sum_{k=K}^{\infty} M_k &= -\sum_{k=K}^{\infty} c\lambda_{\max}^k \log(c\lambda_{\max}^k) \\ &= -c\log c \sum_{k=K}^{\infty} \lambda_{\max}^k - c\log(\lambda_{\max}) \sum_{k=K}^{\infty} k\lambda_{\max}^k \\ &= -c\Big(\frac{\lambda_{\max}^K}{1-\lambda_{\max}} \log(c\lambda_{\max}^K) + \frac{\lambda_{\max}^{K+1}}{(1-\lambda_{\max})^2} \log(\lambda_{\max})\Big). \end{aligned} \tag{5.14}$$

Hence,

$$\sum_{k=1}^{\infty} M_k = \sum_{k=1}^{K-1} M_k + \sum_{k=K}^{\infty} M_k < \infty,$$

which holds for any i and any transition matrix in the compact set $\mathcal{P}^{\epsilon}_{\mathcal{G},\pi}$. By Lemma 9, the series $-\sum_{k=1}^{\infty} F_k(i,i)\log F_k(i,i)$ converges uniformly. Since the the limit of a uniformly convergent series of continuous function is continuous [147, Theorem 7.12], $\mathbb{H}(T_{ii})$ is a continuous function on $\mathcal{P}^{\epsilon}_{\mathcal{G},\pi}$. Finally, $\mathcal{J}(P)$ is a finite weighted sum of continuous functions $\mathbb{H}(T_{ii})$, thus $\mathcal{J}(P)$ is a continuous function.

Regarding (iii), because $\mathcal{J}$ is a continuous function over the compact set $\mathcal{P}^{\epsilon}_{\mathcal{G},\pi}$, the extreme value theorem ensures that Problem 6 admits a global optimum solution (possibly non-unique) and is therefore well-posed. ■

5.3.3 Optimal solution for complete graphs with unitary travel times

We here provide (1) an upper bound for the return time entropy with unitary travel times based on the principle of maximum entropy and (2) the optimal solution to Problem 6 for the complete graph case with unitary travel times.

Lemma 13 *(Maximum achieved return time entropy in a complete graph with unitary weights) Given a strongly connected graph $\mathcal{G}$ with unitary weights and the compact set $\mathcal{P}^{\epsilon}_{\mathcal{G},\pi}$,*

(i) the return time entropy function is upper bounded by

$$\mathcal{J}(P) \leq -\sum_{i=1}^{n}(\pi_i \log \pi_i + (1-\pi_i)\log(1-\pi_i));$$

(ii) when the graph $\mathcal{G}$ is complete, the upper bound is achieved and the transition matrix that maximizes the return time entropy $\mathcal{J}(P)$ is given by $P = \mathbb{1}_n \pi^\top$.

Proof: Regarding (i), by Remark 11, in the case of unitary travel times, we have $\mathbb{E}[T_{ii}] = 1/\pi_i$. Thus, T_{ii} is a discrete random variable with fixed expectation, whose entropy is bounded as shown in Lemma 10. For any transition matrix $P \in \mathcal{P}^{\epsilon}_{\mathcal{G},\pi}$, the return time entropy function $\mathcal{J}(P)$ satisfies

$$\begin{aligned}
\mathcal{J}(P) &= \sum_{i=1}^{n} \pi_i \mathbb{H}(T_{ii}) \leq \sum_{i=1}^{n} \pi_i \max_{T_{ii}}\{\mathbb{H}(T_{ii})\} \\
&= \sum_{i=1}^{n} \pi_i \big(\frac{1}{\pi_i}\log\frac{1}{\pi_i} - (\frac{1}{\pi_i}-1)\log(\frac{1}{\pi_i}-1)\big) \\
&= -\sum_{i=1}^{n} \big(\pi_i \log \pi_i + (1-\pi_i)\log(1-\pi_i)\big),
\end{aligned}$$

where the third line uses (5.1).

Regarding (ii), when the graph is complete and $P = \mathbb{1}_n \pi^\top$, the return time T_{ii} follows the geometric distribution:

$$\mathbb{P}(T_{ii} = k) = \pi_i (1 - \pi_i)^{k-1}.$$

Then by Lemma 10, we obtain the results. ■

5.3.4 Relations with the entropy rate of Markov chains

Given an irreducible Markov chain P with n nodes and stationary distribution π, recall that the *entropy rate of* P is given by

$$\mathcal{H}(P) = -\sum_{i=1}^{n} \pi_i \sum_{j=1}^{n} p_{ij} \log p_{ij}.$$

We next study the relationship between the return time entropy $\mathcal{J}$ with unitary travel times and the entropy rate $\mathcal{H}$.

Theorem 15 *(Relations between the return time entropy with unitary travel times and the entropy rate) For all P in the compact set $\mathcal{P}^{\epsilon}_{\mathcal{G},\pi}$ where $\mathcal{G}$ has unitary travel times, the return time entropy $\mathcal{J}(P)$ and the entropy rate $\mathcal{H}(P)$ satisfy*

$$\mathcal{H}(P) \leq \mathcal{J}(P) \leq n\mathcal{H}(P), \tag{5.15}$$

where n is the number of nodes in the graph $\mathcal{G}$.

We prove this theorem The proof of the following theorem follows from Lemmas 15 and Lemma 16 below.

Remark 12 *Theorem 15 establishes a large gap, possibly of size $O(n)$, between $\mathcal{H}(P)$ and $\mathcal{J}(P)$ and, thereby, optimizing $\mathcal{H}$ and $\mathcal{J}$ are two different matters altogether.*

First, we show that the return time entropy is upper bounded by n times of the entropy rate. As in [67], we define a *Markov trajectory from state i to state j* to be a path with initial state i, final state j, and no intervening state equal to j. Let $\mathcal{T}_{ij}$ be the set of all Markov trajectories from state i to state j. Let $\mathbb{P}[\ell]$ denote the probability of a Markov trajectory $\ell \in \mathcal{T}_{ij}$; clearly $\sum_{\ell \in \mathcal{T}_{ij}} \mathbb{P}[\ell] = 1$. Let L_{ij} be the Markov trajectory random variable that takes value ℓ in $\mathcal{T}_{ij}$ with probability $\mathbb{P}[\ell]$. Finally, we define the entropy of L_{ij} by

$$\mathbb{H}(L_{ij}) = -\sum_{\ell \in \mathcal{T}_{ij}} \mathbb{P}[L_{ij} = \ell] \log \mathbb{P}[L_{ij} = \ell].$$

Lemma 14 *(Entropy of Markov trajectories [67, Theorem 1]) For an irreducible Markov chain with transition matrix P, the entropy $\mathbb{H}(L_{ii})$ of the random Markov trajectory from state i back to state i is given by*

$$\mathbb{H}(L_{ii}) = \frac{\mathcal{H}(P)}{\pi_i}.$$

Through the entropy of the Markov trajectories, we are able to establish the upper bound of the return time entropy in (5.15).

Lemma 15 *(Upper bound of the return time entropy by n times of the entropy rate) Given the compact set $\mathcal{P}^{\epsilon}_{\mathcal{G},\pi}$,*

(i) the return time entropy is upper bounded by

$$\mathcal{J}(P) \leq n\mathcal{H}(P), \quad \text{for all } P \in \mathcal{P}^{\epsilon}_{\mathcal{G},\pi}; \tag{5.16}$$

(ii) the equality in (5.16) *holds if and only if any node of the graph $\mathcal{G}$ has the property that all distinct first return paths have different length, i.e., the return paths are*

distinguishable by their lengths, and in this case,

$$\operatorname*{argmax}_{P\in\mathcal{P}^{\epsilon}_{\mathcal{G},\pi}} \mathcal{J}(P) = \operatorname*{argmax}_{P\in\mathcal{P}^{\epsilon}_{\mathcal{G},\pi}} \mathcal{H}(P).$$

Proof: Regarding (i), the return time random variable T_{ii} is defined by lumping the trajectories in $\mathcal{T}_{ii}$ with the same length,

$$\mathbb{P}\left[T_{ii}=k\right] = \sum_{\ell\in\mathcal{T}_{ii},|\ell|=k} \mathbb{P}\left[L_{ii}=\ell\right], \tag{5.17}$$

where $|\ell|$ denotes the length of the path ℓ. Note that

$$\begin{aligned} -\mathbb{P}\left[T_{ii}=k\right]\log\mathbb{P}\left[T_{ii}=k\right] &= -\big(\sum_{\ell\in\mathcal{T}_{ii},|\ell|=k} \mathbb{P}\left[L_{ii}=\ell\right]\big)\log\big(\sum_{\ell\in\mathcal{T}_{ii},|\ell|=k} \mathbb{P}\left[L_{ii}=\ell\right]\big) \\ &\le -\sum_{\ell\in\mathcal{T}_{ii},|\ell|=k} \mathbb{P}\left[L_{ii}=\ell\right]\log\mathbb{P}\left[L_{ii}=\ell\right], \end{aligned} \tag{5.18}$$

where we used that $(x+y)\log(x+y) \ge x\log x + y\log y$ for $x, y \ge 0$. Since both the return time entropy and the entropy of Markov trajectories are absolutely convergent, we have

$$\begin{aligned} \mathbb{H}(T_{ii}) &= -\sum_{k=1}^{\infty}\mathbb{P}\left[T_{ii}=k\right]\log\mathbb{P}\left[T_{ii}=k\right] \\ &\le -\sum_{k=1}^{\infty}\sum_{\ell\in\mathcal{T}_{ii},|\ell|=k}\big(\mathbb{P}\left[L_{ii}=\ell\right]\log\mathbb{P}\left[L_{ii}=\ell\right]\big) \\ &= \mathbb{H}(L_{ii}), \end{aligned}$$

which along with Lemma 14 imply

$$\mathcal{J}(P) \le n\mathcal{H}(P).$$

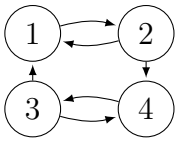

Figure 5.1: An example graph that satisfies the property in Lemma 15(ii)

Regarding (ii), the inequality in (5.16) comes from the inequality (5.18). If any node of the graph $\mathcal{G}$ has the property that all distinct first return paths have different length, then the summation on the right hand side of (5.17) only has one term and the inequality in (5.18) becomes an equality. On the other hand, if for some node of $\mathcal{G}$, there are distinct return paths that have the same length, then one needs to lump the paths with the same length and the inequality in (5.18) becomes strict. Moreover, if the equality holds, then $\mathcal{J}(P)$ is a constant n times of $\mathcal{H}(P)$ and thus they have the same maximizer. ■

Example 2 *For the two-node case in Examples 13(i), the return time entropy is twice the entropy rate. This is not a coincidence since the 2-node complete graph satisfies the property in Lemma 15(ii). Figure 5.1 illustrates a graph with* 4 *nodes that also satisfies the property in Lemma 15(ii).*

In the rest of this subsection, we show that the return time entropy is lower bounded by the entropy rate as shown in (5.15).

