

Markov Processes Beyond Equilibrium and Optimal Control

Theory, Applications, and Examples.



Vom Fachbereich Mathematik und Informatik
der Freien Universität Berlin
zur Erlangung des akademischen Grades
eines Doktors der Naturwissenschaften
genehmigte Dissertation

vorgelegt von

Ralf Banisch

Juli 2015, Freie Universität Berlin

Datum der Disputation: 16. Oktober 2015

Ralf Banisch, *Markov Processes Beyond Equilibrium and Optimal Control*, ©Juli 2015.

Betreuer

Prof. Carsten Hartmann
Freie Universität Berlin
Fachbereich Mathematik und Informatik
Arnimallee 6
14195 Berlin

Gutachter

Prof. Oliver Junge
Zentrum Mathematik - M3
Technische Universität München
85747 Garching bei München

Change your thoughts
and you change the world.
— Norman Vincent Peale

To my family and friends.

Acknowledgements

I would like to thank my supervisor Prof. Carsten Hartmann and my co-supervisor Prof. Christof Schütte for their continuous guidance, support and trust, without which this thesis wouldn't exist. Special thanks to Peter Koltai and Frank Hellmann for many helpful comments that improved the manuscript and to Nataša Conrad for guidance and support in the realm of networks when it was new territory to me. Thanks to Prof. Eric Vanden-Eijnden for helpful discussions on Transition Path Theory. Thanks to Antonia Mey and Nataša Conrad for proofreading the manuscript.

I'd like to thank everybody for bearing with me when writing the thesis took up most of my mind. This applies especially to Antonia Mey, whose patience and support has been invaluable. I also want to thank my friends that made life in Berlin so enjoyable. Thanks to my office mates Anna Djurdjevac, Vikram Sunkara and Patrick Gleiß for having me, and special thanks to Katarzyna Ziółkowska, Marjan Mirahmadi, Nada Cvetkovic, Iliusi Vega, Victor Mireles, Iurii Koznan and many others for awesome office parties. I also wish to thank the Biocomputing Group for their warm welcome three years ago. Finally, all of this would not have been possible without the continuous trust and support of my parents.

This research was founded by the SfB 1114 'Scaling Cascades in Complex Systems' and the Berlin Mathematical School.

Berlin, 01 July 2015

R. B.

Abstract

This thesis is concerned with the long-term dynamics of irreversible Markov processes in discrete and continuous state spaces. In the first part, we study how the long-term dynamics of a reversible Markov process changes if an external force that destroys detailed balance is added. We derive an intuitive and general comparison result in terms of commuting times which indicates that under certain constraints the external driving force will always accelerate the long-term dynamics.

We argue that non-trivial cycles in the probability flow are the key feature of irreversible processes and explain two ways of obtaining cycle decompositions in detail. We study how cycles can be used to construct reversible surrogates of irreversible processes that represent the long-term dynamics more faithfully than simple symmetrization, and apply this to the problem of module detection in directed networks. As a second application, we consider the problem of computing transition pathways between metastable states. This is done by considering the current of reactive trajectories which is computed by Transition Path Theory. We show that this current has cycles if the dynamics is irreversible, and compare two possible Hodge-Helmholtz like splittings of the current into simpler parts. One method is based on a projection, the second is based on cycle decompositions. We show that the second method allows for a computation of the statistics of transition pathways.

In the second part, we study optimal control problems that arise if the external force can be adjusted by a controller who wants to minimize a certain objective function. We focus on linear quadratic (LQ) control problems and show that they are dual to sampling problems which appear e.g. in Molecular Dynamics. A numerical method to approximately solve LQ control problems is derived. The method uses a logarithmic transformation together with a Galerkin projection onto a suitable space of basis functions. The result is a discretization of the entire control problem that replaces the continuous dynamics by a discrete Markov jump process, and preserves the most important structural properties. If the dynamics is metastable, then we propose to utilize the metastable structure by choosing a committor basis, guided by MSM theory. We derive error bounds for this choice which complement standard error bounds from the theory of finite elements. The method is flexible and can also incorporate other choices, e.g. piecewise polynomial or radial basis functions.

Throughout the thesis, we complement theoretical results with careful numerical experiments.

Key words: Irreversible Markov Process, Metastability, Optimal Control, Cycle Decompositions, Transition Path Theory

Zusammenfassung

Diese Arbeit behandelt die Langzeit-Dynamik irreversibler Markovprozesse in diskreten und kontinuierlichen Zustandsräumen. Im ersten Teil wird untersucht wie sich die Langzeit-Dynamik eines reversiblen Prozesses unter dem Einfluss einer externen Kraft ändert, welche die Reversibilität zerstört. Wir geben ein intuitives und allgemeines Resultat mithilfe von commuting times an, welches zeigt dass die externe Kraft unter gewissen Zwangsbedingungen stets für eine Beschleunigung der Langzeit-Dynamik sorgt.

Wir argumentieren, dass Zyklen im Wahrscheinlichkeitsfluss Schlüsselmerkmale irreversibler Prozesse sind und beschreiben zwei Methoden zur Zerlegung des Flusses in Zyklen im Detail. Mit diesen Methoden werden reversible Ersatzprozesse konstruiert, welche die Langzeit-Dynamik irreversibler Prozesse besser abbilden als einfaches Symmetrisieren. Dies wird auf das Problem der Moduldetektion in gerichteten Netzwerken angewendet. Als zweite Anwendung wird das Problem der Berechnung von Übergangspfaden zwischen metastabilen Mengen betrachtet. Dabei betrachten wir den Fluss reaktiver Trajektorien, der durch Transition Path Theory gegeben ist. Wir zeigen, dass dieser Fluss bei irreversibler Dynamik ebenfalls Zyklen enthält, und vergleichen zwei Methoden, um den Fluss in einfachere Bestandteile zu zerlegen. Die erste Methode basiert auf einer Projektion, die zweite auf Zyklenzerlegungen. Wir zeigen, dass die zweite Methode geeignet ist, um die Statistik der Übergangspfade zu bestimmen.

Im zweiten Teil werden Optimalsteuerungsprobleme behandelt, welche entstehen, wenn die externe Kraft von einem externen Agenten gesteuert werden kann. Wir beschränken uns auf linear-quadratische (LQ) Kontrollprobleme und zeigen, dass diese dual zu bestimmten Samplingproblemen z.B. in der Moleküldynamik sind. Eine numerische Methode zur Lösung von LQ Kontrollproblemen wird hergeleitet. Diese Methode benutzt eine logarithmische Transformation zusammen mit einer Galerkin Projektion auf einen geeigneten Unterraum. Das Ergebnis ist eine Diskretisierung des gesamten Kontrollproblems, welches die kontinuierliche Dynamik mit einem diskreten Markovsprungprozess ersetzt und die wesentlichen strukturellen Eigenschaften erhält. Wenn die Dynamik metastabil ist, schlagen wir die Benutzung einer Committorbasis im Sinne der MSM Theorie vor und geben Fehlerschranken an. Diese Fehlerschranken ergänzen Standard-Fehlerschranken aus der Theorie der finiten Elemente. Die Methode ist flexibel und erlaubt die Benutzung anderer Basisfunktionen, z.B. stückweise polynome oder radiale Basisfunktionen.

In der gesamten Arbeit ergänzen wir theoretische Resultate mit sorgfältigen numerischen Experimenten.

Contents

Acknowledgements	i
Abstract (English/Deutsch)	iii
List of figures	ix
List of tables	xi
List of symbols	xiii
Introduction	1
1 Background	9
1.1 Markov Processes	9
1.1.1 Time-reversibility	11
1.1.2 Transfer Operators	12
1.1.3 Markov Diffusions	14
1.1.4 Markov Chains and Markov Jump Processes	17
1.2 Metastability	20
1.2.1 Metastable Sets	20
1.2.2 Metastable Partitions	21
1.2.3 Detecting Metastability	23
1.3 Transition Path Theory	24
1.3.1 Transition Path Theory for Markov Jump Processes	26
1.4 Optimal Control Theory	27
1.4.1 Dynamic Programming	30
1.4.2 Controlled Markov Diffusions	31
1.4.3 Controlled Markov Jump Processes	33
2 Irreversible Markov Chains and Cycles	35
2.1 Cycle Decompositions	35
2.1.1 Coates Graph and Cycle Space	35
2.1.2 Decomposing Flows in the Betti Basis	38
2.1.3 The Circulation Distribution	40
2.1.4 Cycle Decompositions and Reversibility	42

Contents

2.1.5	Cycle Decomposition and Entropy Production	43
2.2	Reversible Surrogates	45
2.2.1	Reversible Surrogates of Unit Cycle Flows	46
2.2.2	Reversible Surrogates of General Markov Chains	49
2.2.3	Reversible Processes are Always Slower than Irreversible Processes	50
2.3	Application I: Modules in Directed Networks	53
2.3.1	Computational Aspects	56
2.3.2	Numerical Examples	59
2.4	Application II: Transition Pathways in Irreversible Markov Processes	65
2.4.1	Hodge-Helmholtz Decomposition	66
2.4.2	Transition Pathways and Cycles	76
3	Optimal Control Theory in Molecular Dynamics	89
3.1	Optimal Control as a Means to Accelerate MD Sampling	90
3.1.1	Duality Between Control and Path Sampling for Jump Processes	95
3.2	MSM Discretizations of LQ Type Control Problems	98
3.2.1	Galerkin Projection Point of View	98
3.2.2	Galerkin Projection of the Dirichlet Problem	99
3.2.3	Interpretation in Terms of a Markov Decision Problem	105
3.2.4	Numerical Results	107
3.3	Multiscale Problems: Averaging	115
4	Conclusion and Outlook	125
A	Appendix A	131
A.1	Proof of theorem 1.1	131
A.2	Proof of theorem 1.7	131
A.3	Proof of theorem 1.16	132
A.4	Expressions for the entropy production rate	133
A.5	Proof of Lemma 2.5	135
A.6	Proof of Theorem 2.6	135
A.7	Proof of Lemma 2.7	136
A.8	The Current of Reactive Trajectories J_{AB}	137
B	Appendix B: Module Detection Methods	139
B.1	MSM Clustering	139
B.2	Alternative Module Detection Methods	140
C	Appendix C	143
C.1	Proof of Theorem 3.4	143
C.2	Best-approximation error bound	145
	Bibliography	155

List of Figures

1	Flowchart	4
2	Introduction, module detection	5
3	Introduction, hodge decomposition	6
4	Introduction, MSM discretization	7
1.1	Partition of double well potential	23
1.2	Reactive trajectories	25
2.1	Oriented graph	37
2.2	Graph with spanning tree	38
2.3	Biased random walk on a cycle	45
2.4	Motifs	47
2.5	Definition of commuting time	52
2.6	Overpartitioning	54
2.7	Barbell graph with spanning tree	59
2.8	Barbell graph metastability	60
2.9	Two or three modules	61
2.10	Two or three modules metastability	62
2.11	Langevin system, potential energy	63
2.12	Langevin system, modules	64
2.13	Earthquake timeseries clustering	65
2.14	Barbell graph hodge decomposition	73
2.15	Transition tube with vortices	73
2.16	Blockmodel, probability distribution of reactive trajectories	74
2.17	Blockmodel, hodge decomposition	75
2.18	Langevin Hodge decomposition	77
2.19	Lumping	78
2.20	Tutorial example.	82
2.21	Transition tube with vortices, productive flow	83
2.22	Barbell graph effective current	84
2.23	Blockmodel, direct pathways	85
2.24	Langevin system, hodge decomposition using cycles.	86
2.25	Langevin system, unproductive cycles.	87

List of Figures

3.1	Discretization diagram	99
3.2	Domain geometry	100
3.3	basis functions	104
3.4	Triple well potential, convergence	108
3.5	Triple well potential, results	109
3.6	Rugged triple well potential, convergence	110
3.7	Rugged triple well potential, results	111
3.8	Alanine Dipeptide	112
3.9	Alanine Dipeptide, implied time scales	114
3.10	Alanine Dipeptide, control results	116
3.11	Blockmodel with averaging	122
3.12	Blockmodel with averaging, results	123
4.1	Approximate policy iteration	130



List of Tables

2.1	The derived chain	41
2.2	Table of simple motifs	48

List of symbols

$(X_t)_t$	a stochastic process in discrete or continuous time
ε_0	best approximation error of the Galerkin method, see section 3.2.2
ε	the $L^2(\mu)$ -error of the Galerkin method, see section 3.2.2
\mathbf{T}	time index set, usually $\mathbf{T} = \mathbb{R}$ or $\mathbf{T} = \mathbb{N}$
\mathcal{B}	Borel- σ -algebra on \mathbb{X}
\mathcal{F}	the σ -algebra of events measurable by $(X_t)_t$.
\mathcal{F}_t	the σ -algebra of events measurable by $\{X_s\}_{0 \leq s \leq t}$; this is the natural Filtration of \mathcal{F} induced by $(X_t)_t$.
$\operatorname{div} f$	divergence of f
Δ	the Laplace operator
$\mathbf{E}_x[\cdot]$	expectation value conditioned on $X_0 = x$
Γ	Betti basis, $\Gamma = \{\gamma_1, \dots, \gamma_b\}$
γ_α	Betti cycle, i.e. basis vector of \mathcal{C}
\mathbb{T}	transition region, $\mathbb{T} = \mathbb{X} \setminus (A \cup B)$
\mathbf{R}	the set of reactive times in TPT, see section 1.3
\mathcal{C}	vector space of cycles in a graph G
\mathcal{C}^*	vector space of cocycles in a graph G
\mathcal{C}_∞	support of ω_∞ , see Lemma 2.4
\mathcal{D}	domain of definition of L
\mathcal{L}	adjoint of L in $L^2(\mathbb{X})$
\mathcal{U}	the space of admissible control strategies $(u_t)_t$

List of symbols

$\text{cap}(A, B)$	capacity of A and B , see (2.26)
μ	invariant measure
\mathbb{N}	the natural numbers
$\nabla^2 u$	the hessian matrix of u
ω_∞	circulation distribution, see Lemma 2.4
∂D	the boundary of the set D
\mathbb{R}	the real numbers
$\sigma(T)$	spectrum of the operator T
τ	a stopping time
\mathbb{U}	the control space
\mathbb{X}	state space
\mathbb{Z}	the integers
A^c	the complement $\mathbb{X} \setminus A$ of A
$C(U)$	the space of continuous functions from U to \mathbb{R}
$C^k(U)$	the functions in $C(U)$ with continuous derivatives up to order k
$C^{k,l}(U \times V)$	the functions $f(u, v) : U \times V \rightarrow \mathbb{R}$ which are C^k w.r.t. $u \in U$ and C^l w.r.t. $v \in V$
$C_0(U)$	the functions in $C(U)$ with compact support
$C_0^k(U)$	the functions in $C^k(U)$ with compact support
D_μ	diagonal matrix with the vector μ on the diagonal
e_P	entropy production rate, see (1.10)
e_{AB}	equilibrium measure, see (2.25)
f_{AB}	Probability current of reactive trajectories, see (1.51)
f_{AB}^+	effective probability current of reactive trajectories, see (1.52)
$G(P)$	Coates graph of P , see definition 2.1
$H(\mu, \nu)$	the relative entropy between μ and ν , see (1.9)
$J_c(x), J_c(x, y)$	passage function of c , see (2.4)
k_{AB}	TPT reaction rate, see section 1.3

L	generator of a Markov process
L^*	adjoint of L in $L^2(\mathbb{X}, \mu)$
$L^p(\mathbb{X}, \mu)$	weighted Banach space, see (1.11)
$N(\mu, \Phi)$	normal distribution with mean μ and covariance matrix Φ
P_∞^s	loop transition matrix, see (2.21)
P_b^s	loop transition matrix based on Betti cycles, see (2.18) and (2.22)
q^+	forward committor, see (1.47)
q^-	backward committor, see (1.47)
t	a fixed time
$T(A, B)$	commuting time between A and B , see (2.27)
T_t	transfer operator, see (1.12)
W_t	Brownian Motion, see section 1.1.3
X_t^-	the time-reversed process $X_t^-(\omega) = X_{-t}(\omega)$.
a.s.	almost surely, i.e. with probability one.
MD	Molecular Dynamics
MJP	Markov jump process
MSM	Markov State Model
ODE	ordinary differential equation
PDE	partial differential equation
SDE	stochastic differential equation
TPT	Transition Path Theory

Introduction

Many real world processes in physics, chemistry, geological sciences, material sciences, social sciences etc. can be modelled by **stochastic Markov processes**. The randomness in these models is often the result of microscopic degrees of freedom that are not explicitly represented, e.g. the various chemical reactions that eventually lead to the formation of a seemingly random decision in an individual, or a bath of water molecules that interact with a larger molecule or colloidal particle. Often the models of interest operate on a vast range of time- and/or length scales, and this proposes a serious challenge: Direct numerical simulations, for example, have to resolve the smallest scales, but the processes of interest typically take place on the largest scales and thus become **rare events**.

Stochastic processes are best understood in equilibrium. We will explain below in more detail how the term 'equilibrium' is used in this thesis, but in order to be in equilibrium, the process must necessarily be (a) autonomous and (b) reversible. In other words, the dynamics that govern the evolution of the process must (a) not change with time and (b) be invariant under time reversal (in a statistical sense, see below). The questions considered in this thesis broadly fit into one of the following two categories:

- How do the long-term dynamics of a reversible Markov process change when non-reversibility is added?
- Can we utilise non-reversibility to accelerate the long-term dynamics?

Over the last years, the long-term dynamics of Markov processes in equilibrium has been well understood. The key here is the self-adjointness of the associated transfer operator, and the resulting spectral properties. We know that the slowest processes are described by the dominant eigenvalues and eigenvectors of the transfer operator [HS06], and that they describe **metastable** transitions between almost invariant sets. With Markov State Models (MSMs) [SFHD99, DHFS00, SNS10, SNL⁺11], we also have an efficient tool for the computation of these objects. By choosing a certain state space discretization, the part of the transfer operator that governs the slowest processes is well approximated by a matrix. This matrix can be sampled by many short trajectories and still allows for the computation of the long-term dynamics that would otherwise require extremely long trajectories. If we move away from equilibrium, however, all of this machinery breaks down.

Introduction

In fact, the non-equilibrium case is vastly more complicated, and a systematic understanding is still missing. Even the term 'non-equilibrium' has different uses in the literature. To appreciate these differences, it is necessary to go from a trajectory-based description of the dynamics to a density-based description. So let $\mu_t(x)$ be the probability density, assuming it exists, to find the system at point x at time t . The typical situation is that $\mu_t \rightarrow \mu_\infty$ in some sense as $t \rightarrow \infty$, and μ_∞ is invariant under the dynamics. We then call the process **ergodic** and μ_∞ a **steady state**. Now take two regions A and B in state space and consider the probability $p(t, A, B)$ to observe the system in region A at time 0 and then in region B at time t . The steady state μ_∞ is an equilibrium state if and only if for any two regions A, B and any $t > 0$, the flow $f(t, A, B) = p(t, A, B) - p(t, B, A)$ is equal to zero when the system is in μ_∞ at time 0. The property $p(t, A, B) = p(t, B, A)$ is called **reversibility** or **detailed balance**. The term 'non-equilibrium' can therefore mean any of the following [CMZ11]:

- (1) **Non-autonomous dynamics:** The dynamics is time-dependent, and a steady state typically does not exist.
- (2) **Transient regime:** A steady state μ_∞ exists and $\mu_t \rightarrow \mu_\infty$, but for the times t considered μ_t is still different from μ_∞ .
- (3) **Non-reversibility:** The dynamics is in a steady state μ_∞ , but some flows $f(t, A, B)$ are non-zero, i.e. detailed balance is violated.

Situation (3) is sometimes called **non-equilibrium steady state (NESS)**, and arises if the stochastic process of our model fails to be reversible. This is the situation we will consider in chapter 2.

At first glance, there seems to be a paradox: The fundamental laws of physics are all invariant under time reversal, so why should our model break this symmetry? The answer is that the scope of any model is limited: We cannot model every part of the system of interest according to fundamental laws. Irreversibility can then enter in many ways, i.e. through coarse graining or effective forces that are the result of an external environment which is not explicitly modelled. We will take the mathematical viewpoint and simply assume that reversibility fails and explore the consequences. One immediate consequence is the following: The flow $f(t, A, B)$ must be **divergence-free**, meaning that for any region A , the total flow into A must equal the total flow out of A . If this were not the case, the probability of being in A would either increase or decrease with time, in contradiction to the system being in the steady state μ_∞ . For discrete state spaces, this means that $f(t, A, B)$ obeys Kirchhoff's first law and must be decomposable into elementary cycles. This observation was first made by Schnakenberg [Sch76]. Many other authors have since then picked up the idea that these elementary cycles encode important information about non-reversible processes [JQQ04, Kal06, Pol15]. In fact, in the light of a widely used analogy between reversible Markov chains in equilibrium and electrostatics, the appearance of cycles in $f(t, A, B)$ when reversibility no longer holds reminds one of magnetostatics, which is hallmarked by cycles in the magnetic field. Cycles will be the primary tool for studying non-reversible Markov chains in chapter 2.

Why should one be interested in studying irreversible Markov processes? The answer is twofold. First, many real world processes are simply best described by a non-reversible model, and it is important to understand the properties of these models. Second, even if the model is reversible, if we are interested in the long term dynamics then simply running the model itself might be a bad idea. If the model is very metastable, then the long-term dynamics is hard to sample from equilibrium simulations, and adding external forces might help to accelerate the sampling. Indeed, a number of results [RS14, HHMS05, SSG10] suggest that a reversible process is always 'slower' than a large class of non-reversible processes that are obtained by adding an external force to the reversible process. Here, 'slower' means e.g. that the second-largest eigenvalue of the transfer operator is closer to one. In this sense, simulations of reversible processes are the worst choice if one is interested in obtaining information about a rare event. We discuss this aspect in chapter 2.

If adding an external force to a reversible process can accelerate the sampling of rare events, it is natural to consider a scenario where an external controller can adjust the external force to minimize a certain objective function. By solving the associated optimization problem, we could then find the best force to add in order to achieve the best improvement in sampling. This is the framework of **Optimal Control Theory**. Indeed, there is a duality principle based on [DPMR96] that establishes a duality between a certain class of optimal control problems and sampling problems in Molecular Dynamics (MD) [HS12]. Solving the optimal control problem exactly provides a zero-variance estimator for the quantity one wanted to sample [HS12]. To exploit this duality, we need to think about methods for solving optimal control problems. This is the topic of chapter 3. One of the first steps in this direction was the seminal work of Bellman [Bel56] who showed that the solution to an optimal control problem can be found via a recursion principle that he called **dynamic programming**. The recursion principle leads to the Bellman equation, which for continuous state spaces takes the form of a nonlinear partial differential equation (PDE). This leads to an immediate problem: For large state spaces, all of the established solution methods for PDEs fail because of the curse of dimensionality. The same is true for other standard solution methods of optimal control problems, such as value iteration and policy iteration, see e.g. [Kus01]. In order to make progress, we need a method for solving optimal control problems approximately that uses the multiscale structure of the problem. We will use MSM theory as a guide to construct such a method in chapter 3.

Organisation of the thesis. Figure 1 shows the structure of the thesis and the dependencies of the different sections. Chapter 1 reviews background material on Markov processes and the different concepts of metastability for such processes, before moving on to the somewhat more specialised topics of Transition Path Theory, which is concerned with the computation of rare event statistics, and Optimal Control. Chapter 2 is concerned with NESS for Markov chains and develops the theory of cycle decompositions for these in section 2.1. This theory is used to construct reversible surrogates of the NESS Markov chain in section 2.2, and we apply these approximations to the problem of module detection in networks in section 2.3. As a second application, we marry Transition Path Theory and cycle decompositions to compute

transition pathways for NESS Markov chains in section 2.4. Chapter 3 begins by reviewing the duality between control problems and sampling. In section 3.2, the MSM-based discretization of the control problems considered in section 3.1 is developed. We close by discussing the averaging method in section 3.3, which helps to reduce the problem size for very large systems.

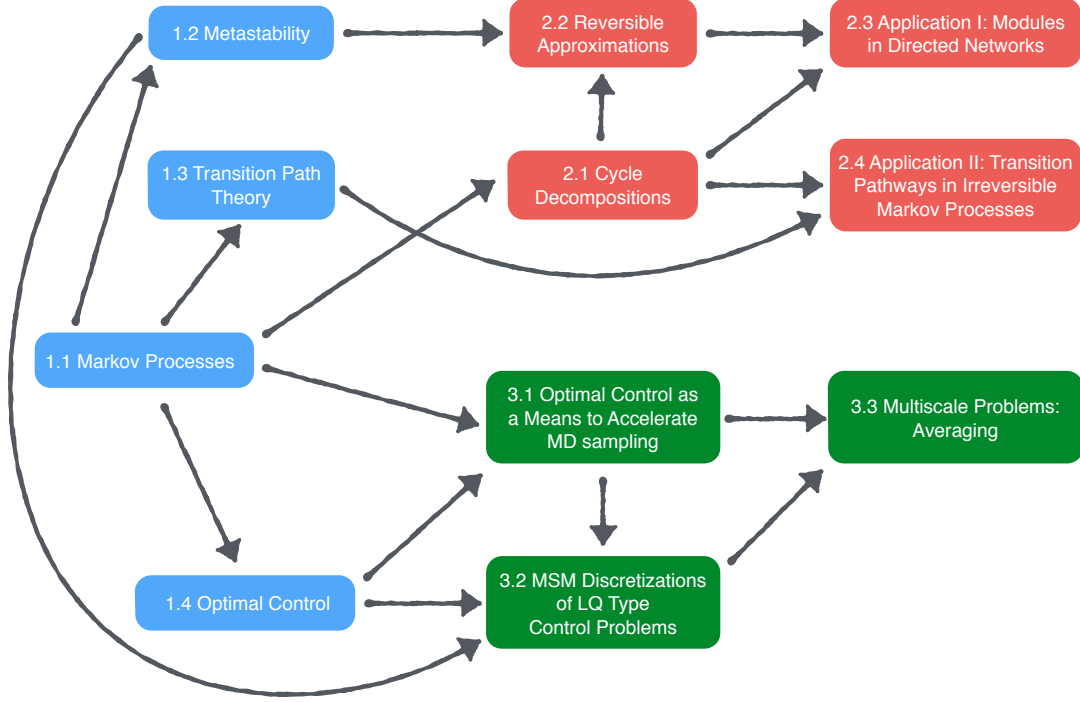


Figure 1 – Organisation of the thesis.

Main results. We highlight the main results. For more details and further references, see the respective sections. Parts of these results are published in [BH14, SBHS13, BC14, HBS⁺14, CBS15, BCS15].

- (1) In section 2.2.3, we prove a general comparison result about **commuting times** for a non-reversible Markov process $(X_t)_t$ with generator L and the reversible process $(X_t^s)_t$ with generator $L^s = \frac{1}{2}(L + L^-)$ obtained by symmetrisation. Here, L^- is the generator of the time-reversed process, see section 1.1.1. Let A and B be two regions in state space and define the commuting time $T(A, B) = \mathbf{E}_A[\tau_B] + \mathbf{E}_B[\tau_A]$ to be the average time to go from A to B and then back to A . The result is that for any disjoint A and B ,

$$T^s(A, B) \geq T(A, B)$$

where $T(A, B)$ and $T^s(A, B)$ are the commuting times for $(X_t)_t$ and $(X_t^s)_t$ respectively. In words: $(X_t^s)_t$ is always slower than $(X_t)_t$. Results in a similar spirit were found in [RS14, HHMS05, SSG10] for Markov diffusions and Markov chains. All of these results

are more technical and use e.g. Large Deviation methods to provide comparisons between the second largest eigenvalue of L and L^s or the asymptotic variance of MCMC estimators based on $(X_t)_t$ and $(X_t^s)_t$. In contrast, the result presented here is in terms of dynamical quantities and thus allows for a clear and direct interpretation. It also tells us what to expect in terms of metastability.

- (2) In section 2.3, we use the theory of cycle decompositions for non-reversible Markov chains that we adapt from [Kal06, JQQ04] to detect **modules** in directed networks. For the purpose of this thesis, modules are metastable sets of a Markov chain that lives on the network. Figure 2 shows an illustration of the structures one can expect to find. One common definition of a module used in the network community is as a set of nodes with high link density. By this definition, the blue, red and green structures in Figure 2 are modules. However, there is a topological difference between the green and the blue and red structures: The shortest cycle connecting any two nodes A, B in the green structure (the yellow ribbon in Figure 2) is very long. As a consequence, a random walker following edge directions cannot spend much time in the green structure, but it can spend long times in the red and blue structures due to the presence of internal cycles.




Figure 2 – Modules in a directed network.

Module detection is a huge area of research [For10, SLB12, SDYB12, DYB10, EL09, New06, NBW06]. In a nutshell, existing topological methods do not pick up on the difference between the green and the blue/red structures, while existing methods based on random walks [DBCS11, SCB⁺14] are confined to reversible processes and thus to undirected networks. Our method is a random walk based method that constructs a suitable reversible surrogate of the non-reversible random walk via cycles. It can differentiate between the green and the blue/red structures in Figure 2 by construction since their metastability is different.

- (3) In section 2.4, we provide a new way to analyse the probability current of reactive trajectories in Transition Path Theory (TPT) for irreversible processes. If identifying

Introduction

metastable sets is the first step to understand long-term dynamics, then understanding the pathways along which the transitions between these sets happen is the second step. This is the domain of TPT [VE06]. One constructs an ensemble of **reactive trajectories**, which are trajectories that start in A and transition to B before returning to A . One interesting quantity to compute is the probability current f_{AB}^+ which is generated by the reactive trajectories. If the dynamics is non-reversible, f_{AB}^+ contains cycles, and it does make sense to look for a decomposition of f_{AB}^+ into simpler parts that allow for a clearer interpretation of the underlying transition mechanism. This is illustrated in Figure 3: It is clear that a transition from $A \in C_1$ to $B \in C_5$ has to go via C_2 , with C_3 and C_4 representing detours. This is not obvious from f_{AB}^+ , but it is possible to find a decomposition of the form $f_{AB}^+ = \text{grad}\Phi + R$ where $\text{grad}\Phi$ represents the transition and R the detours. In section 2.4, we compare decompositions of this type based on Hodge-Helmholtz decompositions with those based on cycles. This leads to new ways of analysing irreversible Markov processes with TPT.




Figure 3 – Decomposition of the probability current f_{AB}^+ into the gradient of a potential Φ and a rotation R .

- (4) In section 3.2, we develop a numerical method to approximately solve **linear-quadratic control problems**. The method uses a logarithmic transformation to turn the Bellman equation (which is a nonlinear PDE) into a linear PDE, and then discretizes the linear

PDE with a Galerkin projection onto a suitable space of basis functions. After the discretization, the logarithmic transform is inverted. The result is a discretization of the entire control problem that replaces the continuous dynamics by a discrete Markov jump process, and preserves the most important structural properties. If the dynamics is metastable, then we propose to utilize the metastable structure by choosing a committor basis, guided by MSM theory. However, the method is flexible and can also incorporate other choices, e.g. piecewise polynomial or radial basis functions [Wen99]. The example in Figure 4 shows that this method reproduces the effective potential, which is the outcome of the control problem, remarkably well with as few as three basis functions, even for a relatively complicated potential energy landscape with many local minima. We give theoretical error bounds in section 3.2.2.