Lemma 16 *(Lower bound of the return time entropy by the entropy rate) Given the compact set* $\mathcal{P}^{\epsilon}_{\mathcal{G},\pi}$,

(i) the return time entropy is lower bounded by

$$\mathcal{J}(P) \geq \mathcal{H}(P), \quad \textit{for all } P \in \mathcal{P}^{\epsilon}_{\mathcal{G},\pi}; \tag{5.19}$$

(ii) the equality in (5.19) *holds if and only if* P *is a permutation matrix.*

Proof: Regarding (i), note that the first hitting time T_{ij} from state i to state j as defined in (5.2) is a random variable , whose entropy is $\mathbb{H}(T_{ij})$. Then by definition, we have in the case of unitary travel times,

$$\begin{aligned}
\mathbb{H}(T_{ij}) &= -\sum_{k=1}^{\infty} \mathbb{P}\left[T_{ij}=k\right] \log \mathbb{P}\left[T_{ij}=k\right] \\
&= -p_{ij}\log p_{ij} - \Big(\sum_{k_1\neq j} p_{ik_1}p_{k_1j}\Big)\log\Big(\sum_{k_1\neq j} p_{ik_1}p_{k_1j}\Big) \\
&\quad - \Big(\sum_{k_1,k_2\neq j} p_{ik_1}p_{k_1k_2}p_{k_2j}\Big)\log\Big(\sum_{k_1,k_2\neq j} p_{ik_1}p_{k_1k_2}p_{k_2j}\Big) \\
&\quad - \cdots \\
&\quad - \Big(\sum_{k_1\cdots k_m\neq j} p_{ik_1}\cdots p_{k_mj}\Big)\log\Big(\sum_{k_1\cdots k_m\neq j} p_{ik_1}\cdots p_{k_mj}\Big) \\
&\quad - \cdots .
\end{aligned}$$

Since $x \mapsto -x\log x$ is a concave function, for $x_i \geq 0$ and for coefficients $\alpha_i \geq 0$ satisfying $\sum_{i=1}^{n}\alpha_i = 1$, we have

$$-\Big(\sum_{i=1}^{n}\alpha_i x_i\Big)\log\Big(\sum_{i=1}^{n}\alpha_i x_i\Big) \geq -\sum_{i=1}^{n}\alpha_i(x_i\log x_i). \tag{5.20}$$

Thus, for $m \geq 1$,

$$\begin{aligned}
-\mathbb{P}\left[T_{ij} = m+1\right] \log \mathbb{P}\left[T_{ij} = m+1\right] &= -\Big(\sum_{k_1 \cdots k_m \neq j} p_{ik_1} \cdots p_{k_m j}\Big) \log\Big(\sum_{k_1 \cdots k_m \neq j} p_{ik_1} \cdots p_{k_m j}\Big) \\
&= -\Big(\sum_{k_1 \neq j} p_{ik_1} \sum_{k_2 \cdots k_m \neq j} p_{k_1 k_2} \cdots p_{k_m j} + p_{ij} \cdot 0\Big) \\
&\quad \cdot \log\Big(\sum_{k_1 \neq j} p_{ik_1} \sum_{k_2 \cdots k_m \neq j} p_{k_1 k_2} \cdots p_{k_m j} + p_{ij} \cdot 0\Big) \\
&\geq -\sum_{k_1 \neq j} p_{ik_1} \Big(\sum_{k_2 \cdots k_m \neq j} p_{k_1 k_2} \cdots p_{k_m j} \\
&\quad \cdot \log\Big(\sum_{k_2 \cdots k_m \neq j} p_{k_1 k_2} \cdots p_{k_m j}\Big)\Big) \\
&= -\sum_{k_1 \neq j} p_{ik_1} \mathbb{P}[T_{k_1 j} = m] \log \mathbb{P}[T_{k_1 j} = m], \qquad (5.21)
\end{aligned}$$

where the inequality uses equation (5.20). Summing both sides of (5.21) over m for $m \geq 1$, we have

$$\begin{aligned}
\mathbb{H}(T_{ij}) &\geq -p_{ij} \log p_{ij} + \sum\nolimits_{k_1 \neq j} p_{ik_1} \mathbb{H}(T_{k_1 j}) \\
&= -p_{ij} \log p_{ij} + \sum\nolimits_{k_1 = 1}^{n} p_{ik_1} \mathbb{H}(T_{k_1 j}) - p_{ij} \mathbb{H}(T_{jj}). \qquad (5.22)
\end{aligned}$$

Let $\mathbb{H}(T)$ be a matrix whose (i,j)-th element is $\mathbb{H}(T_{ij})$. Then equation (5.22) can be put in the matrix form

$$\mathbb{H}(T) \geq -P \circ \log P + P\mathbb{H}(T) - P[\mathbb{H}(T)], \qquad (5.23)$$

where the inequality and the log function are entrywise. Multiplying $\pi^\top$ from the left and $\mathbb{1}_n$ from the right on both sides of (5.23), we have

$$\pi^\top [\mathbb{H}(T)] \mathbb{1}_n \geq -\pi^\top (P \circ \log P) \mathbb{1}_n,$$

which is $\mathcal{J}(P) \geq \mathcal{H}(P)$.

Regarding (ii), if P is a permutation matrix, then $\mathcal{J}(P) = \mathcal{H}(P) = 0$. On the other hand, if P is not a permutation matrix, then there exist 2 or more nonzero elements on at least one row of P. In this case, the inequality in (5.21) is strict for that row for some m, which carries over to (5.22). Thus, $\mathcal{J}(P) > \mathcal{H}(P)$.

■

5.4 Truncated return time entropy and its optimization via gradient descent

We now introduce the truncated and conditional return time entropy and setup a gradient descent algorithm.

5.4.1 The truncated and conditional return time entropies

In practical applications, we may discard events occurring with extremely low probability. In what follows, we study the return time distribution and its entropy conditioned upon the event that the return time is upper bounded. We first introduce a *truncation accuracy* parameter $0 < \eta \ll 1$ that upper bounds the cumulative probabilities of very large return times and we define a *duration* $N_\eta \in \mathbb{Z}_{>0}$ by

$$N_\eta = \left\lceil \frac{w_{\max}}{\eta \pi_{\min}} \right\rceil - 1, \tag{5.24}$$

where $\pi_{\min} = \min_{i \in \{1,\ldots,n\}} \{\pi_i\}$ and $\lceil \cdot \rceil$ is the ceiling function. It is an immediate consequence of the Markov's inequality that, given the fixed stationary distribution π, for all

$i \in \{1, \dots, n\}$,

$$\mathbb{P}[T_{ii} \geq N_\eta + 1] \leq \frac{\mathbb{E}[T_{ii}]}{N_\eta + 1} \leq \frac{w_{\max}}{\pi_i (N_\eta + 1)} \leq \eta,$$

where we used (5.5)

$$\mathbb{E}[T_{ii}] = \frac{\pi^\top (P \circ W) \mathbb{1}_n}{\pi_i} \leq \frac{w_{\max}}{\pi_i}.$$

We now define the conditional return time and its entropy.

Definition 4 *(Conditional return time and its entropy) Given $P \in \mathcal{P}^\epsilon_{\mathcal{G},\pi}$ and a duration N_η, the conditional return time $T_{ii} \,|\, T_{ii} \leq N_\eta$ of state i is defined by*

$$T_{ii} \,|\, T_{ii} \leq N_\eta = \min \Big\{ \sum_{k'=0}^{k-1} w_{X_{k'} X_{k'+1}} \,|\, \sum_{k'=0}^{k-1} w_{X_{k'} X_{k'+1}} \leq N_\eta, \\ X_0 = i, X_k = i, k \geq 1 \Big\}.$$

with probability mass function

$$\mathbb{P}[T_{ii} = k \,|\, T_{ii} \leq N_\eta] = \frac{F_k(i,i)}{\sum_{k=1}^{N_\eta} F_k(i,i)}.$$

Moreover, the conditional return time entropy function $\mathcal{J}_{\text{cond},\eta} : \mathcal{P}^\epsilon_{\mathcal{G},\pi} \mapsto \mathbb{R}_{\geq 0}$ is

$$\begin{aligned} \mathcal{J}_{\text{cond},\eta}(P) &= \sum_{i=1}^n \pi_i \mathbb{H}(T_{ii} \,|\, T_{ii} \leq N_\eta) \\ &= -\sum_{i=1}^n \pi_i \sum_{k=1}^{N_\eta} \frac{F_k(i,i)}{\sum_{k=1}^{N_\eta} F_k(i,i)} \log \frac{F_k(i,i)}{\sum_{k=1}^{N_\eta} F_k(i,i)}. \end{aligned}$$

Given the duration N_η, $\mathcal{J}_{\text{cond},\eta}(P)$ is a finite sum of continuously differentiable functions and thus more tractable than the original return time entropy function $\mathcal{J}(P)$. Next, we introduce a truncated entropy that is even simpler to evaluate.

Definition 5 *(Truncated return time entropy function) Given a compact set $\mathcal{P}^\epsilon_{\mathcal{G},\pi}$ and*

the duration N_η, define the truncated return time entropy function $\mathcal{J}_{\text{trunc},\eta} : \mathcal{P}^{\epsilon}_{\mathcal{G},\pi} \mapsto \mathbb{R}_{\geq 0}$ *by*

$$\mathcal{J}_{\text{trunc},\eta}(P) = -\sum_{i=1}^{n} \pi_i \sum_{k=1}^{N_\eta} F_k(i,i) \log F_k(i,i).$$

The following lemma shows that, for small η, the truncated return time entropy $\mathcal{J}_{\text{trunc},\eta}(P)$ is a good approximation for the conditional return time entropy $\mathcal{J}_{\text{cond},\eta}(P)$. Furthermore, when η is sufficiently small, the truncated return time entropy $\mathcal{J}_{\text{trunc},\eta}(P)$ is also a good approximation for the original return time entropy function $\mathcal{J}(P)$.

Lemma 17 *(Approximation bounds) Given $P \in \mathcal{P}^{\epsilon}_{\mathcal{G},\pi}$ and the truncation accuracy η, we have*

(i) the conditional return time entropy is related to the truncated return time entropy by

$$\mathcal{J}_{\text{trunc},\eta}(P) + \log(1-\eta) < \mathcal{J}_{\text{cond},\eta}(P) < \frac{\mathcal{J}_{\text{trunc},\eta}(P)}{1-\eta}; \tag{5.25}$$

(ii) $\mathcal{J}(P) \geq \mathcal{J}_{\text{trunc},\eta}(P)$ holds trivially and if

$$\eta \leq \frac{w_{\max} \log \lambda_{\max}}{\pi_{\min}(\log \lambda_{\max} - \log c - 1)}, \tag{5.26}$$

then

$$\mathcal{J}(P) - \mathcal{J}_{\text{trunc},\eta}(P) \leq \frac{c \log(\lambda_{\max}^{-1})}{(1-\lambda_{\max})^2}(1+N_\eta)\lambda_{\max}^{N_\eta}, \tag{5.27}$$

where c and $\lambda_{\max}$ are given as in Lemma 12(i);

(iii) $\mathcal{J}(P) = \lim_{\eta \to 0^+} \mathcal{J}_{\text{cond},\eta}(P) = \lim_{\eta \to 0^+} \mathcal{J}_{\text{trunc},\eta}(P)$.

Proof: Regarding (i), for $\mathcal{J}_{\text{cond},\eta}(P)$, we have

$$\begin{aligned}\mathcal{J}_{\text{cond},\eta}(P) &= -\sum_{i=1}^{n}\pi_i \sum_{k=1}^{N_\eta} \frac{F_k(i,i)}{\sum\limits_{k=1}^{N_\eta} F_k(i,i)} \log \frac{F_k(i,i)}{\sum\limits_{k=1}^{N_\eta} F_k(i,i)} \\ &= -\sum_{i=1}^{n}\pi_i \left(\frac{\sum\limits_{k=1}^{N_\eta} F_k(i,i)\log F_k(i,i)}{\sum\limits_{k=1}^{N_\eta} F_k(i,i)} - \log \sum_{k=1}^{N_\eta} F_k(i,i) \right).\end{aligned}$$

On one hand,

$$\begin{aligned}\mathcal{J}_{\text{cond},\eta}(P) &> -\sum_{i=1}^{n}\pi_i \big(\sum_{k=1}^{N_\eta} F_k(i,i)\log F_k(i,i) - \log \sum_{k=1}^{N_\eta} F_k(i,i) \big) \\ &\geq -\sum_{i=1}^{n}\pi_i \sum_{k=1}^{N_\eta} F_k(i,i)\log F_k(i,i) + \log(1-\eta). \end{aligned} \tag{5.28}$$

On the other hand,

$$\begin{aligned}\mathcal{J}_{\text{cond},\eta}(P) &< -\sum_{i=1}^{n}\pi_i \frac{1}{\sum\limits_{k=1}^{N_\eta} F_k(i,i)} \sum_{k=1}^{N_\eta} F_k(i,i)\log F_k(i,i) \\ &\leq -\frac{1}{1-\eta}\sum_{i=1}^{n}\pi_i \sum_{k=1}^{N_\eta} F_k(i,i)\log F_k(i,i).\end{aligned} \tag{5.29}$$

Combining (5.28) and (5.29), we have (5.25).

Regarding (ii), if η satisfies (5.26), we have $c\lambda_{\max}^{N_\eta} \leq e^{-1}$. Then, following the same

arguments as in the proof of Lemma 12(ii) and replacing K in (5.14) with N_η, we have

$$\begin{aligned}
&\mathcal{J}(P) - \mathcal{J}_{\text{trunc},\eta}(P) \\
&\leq -c\left(\frac{\lambda_{\max}^{N_\eta}}{1-\lambda_{\max}}\log(c\lambda_{\max}^{N_\eta}) + \frac{\lambda_{\max}^{N_\eta+1}}{(1-\lambda_{\max})^2}\log(\lambda_{\max})\right) \\
&\leq -\frac{c\lambda_{\max}^{N_\eta}}{(1-\lambda_{\max})^2}\left(N_\eta\log(\lambda_{\max}) + \lambda_{\max}\log(\lambda_{\max}) + \log(c)\right) \\
&\leq -\frac{c\lambda_{\max}^{N_\eta}}{(1-\lambda_{\max})^2}\left(N_\eta\log(\lambda_{\max}) + \log(\lambda_{\max})\right) \\
&= \frac{c\log(\lambda_{\max}^{-1})}{(1-\lambda_{\max})^2}(1+N_\eta)\lambda_{\max}^{N_\eta}.
\end{aligned}$$

Regarding (iii), the results follow from (5.25) and (5.27), respectively. Specifically, in (5.27), since $0 < \lambda_{\max} < 1$, the error $\mathcal{J}(P) - \mathcal{J}_{\text{trunc},\eta}(P)$ goes to 0 exponentially fast as η goes to 0 ($N_\eta \to \infty$). ■

5.4.2 The gradient of the truncated return time entropy

Lemma 17 establishes how $\mathcal{J}_{\text{trunc},\eta}(P)$ is a good approximation to both of $\mathcal{J}(P)$ and $\mathcal{J}_{\text{cond},\eta}(P)$. Since it is also easier to compute $\mathcal{J}_{\text{trunc},\eta}(P)$ than the other two quantities, we focus on optimizing $\mathcal{J}_{\text{trunc},\eta}(P)$ by computing its gradient.

For $k \in \mathbb{Z}_{>0}$, define $G_k = \frac{\partial \operatorname{vec}(F_k)}{\partial \operatorname{vec}(P)} \in \mathbb{R}^{n^2 \times n^2}$ and note

$$G_k = \begin{bmatrix} \frac{\partial \operatorname{vec}(F_k)}{\partial p_{11}} & \frac{\partial \operatorname{vec}(F_k)}{\partial p_{21}} & \cdots & \frac{\partial \operatorname{vec}(F_k)}{\partial p_{(n-1)n}} & \frac{\partial \operatorname{vec}(F_k)}{\partial p_{nn}} \end{bmatrix}. \tag{5.30}$$

Lemma 18 *(Gradient of the truncated return time entropy function) Given $P \in \mathcal{P}^{\epsilon}_{\mathcal{G},\pi}$, the matrix sequence G_k in* (5.30) *satisfies the iteration for $k \in \mathbb{Z}_{>0}$,*

$$G_k = [\mathrm{vec}(\mathbf{1}_{\{k\mathbb{1}_n\mathbb{1}_n^\top = W\}})] + \sum_{i=1}^{w_{\max}} \Phi_i G_{k-i} + \sum_{i=1}^{n}\sum_{j=1}^{n}(E_j F_{k-w_{ij}}^\top \otimes I_n)[\mathrm{vec}(\mathbb{e}_i\mathbb{e}_j^\top)]\mathbf{1}_{\{w_{ij}>0\}}, \tag{5.31}$$

where the initial conditions are $G_k = \mathbb{0}_{n^2\times n^2}$ *for* $k \le 0$. *Moreover, the vectorization of the gradient of* $\mathcal{J}_{\mathrm{trunc},\eta}$ *satisfies*

$$\mathrm{vec}\Big(\frac{\partial \mathcal{J}_{\mathrm{trunc},\eta}(P)}{\partial P}\Big) = -\sum_{i=1}^{n}\pi_i \sum_{k=1}^{N_\eta} \frac{\partial\big(F_k(i,i)\log F_k(i,i)\big)}{\partial F_k(i,i)} G_k^\top \mathbb{e}_{(i-1)n+i}, \tag{5.32}$$

where $\mathbb{e}_{(i-1)n+i} \in \mathbb{R}^{n^2}$ *and*

$$\frac{\partial F_k(i,i)\log F_k(i,i)}{\partial F_k(i,i)} = \begin{cases} 1+\log(F_k(i,i)), & \text{if } F_k(i,i) > 0, \\ 0, & \text{if } F_k(i,i) = 0. \end{cases}$$

Proof: For $k \in \mathbb{Z}_{>0}$, according to (5.6), we have for $p_{uv} > 0$,

$$\frac{\partial \,\mathrm{vec}(F_k)}{\partial p_{uv}} = \mathrm{vec}(\mathbb{e}_u\mathbb{e}_v^\top)\mathbf{1}_{\{k=w_{uv}\}} + (E_v \otimes \mathbb{e}_u\mathbb{e}_v^\top)\,\mathrm{vec}(F_{k-w_{uv}}) + \sum_{i=1}^{n}\sum_{j=1}^{n} p_{ij}(E_j \otimes \mathbb{e}_i\mathbb{e}_j^\top)\frac{\partial\,\mathrm{vec}(F_{k-w_{ij}})}{\partial p_{uv}},$$

where the second term on the right hand side satisfies

$$\begin{aligned}(E_v \otimes \mathbb{e}_u\mathbb{e}_v^\top)\,\mathrm{vec}(F_{k-w_{uv}}) &= \mathrm{vec}(\mathbb{e}_u\mathbb{e}_v^\top F_{k-w_{uv}}E_v) \\ &= (E_v F_{k-w_{uv}}^\top \otimes I_n)\,\mathrm{vec}(\mathbb{e}_u\mathbb{e}_v^\top).\end{aligned}$$

Stacking $\frac{\partial\,\mathrm{vec}(F_k)}{\partial p_{uv}}$'s in a matrix as (5.30), we obtain (5.31).

Since $\mathcal{J}_{\mathrm{trunc},\eta}(P)$ only involves $F_k(i,i)$ for $i = \{1,\dots,n\}$, we only need the corre-

sponding columns in $G_k^\top$ to compute the gradient, which is realized by multiplying the standard unit vector as in (5.32). ■

Remark 13 *Iteration* (5.31) *is an exponentially stable discrete-time delayed linear system subject to and a finite number of impulse inputs and an exponentially vanishing input. Hence, the state* $G_k \to 0$ *exponentially fast as* $k \to \infty$.

5.4.3 Optimizing the truncated entropy via gradient projection

Motivated by the previous analysis, we consider the following problem.

Problem 7 *(Maximization of the truncated return time entropy) Given a strongly connected directed graph* $\mathcal{G}$ *and the stationary distribution* π*, pick a minimum edge weight* $\epsilon > 0$ *and a truncation accurate parameter* $\eta > 0$*, the maximization of the truncated return time entropy function is as follows:*

$$\begin{aligned} &\textit{maximize} \quad \mathcal{J}_{\mathrm{trunc},\eta}(P), \\ &\textit{subject to} \quad P \in \mathcal{P}^{\epsilon}_{\mathcal{G},\pi}. \end{aligned}$$

To solve numerically this nonlinear program, we exploit the results in Lemma 18 and adopt the gradient projection method as presented in [156, Chapter 2.3]:

1: select: minimum edge weight $\epsilon \ll 1$, truncation accuracy $\eta \ll 1$, and initial condition P_0 in $\mathcal{P}^{\epsilon}_{\mathcal{G},\pi}$
2: **for** iteration parameter $s = 0$: (number-of-steps) **do**
3: $\quad \{G_k\}_{k\in\{1,\ldots,N_\eta\}} :=$ solution to iteration (5.31) at P_s
4: $\quad \Delta_s :=$ gradient of $\mathcal{J}_{\mathrm{trunc},\eta}(P_s)$ via equation (5.32)
5: $\quad P_{s+1} := \mathrm{projection}_{\mathcal{P}^{\epsilon}_{\mathcal{G},\pi}}(P_s + (\text{step size}) \cdot \Delta_s)$
6: **end for**

We analyze the computational complexity of this algorithm. To compute step `3:`, we need to evaluate the right-hand side of equation (5.31) by computing three terms. For the first term, we need to do m comparisons, where m is the number of edges in the graph (i.e., the number of variables in the transition matrix), and it takes $O(m)$ elementary operations. For the second term, note that the matrices $\Phi_i \in \mathbb{R}^{n^2 \times n^2}$ introduced in equation (5.13) can be precomputed and is block diagonal with n blocks of size $n \times n$. Also note that $G_k \in \mathbb{R}^{n^2 \times n^2}$ has only m nonzero columns. Thus, we need $O(w_{\max} m n^3)$ operations. For the third term, F_k is updated by equation (5.11), which requires $O(w_{\max} n^3)$ and is the main computational cost. Therefore, it takes $O(w_{\max} m n^3)$ to compute one update of iteration (5.31). Thus, it takes $O(N_\eta w_{\max} m n^3)$ elementary operations to complete step `3:`. In step `5:`, we need to solve a least square problem with linear equalities and inequalities constraints; which requires $O(m^3)$ [157].

5.5 Numerical results

In this section, we provide numerical results on the computation of the maximum return time entropy chain (Subsection 5.5.1) and its application to robotic surveillance problems (Subsection 5.5.2). We compute and compare three chains:

(i) the Markov chain that maximizes the return time entropy (solution of Problem 6), abbreviated as the *MaxReturnEntropy chain*. This chain may be computed for a directed graph with arbitrary integer-valued travel times. Since we do not have a way to solve Problem 6 directly, the MaxReturnEntropy chain is approximated by the solution of Problem 7, which is solved via the gradient projection method. Unless otherwise stated, we choose truncation accuracy $\eta = 0.1$. Note that (5.24) is quite conservative and the actual probabilities being discarded is much less than 0.1.

(ii) the Markov chain that maximizes the entropy rate, as discussed in chapter 4, abbreviated as the *MaxEntropyRate chain*. Note, that this chain can be computed for a directed graph with unitary weights via solving a convex program. If the graph is undirected, the MaxEntropyRate chain can be computed efficiently via the methods described in the previous chapter.

(iii) the Markov chain that minimizes the mean (weighted) hitting time known as the (weighted) Kemeny constant, abbreviated as the *MinKemeny chain*. This chain may be computed for a directed graph with arbitrary travel times via solving a nonlinear nonconvex program [53]. We compute this chain using the solver implemented in the KNITRO/TOMLAB package.

5.5.1 Computation, comparison and intuition

We divide this subsection into two parts. In the first part, we first compare 3 chains on graphs that have unitary travel times. We then summarize several observations in computing the MaxReturnEntropy chain. Finally, we visualize and plot the chains as well as the return time distributions. In the second part, we compare the MaxReturnEntropy chain with the MinKemeny chain on a realistic map taken from [2, Section 6.2] with travel times.

Chains on graphs with unitary travel times

Comparison: We consider 2 simple undirected graphs and solve for the MaxReturnEntropy chain, the MaxEntropyRate chain and the MinKemeny chain for each case. We compare the return time entropy, the entropy rate, and the Kemeny constant of these chains in Table 5.1. The stationary distribution of the ring graph is set to be $\pi = [1/12, 1/6, \ldots, 1/12, 1/6]^\top$, and the stationary distribution of of grid is proportional

x

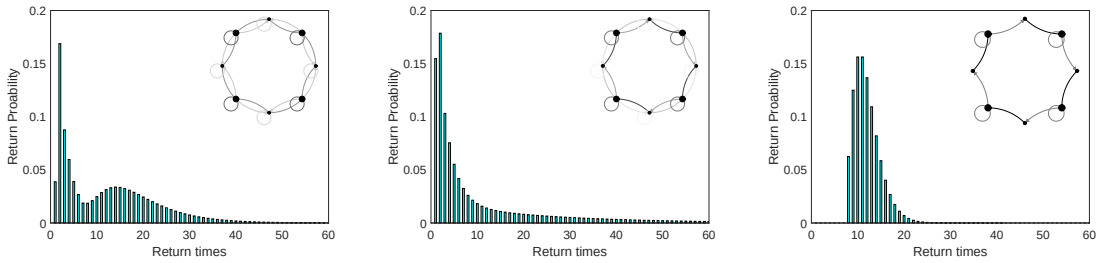

(a) MaxReturnEntropy chain on ring graph (b) MaxEntropyRate chain on ring graph (c) MinKemeny chain on ring graph

Figure 5.2: Return time distributions of node 1 (i.e., top node) on an 8-node ring graph with stationary distribution $\pi = [1/12, 1/6, \ldots, 1/12, 1/6]^\top$. Although the expectations of the first return time distributions in the figure are the same, the histogram is remarkably different for different chains. Specifically, for the nonreversible MaxRetrunEntropy chain, the distribution is bimodal and generates more entropy. The node size is proportional to the stationary distribution.