The most developed method for the discretization of optimal control problems is the Markov chain approximation method (MCA) by Kushner, see [Kus01] and the references therein. It replaces the continuous control problem with a control problem for a Markov chain and can deal with a more general class of control problems. However, the discretization has to use a grid which renders the problem intractable in high dimensions, and only weak convergence results are obtained. Our method includes MCA if we use linear finite elements on a grid as basis functions in the sense that MCA discretizes derivatives by finite differences whereas our method discretizes derivatives by finite elements. In this case, we obtain strong convergence results.




Figure 4 – Left: Potential energy V (blue), effective potential U (green) and MSM approximation (red). Right: Committor basis based on the three main wells.

In terms of linear PDEs, Galerkin-based discretization methods are well known and usually called finite elements in this context [Bra07]. Aside from the structural discretization of the control problem, the novelty of our contribution lies in the identification of the committor basis as an effective basis when metastabilities are present, and the error bounds we compute in section 3.2.2. These complement standard error bounds for finite element methods, which are usually given in terms of piecewise polynomial functions on grids, as well as existing error bounds in the MSM literature [Sar11].

1 Background

1.1 Markov Processes

In this section we give a brief introduction to the mathematical theory of Markov processes. For details see [SS13b].

We denote the **state space** by \mathbb{X} . For us, \mathbb{X} will always be either a countable set or a subset of \mathbb{R}^d and we call \mathbb{X} discrete in the former and continuous in the latter case. Let \mathcal{B} be the σ -algebra of Borel sets on \mathbb{X} . A **stochastic process** on \mathbb{X} is a collection of \mathbb{X} -valued random variables $(X_t)_{t \in \mathbf{T}}$ where the time index set is either $\mathbf{T} = [0, \infty)$ (continuous time) or $\mathbf{T} = \mathbb{N}$ (discrete time). If no specification is necessary, we will often just write $(X_t)_t$. The stochastic process $(X_t)_t$ lives on the probability space $(\Omega, \mathcal{F}, \mathbf{P})$ where $\Omega = \{\omega : \mathbf{T} \rightarrow \mathbb{X}\}$ is the set of \mathbb{X} -valued functions on \mathbf{T} , \mathcal{F} is the σ -algebra generated by the sets $\{X_t^{-1}(B) | t \in \mathbf{T}, B \in \mathcal{B}\}$ and \mathbf{P} is the probability measure generated by the finite-dimensional distributions of $(X_t)_t$, which are the measures μ_{t_1, \dots, t_k} on \mathbb{X}^k defined by

$$\mu_{t_1, \dots, t_k}(B_1 \times B_2 \times \dots \times B_k) = \mathbf{P}[X_{t_1} \in B_1, \dots, X_{t_k} \in B_k], \quad t_i \in \mathbf{T}, \quad B_i \in \mathcal{B}. \quad (1.1)$$

Note that for each fixed $t \in \mathbf{T}$, we have a random variable $\Omega \ni \omega \mapsto X_t(\omega)$. On the other hand, for each fixed $\omega \in \Omega$ we can consider the function $\mathbf{T} \ni t \mapsto X_t(\omega)$ which is sometimes called **sample path, realization or trajectory** of the process.

For every $t \in \mathbf{T}$ we define the **natural filtration** $\mathcal{F}_t \subset \mathcal{F}$ of \mathcal{F} with respect to $(X_t)_t$ to be the σ -algebra generated by the sets $\{X_s^{-1}(B) | s \in \mathbf{T}, s \leq t, B \in \mathcal{B}\}$. That is, \mathcal{F}_t contains all pre-images of \mathcal{B} -measurable subsets of \mathbb{X} under X_s for times $s \leq t$. Intuitively speaking, \mathcal{F}_t contains all events that can be decided given knowledge of $X_s(\omega)$ for $s \leq t$, i.e. the 'past' of X_t . We say that $(X_t)_t$ is a **Markov process** if the so-called Markov property is satisfied:

$$\mathbf{P}[X_t \in B | \mathcal{F}_s] = \mathbf{P}[X_t \in B | X_s] \quad \forall 0 \leq s < t, \forall B \in \mathcal{B}. \quad (1.2)$$

Chapter 1. Background

The Markov property says that the probability of events in the future of $s \in \mathbf{T}$ depends only on the 'present' $X_s(\omega)$ and not on the history of $(X_t)_t$ before time s . It is often important to guarantee that the Markov property holds even if the fixed time t in (1.2) is replaced by a stopping time. A random variable $\tau : \Omega \rightarrow \mathbf{T}$ is called a **stopping time** if

$$\{\tau \leq t\} = \{\omega \in \Omega : \tau(\omega) \leq t\} \in \mathcal{F}_t \quad \forall t \in \mathbf{T}.$$

In other words, τ is a stopping time if we can decide whether or not $\tau \leq t$ has occurred based on the knowledge of the process up to time t only. The process $(X_t)_t$ is said to have the **strong Markov property** if (1.2) still holds with t replaced by a stopping time τ .

A Markov process is called **homogeneous** if $\mathbf{P}[X_{t+h} \in B | X_t] = \mathbf{P}[X_h \in B | X_0]$ for all $t, h \in \mathbf{T}$ and all $B \in \mathcal{B}$ holds. The motion of a homogeneous Markov process is completely described by its **transition function** $p : \mathbf{T} \times \mathbb{X} \times \mathcal{B} \rightarrow [0, 1]$ according to

$$p(t, x, B) = \mathbf{P}[X_t \in B | X_0 = x]. \quad (1.3)$$

We recall the most important properties of transition functions:

1. $x \mapsto p(t, x, B)$ is measurable for fixed $t \in \mathbf{T}$ and fixed $B \in \mathcal{B}$.
2. $B \mapsto p(t, x, B)$ is a probability measure for fixed $t \in \mathbf{T}$ and fixed $x \in \mathbb{X}$.
3. We have the **Chapman-Kolmogorov equation**:

$$p(t+s, x, B) = \int_{\mathbb{X}} p(t, x, dy) p(s, y, B) \quad \forall t, s \in \mathbf{T}, x \in \mathbb{X}, B \in \mathcal{B}. \quad (1.4)$$

This follows from the Markov property (1.2).

Next we discuss invariant measures and ergodicity. A probability measure μ is said to be **invariant** w.r.t. $(X_t)_t$ iff

$$\int_{\mathbb{X}} p(t, x, B) \mu(dx) = \mu(B) \quad \forall t \in \mathbf{T}, B \in \mathcal{B}. \quad (1.5)$$

In this thesis, μ will always denote a probability measure over \mathbb{X} , that is, $\mu(\mathbb{X}) = 1$. For a measurable function $u : \mathbb{X} \rightarrow \mathbb{R}$ we denote the expectation value of u w.r.t. the measure μ by

$$\mathbf{E}_\mu[u] = \int_{\mathbb{X}} u(x) \mu(dx).$$

The space of measurable functions $u : \mathbb{X} \rightarrow \mathbb{R}$ with $\mathbf{E}[|u|] < \infty$ is denoted by $L^1(\mathbb{X}, \mu)$. Then the Markov process $(X_t)_t$ is said to be **ergodic** w.r.t. μ if [Wal82]

$$\mathbf{E}_\mu[u] = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T u(X_t) dt \quad \forall u \in L^1(\mathbb{X}, \mu) \quad (1.6)$$

holds for almost all initial values X_0 . If $\mathbf{T} = \mathbb{N}$, then the integral on the RHS is replaced by a sum. For an ergodic Markov process, averages over μ can be replaced by long-time averages, which is the foundation of Monte Carlo sampling and similar numerical methods. The existence and uniqueness of an invariant measure μ is a necessary, but not a sufficient condition for $(X_t)_t$ to be ergodic, and for many real systems both ergodicity and existence and uniqueness of μ are very difficult to prove. In this thesis we will mostly deal with ergodic Markov processes.

1.1.1 Time-reversibility

We now discuss the very important concept of time-reversibility. Since we need the concept of forward and backward dynamics to do so we need \mathbf{T} to be symmetric around $t = 0$, so let $\mathbf{T} = \mathbb{R}$ or $\mathbf{T} = \mathbb{Z}$ if the time horizon is infinite, or let $\mathbf{T} = [-T, T]$ or $\mathbf{T} = \mathbb{Z} \cup [-T, T]$ for some $T > 0$ if the time horizon is finite. Let the process $(X_t)_t$ be homogeneous and ergodic with unique invariant measure μ , and let $X_t \sim \mu$ for all $t \in \mathbf{T}$. We define the time reversal transformation $R : (\Omega, \mathcal{F}) \rightarrow (\Omega, \mathcal{F})$ by $(R\omega)(t) = \omega(-t)$ for all $t \in \mathbf{T}$. Obviously R is invertible and $R^{-1} = R$. Then the **time-reversed process** $(X_t^-)_t$ is defined by¹ $X_t^-(\omega) = X_{-t}(\omega)$, and we denote the law of $(X_t^-)_t$ by $\mathbf{P}^- = \mathbf{P} \circ R$. We often call $(X_t)_t$ the forward and $(X_t^-)_t$ the backward process. A Markov process is called **time-reversible** or simply **reversible** if $\mathbf{P} = \mathbf{P}^-$, i.e. if the laws of the forward and backward process are the same. In particular, we have

$$\mathbf{P}[X_t \in A, X_0 \in B] = \mathbf{P}^- [X_t \in A, X_0 \in B] = \mathbf{P}[X_t^- \in A, X_0^- \in B] = \mathbf{P}[X_0 \in A, X_t \in B] \quad (1.7)$$

for all $A, B \in \mathcal{B}$ and all $t \in \mathbf{T}$, where we used the definition of \mathbf{P}^- , X_t^- and the homogeneity of X_t . The equation (1.7) can be written in terms of the transition function by using

$$\mathbf{P}[X_t \in A, X_0 \in B] = \int_B \mathbf{P}[X_t \in A | X_0 = x] \mu(dx) = \int_B p(t, x, A) \mu(dx),$$

which holds since $X_0 \sim \mu$. Then (1.7) can be written as

$$\int_B p(t, x, A) \mu(dx) = \int_A p(t, x, B) \mu(dx). \quad (1.8)$$

If the transition function is absolutely continuous w.r.t. the Lebesgue measure² dx on \mathbb{X} , that is, if a density $p(t, x, y)$ exists such that

$$p(t, x, A) = \int_A p(t, x, y) dy \quad \forall A \in \mathcal{B},$$

then (1.8) simply reads $\mu(dx) p(t, x, y) = \mu(dy) p(t, y, x)$ for all $t \in \mathbf{T}$ and μ -a.e. $x, y \in \mathbb{X}$. This equation is often called **detailed balance**. It turns out that detailed balance is both necessary

¹Another definition often used in the literature is $X_t^- = X_{T-t}$ for $\mathbf{T} = [0, T]$, which reduces to our definition for homogeneous processes after a shift $t \mapsto t - \frac{T}{2}$.

²If \mathbb{X} is discrete, then the Lebesgue measure dx is replaced by the counting measure.

Chapter 1. Background

and sufficient for $(X_t)_t$ to be reversible³.

Finally, we introduce an information-theoretic characterization of reversibility. We define the **relative entropy** of two probability measures μ and ν as

$$H(\mu, \nu) = \int_{\mathbb{X}} \log \frac{d\mu}{d\nu}(x) \mu(dx) \quad \text{if } \mu \ll \nu, \quad H(\mu, \nu) = +\infty \quad \text{otherwise.} \quad (1.9)$$

Let $\mathcal{F}_{[0,T]}$ be the σ -algebra generated by $\{X_t^{-1}(B) \mid t \in \mathbf{T}, 0 \leq t \leq T, B \in \mathcal{B}\}$ and denote by $\mathbf{P}_{[0,T]}$ and $\mathbf{P}_{[0,T]}^-$ the laws \mathbf{P} and \mathbf{P}^- restricted to $\mathcal{F}_{[0,T]}$, meaning that e.g. $\mathbf{P}_{[0,T]}$ is the probability measure associated to the paths of finite length $[0, T] \ni t \mapsto X_t(\omega)$. Then the **entropy production rate** e_P is defined as

$$e_P = \lim_{T \rightarrow \infty} \frac{1}{T} H(\mathbf{P}_{[0,T]}, \mathbf{P}_{[0,T]}^-). \quad (1.10)$$

That is, e_P is the rate of growth of relative entropy between \mathbf{P} and \mathbf{P}^- as the time interval $[0, T]$ considered gets larger. Being a relative entropy, e_P is always positive and $e_P = 0$ if and only if $\mathbf{P} = \mathbf{P}^-$ a.s.. In fact since $H(\mu, \nu)$ is a distance measure⁴ for the probability distributions μ and ν , the entropy production rate can be seen as a measure for the distance between X_t and X_t^- and thus as a measure for the degree of irreversibility of $(X_t)_t$. In summary, we have the following equivalent characterizations of reversibility for homogeneous Markov processes:

1. *Symmetry under time reversal*: The laws of forward and backward process are the same, that is $\mathbf{P} = \mathbf{P}^-$,
2. *Symmetry of two-point correlations*: The relation $\mathbf{P}[X_t \in A, X_0 \in B] = \mathbf{P}[X_t \in B, X_0 \in A]$ holds for all $t \in \mathbf{T}$ and all $A, B \in \mathcal{B}$,
3. *Detailed balance*: $\mu(dx)p(t, x, y) = \mu(dy)p(t, y, x)$ holds for all $t \in \mathbf{T}$ and μ -a.e. $x, y \in \mathbb{X}$,
4. *Zero entropy production*: $e_P = 0$.

1.1.2 Transfer Operators

From here on we will need the weighted Banach spaces

$$L^p(\mathbb{X}, \mu) = \left\{ u : \mathbb{X} \rightarrow \mathbb{R} : \int_{\mathbb{X}} |u(x)|^p \mu(dx) < \infty \right\}, \quad 1 \leq p < \infty \quad (1.11)$$

and

$$L^\infty(\mathbb{X}, \mu) = \left\{ u : \mathbb{X} \rightarrow \mathbb{R} : \mu - \text{ess sup}_{x \in \mathbb{X}} |u(x)| < \infty \right\}$$

³From detailed balance, it is not hard to show that the symmetry of the two-point correlations (1.7) also holds for all finite-dimensional distributions (1.1) provided $(X_t)_t$ is homogeneous, which implies $\mathbf{P} = \mathbf{P}^-$.

⁴The relative entropy is not symmetric, $H(\mu, \nu) \neq H(\nu, \mu)$. Thus $H(\mu, \nu)$ is not a metric, but it still provides a useful notion of distance between probability measures.

with the corresponding norms $\|u\|_p$ and $\|u\|_\infty$ respectively. Due to Hölder's inequality $L^q(\mathbb{X}, \mu) \subset L^p(\mathbb{X}, \mu)$ for all $1 \leq p \leq q \leq \infty$.