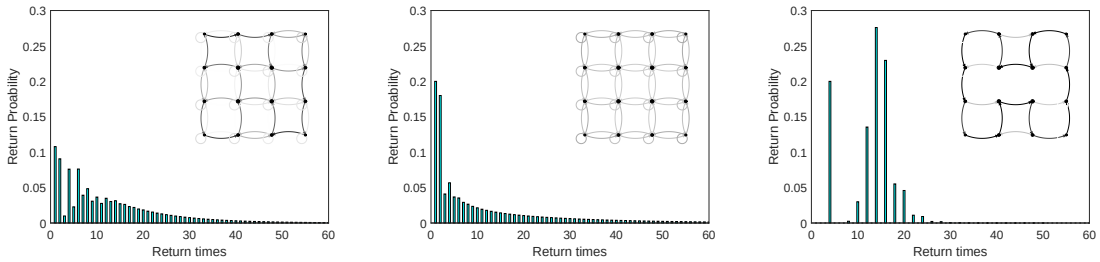

(a) MaxReturnEntropy chain on 4×4 grid (b) MaxEntropyRate chain on 4×4 grid (c) MinKemeny chain on 4×4 grid

Figure 5.3: Return time distributions of node 6 (i.e., second node on the second row) on a 4×4 grid with stationary distribution π proportional to the node degree and unitary travel times. The node size is proportional to the stationary distribution.

to the degree of nodes. To evaluate the value of $\mathcal{J}(P)$, we set $\eta = 10^{-2}$. From the table, we notice that the MaxReturnEntropy chain has the highest value of the return time entropy in both cases. It also has relatively good performance in terms of the entropy rate and the Kemeny constant, which indicates that the MaxReturnEntropy chain is potentially a good combination of speed (expected traversal time) and unpredictability. Furthermore, it is clear that (5.15), which characterizes the relationship between the entropy rate and the return time entropy, holds.

Table 5.1: Comparison between different chains on different graphs

Graph	Markov chains	$\mathcal{J}(P)$	$\mathcal{H}(P)$	Kemeny constant
8-node ring	MaxReturnEntropy	2.4927	0.8698	10.0479
	MaxEntropyRate	2.3510	0.9883	19.5339
	MinKemeny	1.9641	0.4621	6.1667
4-by-4 grid	MaxReturnEntropy	3.6539	0.9491	16.3547
	MaxEntropyRate	3.2844	1.4021	30.8661
	MinKemeny	2.0990	0.2188	10.0938

Observations: In computing the MaxReturnEntropy chain, we observe some interesting properties of our problem. First, when solving Problem 7 by the gradient projection method with different initial conditions, we found different optimal solutions, and they have slightly different optimal values. This suggests that Problem 6 is unlikely to be a convex problem. Secondly, the global optimal solution to Problem 6 is possibly not unique in general. For instance, for an undirected ring graph with even number of nodes and certain stationary distribution, exchanging the probability of going right and that of going left for all nodes does not change the return time entropy. Thirdly, the optimal solution to Problem 6 is likely to be nonreversible because none of the approximate optimal solutions we have encountered are reversible. This again indicates that the MaxReturnEntropy chain is a good combination of unpredictability and speed. Fourth, even if we set the edge weight $\epsilon = 0$, the MaxReturnEntropy chain is always irreducible.

Intuition: In order to provide intuition for the maximization of the return time entropy, we compare and plot the chains as well as the return time distribution of a same node on the 8-node ring graph and the 4×4 grid graph in Fig. 5.2 and Fig. 5.3, respectively. Since the stationary distribution is fixed and identical for all chains in each case, the expectations of the probability mass functions in each figure are the same. From the figures, we note that for the MaxReturnEntropy chain, the return time distribution is reshaped so that the distribution is more spread out and it is more difficult to predict the return time. In contrast, the return time distribution for the MinKemeny chain has a predictable pattern and the return time probability is constantly 0 for some time intervals. Moreover, from the visualization of the chains, we notice that the MaxReturnEntropy chain has a net flow on the graph, which again indicates its nonreversibility.

MaxReturnEntropy and MinKemeny on a realistic map

In this part, we compare the MaxReturnEntropy chain with the MinKemeny chain on a realistic map with travel times. The problem data is taken from [2, Section 6.2]: a small area in San Francisco (SF) is modeled by a fully connected directed graph with 12 nodes and by-car travel times on edges measured in seconds. The map is shown in Fig. 5.4. The importance of the a location (node) is characterized by the the number of crimes recorded at that place during a specific period, and the surveillance agent should visit the places with higher crime rate more often. The visit frequency is set to be $[\frac{133}{866}, \frac{90}{866}, \frac{89}{866}, \frac{87}{866}, \frac{83}{866}, \frac{83}{866}, \frac{74}{866}, \frac{64}{866}, \frac{48}{866}, \frac{43}{866}, \frac{38}{866}, \frac{34}{866}]^\top$. For simplicity, we quantize the travel times by treating a minute as one unit of time, i.e., dividing the travel times by 60 and round the result to the smallest integer that is larger than it, and by doing so, we have $w_{\max} = 9$. The pairwise travel times are recorded in Table 5.2.

First, we compare three key metrics of the MaxReturnEntropy chain and MinKemeny chain. The results are reported in Table 5.3. It can be observed that the MaxRetur-

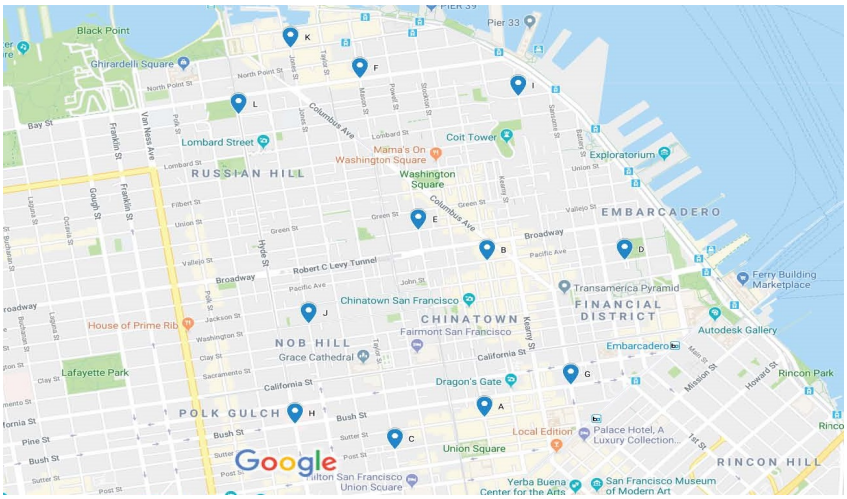

Figure 5.4: San Francisco (SF) crime map from [2, Section 6.2].

nEntropy chain is much better than the MinKemeny chain regarding the return time entropy and the entropy rate. This better performance in terms of the unpredictability is obtained at the cost of being slower as indicated by the larger weighted Kemeny constant.

We also plot the return time distribution of location A in Fig. 5.5. Apparently, the MaxReturnEntropy chain spreads the return time probabilities over the possible return times and it is hard to predict the exact time the surveillance agent comes back to the location. In contrast, the MinKemeny chain tries to achieve fast traversal on the graph and the return times distribute over a few intervals.

5.5.2 Application to the Robotic Surveillance Problem

In this subsection, we provide simulation results in the application of robotic surveillance.

Setup: Consider the scenario where a single agent performs the surveillance task by moving randomly according to a Markov chain on the road map. The intruder is able to observe the local behaviors of the surveillance agent, e.g., presence/absence and duration

Table 5.2: The quantized pairwise by-car travel times on SF crime map

Location	A	B	C	D	E	F	G	H	I	J	K	L
A	1	3	3	5	4	6	3	5	7	4	6	6
B	3	1	5	4	2	4	4	5	5	3	5	5
C	3	5	1	7	6	8	3	4	9	4	8	7
D	6	4	7	1	5	6	4	7	5	6	6	7
E	4	3	6	5	1	3	5	5	6	3	4	4
F	6	4	8	5	3	1	6	7	3	6	2	3
G	2	5	3	5	6	7	1	5	7	5	7	8
H	3	5	2	7	6	7	3	1	9	3	7	5
I	8	6	9	4	6	4	6	9	1	8	5	7
J	4	3	4	6	3	5	5	3	7	1	5	3
K	6	4	8	6	4	2	6	6	4	5	1	3
L	6	4	6	6	3	3	6	4	5	3	2	1

Table 5.3: Comparison between different chains on SF crime map

Markov chains	$\mathcal{J}(P)$	$\mathcal{H}(P)$	Weighted Kemeny constant
MaxReturnEntropy	5.0078	1.7810	63.6007
MinKemeny	2.4678	0.6408	24.2824

between visits, and he/she plans and decides the time of attack so as to avoid being captured. It takes a certain amount of time for the intruder to complete an attack, which is called the *attack duration* of the intruder. A successful detection/capture happens when the surveillance agent and the intruder are at the same location and the intruder is attacking.

Intruder model (success probability maximizer with bounded patience): Consider a rational intruder that exploits the return time statistics of the Markov chains and chooses an optimal attack time so as to minimize the probability of being captured. The intruder picks a node i to attack randomly according to the stationary distribution, and it collects and learns the probability distribution of node i's first return time. Suppose the intruder and the surveillance agent are at the same node i at the beginning and the attack duration of the intruder is τ. If the intruder observes that the surveillance agent leaves the node

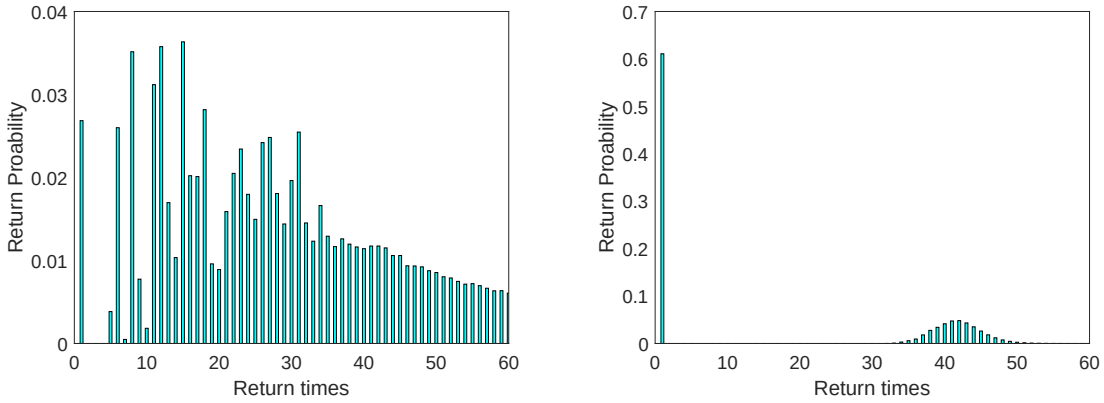

(a) MaxReturnEntropy chain on SF crime map (b) MinKemeny chain on SF crime map

Figure 5.5: Return time distributions of location A on SF crime map. Note that the scales of the vertical axes are different in the two figures.

and does not come back for s periods, he/she can attack with the probability of being captured given by

$$\sum_{k=1}^{\tau} \mathbb{P}\left[T_{ii} = s + k \,|\, T_{ii} > s\right]. \tag{5.33}$$

Mathematically speaking, (5.33) is the conditional cumulative return probability for the surveillance agent. Specifically for $s = 0$, (5.33) is the capture probability when the intruder attacks immediately after the agent leaves the node. Then, the optimal time of attack s_i for the intruder is given by

$$s_i = \text{argmin}_{0 \le s \le S_i} \{ \sum_{k=1}^{\tau} \mathbb{P}\left[T_{ii} = s + k \,|\, T_{ii} > s\right] \}. \tag{5.34}$$

The reason there is an upper bound S_i on s is that the event $T_{ii} > s$ happens with very low probability when s is large, and the intruder may be unwilling to wait for such an event to happen. Let $\delta \in (0, 1)$ be the *degree of impatience* of the intruder, then S_i can

be chosen as the minimal positive integer such that the following holds,

$$\mathbb{P}\left[T_{ii} \geq S_i\right] \leq \delta,$$

where a smaller δ implies a larger S_i and a more patient intruder. In other words, when δ is small, the intruder is willing to wait for a rare event to happen. Note that the value of S_i is also dependent on the node i that the intruder chooses to attack, and thus the argmin in (5.34) is over different ranges when the intruder attacks different nodes. In summary, the intruder is dictated by two parameters: the attack duration τ and the degree of impatience δ, and the strategy for the intruder is as follows: waits until the event that the surveillance agent leaves and does not come back for the first s_i steps happens, then attacks immediately.

From the surveillance point of view, the probability of capturing the rational intruder when he/she attacks node i is

$$\mathbb{P}_i\left[\text{Capture}\right] = \sum_{k=1}^{\tau} \mathbb{P}\left[T_{ii} = s_i + k \,|\, T_{ii} > s_i\right],$$

and the performance of the Markov chains can be evaluated by the total probability of capture as follows

$$\mathbb{P}\left[\text{Capture}\right] = \sum_{i=1}^{n} \pi_i \mathbb{P}_i\left[\text{Capture}\right]. \tag{5.35}$$

Simulation results: Designing an optimal defense mechanism for the rational intruder is an interesting yet challenging problem in its own. Instead, we use the MaxReturnEntropy chain as a heuristic solution and compare its performance with other chains. In the following, we consider two types of graphs: the grid graph and the SF crime map. The degree of impatience of the intruder is set to be $\eta = 0.1$ in this part.

We first consider a 4×4 grid and plot the probability of capture defined by (5.35)

for the chains in comparison in Fig. 5.6. It can be observed that, when defending against the rational intruder described above, the MaxReturnEntropy chain outperforms all other chains when the attack duration of the intruder is small or moderate. The unpredictability in the return time prevents the rational intruder from taking advantage of the visit statistics learned from the observations. The MinKemeny chain, which emphasizes a faster traversal, has a hard time capturing the intruder when the attack duration of the intruder is small. This is because the agent moves in a relatively more predictable way, and the return time statistics may have a pattern that could be exploited. The MaxEntropyRate chain has the in-between performance.

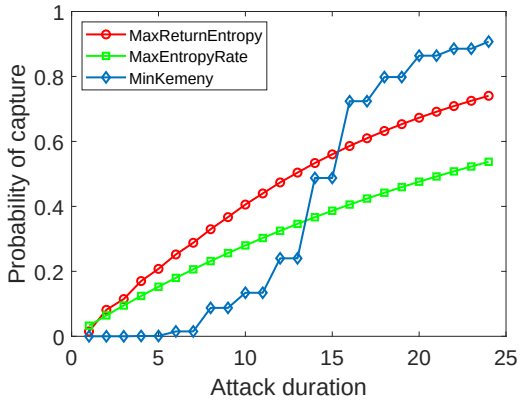

Figure 5.6: Performance of different chains on a 4×4 grid.

For the SF crime map, we use the same problem data as described in Subsection 5.5.1. Since the MaxEntropyRate chain does not generalize to the case when there are travel times, we compare the performance of the MaxReturnEntropy chain and the MinKemny chain. Again, The MaxReturnEntropy chain outperforms the MinKemeny chain when the attack duration of the intruder is relatively small.

Conclusion: The simulation results presented in this subsection demonstrate that the MaxReturnEntropy chain is an effective strategy against the intruder with reasonable

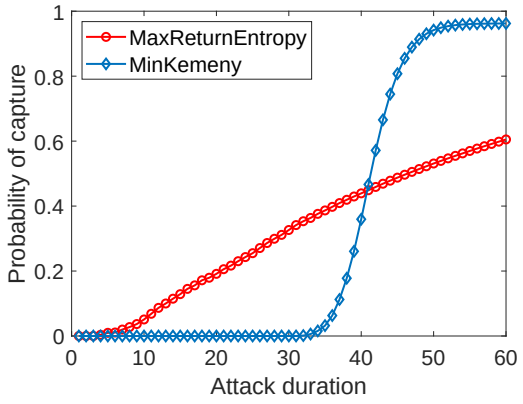

Figure 5.7: Performance of different chains on the SF crime map.

amount of knowledge and level of intelligence, particularly when the attack duration of the intruder is small or moderate. With the property of both unpredictability and speed, the MaxReturnEntropy chain should also work well in a much more broader range of scenarios.

5.6 Summary

In this chapter, we proposed and optimized a new metric that quantifies the unpredictability of Markov chains over a directed strongly connected graph with travel times, i.e., the return time entropy. We characterized the return time probabilities and showed that optimizing the return time entropy is a well-posed problem. For the case of unitary travel times, we established an upper bound for the return time entropy by using the maximum entropy principle and obtained an analytic solution for the complete graph. We connected the return time entropy with the well-known entropy rate of Markov chains and showed that the return time entropy is lower bounded by the entropy rate and upper bounded by n times the entropy rate. In order to solve the optimization problem numerically, we approximated the return time entropy as well as a practically useful conditional

return time entropy by the truncated return time entropy. We derived the gradient of the truncated return time entropy and proposed to solve the problem by the gradient projection method. We applied our results to the robotic surveillance problem and found that the chain with maximum return time entropy is a good trade-off between speed and unpredictability, and it performs better than several existing chains against a rational intruder.

Chapter 6

Conclusions and Future Work

Before we conclude this thesis and present future directions, it is interesting to note the role that unpredictability plays in both endeavors: biological and robotic. In the design of stochastic surveillance strategies, the role of unpredictability is quite apparent as some of the solutions presented directly try to maximize entropy. In the description of collective cell migration, randomness plays a much subtler role: a Langevin-type noise models the formation of ruffling lamellipodia in the theoretical description. These ruffling behaviors are key to breaking the symmetry of the system, initiating persistent motion in groups of cells, and forming trends in traction and stress profiles in colonies. All of these behaviors are salient features of collective migration in groups of cells, which are observed in experiments and would not be observed in simulations if not for the addition of noise.

6.1 Summary

In chapter 2, we presented a theoretical description of cell migration that accounts for known individual cell behaviors, such as contact inhibition of locomotion and force-induced repolarization, and is able to reproduce the motion of a single cell, two cell collisions, small groups of cells and large colonies. This description provides a unified

framework to connect a large number of experiments in different conditions and with different cell types. Moreover, it allows a direct connection between specific molecular perturbations in cell adhesion, cell polarization, the generation of traction forces and mechanical feedback, and their effect on collective cell migration. The theory presented in this chapter provides a framework within which the results of multiple experiments can be understood and future experiments organized.

In chapter 3, we studied the meeting time of multiple random walkers on a graph and presented necessary and sufficient conditions for finiteness and novel closed-form expressions for the expected time to meeting between a single pursuer and a single evader, multiple pursuers and multiple evaders, and extended the treatment to continuous-time chains. We also provide sufficient conditions for certain pairs (or tuples) of Markov chains that satisfy conditions on their absorbing classes to have finite meeting times. Finally, we discuss connections to other metrics relevant to Markov chains such as the hitting time.

In chapter 4, we considered the optimization problem of maximizing the entropy rate of a Markov chain with prescribed stationary distribution. We showed this problem is strictly convex with a unique global optimizer. We provided a fast iterative algorithm with rigorous convergence guarantees to compute the so-called entropic vector; as a function of this entropic vector, we provide a closed-form formula for the maximum entropy Markov chain with prescribed stationary distribution. We then characterized several properties of maxentropic chains. The interest for Markov chains with maximum entropy rate and prescribed stationary distributions arises naturally in robotic surveillance; accordingly we showed some realizations of optimal chains for prototypical robotic roadmaps.

In chapter 5, we proposed and optimized a different notion of unpredictability of Markov chains over a directed strongly connected graph with travel times based on the

entropy of the random variable associated with the return time to nodes. We characterized the return time probabilities and showed that optimizing the return time entropy is a well-posed problem. For the case of unitary travel times, we show that the return time entropy is lower bounded by the entropy rate and upper bounded by n times the entropy rate. In order to solve the optimization problem numerically, we approximated the return time entropy as well as a practically useful conditional return time entropy by the truncated return time entropy. We applied our results to the robotic surveillance problem and found that the chain with maximum return time entropy is a good trade-off between speed and unpredictability, and it performs better than several existing chains against a rational intruder.

In general, it is noted that the maximum entropy rate Markov chain and the maximum return time entropy chain are both good solutions when the intruder is capable of observing the motion of the surveillance agent. The maximum return time entropy chain formulation has three main advantages.

(i) It is more reasonable to expect that an intruder would have access to the return times of the agent than have knowledge of the sequence of locations of the agent. (For e.g., an intruder hovering just outside the visibility range of the surveillance agent at a node). Thus the entropy of a more relevant random variable is being maximized.

(ii) This chain is naturally non-reversible and comparison of hitting times indicate that in general, the maximum return time entropy chain has lower travel times than the maximum entropy rate Markov chain.

(iii) The formulation can tackle directed, weighted graphs.

The maximum entropy rate Markov chain presents the following advantages.

(i) The computation of maximum entropy rate Markov chains is significantly more computationally efficient than the maximum return time entropy chain enabling its use on much larger graphs and potential recomputation in online scenarios, i.e., where the stationary distribution is changing with time.

(ii) Maximum entropy rate Markov chains optimize a more well-studied quantity, i.e. the entropy rate, and hence is more widely applicable as mentioned in subsection 4.1.2.

The above discussion indicates that for applications to robotic surveillance where the strategy is precomputed and computational costs are not a concern, the maximum return time entropy chain should be the preferred strategy.

6.2 Future Work

The work described in this thesis leaves open several directions for future research. We present future directions for each strand of research discussed.

Collective cell migration: The theoretical description presented in chapter 2 shows that the collective migration of small groups of cells can be understood within the same framework as single cell migration and the expansion of large colonies. Extensions of this work to 2D and 3D systems [158], as well as the consideration of cell shapes or biochemical signaling, will help further elucidate how different modes of collective migration emerge in developing embryos.

Aside from theoretical directions in which this work can be extended, the theory bears several predictions which could be used to motivate experiments on cell motility. Our results predict that the diffusion constant of the motion of a single cell depends quadratically with the characteristic traction force of the cell type. We established that collision

dynamics between cells are strongly dependent on parameters such as the traction force and inter-cellular adhesion both of which can be targeted experimentally. For developing persistence in groups of cells, we see that neighbor-enabled repolarization (NER) plays a crucial role, hence, experiments which alter the polarization dynamics of cells should measurably affect the persistence characteristics observed in clusters. Further our description predicts that absence of cell proliferation can lead to breakage of cell colonies when the intercellular adhesion is weak. This can be probed using experimental techniques which limit proliferation [159] and alter inter-cellular adhesion [160].

Meeting times of random walks: Several future directions of interest are left open by the work exploring the meeting times of random walks. Though we provide closed-form expressions here, the complexity involved in the calculation makes the computation expensive for large number of agents and on graphs with large number of nodes. It would be of practical interest to devise a formulation which has lower complexity. The literature on computationally efficient methods to calculate the SimRank of nodes on a graph might provide alternative formulations [79]. An extension of the work discussed here would be to consider walkers moving with travel times similar to the case of doubly weighted graphs described in [52].