Unless stated otherwise we will let $\mathbf{T} = [0, \infty)$. The transition function introduced in (1.3) is one way to completely characterize the dynamics of a homogeneous Markov process. Another is via the family of **Transfer Operators**⁵ $T_t : L^1(\mathbb{X}, \mu) \rightarrow L^1(\mathbb{X}, \mu)$, which are defined for $t \geq 0$ as

$$(T_t u)(x) = \mathbf{E}_x[u(X_t)] = \int_{\mathbb{X}} u(y) p(t, x, dy), \quad (1.12)$$

where $\mathbf{E}_x[u(X_t)]$ is the expectation value of $u(X_t)$ conditioned on $X_0 = x$. The operator T_0 is the identity, and from the Chapman-Kolmogorov equation (1.4) we get the semigroup property

$$T_t T_s = T_s T_t = T_{s+t} \quad \forall s, t \in \mathbf{T}, s, t \geq 0.$$

The invariance of the transition functions w.r.t. μ (1.5) guarantees that $\|T_t\|_1 = 1$, thus the family of Transfer operators is a one-parameter family of contraction operators, and as such it can be written as $T_t = \exp(tL)$ for an operator $L : \mathcal{D} \subset L^1(\mathbb{X}, \mu) \rightarrow L^1(\mathbb{X}, \mu)$ which is defined as

$$(Lu)(x) = \lim_{t \downarrow 0} \frac{T_t u(x) - u(x)}{t}, \quad (1.13)$$

and $u \in \mathcal{D}$ if the limit in (1.13) exists. The operator L is called the **generator** of the Markov process $(X_t)_t$. If $\mathbf{T} = \mathbb{N}$, we simply set $L = T_1 - I$. The operators T_t and L describe the evolution of observables with time. This is summarized in the following theorem (for $\mathbf{T} = \mathbb{R}$):

Theorem 1.1 *Let $f : \mathbb{X} \rightarrow \mathbb{R}$ be continuous and bounded and define $u : \mathbf{T} \times \mathbb{X} \rightarrow \mathbb{R}$ as $u(t, x) = \mathbf{E}_x[f(X_t)]$. Suppose that $u(t, \cdot) \in \mathcal{D}$ for each $t \in \mathbf{T}$ and that the map $t \mapsto u(t, x)$ is differentiable. Then u satisfies the **backward Kolmogorov equation***

$$\frac{\partial u}{\partial t} = Lu \quad \text{on } (0, \infty) \times \mathbb{X}, \quad u(0, x) = f(x). \quad (1.14)$$

The theorem follows from (1.13) and properties of the semigroup T_t . We prove it in appendix A. One typical situation where the assumptions of the theorem are satisfied is $u \in C^{1,2}(\mathbf{T} \times \mathbb{X})$ in the context of diffusions, see below and [Øks03]. The counterpart of Theorem 1.1 is concerned with the evolution of probability densities:

Theorem 1.2 *Suppose that the law of $(X_t)_t$ has a density v w.r.t. the Lebesgue measure, that is $\mathbf{P}(X_t \in B) = \int_B v(x, t) dx$. Let $X_0 \sim v_0$. Then under some regularity assumptions, v solves the **forward Kolmogorov equation***

$$\frac{\partial v}{\partial t} = \mathcal{L}v \quad \text{on } (0, \infty) \times \mathbb{X}, \quad v(0, x) = v_0(x) \quad (1.15)$$

⁵A priori, T_t is defined on $L^\infty(\mathbb{X}, \mu)$, but it can be extended to $L^p(\mathbb{X}, \mu)$ for all $p \geq 1$.

Chapter 1. Background

where \mathcal{L} is the formal L^2 -adjoint of L , that is $\int_{\mathbb{X}} v(x)(Lu)(x)dx = \int_{\mathbb{X}} (\mathcal{L}v)(x)u(x)dx$.

Equation (1.15) is also known as the **Fokker-Planck-equation**. One usually needs stronger assumptions on the coefficients of \mathcal{L} in order for (1.15) to hold than one needs for (1.14). For example, in the context of diffusions (1.15) holds if $v \in C^{1,2}(\mathbf{T} \times \mathbb{X})$ and additionally the coefficients of L , $b(t, \cdot)$ and $\sigma(t, \cdot)$, are in $C^1(\mathbb{X})$ and $C^2(\mathbb{X})$ for every $t \in \mathbf{T}$ respectively. Cf. below and [LM94]. It is therefore typically better to study the backward equation instead of the forward equation.

The transfer operator point of view is most powerful if T_t and L are regarded as operators on the Hilbert space $L^2(\mathbb{X}, \mu)$ equipped with the scalar product

$$\langle u, v \rangle_{\mu} := \int_{\mathbb{X}} u(x)v(x)\mu(dx). \quad (1.16)$$

Then the adjoint T_t^* of T_t in $L^2(\mathbb{X}, \mu)$ is equal to T_t^- , the transfer operator of the backward process X_t^- :

$$\langle u, T_t v \rangle_{\mu} = \mathbf{E}_{\mu}[u(X_0)v(X_t)] = \mathbf{E}_{\mu}[u(X_0^-)v(X_t^-)] = \mathbf{E}_{\mu}[u(X_t^-)v(X_0^-)] = \langle T_t^- u, v \rangle_{\mu}$$

by using the definition of T_t , the fact that $X_t = X_t^-$ the homogeneity of X_t^- and the invariance of μ . Similarly, $L^* = L^-$, L^- being the generator of the backward process X_t^- . If $(X_t)_t$ is reversible, then T_t and L are essentially self-adjoint operators on $L^2(\mathbb{X}, \mu)$, and consequently $\sigma(T_t) \subset [-1, 1]$ and $\sigma(L) \subset (-\infty, 0]$. Furthermore, if $(X_t)_t$ is ergodic then the essential spectral radius of T_t is bounded away from one. Spectral properties of T_t have important connections to metastability, see section 1.2.

1.1.3 Markov Diffusions

In what follows $\mathbb{X} = \mathbb{R}^d$ and $\mathbf{T} = \mathbb{R}$. Markov diffusions are generalizations of **Brownian Motion**, the prototypical stochastic process named after the erratic movement of pollen particles observed by the botanist Robert Brown in 1827. A Brownian Motion or **Wiener Process** W_t on \mathbb{R}^d is a homogeneous Markov Process with independent gaussian increments $W_t - W_s \sim N(0, (t-s)I)$ and a.s. continuous sample paths. Here $N(\mu, \Phi)$ denotes the normal distribution with mean μ and covariance matrix Φ .

To describe diffusions we shall use the theory of Stochastic Differential Equations (SDEs) in the Itô sense, for an introduction to SDEs and Itô calculus see [Øks03]. A Markov diffusion $(X_t)_t$ is defined as the solution to the SDE

$$dX_t = b(X_t, t)dt + \sigma(X_t, t)dW_t, \quad X_0 = x_0. \quad (1.17)$$

with W_t being Brownian motion on \mathbb{R}^m , the drift vector field $b: \mathbb{R}^d \times \mathbf{T} \rightarrow \mathbb{R}^d$ and the diffusion coefficient $\sigma: \mathbb{R}^d \times \mathbf{T} \rightarrow \mathbb{R}^{d \times m}$. The matrix-valued function $a(X_t, t) = \frac{1}{2}\sigma(X_t, t)\sigma^T(X_t, t)$ is

called the diffusion matrix. Existence and uniqueness of a solution to (1.17) is guaranteed under mild regularity assumptions on b and σ , see [Øks03]. We will always work with diffusions where existence and uniqueness is guaranteed, and furthermore where $(X_t)_t$ is ergodic and admits a unique invariant measure μ .

In the following we shall assume that b and σ do not depend explicitly on time, i.e. $b(x, t) = b(x)$ and $\sigma(x, t) = \sigma(x)$. Then the generator of the diffusion (1.17) acts on $u \in C_0^2(\mathbb{R}^d)$ as

$$Lu(x) = \sum_i b_i(x) \frac{\partial u}{\partial x_i} + \sum_{i,j} a_{ij}(x) \frac{\partial^2 u}{\partial x_i \partial x_j} = b(x) \cdot \nabla u + a : \nabla^2 u \quad (1.18)$$

where $\nabla^2 u$ is the Hessian matrix of u and $a : \nabla^2 u$ denotes the trace of the matrix product between a and $\nabla^2 u$. If the matrix a is **positive-definite**, that is

$$\sum_{i,j} a_{ij}(x) \xi_i \xi_j > 0 \quad \forall \xi \in \mathbb{R}^d,$$

then L is **elliptic**. If a is positive-semidefinite, i.e. the above inequality is not strict, then L is said to be **semi-elliptic**. If a constant $\theta > 0$ exists such that

$$\sum_{i,j} a_{ij}(x) \xi_i \xi_j > \theta \|\xi\|^2 \quad \forall \xi \in \mathbb{R}^d, \quad (1.19)$$

holds for all $x \in D \subset \mathbb{R}^d$, then L is said to be **uniformly elliptic** on the domain D . If $(X_t)_t$ has invariant measure μ and this measure admits a density w.r.t. the Lebesgue measure on \mathbb{R}^d which we also denote by μ , then the time-reversed diffusion $(X_t^-)_t$ satisfies the SDE [HP86, Met07]

$$dX_t^- = b^-(X_t^-) dt + \sigma(X_t^-) dW_t, \quad b_i^-(x) = -b_i(x) + \frac{2}{\mu(x)} \nabla^j (a_{ij}(x) \mu(x)). \quad (1.20)$$

Since the diffusion coefficient σ is unchanged under time reversal, a Markov diffusion is reversible iff $b = b^-$.

The solutions of stochastic differential equations can be used to describe solutions of certain PDEs. The backward Kolmogorov equation (1.14) is the first example of this interplay, and one might wonder if this can be generalized. This is indeed the case, and the most classical example of this is given by the famous **Feynman-Kac formula** [Øks03]:

Theorem 1.3 (The Feynman-Kac formula.) *Let $g \in C_0^2(\mathbb{R}^d)$ and $f \in C(\mathbb{R}^d)$, and assume that f is bounded from below. Put*

$$u(t, x) = \mathbf{E}_x \left[\exp \left(- \int_0^t f(X_s) ds \right) g(X_t) \right]. \quad (1.21)$$

Chapter 1. Background

Then u satisfies the PDE

$$\frac{\partial u}{\partial t} = Lu - fu \quad \text{on } (0, \infty) \times \mathbb{R}^d, \quad u(0, x) = g(x), \quad x \in \mathbb{R}^d. \quad (1.22)$$

Moreover, if $w \in C^{1,2}(\mathbb{R} \times \mathbb{R}^d)$ is bounded on $K \times \mathbb{R}^d$ for each compact $K \subset \mathbb{R}$ and w solves (1.22), then $w = u$.

The backward Kolmogorov equation and the Feynman-Kac formula both consider parabolic PDEs. Similar results exist for elliptic boundary value problems: Suppose we are interested in the solution h of a PDE on a domain $D \subset \mathbb{R}^d$. Let $(X_t)_t$ be a Markov diffusion with semi-elliptic generator L and initial condition inside D , and let τ_D be the time X_t hits the boundary ∂D . Then τ_D is a stopping time, and we assume that τ_D is a.s. finite for all initial conditions $x \in D$. Then the following theorem holds [Øks03] (see [Kai12] or [FS06, App. D] for a detailed proof):

Theorem 1.4 (A Feynman-Kac formula for boundary value problems.) *Let $f \geq 0$ be a continuous function on \mathbb{R}^d . Let $g \in C(\partial D)$ be bounded. Consider the boundary value problem*

$$\begin{aligned} Lh(x) - f(x)h(x) &= 0 && \text{on } D, \\ \lim_{x \rightarrow y} h(x) &= g(y) && y \in \partial D. \end{aligned} \quad (1.23)$$

Then if a bounded solution h to (1.23) exists, it has the form

$$h(x) = \mathbf{E}_x \left[\exp \left(- \int_0^{\tau_D} f(X_s) ds \right) g(X_{\tau_D}) \right]. \quad (1.24)$$

We now give two examples for SDEs which occur frequently in Molecular Dynamics applications.

1. The *overdamped Langevin equation* is given by

$$dX_t = -\nabla V(X_t) dt + \sqrt{2\varepsilon} dW_t \quad (1.25)$$

where $V \in C^1(\mathbb{R}^d)$ is the potential and $\varepsilon > 0$ the temperature. If V grows faster than linear as $\|x\| \rightarrow \infty$, then the dynamics (1.25) is ergodic with invariant measure

$$\mu(x) = Z^{-1} e^{-\varepsilon^{-1} V(x)}, \quad Z = \int e^{-\varepsilon^{-1} V(x)} dx. \quad (1.26)$$

The quantity Z is called the partition function. The generator of (1.25) acts on $u \in C^2(\mathbb{R}^d)$ as $Lu = \mu^{-1} \nabla \cdot (\mu \nabla u)$. The overdamped Langevin equation is an example of a reversible dynamics.

2. The second example is the famous *Langevin equation*

$$\begin{aligned} dX_t &= m^{-1}P_t dt \\ dP_t &= -(\nabla V(X_t) + \gamma m^{-1}P_t) dt + \sqrt{2\gamma\epsilon} dW_t \end{aligned} \quad (1.27)$$

where $X_t, P_t \in \mathbb{R}^{3d}$ denote the position and momentum of d particles at time t , m is the mass of the particles⁶ and γ the friction coefficient. The overdamped Langevin equation (1.25) arises as the high-friction limit of (1.27). Langevin dynamics is ergodic with invariant measure

$$d\mu(x, p) = e^{-\epsilon^{-1}H(x, p)} dx dp, \quad H(x, p) = V(x) + \frac{1}{2m} p^T p. \quad (1.28)$$

The function $H: \mathbb{R}^{3d} \times \mathbb{R}^{3d} \rightarrow \mathbb{R}$ is called the **Hamiltonian**. Langevin dynamics is non-reversible. The generator L of (1.27) acts on functions $u \in C^2(\mathbb{R}^{3d} \times \mathbb{R}^{3d})$ like

$$Lu = \epsilon\gamma\Delta_p u + m^{-1}p \cdot \nabla_x u - \nabla_x V \cdot \nabla_p u - \gamma m^{-1}p \cdot \nabla_p u.$$

1.1.4 Markov Chains and Markov Jump Processes

If the state space \mathbb{X} is discrete and \mathbf{T} is continuous, we are in the realm of Markov jump processes (MJPs). If \mathbf{T} is also discrete, we are in the realm of Markov chains. In this section, we will discuss the case where $\mathbb{X} = \{1, \dots, d\}$ is a finite set. This has the advantage that the operators in question become finite matrices and thus many technical complications are avoided. The function spaces $L^p(\mathbb{X}, \mu)$ for $1 \leq p \leq \infty$ all coincide with \mathbb{R}^d , and the transition function (1.3) always has a density

$$p(t, x, y) = \mathbf{P}[X_t = y | X_0 = x]. \quad (1.29)$$

The equation (1.5) for the stationary distribution μ becomes

$$\sum_{x \in \mathbb{X}} p(t, x, y) \mu(x) = \mu(y). \quad (1.30)$$

As before, μ gives rise to a scalar product $\langle u, v \rangle_\mu = \sum_{x \in \mathbb{X}} \mu(x) u(x) v(x)$. The action of the transfer operator (1.12) can be written as

$$T_t u(x) = \mathbf{E}[u(X_t) | X_0 = x] = \sum_{y \in \mathbb{X}} u(y) \mathbf{P}[X_t = y | X_0 = x] = \sum_{y \in \mathbb{X}} u(y) p(t, x, y), \quad (1.31)$$

⁶It is possible to consider different masses and friction coefficients for all particles, which we won't do here to keep notation simple.

Chapter 1. Background

hence T_t acts as a matrix on $u \in \mathbb{R}^d$, and the components of T_t are the transition probabilities (1.29) which we shall always assume to be continuous at $t = 0$, that is

$$\lim_{t \downarrow 0} p(t, x, y) = \delta_{xy} \quad \forall x, y \in \mathbb{X}. \quad (1.32)$$

This guarantees that the trajectories of $(X_t)_t$ are right continuous functions with left limits, so-called **càdlàg functions**. Trajectories of $(X_t)_t$ are piecewise constant with distinct jumps at random times. From (1.30) we see that $\mu^T T_t = \mu^T$, i.e. μ is a left eigenvector to the eigenvalue 1 of T_t . The semigroup of transfer operators T_t becomes a semigroup of stochastic matrices, i.e. matrices with row sum one, with generator

$$L = \lim_{t \downarrow 0} \frac{T_t - I}{t}. \quad (1.33)$$

The generator matrix L has row sum zero, and $\mu^T L = 0$, i.e. μ is a left eigenvector of L to the eigenvalue 0. The off-diagonal entries of L are all nonnegative and the diagonal entries are given by $l_{xx} = -\sum_{y \neq x} l_{xy}$. The reversed-time generator $L^- = L^*$ is given by $L^- = D_\mu^{-1} L^T D_\mu$ where $D_\mu = \text{diag}(\mu)$, this follows from $\langle \cdot, L(\cdot) \rangle_\mu = \langle L^-(\cdot), \cdot \rangle_\mu$. In components,

$$l_{xy}^- = \frac{\mu(y)}{\mu(x)} l_{yx} \quad \forall x, y \in \mathbb{X}. \quad (1.34)$$

Both L and L^- have the same stationary distribution μ .