The hitting time of a Markov chain is intimately connected with the fundamental matrix of a Markov chain [57]. It would be interesting to see whether meeting times between Markov chains can be expressed in terms of their fundamental matrices. Algebraic techniques used to compute the pseudo-inverse could potentially aid in the discovery of such an expression [161]. The notion of random walks on Kronecker product graphs, which is extensively used in chapter 3, can potentially be utilized to provide closed-form formulations of other quantities such as the expected time to capture of all evaders or the coalescence time of multiple random walkers [130].

In the robotic context, the expression for meeting times can be utilized to setup

pursuit-evasion games in which policies correspond to choice of Markov chains. It would be of interest to study optimal strategies for when the evader assumes certain behaviors (such as a deterministic walk or a random walk such as the equal neighbor Markov chain) and whether Nash equilibria exist for such games. Also, the computational tools provided here can be employed in the design of random walks with meeting time constraints. This could be useful in scenarios where communication is limited by distance and the agents need to communicate regarding events occurring on the graph (such as detection of anomalies) by being in close range.

Maximum entropy surveillance: The work on designing maximum entropy rate surveillance strategies was primarily targeted at the single- agent case on an undirected graph with self-loops at every node. As such, it is potentially important to extend the analysis in chapter 4 to more general graph settings, including graphs without a complete set of self-loops and directed graphs with asymmetric adjacency matrices. For graphs without a complete set of self-loops, we present a mathematical conjecture inspired by our results on the maxentropic matrix and vector maps.

Conjecture: Given a connected graph G with binary adjacency matrix A, the set of stationary distributions for all irreducible Markov chains over G is $\{[x]Ax/\|[x]Ax\|_1 \mid x \in \mathbb{R}^n_{>0}\}$.

Numerical simulations indicate that the set of feasible stationary distributions over sparse graphs without self-loops is of measure zero (for an appropriately defined measure).

Several open directions remain in the work on designing maximum return time entropy surveillance strategies. First of all, a simple closed- form expression for the return time entropy would enable us to establish more properties of the objective function and thus make the optimization problem more tractable. Second, it is interesting to design an optimal Markov chain that maximizes the probability of capture of the intruder model proposed in this chapter. Third, we believe there are more application scenarios for

Markov chains where the return time entropy is an appropriate quantity to optimize.

Most importantly, appropriate notions of unpredictability for the multi-agent case are yet to be developed for both approaches and could lead to the design of effective strategies against intruders with advanced planning and sensing capabilities. It should be possible to generate more entropy rate as well as return time entropy using multiple agents. With the help of an appropriately defined notion of group entropy rate or group return time entropy, one could utilize this quantity to further increase the unpredictability of surveillance strategies. This could also lead to strategies that trade-off between unpredictability and speed more effectively, e.g., on a shared graph some of the agents could optimize for entropy while others could be optimized for speed. It is also of interest to combine notions of unpredictability with speed of traversal of graphs; see recent related work in [52].

Bibliography

[1] R. A. Desai, S. B. Gopal, S. Chen, and C. S. Chen, *Contact inhibition of locomotion probabilities drive solitary versus collective cell migration*, *Journal of The Royal Society Interface* **10** (2013), no. 88 20130717.

[2] S. Alamdari, E. Fata, and S. L. Smith, *Persistent monitoring in discrete environments: Minimizing the maximum weighted latency between observations*, *Int J Robotic Research* **33** (2014), no. 1 138–154.

[3] U. Lopez, J. Gautrais, I. D. Couzin, and G. Theraulaz, *From behavioural analyses to models of collective motion in fish schools*, *Interface focus* **2** (2012), no. 6 693–707.

[4] M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, M. M. Viale, and V. Zdravkovic, *Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study*, *Proceedings of the National Academy of Sciences* **105** (2008), no. 4 1232–1237, [http://www.pnas.org/content/105/4/1232.full.pdf].

[5] M. Turpin, N. Michael, and V. Kumar, *Trajectory design and control for aggressive formation flight with quadrotors*, *Autonomous Robots* **33** (Aug, 2012) 143–156.

[6] W. Yu, W. Ren, W. X. Zheng, G. Chen, and J. L, *Distributed control gains design for consensus in multi-agent systems with second-order nonlinear dynamics*, *Automatica* **49** (2013), no. 7 2107 – 2115.

[7] S. Spring, V. Kumar, and R. Nagpal, *Design of control policies for spatially inhomogeneous robot swarms with application to commercial pollination*, in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pp. 378–385, IEEE, 2011.

[8] R. H. Pritchard, Y. Y. S. Huang, and E. M. Terentjev, *Mechanics of biological networks: from the cell cytoskeleton to connective tissue*, *Soft matter* **10** (2014), no. 12 1864–1884.

[9] T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet, *Novel type of phase transition in a system of self-driven particles*, *Physical Review Letters* **75** (1995), no. 6-7 1226–1229.

[10] S. Martínez, J. Cortés, and F. Bullo, *Motion coordination with distributed information*, *IEEE Control Syst. Mag.* **27** (2007), no. 4 75–88.

[11] H. Oh, A. R. Shirazi, C. Sun, and Y. Jin, *Bio-inspired self-organising multi-robot pattern formation: A review*, *Robotics and Autonomous Systems* **91** (2017) 83 – 100.

[12] J. Halloy, G. Sempo, G. Caprari, C. Rivault, M. Asadpour, F. Tâche, I. Saïd, V. Durier, S. Canonge, J. M. Amé, C. Detrain, N. Correll, A. Martinoli, F. Mondada, R. Siegwart, and J. L. Deneubourg, *Social integration of robots into groups of cockroaches to control self-organized choices*, *Science* **318** (2007), no. 5853 1155–1158.

[13] V. Schaller, C. Weber, C. Semmrich, E. Frey, and A. R. Bausch, *Polar patterns of driven filaments*, *Nature* **467** (2010), no. 7311 73.

[14] D. J. Montell, *Border-cell migration: the race is on*, *Nature reviews Molecular cell biology* **4** (2003), no. 1 13.

[15] M. Moussaïd, D. Helbing, and G. Theraulaz, *How simple rules determine pedestrian behavior and crowd disasters*, *Proceedings of the National Academy of Sciences* **108** (2011), no. 17 6884–6888.

[16] J. L. Deneubourg, S. Aron, S. Goss, and J. M. Pasteels, *The self-organizing exploratory pattern of the argentine ant*, *Journal of Insect Behavior* **3** (Mar, 1990) 159–168.

[17] S. Bazazi, J. Buhl, J. J. Hale, M. L. Anstey, G. A. Sword, S. J. Simpson, and I. D. Couzin, *Collective motion and cannibalism in locust migratory bands*, *Current Biology* **18** (2008), no. 10 735–739.

[18] R. N. Smith, M. Schwager, S. L. Smith, B. H. Jones, D. Rus, and G. S. Sukhatme, *Persistent ocean monitoring with underwater gliders: Adapting sampling resolution*, *Journal of Field Robotics* **28** (2011), no. 5 714–741.

[19] A. Macwan, G. Nejat, and B. Benhabib, *Optimal deployment of robotic teams for autonomous wilderness search and rescue*, in *Proc IROS*, (San Francisco, CA, USA), pp. 4544–4549, Sept., 2011.

[20] J. Grace and J. Baillieul, *Stochastic strategies for autonomous robotic surveillance*, in *Proc CDC-ECC*, (Seville, Spain), pp. 2200–2205, Dec., 2005.

[21] P. Friedl and D. Gilmour, *Collective cell migration in morphogenesis, regeneration and cancer*, *Nature reviews Molecular cell biology* **10** (2009), no. 7 445.

[22] C. Carmona-Fontaine, H. K. Matthews, S. Kuriyama, M. Moreno, G. A. Dunn, M. Parsons, C. D. Stern, and R. Mayor, *Contact inhibition of locomotion in vivo controls neural crest directional migration*, *Nature* **456** (2008), no. 7224 957.

[23] V. Lecaudey and D. Gilmour, *Organizing moving groups during morphogenesis*, *Current opinion in cell biology* **18** (2006), no. 1 102–107.

[24] P. Haas and D. Gilmour, *Chemokine signaling mediates self-organizing tissue migration in the zebrafish lateral line*, *Developmental cell* **10** (2006), no. 5 673–680.

[25] K. J. Cheung and A. J. A. J. Ewald, *A collective route to metastasis: Seeding by tumor cell clusters*, *Science* **352** (2016), no. 6282 167–169.

[26] A. J. Ridley, M. A. Schwartz, K. Burridge, R. A. Firtel, M. H. Ginsberg, G. Borisy, J. T. Parsons, and A. R. Horwitz, *Cell migration: integrating signals from front to back*, *Science* **302** (2003), no. 5651 1704–1709.

[27] M. Prasad and D. J. Montell, *Cellular and molecular mechanisms of border cell migration analyzed using time-lapse live-cell imaging*, *Developmental cell* **12** (2007), no. 6 997–1005.

[28] N. B. David, D. Sapède, L. Saint-Etienne, C. Thisse, B. Thisse, C. Dambly-Chaudière, F. M. Rosa, and A. Ghysen, *Molecular basis of cell migration in the fish lateral line: role of the chemokine receptor cxcr4 and of its ligand, sdf1*, *Proceedings of the National Academy of Sciences* **99** (2002), no. 25 16297–16302.

[29] C. W. Reynolds, *Flocks, herds, and schools: A distributed behavioral model*, *Computer Graphics* **21** (1987), no. 4 25–34.

[30] B. Szabo, G. J. Szöllösi, B. Gönci, Z. Jurányi, D. Selmeczi, and T. Vicsek, *Phase transition in the collective migration of tissue cells: experiment and model*, *Physical Review E* **74** (2006), no. 6 061908.

[31] L. Conradt and T. J. Roper, *Consensus decision making in animals*, *Trends in ecology & evolution* **20** (2005), no. 8 449–456.

[32] A. J. W. Ward, J. E. Herbert-Read, D. J. T. Sumpter, and J. Krause, *Fast and accurate decisions through collective vigilance in fish shoals*, *Proceedings of the National Academy of Sciences* **108** (2011), no. 6 2312–2315.

[33] A. Bricard, J. B. Caussin, N. Desreumaux, O. Dauchot, and D. Bartolo, *Emergence of macroscopic directed motion in populations of motile colloids*, *Nature* **503** (2013), no. 7474 95.

[34] F. Ginelli, F. Peruani, M. Bär, and H. Chaté, *Large-scale collective properties of self-propelled rods*, *Physical review letters* **104** (2010), no. 18 184502.

[35] J. Galanis, D. Harries, D. L. Sackett, W. Losert, and R. Nossal, *Spontaneous patterning of confined granular rods*, *Physical review letters* **96** (2006), no. 2 028002.

[36] S. R. K. Vedula, A. Ravasio, C. T. Lim, and B. Ladoux, *Collective cell migration: a mechanistic perspective*, *Physiology* **28** (2013), no. 6 370–379.

[37] T. E. Angelini, E. Hannezo, X. Trepat, M. Marquez, J. J. Fredberg, and D. A. Weitz, *Glass-like dynamics of collective cell migration*, *Proceedings of the National Academy of Sciences* **108** (2011), no. 12 4714–4719.

[38] X. Trepat, M. R. Wasserman, T. E. Angelini, E. Millet, D. A. Weitz, J. P. Butler, and J. J. Fredberg, *Physical forces during collective cell migration*, *Nature physics* **5** (2009), no. 6 426.

[39] M. Poujade, E. Grasland-Mongrain, A. Hertzog, J. Jouanneau, P. Chavrier, B. Ladoux, A. Buguin, and P. Silberzan, *Collective migration of an epithelial monolayer in response to a model wound*, *Proceedings of the National Academy of Sciences* **104** (2007), no. 41 15988–15993.

[40] D. Bi, X. Yang, M. C. Marchetti, and M. L. Manning, *Motility-driven glass and jamming transitions in biological tissues*, *Physical Review X* **6** (2016), no. 2 021011.

[41] S. Garciad, E. Hannezo, J. Elgeti, J.-F. Joanny, P. Silberzan, and N. S. Gov, *Physics of active jamming during collective cellular motion in a monolayer*, *Proceedings of the National Academy of Sciences* **112** (2015), no. 50 15314–15319.

[42] N. Sepúlveda, L. Petitjean, O. Cochet, E. Grasland-Mongrain, P. Silberzan, and V. Hakim, *Collective cell motion in an epithelial sheet can be quantitatively described by a stochastic interacting particle model*, *PLoS computational biology* **9** (2013), no. 3 e1002944.

[43] T. Vicsek and A. Zafeiris, *Collective motion*, *Physics Reports* **517** (2012), no. 3-4 71–140.

[44] S. Henkes, Y. Fily, and M. C. Marchetti, *Active jamming: Self-propelled soft particles at high density*, *Physical Review E* **84** (2011), no. 4 040301.

[45] M. C. Marchetti, J.-F. Joanny, S. Ramaswamy, T. B. Liverpool, J. Prost, M. Rao, and R. A. Simha, *Hydrodynamics of soft active matter*, *Reviews of Modern Physics* **85** (2013), no. 3 1143.

[46] A. Machado, G. Ramalho, J. D. Zucker, and A. Drogoul, *Multi-agent patrolling: An empirical analysis of alternative architectures*, in *Multi-Agent-Based Simulation II*, Lect. Notes Comp. Science, pp. 155–170. Springer Verlag, 2003.

[47] H. Santana, G. Ramalho, V. Corruble, and B. Ratitch, *Multi-agent patrolling with reinforcement learning*, in *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 3*, pp. 1122–1129, IEEE Computer Society, 2004.

[48] H. N. Chu, A. Glad, O. Simonin, F. Sempe, A. Drogoul, and F. Charpillet, *Swarm approaches for the patrolling problem, information propagation vs. pheromone evaporation*, in *Tools with Artificial Intelligence, 2007. ICTAI 2007. 19th IEEE International Conference on*, vol. 1, pp. 442–449, IEEE, 2007.

[49] Y. Elmaliach, N. Agmon, and G. A. Kaminka, *Multi-robot area patrol under frequency constraints*, in *Proc ICRA*, (Roma, Italy), pp. 385–390, Apr., 2007.

[50] K. Srivastava, D. M. Stipanovìc, and M. W. Spong, *On a stochastic robotic surveillance problem*, in *Proc CDC*, (Shanghai, China), pp. 8567–8574, Dec., 2009.

[51] T. Sak, J. Wainer, and S. Goldenstein, *Probabilistic multiagent patrolling*, in *Brazilian Symposium on Artificial Intelligence* (G. Zaverucha and A. L. da Costa, eds.), vol. 5249 of *Lect. Notes Comp. Science*, pp. 124–133. Springer Verlag, 2008.

[52] R. Patel, P. Agharkar, and F. Bullo, *Robotic surveillance and Markov chains with minimal weighted Kemeny constant*, *IEEE Trans. Autom. Control* **60** (2015), no. 12 3156–3157.

[53] R. Patel, A. Carron, and F. Bullo, *The hitting time of multiple random walks*, *SIAM Journal on Matrix Analysis and Applications* **37** (2016), no. 3 933–954.

[54] A. B. Asghar and S. L. Smith, *Stochastic patrolling in adversarial settings*, in *Proc ACC*, (Boston, USA), pp. 6435–6440, July, 2016.

[55] P. Agharkar and F. Bullo, *Quickest detection over robotic roadmaps*, *IEEE Trans Robotics* **32** (2016), no. 1 252–259.

[56] G. Cannata and A. Sgorbissa, *A minimalist algorithm for multirobot continuous coverage*, *IEEE Trans Robotics* **27** (2011), no. 2 297–312.

[57] J. G. Kemeny and J. L. Snell, *Finite Markov Chains*. Springer Verlag, 1976.

[58] A. Z. Broder and A. R. Karlin, *Bounds on the cover time*, *Journal of Theoretical Probability* **2** (1989), no. 1 101–120.

[59] J. J. Hunter, *The role of Kemeny's constant in properties of Markov chains*, *Communications in Statistics-Theory and Methods* **43** (2014), no. 7 1309–1321.

[60] M. Levene and G. Loizou, *Kemeny's constant and the random surfer*, *American Mathematical Monthly* **109** (2002), no. 8 741–745.

[61] S. Kirkland, *Fastest expected time to mixing for a Markov chain on a directed graph*, *Linear Algebra and its Applications* **433** (2010), no. 11-12 1988–1996.

[62] R. Elsässer and T. Sauerwald, *Tight bounds for the cover time of multiple random walks*, *Theoretical Computer Science* **412** (2011), no. 24 2623–2641.

[63] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. John Wiley & Sons, 2012.

[64] J. Justesen and T. Hoholdt, *Maxentropic Markov chains*, *IEEE Trans Information Theory* **30** (1984), no. 4 665–667.

[65] Z. Burda, J. Duda, J. M. Luck, and B. Waclaw, *Localization of the maximal entropy random walk*, *Physical Review Letters* **102** (2009) 160602.

[66] Y. Chen, T. Georgiou, M. Pavon, and A. Tannenbaum, *Efficient-robust routing for single commodity network flows*, *IEEE Transactions on Automatic Control* (2017).

[67] L. Ekroot and T. M. Cover, *The entropy of Markov trajectories*, *IEEE Trans Information Theory* **39** (1993) 1418–1421.

[68] N. Demir and B. Açıkmeşe, *Probabilistic density control for swarm of decentralized on-off agents with safety constraints*, in *American Control Conference (ACC), 2015*, pp. 5238–5244, IEEE, 2015.

[69] B. Açikmeşe and S. D. Bayard, *A markov chain approach to probabilistic swarm guidance*, in *American Control Conference (ACC), 2012*, pp. 6300–6307, IEEE, 2012.

[70] V. Deshmukh, K. Elamvazhuthi, S. Biswal, Z. Kakish, and S. Berman, *Mean-field stabilization of markov chain models for robotic swarms: Computational approaches and experimental results*, *IEEE Robotics and Automation Letters* (2018).

[71] A. Prorok, M. A. Hsieh, and V. Kumar, *Fast redistribution of a swarm of heterogeneous robots*, in *Proceedings of the 9th EAI International Conference on Bio-inspired Information and Communications Technologies (formerly BIONETICS)*, pp. 249–255, ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2016.

[72] P. G. Doyle and J. L. Snell, *Random Walks and Electric Networks*. Mathematical Association of America, 1984.

[73] P. Tetali, *Random walks and the effective resistance of networks*, *Journal of Theoretical Probability* **4** (1991), no. 1 101–109.

[74] L. Backstrom and J. Leskovec, *Supervised random walks: predicting and recommending links in social networks*, in *ACM International Conference on Web Search and Data Mining*, (New York, USA), pp. 635–644, 2011.

[75] C. Gkantsidis, M. Mihail, and A. Saberi, *Random walks in peer-to-peer networks*, in *Annual Joint Conference of the IEEE Computer and Communications Societies*, (Hong Kong, China), pp. 120–130, Mar., 2004.

[76] Q. Lv, P. Cao, E. Cohen, K. Li, and S. Shenker, *Search and replication in unstructured peer-to-peer networks*, in *International Conference on Supercomputing*, (New York, USA), pp. 84–95, 2002.

[77] M. S. Squillante, D. Towsley, and S. Barker, *Improving the scalability of search in networks through multiple random walks*, *ACM SIGMETRICS Performance Evaluation Review* **42** (2014), no. 2 73–75.

[78] T. Ohwa, *Exact computation for meeting times and infection times of random walks on graphs*, *Pacific Journal of Mathematics for Industry* **7** (2015), no. 1 1–9.

[79] C. Li, J. Han, G. He, X. Jin, Y. Sun, Y. Yu, and T. Wu, *Fast computation of SimRank for static and dynamic information networks*, in *International Conference on Extending Database Technology*, (New York, USA), pp. 465–476, 2010.

[80] B. Lin, T. Yin, Y. I. Wu, T. Inoue, and A. Levchenko, *Interplay between chemotaxis and contact inhibition of locomotion determines exploratory cell migration*, *Nature communications* **6** (2015) 6619.

[81] E. Scarpa, A. Roycroft, E. Theveneau, E. Terriac, M. Piel, and R. Mayor, *A novel method to study contact inhibition of locomotion using micropatterned substrates*, *Biology open* **2** (2013), no. 9 901–906.

[82] B. Stramer and R. Mayor, *Mechanisms and in vivo functions of contact inhibition of locomotion*, *Nature Reviews Molecular Cell Biology* **18** (2017), no. 1.

[83] R. Mayor and C. Carmona-Fontaine, *Keeping in touch with contact inhibition of locomotion*, *Trends in cell biology* **20** (2010), no. 6 319–328.

[84] M. Abercrombie and J. E. M. Heaysman, *Observations on the social behaviour of cells in tissue culture: I. speed of movement of chick heart fibroblasts in relation to their mutual contacts*, *Experimental cell research* **5** (1953), no. 1 111–131.

[85] P. Bun, J. Liu, H. Turlier, Z. Liu, K. Uriot, J.-F. Joanny, and M. Coppey-Moisan, *Mechanical checkpoint for persistent cell polarization in adhesion-naive fibroblasts*, *Biophysical journal* **107** (2014), no. 2 324–335.

[86] G. F. Weber, M. A. Bjerke, and D. W. DeSimone, *A mechanoresponsive cadherin-keratin complex directs polarized protrusive behavior and collective cell migration*, *Developmental cell* **22** (2012), no. 1 104–115.

[87] J. R. Davis, A. Luchici, F. Mosis, J. Thackery, J. A. Salazar, Y. Mao, G. A. Dunn, T. Betz, M. Miodownik, and B. M. Stramer, *Inter-cellular forces orchestrate contact inhibition of locomotion*, *Cell* **161** (2015), no. 2 361–373.

[88] E. Scarpa, A. Szabó, A. Bibonne, E. Theveneau, M. Parsons, and R. Mayor, *Cadherin switch during emt in neural crest cells leads to contact inhibition of locomotion via repolarization of forces*, *Developmental cell* **34** (2015), no. 4 421–434.

[89] M. L. Woods, C. Carmona-Fontaine, C. P. Barnes, I. D. Couzin, R. Mayor, and K. M. Page, *Directional collective cell migration emerges as a property of cell interactions*, *PLoS One* **9** (2014), no. 9 e104969.

[90] B. A. Camley, Y. Zhang, Y. Zhao, B. Li, E. Ben-Jacob, H. Levine, and W.-J. Rappel, *Polarity mechanisms such as contact inhibition of locomotion regulate persistent rotational motion of mammalian cells on micropatterns*, *Proceedings of the National Academy of Sciences* **111** (2014), no. 41 14770–14775.

[91] B. Smeets, R. Alert, J. Pešek, I. Pagonabarraga, H. Ramon, and R. Vincent, *Emergent structures and dynamics of cell colonies by contact inhibition of locomotion*, *Proceedings of the National Academy of Sciences* **113** (2016), no. 51 14621–14626.

[92] G. Reig, E. Pulgar, and L. M. Concha, *Cell migration: from tissue culture to embryos*, *Development* **141** (2014), no. 10 1999–2013.

[93] S. Etienne-Manneville, *Polarity proteins in migration and invasion*, *Oncogene* **27** (2008), no. 55 6970.

[94] R. J. Petrie, A. D. Doyle, and K. M. Yamada, *Random versus directionally persistent cell migration*, *Nature reviews Molecular cell biology* **10** (2009), no. 8 538.

[95] T. Das, K. Safferling, S. Rausch, N. Grabe, H. Boehm, and J. P. Spatz, *A molecular mechanotransduction pathway regulates collective migration of epithelial cells*, *Nature cell biology* **17** (2015), no. 3 276.

[96] R. Farooqui and G. Fenteany, *Multiple rows of cells behind an epithelial wound edge extend cryptic lamellipodia to collectively drive cell-sheet movement*, *J Cell Sci* **118** (2005), no. 1 51–63.

[97] R. Mayor and S. Etienne-Manneville, *The front and rear of collective cell migration*, *Nature reviews Molecular cell biology* **17** (2016), no. 2 97.