As we saw in section 1.1.2, the evolution of conditional expectations $u(t, x) = \mathbf{E}_x[f(X_t)]$ is governed by the backward Kolmogorov equation (1.14). For Markov jump processes equation (1.14) has the same form, but becomes a system of d coupled linear ODEs instead of a PDE. The same is true for the Fokker-Planck equation (1.15). The operator appearing in (1.15) is simply $\mathcal{L} = L^T$. The Fokker-Planck-equation itself becomes

$$\frac{\partial v(t, x)}{\partial t} = (L^T v)(t, x) = \sum_{y \neq x} l_{yx} v(t, y) + l_{xx} v(t, x) \quad (1.35)$$

and can be interpreted as a balance equation for the probability $v(t, x) = \mathbf{P}(X_t = x)$. Therefore, l_{yx} for $y \neq x$ is the rate at which $(X_t)_t$ jumps from y to x , and $l_{xx} = -\sum_{y \neq x} l_{xy}$ is the **escape rate** of state x .

To further illustrate this, let $0 = t_0 < t_1 < t_2 < \dots < t_n < \dots$ denote the sequence of jump times of $(X_t)_t$. Then $X_n := X(t_n)$ defines a Markov chain which is called the **embedded chain** of the jump process $(X_t)_t$. The Markov jump process is then completely characterized by the chain $(X_n)_n$ of states visited and the sequence of **waiting times** $\tau_i = t_{i+1} - t_i$. The embedded chain $(X_n)_n$ has a transition matrix Q with entries

$$q_{xy} = \mathbf{P}[X_{n+1} = y | X_n = x] = \begin{cases} \frac{l_{xy}}{-l_{xx}} & x \neq y \\ 0 & x = y. \end{cases} \quad (1.36)$$

The waiting times τ_i are exponentially distributed with parameter $w_x = \mathbf{E}[\tau_i | X_i = x] = -1/l_{xx}$.

The ergodicity properties of Markov chains and Markov jump processes on finite state spaces can be completely characterized by some simple structural properties. For a Markov chain with transition matrix Q we say that $x, y \in \mathbb{X}$ **communicate** if there are $m, m' \in \mathbb{N}$ such that $(Q^m)_{xy} > 0$ and $(Q^{m'})_{yx} > 0$, i.e. y can eventually be reached from x and vice versa. Communication is an equivalence relation and divides \mathbb{X} into communication classes. If there is only one class, then the Markov chain is called **irreducible**. A MJP is called irreducible if its embedded chain is irreducible. An irreducible Markov chain on finite state space \mathbb{X} is necessarily also **positive recurrent**, meaning that for every $x \in X$, $\mathbf{E}_x[\tau_x] < \infty$ where $\tau_x = \inf\{n > 0 : X_n = x\}$ is the first return time to x . The state $x \in \mathbb{X}$ of a Markov chain is called **periodic** with period p if returns to x must occur in multiples of p time steps. Formally, the period p of $x \in \mathbb{X}$ is defined as

$$p = \gcd\{n : (Q^n)_{xx} > 0\}.$$

If $p = 1$, then x is said to be **aperiodic**. If all $x \in \mathbb{X}$ are aperiodic then the chain $(X_n)_n$ is said to be aperiodic. For Markov chains, we have the very important **ergodic theorem** [Bré99]:

Theorem 1.5 (Ergodic theorem for Markov chains.)

- (a) Let $(X_n)_n$ be an irreducible positive recurrent Markov chain on \mathbb{X} with stationary distribution μ and transition matrix P , and let $f : \mathbb{X} \rightarrow \mathbb{R}$ be such that $\mathbf{E}_\mu[f] < \infty$. Then for any initial distribution ν_0 , \mathbf{P}_{ν_0} -a.s.

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{n=1}^T f(X_n) = \mathbf{E}_\mu[f].$$

- (b) If additionally $(X_n)_n$ is aperiodic, then for any initial distribution ν_0 ,

$$\lim_{n \rightarrow \infty} \|\nu_0^T P^n - \mu^T\| = 0$$

with $\|\mu\| := \sum_{x \in \mathbb{X}} |\mu(x)|$.

Theorem 1.5 is the reason one calls a Markov chain **ergodic** if it is irreducible, positive recurrent and aperiodic. This implies both (a) and the stronger property (b). MJPs cannot be periodic, thus a MJP is called ergodic if it is irreducible and positive recurrent. The transition matrix Q of an ergodic Markov chain has a simple eigenvalue 1 with the stationary distribution μ being the unique left eigenvector, $\mu^T Q = \mu^T$. The generator L of an ergodic MJP has a simple eigenvalue 0 with the stationary distribution μ being the unique left eigenvector, $\mu^T L = 0$.

We finish this section by noting that, like the forward and backward Kolmogorov equations, the Theorems 1.3 and 1.4 discussed in section 1.1.3 also have analogues for MJPs. We first present the analogue of Theorem 1.3:

Theorem 1.6 (The discrete Feynman-Kac formula.) *Let $f, g \in \mathbb{R}^d$. Put*

$$u(t, x) = \mathbf{E}_x \left[\exp \left(- \int_0^t f(X_s) ds \right) g(X_t) \right]. \quad (1.37)$$

Then u satisfies the ODE

$$\frac{\partial u}{\partial t} = Lu - fu \quad \text{on } (0, \infty) \times \mathbb{X}, \quad u(0, x) = g(x), \quad x \in \mathbb{X}. \quad (1.38)$$

Moreover, if $w \in C^{1,0}(\mathbb{R} \times \mathbb{X})$ solves (1.38), then $w = u$.

As for the Kolmogorov equations, in the discrete Feynman-Kac formula the parabolic PDE (1.22) is replaced by a system of d coupled linear ODEs for which existence and uniqueness is simple. Next we present the analogue of Theorem 1.4:

Theorem 1.7 (A discrete Feynman-Kac formula for boundary value problems.) *Let $D \subset \mathbb{X}$ and $f, g : \mathbb{X} \rightarrow \mathbb{R}$ with $f \geq 0$. Consider the boundary value problem*

$$\begin{aligned} Lh(x) - f(x)h(x) &= 0 & x \in D \\ h(x) &= g(x), & x \in \mathbb{X} \setminus D. \end{aligned} \quad (1.39)$$

Then the unique solution to (1.23) is

$$h(x) = \mathbf{E}_x \left[\exp \left(- \int_0^{\tau_D} f(X_s) ds \right) g(X_{\tau_D}) \right]. \quad (1.40)$$

Now the equation (1.39) is a linear system of equations instead of a PDE. We prove Theorem 1.7 in Appendix A. It also holds for Markov chains if we set $L := Q - I$.

1.2 Metastability

The concept of metastability is of central importance for this thesis, and it is a concept that is particularly elusive. Although the intuition behind it is usually clear, the mathematical definition of metastability is not. Several different aspects of metastability exist, and accordingly one finds several different definitions in the literature [BEGK02, BGK05, HS06, DSS12]. Here we review some of them. A good overview can be found in [SS13b].

1.2.1 Metastable Sets

Perhaps the simplest aspect of metastability is that of an almost invariant set. A set $A \subset \mathbb{X}$ is called **invariant** for the timescale T if $\mathbf{P}[X_t \in A | X_0 \in A] = 1$ for all $0 \leq t \leq T$, that is the process $(X_t)_t$ is certain to remain in A over the timespan T . The set A is called invariant if it is

invariant for $T = \infty$. Invariance is a very strict concept. For example, for an ergodic process an invariant set A must necessarily have $\mu(A) \in \{0, 1\}$, therefore invariant sets do not allow for a decomposition of \mathbb{X} into multiple sets of finite measure. It is therefore natural to relax the condition of invariance to define metastable sets:

Definition 1.8 (metastable sets) *The set $A \subset \mathbb{X}$ is called **metastable** for the timescale T if $p_t(A, A) := \mathbf{P}[X_t \in A | X_0 \in A] \approx 1$ for all $0 \leq t \leq T$.*

This defines metastable sets as sets which are *almost invariant* with respect to the timescale T chosen. Obviously $p_T(A, A) \approx 1$ is a vague statement. However, the larger $p_T(A, A)$ is, the more metastable A is for the timescale T , and we are often interested in the most metastable sets. Since metastability is defined relative to a timescale T , an obvious question is how to choose T . Definition 1.8 becomes more restrictive the larger T is, so we will have fewer and fewer metastable sets as T is chosen larger and larger. This hierarchical structure is natural and reflects the multiscale structure of many dynamical processes. Usually we will select an appropriate reference timescale t_f and we are then interested in finding sets which are metastable w.r.t. $T \gg t_f$. There are two alternative ways to characterize the metastability of a set:

1. **Large hitting times.** Denote by $\tau_D = \inf\{t \geq 0 : X_t \in D\}$ the hitting time of a set $D \subset \mathbb{X}$. Suppose A is a metastable set in the sense of definition 1.8 for some $T \gg t_f$, and let $D = A^c$ be the complement of A . Since metastable sets are hard to leave, we expect the mean time it takes to hit D from any $x \in A$ to be large, i.e. $\mathbf{E}_x[\tau_D] \gg t_f$ for all $x \in A$. Of course $\mathbf{E}_x[\tau_D]$ depends on x . In some ideal cases, $\mathbf{E}_x[\tau_D]$ will be almost constant for all $x \in A$, but often this is only true for x sufficiently far away from ∂A or for very low temperatures.
2. **Small exit rates.** With the same setting as above, suppose that we are in the situation that $\mathbf{E}_x[\tau_D]$ is almost constant for all $x \in A$. Then heuristically, the exit rate of A can be defined as $\Gamma(A) = 1/\mathbf{E}_x[\tau_D]$ and we expect $\Gamma(A) \ll t_f^{-1}$. More rigorously, it can be shown [BBI12] that the distribution of the random variable τ_D conditioned on $x \in A$ is asymptotically independent of x and exponential, hence $\mathbf{P}[\tau_D = t | X_0 = x] \sim \exp(-\Gamma(A)t)$, which defines the rate $\Gamma(A)$, and indeed one has $\Gamma(A) \ll t_f^{-1}$.

1.2.2 Metastable Partitions

The notion of metastable sets can be used to partition \mathbb{X} . This is tied to the construction of Markov State Models (MSMs) [Sar11]. A **full partition** of \mathbb{X} is a collection of sets A_1, \dots, A_m with $A_i \cap A_j = \emptyset$ and $\bigcup_{i=1}^m A_i = \mathbb{X}$. We can use definition 1.8 to evaluate the metastability of a full partition:

Chapter 1. Background

Definition 1.9 The metastability index $D_T(A_1, \dots, A_m)$ of the full partition A_1, \dots, A_m of \mathbb{X} for the timescale T is defined as

$$D_T(A_1, \dots, A_m) = \sum_{k=1}^m p_T(A_k, A_k). \quad (1.41)$$

The partition A_1, \dots, A_m is called metastable for the timescale T if $D_T(A_1, \dots, A_m) \approx m$.

The metastability index can be used as an optimization criterion to find the best metastable partition $(A_1^*, \dots, A_m^*) = \operatorname{argmax}_{A_1, \dots, A_m} D_T(A_1, \dots, A_m)$ among all full partitions into m sets for a given timescale T .

Full partitions suffer from a problem which is illustrated in Figure 1.1. Suppose the dynamics $(X_t)_t$ is a Markov diffusion given by (1.27) with the potential V shown in Figure 1.1. For such a dynamics, deep wells are obvious metastable sets. But the assumption that the whole energy landscape can be decomposed into deep wells is usually not valid, and in fact the exact location of the boundary between A_1 and A_2 in Figure 1.1 is largely arbitrary. In this situation it is more meaningful to search for an **incomplete partition** of \mathbb{X} , that is a collection C_1, \dots, C_m of disjoint sets such that the complement of $C = \bigcup_{i=1}^m C_i$ is nonempty. The metastability of an incomplete partition can be characterized via the mean hitting times in the following way [BEGK02]:

Definition 1.10 (Metastability of incomplete partitions) For an incomplete partition C_1, \dots, C_m of \mathbb{X} , let $C = \bigcup_{i=1}^m C_i$. We define the metastability of the set C_i to be

$$\rho_i = \frac{\sup_{x \in \mathbb{X} \setminus C} \mathbf{E}_x[\tau_C]}{\inf_{x \in C_i} \mathbf{E}_x[\tau_{C \setminus C_i}]}. \quad (1.42)$$

The partition C_1, \dots, C_m is called ρ -metastable if $\rho := \max_i \rho_i \ll 1$.

In this definition, the time $t_R = \sup_{x \in \mathbb{X} \setminus C} \mathbf{E}_x[\tau_C]$ is the maximum return time to one of the metastable sets from outside and $T_i = \inf_{x \in C_i} \mathbf{E}_x[\tau_{C \setminus C_i}]$ is the minimum time it takes to escape C_i and jump into any other metastable set C_j . The metastability ρ_i is the degree of separation between the two timescales t_R and T_i . A large degree of metastability in the sense of Definition 1.10, i.e. a small ρ_i , requires t_R to be small, so the metastable sets have to be easy to reach, and T_i to be large, so transitions between different sets have to be hard. The key difference between full and incomplete metastable partitions is the following: For the incomplete partition shown in Figure 1.1, a transition from C_1 to C_2 is necessarily a transition from the left to the right well. But in terms of the full partition shown in Figure 1.1, a transition from A_1 to A_2 is any crossing of the boundary between A_1 and A_2 , and such boundary crossings may happen many times without any true transition from the left to the right well. In other words: The sets C_1, \dots, C_m define the locations of the 'cores' of metastable regions. For this reason they are also called core sets in this context [DSS12, Sar11].




Figure 1.1 – A metastable partition A_1, A_2 of the double well potential V . Because of the extended transition region, the exact position of the boundary between A_1 and A_2 is largely arbitrary.

1.2.3 Detecting Metastability

With the Definitions 1.8, 1.9 and 1.10 of metastable sets and metastable partitions, some obvious questions remain: How does one detect a priori if a system is metastable or not? How does one find metastable sets and partitions? How does one determine the number of metastable sets in a partition?

Definition 1.10 indicates one possible answer to these questions via the hitting times: Select some test set $A \subset \mathbb{X}$, and compute the hitting times $\mathbf{E}_x[\tau_A]$ for all $x \in \mathbb{X} \setminus A$, this can be done by solving a linear system of the form (1.23). Then for the metastable decomposition C_1, \dots, C_m in Definition 1.10, we expect that $\mathbf{E}_x[\tau_A]$ is large and almost constant for all $x \in C_i$. Moreover, if $x_i \in C_i$ and $x_j \in C_j$ with $i \neq j$, we expect that generically $\mathbf{E}_{x_i}[\tau_A]$ is very different from $\mathbf{E}_{x_j}[\tau_A]$. Generically here means that this property holds for many, but not necessarily all test sets $A \subset \mathbb{X}$. This way, hitting time distributions for different test sets may be used to detect the metastable sets C_1, \dots, C_m , see [SS14].

Metastability is best understood if the dynamics is reversible. For reversible Markov diffusions and MJPs it is known [BEGK02, BGK05] that if a metastable decomposition in the sense of Definition 1.10 exists, then the spectrum of the generator L has m dominant eigenvalues $0 = \lambda_1 > \lambda_2 > \dots > \lambda_m$, and the inverses of the dominant eigenvalues are asymptotically close to the jumping timescales $T_i = \inf_{x \in C_i} \mathbf{E}_x[\tau_{C \setminus C_i}]$ from Definition 1.10 in a certain ordering. Moreover, the location of the metastable sets is encoded in the eigenvectors of the m dominant eigenvalues. In other words, for a reversible Markov diffusion or MJP, metastability can be detected by examining the dominant spectrum of L .