[98] J. R. Davis, C.-Y. Huang, J. Zanet, S. Harrison, E. Rosten, S. Cox, D. Y. Soong, G. A. Dunn, and B. M. Stramer, *Emergence of embryonic pattern through contact inhibition of locomotion*, *Development* **139** (2012), no. 24 4555–4560.

[99] V. Villar-Cerviño, M. Molano-Mazón, T. Catchpole, M. Valdeolmillos, M. Henkemeyer, L. M. Martínez, V. Borrell, and O. Marín, *Contact repulsion controls the dispersion and final distribution of cajal-retzius cells*, *Neuron* **77** (2013), no. 3 457–471.

[100] A. Szabó, M. Melchionda, G. Nastasi, M. L. Woods, S. Campo, R. Perris, and R. Mayor, *In vivo confinement promotes collective migration of neural crest cells*, *J Cell Biol* (2016) jcb–201602083.

[101] S. J. Streichan, C. R. Hoerner, T. Schneidt, D. Holzer, and L. Hufnagel, *Spatial constraints control cell proliferation in tissues*, *Proceedings of the National Academy of Sciences* **111** (2014), no. 15 5586–5591.

[102] J. Ranft, M. Basan, J. Elgeti, J.-F. Joanny, J. Prost, and F. Jülicher, *Fluidization of tissues by cell division and apoptosis*, *Proceedings of the National Academy of Sciences* **107** (2010), no. 49 20863–20868.

[103] D. J. Higham, *An algorithmic introduction to numerical simulation of stochastic differential equations*, *SIAM review* **43** (2001), no. 3 525–546.

[104] I. Pushkarsky, Y. Liu, W. Weaver, T.-W. Su, O. Mudanyali, A. Ozcan, and D. D. Carlo, *Automated single-cell motility analysis on a chip using lensfree microscopy*, *Scientific reports* **4** (2014) 4717.

[105] M. H. Gail and C. W. Boone, *The locomotion of mouse fibroblasts in tissue culture*, *Biophysical journal* **10** (1970), no. 10 980–993.

[106] B. M. Stramer, G. A. Dunn, J. R. Davis, and R. Mayor, *Rediscovering contact inhibition in the embryo*, *Journal of microscopy* **251** (2013), no. 3 206–211.

[107] R. Vidal, O. Shakernia, H. J. Kim, D. H. Shim, and S. Sastry, *Probabilistic pursuit-evasion games: theory, implementation, and experimental evaluation*, *IEEE Trans Robotics & Automation* **18** (2002), no. 5 662–669.

[108] C. Cooper, A. Frieze, and T. Radzik, *Multiple random walks in random regular graphs*, *SIAM Journal on Discrete Mathematics* **23** (2009), no. 4 1738–1761.

[109] A. Bonato, P. Golovach, G. Hahn, and J. Kratochvíl, *The capture time of a graph*, *Discrete Mathematics* **309** (2009), no. 18 5588–5595.

[110] S. Alpern, R. Fokkink, R. Lindelauf, and G. Olsder, *The "princess and monster" game on an interval*, *SIAM Journal on Control and Optimization* **47** (2008), no. 3 1178–1190.

[111] Y. Elmaliach, A. Shiloni, and G. A. Kaminka, *A realistic model of frequency-based multi-robot polyline patrolling*, in *International Conference on Autonomous Agents*, (Estoril, Portugal), pp. 63–70, May, 2008.

[112] F. Pasqualetti, A. Franchi, and F. Bullo, *On cooperative patrolling: Optimal trajectories, complexity analysis and approximation algorithms*, *IEEE Trans Robotics* **28** (2012), no. 3 592–606.

[113] T. H. Blackwell and J. S. Kaufman, *Response time effectiveness: Comparison of response time and survival in an urban emergency medical services system*, *Academic Emergency Medicine* **9** (2002), no. 4 288–295.

[114] M. Papageorgiou, C. Diakaki, V. Dinopoulou, A. Kotsialos, and Y. Wang, *Review of road traffic control strategies*, *Proc. IEEE* **91** (2003), no. 12 2043–2067.

[115] S. Susca, S. Martínez, and F. Bullo, *Monitoring environmental boundaries with a robotic sensor network*, *IEEE Trans Control Systems Technology* **16** (2008), no. 2 288–296.

[116] S. R. Etesami and T. Başar, *Convergence time for unbiased quantized consensus over static and dynamic networks*, *IEEE Trans. Autom. Control* **61** (2016), no. 2 443–455.

[117] A. Israeli and M. Jalfon, *Token management schemes and random walks yield self-stabilizing mutual exclusion*, in *ACM Symposium on Principles of Distributed Computing*, (New York, USA), pp. 119–131, 1990.

[118] G. Jeh and J. Widom, *SimRank: a measure of structural-context similarity*, in *ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD)*, (New York, USA), pp. 538–543, 2002.

[119] P. Tetali and P. Winkler, *On a random walk problem arising in self-stabilizing token management*, in *Symposium on Principles of Distributed Computing*, pp. 273–280, 1991.

[120] D. Coppersmith, P. Tetali, and P. Winkler, *Collisions among random walks on a graph*, *SIAM Journal on Discrete Mathematics* **6** (1993), no. 3 363–374.

[121] N. H. Bshouty, L. Higham, and J. Warpechowska-Gruca, *Meeting times of random walks on graphs*, *Information Processing Letters* **69** (1999), no. 5 259–265.

[122] D. J. Aldous, *Meeting times for independent Markov chains*, *Stochastic Processes and their Applications* **38** (1991), no. 2 185–193.

[123] D. Aldous and J. A. Fill, *Reversible Markov Chains and Random Walks on Graphs*, 2002. Unfinished monograph, recompiled 2014, available at `http://www.stat.berkeley.edu/~aldous/RWG/book.html`.

[124] C. D. Meyer, *Matrix Analysis and Applied Linear Algebra*. SIAM, 2001.

[125] J. R. Norris, *Markov Chains*. Cambridge University Press, 1997.

[126] R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis*. Cambridge University Press, 1994.

[127] J. Leskovec, D. Chakrabarti, J. Kleinberg, C. Faloutsos, and Z. Ghahramani, *Kronecker graphs: An approach to modeling networks*, *Journal of Machine Learning Research* **11** (Mar., 2010) 985–1042.

[128] P. M. Weichsel, *The Kronecker product of graphs*, *Proceedings of the American Mathematical Society* **13** (1962), no. 1 47–52.

[129] F. Harary and C. A. Trauth Jr., *Connectedness of products of two directed graphs*, *SIAM J Applied Mathematics* **14** (1966), no. 2 250–254.

[130] C. Cooper, R. Elsasser, H. Ono, and T. Radzik, *Coalescing random walks and voting on connected graphs*, *SIAM Journal on Discrete Mathematics* **27** (2013), no. 4 1748–1758.

[131] L. Lovász, *Random walks on graphs: A survey*, in *Combinatorics: Paul Erdős is Eighty* (T. S. D. Miklós, V. T. Sós, ed.), vol. 2, pp. 353–398. János Bolyai Mathematical Society, 1993.

[132] P. Burgisser, M. Clausen, and A. Shokrollahi, *Algebraic Complexity Theory*. Springer Verlag, 1997.

[133] J. R. Gilbert, C. Moler, and R. Schreiber, *Sparse matrices in MATLAB: Design and implementation*, *SIAM Journal on Matrix Analysis and Applications* **13** (1992), no. 1 333–356.

[134] D. M. Foster and J. A. Jacquez, *Multiple zeros for eigenvalues and the multiplicity of traps of a linear compartmental system*, *Mathematical Biosciences* **26** (1975), no. 1 89–97.

[135] R.-H. Li, J. X. Yu, and J. Liu, *Link prediction: The power of maximal entropy random walk*, in *ACM Int. Conf. on Information and Knowledge Management*, (Glasgow, UK), pp. 1147–1156, Oct., 2011.

[136] J. K. Ochab and Z. Burda, *Maximal entropy random walk in community detection*, *The European Physical Journal Special Topics* **216** (2013), no. 1 73–81.

[137] J.-G. Yu, J. Zhao, J. Tian, and Y. Tan, *Maximal entropy random walk for region-based visual saliency*, *IEEE Transactions on Cybernetics* **44** (2014), no. 9 1661–1672.

[138] L. Wang, J. Zhao, X. Hu, and J. Lu, *Weakly supervised object localization via maximal entropy random walk*, in *IEEE Int. Conf. on Image Processing*, (Paris, France), pp. 1614–1617, Oct., 2014.

[139] P. Korus and J. Huang, *Improved tampering localization in digital image forensics based on maximal entropy random walk*, *IEEE Signal Processing Letters* **23** (2016), no. 1 169–173.

[140] X. Peng and Z. Zhang, *Maximal entropy random walk improves efficiency of trapping in dendrimers*, *The Journal of Chemical Physics* **140** (2014), no. 23 234104.

[141] Y. Fan and H. Liu, *Mitigating congestion in complex transportation networks via maximum entropy*, 2017.

[142] S. Bandyopadhyay, S. J. Chung, and F. Y. Hadaegh, *Probabilistic and distributed control of a large-scale swarm of autonomous agents*, *IEEE Transactions on Robotics* **33** (2017), no. 5 1103–1123.

[143] R. Abraham, J. E. Marsden, and T. S. Ratiu, *Manifolds, Tensor Analysis, and Applications*, vol. 75 of *Applied Mathematical Sciences*. Springer Verlag, 2 ed., 1988.

[144] S. G. Krantz and H. R. Parks, *The Implicit Function Theorem*. Birkhäuser, 2013.

[145] M. Grant and S. Boyd, *CVX: Matlab software for disciplined convex programming, version 2.1*, Oct., 2016.

[146] A. F. Karr, *Markov chains and processes with a prescribed invariant measure*, *Stochastic Processes and their Applications* **7** (1978), no. 3 277–290.

[147] W. Rudin, *Principles of Mathematical Analysis*. International Series in Pure and Applied Mathematics. McGraw-Hill, 3 ed., 1976.

[148] Y. Nesterov and A. Nemirovskii, *Interior-Point Polynomial Algorithms in Convex Programming*. SIAM, 1994.

[149] L. N. Trefethen and D. Bau III, *Numerical Linear Algebra*. SIAM, 1997.

[150] L. Vandenberghe and S. Boyd, *Semidefinite programming*, *SIAM Review* **38** (Mar., 1996) 49–95.

[151] A. Ben-Tal and A. Nemirovski, *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications*. SIAM, 2001.

[152] K. Jung, D. Shah, and J. Shin, *Distributed averaging via lifted Markov chains*, *IEEE Trans Information Theory* **56** (2010), no. 1 634–647.

[153] M. H. A. Davis and R. B. Vinter, *Stochastic Modelling and Control*. Springer Verlag, 1985.

[154] S. Guiasu and A. Shenitzer, *The principle of maximum entropy*, *The Mathematical Intelligencer* **7** (1985), no. 1 42–48.

[155] A. Hmamed and E. Tissir, *Further results on the stability of discrete-time matrix polynomials*, *International Journal of Systems Science* **29** (1998), no. 8 819–821.

[156] D. P. Bertsekas, *Nonlinear Programming*. Athena Scientific, 3 ed., 2016.

[157] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.

[158] D. Cai, W. Dai, M. Prasad, J. Luo, N. S. Gov, and D. J. Montell, *Modeling and analysis of collective cell migration in an in vivo three-dimensional environment*, *Proceedings of the National Academy of Sciences* **113** (2016), no. 15 E2134–E2141.

[159] M. A. Jordan, R. J. Toso, D. Thrower, and L. Wilson, *Mechanism of mitotic block and inhibition of cell proliferation by taxol at low concentrations*, *Proceedings of the National Academy of Sciences* **90** (1993), no. 20 9552–9556.

[160] I. C. Mitchell, T. S. Brown, L. S. Terada, J. F. Amatruda, and F. E. Nwariaku, *Effect of vascular cadherin knockdown on zebrafish vasculature during development*, *PLOS ONE* **5** (01, 2010) 1–10.

[161] J. J. Hunter, *Generalized inverses and their application to applied probability problems*, *Linear Algebra and its Applications* **45** (1982) 157–198.