A similar relation between metastable full partitions in the sense of Definition 1.9 and the spectrum of the transfer operator T_t is known [HS06, Hui01]:

Theorem 1.11 *Let A_1, \dots, A_m be a full partition of \mathbb{X} . Let T_t be reversible and $1 = \lambda_1 > \lambda_2 > \dots > \lambda_m$ be the m dominant eigenvalues of T_t with corresponding eigenfunctions u_j . Furthermore, let $a > 0$ be such that $\sigma(T_t) \subset [-a, a] \cup \{\lambda_m\} \cup \dots \cup \{\lambda_1\}$. Then the metastability index $D_t(A_1, \dots, A_m)$*

can be bounded by

$$\kappa_1 \lambda_1 + \dots \kappa_m \lambda_m + c \leq D_t(A_1, \dots, A_m) \leq \lambda_1 + \dots \lambda_m \quad (1.43)$$

where $\kappa_i = \|Qu_i\|_2^2$ and $c = a(1 - \kappa_1) \dots (1 - \kappa_m)$. Q is the orthogonal projection onto the space

$$\mathcal{D} = \left\{ u \in L^2(\mathbb{X}, \mu) : u = \sum_{i=1}^m c_i \chi_{A_i}, c_1, \dots, c_m \in \mathbb{R} \right\}.$$

Theorem 1.11 establishes a relation between the metastability index $D_t(A_1, \dots, A_m)$ of the full partition A_1, \dots, A_m for the timescale t and the dominant spectrum of the transfer operator T_t for the same timescale. Theorem 1.11 can be read in the following way: If T_t has m dominant eigenvalues close to 1, then we can find a partition into m sets which has a metastability index close to m , namely the partition that minimizes the projection errors $\kappa_i = \|Qu_i\|_2^2$ of the eigenfunctions u_i onto the space \mathcal{D} of functions which are constant on the sets A_i . In other words, the sets A_i should be selected so that the eigenfunctions u_j are almost constant on the A_i .

1.3 Transition Path Theory

As we saw in the last section, for a metastable Markov process several regions in state space \mathbb{X} might exist in which the process stays for a long time, such that transitions between these regions become rare events. Often it is precisely those transitions one is interested in. This is where Transition Path Theory (TPT) steps in as an exact theory to describe and bias on the statistical ensemble governing those transitions. An introduction to TPT can be found in [VE06, Met07]. Here we only give a short overview of the objects one is dealing with in TPT.

Fix two subsets $A, B \subset \mathbb{X}$ with smooth boundaries, nonzero measure and $A \cap B = \emptyset$. In TPT, A is the **reactant state** and B the **product state**, and we are interested in reactions from A to B . We usually think of A and B as well-defined metastable sets which are 'far apart' in the sense that the dynamical process $(X_t)_t$ we are interested in takes a long time to make transitions between A and B , cf. Definition 1.10. However, TPT is still exact if this is not the case. Suppose that $\mathbf{T} = \mathbb{R}$ and that we have an infinitely long equilibrium trajectory $\mathbf{T} \ni t \mapsto X_t(\omega)$, which we denote by $X(t)$ for convenience. Now we define the stopping times

$$\tau_{AB}^+(t) = \inf\{t' \geq t : X(t') \in A \cup B\}, \quad \tau_{AB}^-(t) = \sup\{t' \leq t : X(t') \in A \cup B\}. \quad (1.44)$$

That is $\tau_{AB}^+(t)$ is the first time after time t when $X(t)$ hits $A \cup B$ and $\tau_{AB}^-(t)$ is the last time before time t when $X(t)$ was in $A \cup B$. For the realization $X(t)$ we define the set of reactive times \mathbf{R} as

$$t \in \mathbf{R} \Leftrightarrow X(t) \notin A \cup B, \quad X(\tau_{AB}^+(t)) \in B, \quad X(\tau_{AB}^-(t)) \in A. \quad (1.45)$$




Figure 1.2 – An illustration of the ensemble of reactive trajectories in TPT. For the realization $X(t)$ shown, the part in red is a reactive piece.

In other words $t \in \mathbf{R}$ if (1) $X(t)$ is neither in A nor in B , (2) the trajectory $X(t)$ followed into the past starting at time t hits A before B , and (3) the trajectory $X(t)$ followed into the future starting at time t hits B before A , see Figure 1.2 for an illustration. \mathbf{R} is a disjoint union of intervals, each interval representing one reactive piece. It is important to notice that $\tau_{AB}^+(t)$, $\tau_{AB}^-(t)$ and \mathbf{R} depend on the realization $X(t)$ and are therefore random variables themselves. In TPT one is interested in computing the following objects:

1. The **probability density of reactive trajectories** $\mu_R(x)$. Assuming that it exists, $\mu_R(x)$ is the probability density of finding X_t at x conditioned on X_t being reactive (that is, t being in \mathbf{R}). Formally,

$$\int_C \mu_R(x) dx = \mathbf{P}[X_t \in C | t \in \mathbf{R}] \quad \forall C \in \mathcal{B}. \quad (1.46)$$

2. The **probability current of reactive trajectories** J_{AB} . This probability current is a vector field on \mathbb{X} with the property that if J_{AB} is integrated over a surface ∂S which is the boundary of a region $S \subset \mathbb{X} \setminus (A \cup B)$, we get the probability flux of reactive trajectories across that surface. To make this definition precise we need to know what vector fields on \mathbb{X} are. We reserve that for later when we specialize to \mathbb{X} being discrete.
3. The **reaction rate** k_{AB} . As stated above, the set \mathbf{R} of reactive times is a disjoint union of intervals where each interval represents one reaction. Let N_T be the number of reactions or disjoint intervals in $\mathbf{R} \cap [-T, T]$. Then the reaction rate is defined as the limit

$$k_{AB} = \lim_{T \rightarrow \infty} \frac{N_T}{2T}.$$

1.3.1 Transition Path Theory for Markov Jump Processes

We now specialize to the setting of Section 1.1.4, that is $\mathbb{X} = \{1, \dots, d\}$ is finite and $(X_t)_t$ is a Markov jump process on \mathbb{X} with generator L . TPT for MJPs has been studied extensively in [Met07]. A prominent role is played by the committor functions:

Definition 1.12 (Forward and backward committors.) *Let $\tau = \inf\{t \geq 0 : X_t \in A \cup B\}$ and $\tau^- = \inf\{t \geq 0 : X_t^- \in A \cup B\}$. The forward and backward committor functions q^+ and q^- are defined as*

$$q^+(x) = \mathbf{P}[X_\tau \in B | X_0 = x], \quad q^-(x) = \mathbf{P}[X_{\tau^-}^- \in A | X_0 = x]. \quad (1.47)$$

By applying theorem 1.7 with $f = 0$ and $g = \mathbf{1}_B$, we see immediately that the forward committor is the unique solution of the linear system

$$\begin{aligned} (Lq^+)(x) &= 0, & x \in \mathbb{X} \setminus (A \cup B) \\ q^+(x) &= 0, & x \in A \\ q^+(x) &= 1, & x \in B. \end{aligned} \quad (1.48)$$

Likewise, by applying theorem 1.7 to the reversed process X_t^- with $f = 0$ and $g = \mathbf{1}_A$, we see that the backward committor is the unique solution of the linear system

$$\begin{aligned} (L^*q^-)(x) &= 0, & x \in \mathbb{X} \setminus (A \cup B) \\ q^-(x) &= 1, & x \in A \\ q^-(x) &= 0, & x \in B. \end{aligned} \quad (1.49)$$

Since \mathbb{X} is discrete, the probability density of reactive trajectories μ_R exists, and we are now in a position to compute it. First, observe that $\mu_R(x) = 0$ for $x \in A \cup B$, hence we may assume that $x \in \mathbb{X} \setminus (A \cup B)$. Now we perform a Bayesian inversion:

$$\mu_R(x) = \mathbf{P}[X_t = x | t \in \mathbf{R}] = \frac{\mathbf{P}[t \in \mathbf{R} | X_t = x]}{\mathbf{P}[t \in \mathbf{R}]} \mathbf{P}[X_t = x].$$

But by the definition (1.45) of \mathbf{R} and the homogeneity of $(X_t)_t$,

$$\mathbf{P}[t \in \mathbf{R} | X_t = x] = \mathbf{P}[0 \in \mathbf{R} | X_0 = x] = \mathbf{P}[X_\tau \in B, X_{\tau^-}^- \in A | X_0 = x]$$

with τ and τ^- as in definition 1.12. By the same definition and the Markov property we finally get $\mathbf{P}[t \in \mathbf{R} | X_t = x] = q^+(x)q^-(x)$ and

$$\mu_R(x) = Z_R^{-1} q^+(x) \mu(x) q^-(x), \quad Z_R = \mathbf{P}[t \in \mathbf{R}] = \sum_{x \in \mathbb{X}} q^+(x) \mu(x) q^-(x). \quad (1.50)$$

Now we come to the probability current of reactive trajectories. The discrete analogue to vector fields are **flows**, that is, antisymmetric functions on $\mathbb{X} \times \mathbb{X}$. We first define the probability current of reactive trajectories between states $x, y \in \mathbb{X}$:

$$f_{AB}(x, y) := \lim_{s \downarrow 0} \frac{1}{s} \mathbf{P}[X_t = x, X_{t+s} = y, [t, t+s] \in \mathbf{R}] = q^-(x) \mu(x) l_{xy} q^+(y). \quad (1.51)$$

For a derivation of the last identity see [Met07]. The current f_{AB} is antisymmetrized to give the **effective probability current**

$$f_{AB}^+(x, y) = \frac{1}{2} (f_{AB}(x, y) - f_{AB}(y, x)). \quad (1.52)$$

The effective probability current plays the role of J_{AB} in discrete TPT. The fact that f_{AB}^+ can be negative does not mean that probability currents can be negative, rather the sign of f_{AB}^+ indicates the direction in which *more* reactive trajectories flow. Both f_{AB} and f_{AB}^+ have interesting conservation properties. For any function $F: \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$, we define the **divergence** of F by

$$\operatorname{div} F: \mathbb{X} \rightarrow \mathbb{R}, \quad \operatorname{div} F(x) = \sum_y F(x, y) - \sum_y F(y, x). \quad (1.53)$$

The divergence of F at x measures the difference between the amount of flow transported by F into x and the amount of flow transported out of x . Since divergence is unchanged under antisymmetrization, $\operatorname{div} f_{AB} = \operatorname{div} f_{AB}^+$ and anything we say about the divergence of f_{AB} also holds for f_{AB}^+ . The current f_{AB} is divergence-free on $\mathbb{X} \setminus (A \cup B)$:

$$\operatorname{div} f_{AB}(x) = 0 \quad \forall x \in \mathbb{X} \setminus (A \cup B). \quad (1.54)$$

This follows from the committor equations (1.48) and (1.49). The fact that f_{AB} is divergence-free on $\mathbb{X} \setminus (A \cup B)$ reflects the fact that a reactive trajectory entering $x \in \mathbb{X} \setminus (A \cup B)$ must also leave x and stay reactive at the same time. This is not so on $A \cup B$, in fact A acts as a source and B as a sink of reactive trajectories, and it can be shown [Met07] from (1.54) and (1.51) that the reaction rate can be expressed as

$$\begin{aligned} k_{AB} &= \sum_{x \in A} \operatorname{div} f_{AB}(x) = - \sum_{x \in B} \operatorname{div} f_{AB}(x) \\ &= \sum_{x \in A, y \in \mathbb{X}} f_{AB}(x, y) = \sum_{x \in \mathbb{X}, y \in B} f_{AB}(x, y). \end{aligned} \quad (1.55)$$

The reaction rate can also be expressed with f_{AB}^+ in the same way.

1.4 Optimal Control Theory

In this section we give a short and nontechnical introduction to stochastic optimal control theory, following [FS06]. A complete theory of stochastic optimal control needs the theory

Chapter 1. Background

of viscosity solutions which we omit, for details see also [FS06]. Optimal control theory is a mathematical framework which models the following situation: Given a dynamical system which can be influenced by an external controller, we seek to minimize an objective function over all admissible controls. In stochastic optimal control, the dynamical system is a stochastic process. In the following we describe the controlled processes and objective functions used in this thesis in more detail and then describe the dynamic programming approach for solving the resulting optimization problem formally. We finish by stating several verification theorems.

Let the **controlled process** $(X_t^u)_t$ be a stochastic process on the state space \mathbb{X} whose evolution is influenced by another stochastic process $(u_t)_t$ which takes values in the **control space** \mathbb{U} . We assume that \mathbb{U} is a complete separable metric space. Often $\mathbb{U} = \mathbb{R}^k$ for some $k > 0$. We call the process $(u_t)_t$ an **admissible control strategy** if

1. u_t is \mathcal{F}_t -adapted⁷,
2. $u_t(\omega) \in \mathbb{U}$ for every $\omega \in \Omega$ and every $t \in \mathbf{T}$,
3. The evolution equation for X_t^u has a unique strong⁸ solution.

The second condition says that we are only allowed to select controls from the control space \mathbb{U} , while the first condition says that we are only allowed to base our choice of u_s at time s on the history of $(X_t^u)_t$ from time 0 to time s . The third condition becomes meaningful once we specify an evolution equation for X_t^u . The space of all admissible control strategies is denoted by \mathcal{U} . We call $u \in \mathcal{U}$ a **Markov control** if u is of the form $u_t = \alpha(t, X_t^u)$ for some function $\alpha : \mathbf{T} \times \mathbb{X} \rightarrow \mathbb{U}$.

We have to specify how the control u_t influences the evolution of X_t^u . In specific examples, this is done by specifying the evolution equation for X_t^u , but the following construction can serve as a guiding principle: Suppose that we have a family of infinitesimal generators L^α indexed by $\alpha \in \mathbb{U}$. Then for a constant control strategy $u \equiv \alpha \in \mathbb{U}$, we let (X_t^u) be a Markov process with generator L^α . If $(u_t)_t$ is not constant, then the choice of generator at time s depends on u_s . We now give evolution equations for X_t^u for the systems of interest in this thesis:

- **Controlled Markov Diffusions.** Let $\mathbb{X} = \mathbb{R}^n$. The controlled process X_t^u satisfies the controlled SDE

$$dX_t^u = b(t, X_t^u, u_t)dt + \sigma(t, X_t^u, u_t)dW_t, \quad X_0^u = x \quad (1.56)$$

where W_t is Brownian motion on \mathbb{R}^m , and $b : \mathbf{T} \times \mathbb{R}^d \times \mathbb{U} \rightarrow \mathbb{R}^d$ and $\sigma : \mathbf{T} \times \mathbb{R}^d \times \mathbb{U} \rightarrow \mathbb{R}^{d \times m}$ are the control-dependent drift vector field and diffusion coefficient.

⁷A stochastic process $(u_t)_t$ is called \mathcal{F}_t -adapted if the random variable u_t is \mathcal{F}_t -measurable for every $t \in \mathbf{T}$.

⁸A strong solution to the SDE (1.17) is a stochastic process $(X_t)_t$ that satisfies (1.17) with probability one, with the Brownian motion W_t given in advance. See [Øks03] for an account on weak and strong solutions. If $(X_t^u)_t$ is a MJP, the word 'strong' is omitted.

- **Controlled MJP.** Let $\mathbb{X} = \{1, \dots, d\}$ be finite and let L^α be a generator matrix on \mathbb{X} with entries $l(x, y; \alpha)$ for every $\alpha \in \mathbb{U}$. If the control $u_s = \alpha$ is used at time s , then the jump rate from $x \in \mathbb{X}$ to $y \in \mathbb{X}$ at time s is given by $l(x, y; \alpha)$. The probability $v^u(x, t) = \mathbf{P}(X_t^u = x)$ evolves according to the controlled Fokker-Planck equation

$$\frac{\partial v^u(x, t)}{\partial t} = \sum_{y \neq x} l(x, y; u_t) v^u(x, t) + l(x, x; u_t) v^u(x, t).$$

In both examples, X_t^u is not necessarily a Markov process. However, if $(u_t)_t$ is a Markov control policy, then X_t^u is a Markov process under some additional non-degeneracy assumptions, see [FS06]. The goal of the controller is to minimize the **objective function** J^u over all admissible controls $u \in \mathcal{U}$. The most common objective functions are

1. *finite time horizon:* On a finite time horizon $[0, T]$, the objective function J^u with initial data (t, x) has the form

$$J^u(t, x) = \mathbf{E}_{tx} \left[\int_t^T f(s, X_s^u, u_s) ds + z(X_T^u) \right] \quad (1.57a)$$

where $\mathbf{E}_{tx}[\dots]$ denotes expectation conditioned on $X_t^u = x$. The function $f: [0, T] \times \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{R}$ is called the **running cost**, $z: \mathbb{X} \rightarrow \mathbb{R}$ the **terminal cost** and $T > 0$ the terminal time. Both f and z are required to be measurable functions.

2. *indefinite time horizon:* Let $A \subset \mathbb{X}$ and let the stopping time τ be the first exit time of X_t^u from A (of course τ depends on u). Then J^u has the form

$$J^u(x) = \mathbf{E}_x \left[\int_0^\tau f(X_s^u, u_s) ds + z(X_\tau^u) \right] \quad (1.57b)$$

with running cost $f: \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{R}$ and terminal cost $z: \partial A \rightarrow \mathbb{R}$.

3. *infinite time horizon:* Here we use either a discounted cost criterion

$$J_\lambda^u(x) = \mathbf{E}_x \left[\int_0^\infty e^{-\lambda s} f(X_s^u, u_s) ds \right] \quad (1.57c)$$

with discounting parameter $\lambda > 0$ or an averaged cost criterion

$$J^u(x) = \limsup_{T \rightarrow \infty} \mathbf{E}_x \left[\frac{1}{T} \int_0^T f(X_s^u, u_s) ds \right]. \quad (1.57d)$$

The solution of an optimal control problem is encoded in the following two key quantities: The optimal cost-to-go or **value function** V and, if it exists, the **optimal control policy** u^* :

$$V(t, x) = \inf_{u \in \mathcal{U}} J^u(t, x), \quad u^* = \operatorname{argmin}_{u \in \mathcal{U}} J^u. \quad (1.58)$$

Note that the optimization in (1.58) is carried out over all admissible control policies, which is

an extremely difficult optimization problem. If one knows a priori that u^* exists and is a Markov control policy, then one can restrict the optimization to the much smaller space of Markov controls, but this is not always possible. With dynamic programming which is described next, the optimization over \mathcal{U} can be reduced to a much simpler pointwise optimization.

1.4.1 Dynamic Programming

We review the dynamic programming principle for the case (1.57a) of a finite time horizon problem, which associates the value function (1.58) of the optimal control problem with a non-linear PDE. The discussion here is completely formal, we leave all mathematical rigour to the presentation of the verification theorems in the next section, which state that solutions to certain PDEs coincide with the value function of an optimal control problem.

Suppose that an optimal policy u^* exists, and take $u_s = \alpha \in \mathbb{U}$ for $t \leq s < t+h$ and $u_s = u_s^*$ for $t+h \leq s \leq T$. Then

$$V(t, x) \leq J^u(t, x) = \mathbf{E}_{tx} \left[\int_t^{t+h} f(s, X_s^u, \alpha) ds + V(t+h, X_{t+h}^u) \right] \quad (1.59)$$

Because $u_s = \alpha$ is constant on $[t, t+h]$, $(X_s^u)_s$ is a Markov process with generator L^α on $[t, t+h]$. Suppose that all the generators L^α have a common domain of definition \mathcal{D} and that $V(t, \cdot) \in \mathcal{D}$. Then by (1.13),

$$L^\alpha V(t, x) = \lim_{h \downarrow 0} \frac{1}{h} \left(\mathbf{E}_{tx} [V(t, X_{t+h}^u)] - V(t, x) \right).$$

Furthermore, if $V(\cdot, x) \in C^1$ then

$$\lim_{h \downarrow 0} \frac{1}{h} \left(\mathbf{E}_{tx} [(V(t+h, X_{t+h}^u))] - \mathbf{E}_{tx} [(V(t, X_{t+h}^u))] \right) = \frac{\partial V(t, x)}{\partial t}$$

and

$$\lim_{h \downarrow 0} \frac{1}{h} \mathbf{E}_{tx} \left[\int_t^{t+h} f(s, X_s^u, \alpha) ds \right] = f(t, x, \alpha).$$

Using all of this, we can subtract $V(t, x)$ on both sides in (1.59), divide by h and take the limit $h \rightarrow 0$, which gives after collecting all the terms:

$$0 \leq f(t, x, \alpha) + \frac{\partial V(t, x)}{\partial t} + L^\alpha V(t, x). \quad (1.60)$$

On the other hand we expect that equality in (1.59) holds if $u = u^*$ on $[t, T]$. Let $\alpha^* = \mathbf{E}[u_t^* | \mathcal{F}_t]$ and note that $\alpha^* \in \mathbb{U}$ is a.s. constant since u_t^* is \mathcal{F}_t -adapted. Then if u^* is continuous, we have

$$0 = f(t, x, \alpha^*) + \frac{\partial V(t, x)}{\partial t} + L^{\alpha^*} V(t, x).$$

This equality together with the inequality (1.60) yields the **Bellman equation**

$$0 = \min_{\alpha \in \mathbb{U}} \left\{ f(t, x, \alpha) + \frac{\partial V(t, x)}{\partial t} + L^\alpha V(t, x) \right\} \quad 0 \leq t \leq T, x \in \mathbb{X}. \quad (1.61)$$

For Markov diffusions, equation (1.61) is a nonlinear parabolic PDE in $[0, T] \times \mathbb{X}$ together with the terminal condition $V(T, x) = z(x)$ for all $x \in \mathbb{X}$. For MJPs, (1.61) is a set of coupled nonlinear evolution equations. The advantage of the Bellman equation is that the minimization over \mathcal{U} in (1.58) has been reduced to a pointwise minimization over \mathbb{U} , which is a lot more feasible computationally. Still, a direct solution of (1.61) is not feasible if \mathbb{X} is high dimensional. The Bellman equation also suggests that an optimal Markov control policy $u_t^* = \alpha^*(t, X_t^u)$ should exist, where the function α^* should satisfy

$$\alpha^*(t, x) = \operatorname{argmin}_{\alpha \in \mathbb{U}} \{ f(t, x, \alpha) + L^\alpha V(t, x) \} \quad 0 \leq t \leq T, x \in \mathbb{X}. \quad (1.62)$$

The derivation of (1.61) is formal in the sense that we made several strong and unjustified assumptions, e.g. the existence of an optimal control, sufficient smoothness of V and the assumption that $V(t, \cdot) \in \mathcal{D}$ for all $0 \leq t \leq T$. With some technical effort these assumptions can be relaxed and the derivation made rigorous, see [FS06].

1.4.2 Controlled Markov Diffusions

We now specialise to the diffusion setting. Let $(X_t^u)_t$ be a controlled Markov diffusion in \mathbb{R}^d evolving according to the SDE (1.56). We let $\mathbb{U} = \mathbb{R}^k$ and recall that for any admissible u the SDE (1.56) is required to have a unique strong solution. For any $\alpha \in \mathbb{U}$, define the action of the generator L^α on $g \in C_0^{1,2}(\mathbb{T} \times \mathbb{R}^d)$ as

$$L^\alpha g(t, x) = b(t, x, \alpha) \cdot \nabla g + a(t, x, \alpha) : \nabla^2 g \quad (1.63)$$

where $a(t, x, \alpha) = \frac{1}{2} \sigma(t, x, \alpha) \sigma^T(t, x, \alpha)$. We now state a verification theorem that establishes a link between solutions to the Bellman equation (1.61) and solutions to the corresponding optimal control problem [vH07].

Theorem 1.13 (Verification, finite time horizon.) *Let $O = [0, T] \times \mathbb{X}$. Suppose $V \in C^{1,2}(O)$ is a solution to*

$$\begin{aligned} 0 &= \frac{\partial V(t, x)}{\partial t} + \min_{\alpha \in \mathbb{U}} \{ L^\alpha V(t, x) + f(t, x, \alpha) \}, & (t, x) \in O, \\ V(T, x) &= z(x), & x \in \mathbb{X} \end{aligned} \quad (1.64)$$

and assume that $|\mathbf{E}[V(0, X_0)]| < \infty$. Denote by $\mathcal{U}_0 \subset \mathcal{U}$ the class of admissible control strategies u such that

$$\int_0^t \nabla V(s, X_s^u) \cdot \sigma(s, X_s^u, u_s) dW_s$$

Chapter 1. Background

is a martingale (rather than a local martingale). Then $V(t, x) \leq J^u(t, x)$ for every $u \in \mathcal{U}_0$ and all $(t, x) \in O$, where the cost function J^u is given by (1.57a). If further a minimum

$$\alpha^*(t, x) = \operatorname{argmin}_{\alpha \in \mathbb{U}} \{L^\alpha V(t, x) + f(t, x, \alpha)\}, \quad (t, x) \in O \quad (1.65)$$

exists and $u_t^* = \alpha^*(t, X_t^u)$ is in \mathcal{U}_0 , then $V(t, x) = J^{u^*}(t, x)$ for all $(t, x) \in O$, and u^* is an optimal Markov control strategy for the control problem (1.57a).

With the verification Theorem 1.13 at hand, we know that once we find a solution to the Bellman equation (1.64), we know that this solution coincides with the value function of the finite time horizon optimal control problem (1.57a). Furthermore, if the minimum in (1.65) exists and defines an admissible control strategy, then we know that an optimal control for problem (1.57a) is given by the Markov control $u_t^* = \alpha^*(t, X_t^u)$. This approach is very practical and suffices for this thesis. There is also a verification theorem for the case of an indefinite time horizon [vH07]:

Theorem 1.14 (Verification, indefinite time horizon.) *Let $A \subset \mathbb{X}$ be open and bounded with closure \bar{A} . Suppose $V \in C^2(\bar{A})$ is a solution to*

$$\begin{aligned} 0 &= \min_{\alpha \in \mathbb{U}} \{L^\alpha V(x) + f(x, \alpha)\}, & x \in A, \\ V(x) &= z(x), & x \in \partial A. \end{aligned} \quad (1.66)$$

Let $X_0^u \in A$ a.s. and let τ be the first exit time of X_t^u from A . Denote by $\mathcal{U}_0 \subset \mathcal{U}$ the class of admissible control strategies u such that $\tau < \infty$ a.s. and

$$\mathbf{E} \left[\int_0^\tau \nabla V(s, X_s^u) \cdot \sigma(s, X_s^u, u_s) dW_s \right] = 0.$$

Then $V(x) \leq J^u(x)$ for every $u \in \mathcal{U}_0$ and all $x \in A$, where the cost function J^u is given by (1.57b). If further a minimum

$$\alpha^*(x) = \operatorname{argmin}_{\alpha \in \mathbb{U}} \{L^\alpha V(x) + f(x, \alpha)\}, \quad x \in A \quad (1.67)$$

exists and $u_t^* = \alpha^*(X_t^u)$ is in \mathcal{U}_0 , then $V(x) = J^{u^*}(x)$ for all $x \in A$, and u^* is an optimal Markov control strategy for the control problem (1.57b).

Results similar to Theorem (1.13) and (1.14) exist for both the discounted cost and averaged cost criteria for infinite time horizons, but these results are omitted here. They can be found e.g. in [FS06].

1.4.3 Controlled Markov Jump Processes

We now specialize to controlled MJPs. Let $\mathbb{X} = \{1, \dots, d\}$ and let L^α be a generator matrix on \mathbb{X} with entries $l(x, y; \alpha)$ for every $\alpha \in \mathbb{U}$. Denote by $\mathcal{U}_0 \subset \mathcal{U}$ the class of admissible control strategies such that for all functions $\Phi: \mathbb{X} \rightarrow \mathbb{R}$ and for any stopping time τ the Dynkin formula holds:

$$\mathbf{E}_x[\Phi(X_\tau^u)] = \Phi(x) + \mathbf{E}_x \left[\int_0^\tau L^{u_s} \Phi(X_s^u) ds \right]. \quad (1.68)$$

In [FS06], the term 'admissible control system' is used for the class \mathcal{U}_0 . Essentially, the martingale properties used to define \mathcal{U}_0 in the previous section together with the smoothness assumptions on V are equivalent to (1.68). This allows us to carry the verification theorems (1.13) and (1.14) over to the discrete case. Specifically, we have the following:

Theorem 1.15 (Verification, finite time horizon.) *Let $O = [0, T) \times \mathbb{X}$. Suppose $V \in C^{1,0}(O)$ is a solution to*

$$\begin{aligned} 0 &= \frac{\partial V(t, x)}{\partial t} + \min_{\alpha \in \mathbb{U}} \{L^\alpha V(t, x) + f(t, x, \alpha)\}, & (t, x) \in O, \\ V(T, x) &= z(x), & x \in \mathbb{X}. \end{aligned} \quad (1.69)$$

Then $V(t, x) \leq J^u(t, x)$ for every $u \in \mathcal{U}_0$ and all $(t, x) \in O$, where the cost function J^u is given by (1.57a). If further a minimum

$$\alpha^*(t, x) = \operatorname{argmin}_{\alpha \in \mathbb{U}} \{L^\alpha V(t, x) + f(t, x, \alpha)\}, \quad (t, x) \in O \quad (1.70)$$

exists and $u_t^ = \alpha^*(t, X_t^u)$ is in \mathcal{U}_0 , then $V(t, x) = J^{u^*}(t, x)$ for all $(t, x) \in O$, and u^* is an optimal Markov control strategy for the control problem (1.57a).*

And for the indefinite time horizon case:

Theorem 1.16 (Verification, indefinite time horizon.) *Let $A \subset \mathbb{X}$. Suppose $V: \mathbb{X} \rightarrow \mathbb{R}$ is a solution to*

$$\begin{aligned} 0 &= \min_{\alpha \in \mathbb{U}} \{L^\alpha V(x) + f(x, \alpha)\}, & x \in A, \\ V(x) &= z(x), & x \in \mathbb{X} \setminus A. \end{aligned} \quad (1.71)$$

Let $X_0^u \in A$ a.s. and let τ be the first exit time of X_t^u from A and assume that τ is a.s. finite for all $u \in \mathcal{U}_0$. Then $V(x) \leq J^u(x)$ for every $u \in \mathcal{U}_0$ and all $x \in A$, where the cost function J^u is given by (1.57b). If further a minimum

$$\alpha^*(x) = \operatorname{argmin}_{\alpha \in \mathbb{U}} \{L^\alpha V(x) + f(x, \alpha)\}, \quad x \in A \quad (1.72)$$

Chapter 1. Background

exists and $u_t^* = \alpha^*(X_t^u)$ is in \mathcal{U}_0 , then $V(x) = J^{u^*}(x)$ for all $x \in A$, and u^* is an optimal Markov control strategy for the control problem (1.57b).

We prove Theorem (1.16) in appendix A, the proof of Theorem (1.15) is analogous.

2 Irreversible Markov Chains and Cycles

This chapter studies irreversible Markov chains as a prototype for non-equilibrium systems. We use graph-theoretical representations of the transition matrix P of the Markov chain, and the notion of cycle decompositions in particular, as our main tools. Both are introduced in detail in section 2.1. We first give a purely algebraic construction to arrive at a cycle decomposition in section 2.1.2, then we present a stochastic construction in section 2.1.3. We then discuss the connections between cycle decompositions and reversibility (or irreversibility) in section 2.1.4 and give connections to the entropy production rate functional in section 2.1.5.

The second part of the chapter is devoted to the construction of reversible surrogates to irreversible Markov processes. We formulate the problem in section 2.2 and classify possible solutions to a simple subproblem in section 2.2.1. The solution to the general problem via cycle decompositions is constructed in section 2.2.2. Finally, a theoretical result that shows that commuting times for reversible processes are always accelerated if a non-reversible term is added is shown in section 2.2.3.

The third part of the chapter is concerned with applications and numerical examples. We discuss two main applications: Module detection in section 2.3 and dealing with loops in transition paths in section 2.4. Section 2.3 uses the ideas developed in sections 2.1 and 2.2. Section 2.4 also uses TPT, which was introduced in section 1.3.

2.1 Cycle Decompositions

2.1.1 Coates Graph and Cycle Space

Here and in the rest of the chapter, we assume the state space \mathbb{X} to be finite. This allows us to give a graph-theoretical characterization of irreversibility in terms of cycles in a directed graph which we associate to our Markov process. To set this up, first note that we are free to think of any linear operator L on $\mathbb{R}^{\mathbb{X}}$ as a bilinear form $L : \mathbb{R}^{\mathbb{X}} \times \mathbb{R}^{\mathbb{X}} \rightarrow \mathbb{R}$ acting as $L(v, w) = \langle v, Lw \rangle$.

Chapter 2. Irreversible Markov Chains and Cycles

Here, $\langle v, w \rangle = v^T w$ is the canonical scalar product on $\mathbb{R}^{\mathbb{X}}$. By restricting this bilinear form to the canonical basis on $\mathbb{R}^{\mathbb{X}}$, we obtain a map $L: \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$. The values $L(x, y) \equiv l_{xy}$ for $x, y \in \mathbb{X}$ are simply the entries of L viewed as a matrix. We will freely switch between these different points of view. Now we make the following definition:

Definition 2.1 *Let L be a linear operator on $\mathbb{R}^{\mathbb{X}}$. The **Coates graph** of L is the directed graph $G(L)$ with vertex set \mathbb{X} and edge set $E = \{(x, y) \in \mathbb{X} \times \mathbb{X} : L(x, y) > 0\}$.*

Let $(X_n)_n$ be a Markov chain on \mathbb{X} with transition matrix P and associated Coates graph $G(P)$. We now develop graph-theoretical interpretations of several properties of $(X_n)_n$ in terms of $G(P)$. These results generalize immediately to continuous time MJPs by replacing $G(P)$ with $G(L)$. A first observation is that ergodicity translates into a form of connectivity: A **directed path** in G is a sequence (x_1, \dots, x_s) of nodes in G where every (x_i, x_{i+1}) is an edge in G for $i = 1, \dots, s-1$. We call the nodes x, y in G **strongly connected** if there is a directed path from x to y and a (possibly different) directed path from y to x . x and y are called **weakly connected** in $G(P)$ if they are strongly connected in $G(P + P^T)$. Strong (weak) connectivity is an equivalence relation and partitions G into equivalence classes which we call strongly (weakly) connected components. We call G strongly (weakly) connected if it has only one strongly (weakly) connected component (which then must be all of \mathbb{X}). Then we have

Lemma 2.2 *The communication classes of P equal the strongly connected components of the Coates graph $G(P)$. In particular, if P is ergodic then $G(P)$ is strongly connected.*

Proof: Two states x, y are said to communicate if there are natural numbers r, s such that $(P^r)_{xy} > 0$ and $(P^s)_{yx} > 0$. But then $p_{xx_2} p_{x_2 x_3} \dots p_{x_{r-1} y} > 0$ for some sequence of nodes $x, x_2, \dots, x_{r-1}, y$, which implies that there is a directed path of length r from x to y in $G(P)$. For the same reasons, there is also a directed path of length s from y to x in $G(P)$. This shows that x and y are strongly connected if they communicate. The converse is trivial, hence communication and strong connectivity are logically equivalent. If P is ergodic, then P must be irreducible, hence it must have only one communication class, therefore $G(P)$ has only one strongly connected component. ■

The study of non-ergodic Markov chains essentially reduces to the study of each communication class separately. We therefore assume ergodicity of $(X_n)_n$ and thus strong connectivity of $G(P)$ from here on. In the rest of this section, we introduce the cycle and cocycle vector spaces of general oriented graphs, following [Kal06] and [Pol15]. Cycles in the Coates graph $G(P)$ will play a prominent role later on if P is irreversible.

Let $G = (V, E)$ be an oriented graph with edges $E = \{e_1, \dots, e_{|E|}\}$. The **edge vector space** $\mathbb{R}^{|E|}$ is the space of formal linear combinations of edges (in an algebraic topology context, $\mathbb{R}^{|E|}$ would be called the space of one-chains). We endow $\mathbb{R}^{|E|}$ with the standard Euclidean scalar product $\langle \cdot, \cdot \rangle$. Since edges are oriented, every edge e has a unique **source** $s(e) \in V$ and a unique

target $t(e) \in V$. We identify $s(e)$ and $t(e)$ with the associated canonical basis vectors in $\mathbb{R}^{|V|}$ and denote the boundary of an edge e by $\partial e = t(e) - s(e)$. By linear extension this gives us a boundary operator $\partial: \mathbb{R}^{|E|} \rightarrow \mathbb{R}^{|V|}$. We define the vector space of **one-cycles** as $\mathcal{C} = \ker \partial$, and the space \mathcal{C}^* of **cocycles** as the row space of ∂ . We reserve the term **cycles** for elements of \mathcal{C} of the form $c = \sum_i s_i e_i$ with $s_i \in \{0, \pm 1\}$. The coefficient s_i is positive if the orientations of c and e_i agree, negative if they are opposite and zero if e_i is not present in c . Graphically, cycles are sequences of edges such that for each incoming edge at a vertex x , there is an outgoing one. The rows of ∂ correspond to simple cocycles, which are minimal sets of edges that disconnect G into two subgraphs. An example graph with some cycles and a cocycle is shown in Figure 2.1.




Figure 2.1 – Example of an oriented graph G . Some cycles are $c_1 = e_2 + e_3$, $c_2 = e_1 + e_5 + e_6 + e_7$, $c_3 = e_2 + e_4 - e_5$ and $c_4 = -e_3 + e_4 - e_5$. The dotted edges denote the cocycle $c^* = e_4 + e_5 - e_6$ whose removal disconnects the red vertex from the rest of G .

Notice that we distinguish between the cycles $c_3 = e_2 + e_4 - e_5$ and $c_4 = -e_3 + e_4 - e_5$. This is different from e.g. the treatment in [Sch76] and [Pol15], where $-e$ denotes e with opposite orientation, and e and $-e$ cannot both be present in E . For us $-e$ means that the edge e is traversed in the opposite orientation, and there may or may not be another distinct edge $e' \in E$ which is e with that reversed orientation. This has the effect of regarding e.g. $c_1 = e_2 + e_3$ not as identical to zero, but as a cycle of length 2.

The dimension of \mathcal{C}^* is easily shown¹ to be $|V| - 1$. By the rank-nullity theorem the dimension of \mathcal{C} is equal to the Betti number $b = |E| - |V| + 1$. There is a canonical way to construct a basis $\{\gamma_1, \dots, \gamma_b\}$ of \mathcal{C} [Sch76]: Let $T \subset G$ be a spanning tree of G (that is, a subgraph of G that weakly connects all vertices of G , but contains no cycles). See Figure 2.2 for an illustration. Edges of G which belong to T are called **cochords** and denoted by $\{e_\mu\}_{\mu=1}^{|V|-1}$. Edges of G which do not belong to T are called **chords** and denoted by $\{e_\alpha\}_{\alpha=1}^b$. Adding a chord e_α to T encloses one

¹The columns of ∂ correspond to edges and every column has exactly one $+1$ (for its target) and one -1 (for its source). Therefore, the rows of ∂ sum to zero. Now let us remove the row corresponding to $x_0 \in \mathbb{X}$. We have to show that the remaining $|V| - 1$ rows are linearly independent, i.e. that $0 = \sum_{x \neq x_0} \alpha_x \partial_{x e}$ implies all $\alpha_x = 0$. Pick any $x^* \in \mathbb{X}$. Because G is strongly connected, there is a directed path $(x_0, x_1, \dots, x_n, x^*)$ from x_0 to x^* with $e_i := (x_i, x_{i+1}) \in E$. The column corresponding to e_0 only has one $+1$ and no -1 because $x_0 = s(e_0)$ has been removed. Therefore, $\alpha_{x_1} = 0$. By induction, all $\alpha_{x_i} = 0$ and finally $\alpha_{x^*} = 0$.

cycle γ_α which we orient in the direction of e_α . The set $\Gamma = \{\gamma_1, \dots, \gamma_b\}$ of cycles obtained in this way is linearly independent because $\langle e_\alpha, \gamma_{\alpha'} \rangle = \delta_{\alpha\alpha'}$. Since $\dim \mathcal{C} = b$, Γ is a basis of \mathcal{C} . We term Γ the **Betti basis** of \mathcal{C} .




Figure 2.2 – Oriented graph G with spanning tree T (dark edges) and chords e_1, e_2, e_3 (red edges). The corresponding Betti cycles are $\gamma_1 = e_1 + e_5 + e_6 + e_7$, $\gamma_2 = e_2 + e_4$ and $\gamma_3 = e_3 - e_5 - e_4$.

2.1.2 Decomposing Flows in the Betti Basis

Let $F : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$ be a flow². Recall the definition of divergence in section 1.3. We call F **divergence-free** if

$$\operatorname{div} F(x) = \sum_y F(x, y) - \sum_y F(y, x) = 0. \quad (2.1)$$

In physical terms, (2.1) expresses a conservation law: Since the flow into x balances the flow out of x , the type of mass which is transported by F is conserved at x . In the context of electrical networks, F would be an electrical current and (2.1) is called **Kirchhoff's current law**. Let $\mathbb{R}^{|\mathbb{E}|}$ be the space of one-chains and \mathcal{C} be the space of one-cycles of $G(F)$. We can associate a one-chain \underline{F} to F by putting $\underline{F} = \sum_{x,y} F(x, y) e_{xy}$ where e_{xy} denotes the edge with source x and target y . From (2.1) we get that

$$\partial \underline{F} = \sum_{x,y} F(x, y) \partial e_{xy} = \sum_{x,y} F(x, y) (y - x) = \sum_x x \left(\sum_y -F(x, y) + \sum_y F(y, x) \right) = 0.$$

Hence $\underline{F} \in \mathcal{C}$, and we can use the Betti basis to decompose \underline{F} in terms of the Betti basis [Kal06]:

Lemma 2.3 (Cycle decomposition, algebraic version.) *Let $F : \mathbb{X} \times \mathbb{X} \rightarrow [0, \infty)$ be divergence-free, and let $\Gamma = \{\gamma_1, \dots, \gamma_b\}$ be the Betti basis of the space of one-cycles in $G(F)$. Then there are*

²In section 1.3, we required flows to be antisymmetric functions on $\mathbb{X} \times \mathbb{X}$, which is standard. We drop this requirement for the moment to be slightly more general.

weights $\omega(\gamma_\alpha) > 0$ such that

$$F(x, y) = \sum_{\alpha=1}^b \omega(\gamma_\alpha) \langle e_{xy}, \gamma_\alpha \rangle \quad (2.2)$$

holds for all $x, y \in \mathbb{X}$. The weights are given by $\omega(\gamma_\alpha) = F(s(e_\alpha), t(e_\alpha))$ where e_α is the chord corresponding to γ_α .

Proof: We associate the one-chain $\underline{F} = \sum_{x,y} F(x, y) e_{xy}$ to F . Since F is divergence-free, $\underline{F} \in \mathcal{C}$ and can be expressed in terms of the Betti basis as

$$\underline{F} = \sum_{\alpha=1}^b \omega(\gamma_\alpha) \gamma_\alpha \quad (2.3)$$

with real numbers $\omega(\gamma_\alpha)$. We get (2.2) from (2.3) by taking the scalar product with e_{xy} on both sides. If we take the scalar product with e_α on both sides, we get $\omega(\gamma_\alpha) = \langle e_\alpha, \underline{F} \rangle = F(s(e_\alpha), t(e_\alpha))$ since $\langle e_\alpha, \gamma_{\alpha'} \rangle = \delta_{\alpha\alpha'}$, and $\langle e_\alpha, \underline{F} \rangle > 0$ because $\underline{F} > 0$. ■

Example. Consider the graph shown in Figure (2.2). The flow $F(x, y) = \mu_x p_{xy}$ of the simple random walk on this graph is given by

$$F = \frac{1}{15} (3e_1 + 2e_2 + e_3 + e_4 + 2e_5 + 3e_6 + 3e_7).$$

The decomposition of F into Betti cycles shown in Figure 2.2 reads $F = \frac{1}{15} (3\gamma_1 + 2\gamma_2 + \gamma_3)$, this is (2.3). Note that the weights in (2.2) are positive, but the scalar products $\langle e_{xy}, \gamma_\alpha \rangle$ can be either +1 (if the orientations of e_{xy} and γ_α agree) or -1 (if they disagree). For example, the flow through the edge e_5 is $\langle e_5, F \rangle = \frac{1}{15} (3\langle e_5, \gamma_1 \rangle + 2\langle e_5, \gamma_2 \rangle + \langle e_5, \gamma_3 \rangle) = \frac{1}{15} (3 - 1)$.

Computational aspects. To compute the cycle decomposition (2.2), we need to perform the following steps:

1. Find a spanning tree $T \subset G(F)$. Identify the chords $\{e_\alpha, \alpha = 1 \dots b\}$.
2. For each chord e_α , find the sequence $\lambda_\alpha = \sum_{\mu=1}^{|V|-1} s_\mu e_\mu$ of cochords that connects $t(e_\alpha)$ with $s(e_\alpha)$. Then set $\gamma_\alpha := e_\alpha + \lambda_\alpha$.
3. For each chord e_α , compute the corresponding weight $\omega(\gamma_\alpha) = \langle e_\alpha, F \rangle$.

These steps are fairly easy to perform. The computationally most demanding part for large graphs is the identification of the spanning tree T , which can be done e.g. by the Prim [Pri57] algorithm in time $\mathcal{O}(|E| \log |\mathbb{X}|)$. There are two major disadvantages of (2.2): (i) The Betti cycles γ_α and the weights $\omega(\gamma_\alpha)$ are not unique; they depend on the choice of spanning tree T . (ii) Since $\omega(\gamma_\alpha) > 0$, we can interpret ω as a probability distribution over Γ . But the Betti cycles γ_α

themselves are algebraic objects that do not allow for a stochastic interpretation. For example, it is not possible for the Markov chain $(X_n)_n$ to pass through all the edges of γ_3 in Figure 2.2 in the direction which is induced by γ_3 since e_4 and e_5 can only be passed through in the opposite direction. To remedy this, we now present a stochastic version of (2.2) due to Qian [JQQ04].

2.1.3 The Circulation Distribution

Let \mathcal{C} be the space of one-cycles of $G(P)$. We define a subset $\mathcal{C}_0 \subset \mathcal{C}$ of cycles that can be given a stochastic meaning by

$$\mathcal{C}_0 = \left\{ c \in \mathcal{C} : c \text{ is simple, } c = \sum_i s_i e_i \text{ with } s_i = \{0, 1\} \right\}.$$

Here **simple** means that c has no self-intersections, i.e. there are no two edges in c that have the same source or the same target nodes. We can denote a cycle $c \in \mathcal{C}_0$ unambiguously by the sequence of edges that c visits, i.e. $c = [x_1, \dots, x_p]$ is the cycle with edges $(x_1, x_2), \dots, (x_{p-1}, x_p), (x_p, x_1)$. The orientation of $c \in \mathcal{C}_0$ always agrees with the orientation of its edges, therefore we may introduce the **passage functions** that describe which nodes and edges are part of c as follows:

$$J_c(x) = \begin{cases} 1 & \text{if } x \text{ is a node in } c, \\ 0 & \text{otherwise} \end{cases}; \quad J_c(x, y) = \begin{cases} 1 & \text{if } (x, y) \text{ is an edge in } c, \\ 0 & \text{otherwise} \end{cases}. \quad (2.4)$$

The two passage functions are related by the balance property

$$\sum_y J_c(x, y) = \sum_y J_c(y, x) = J_c(x). \quad (2.5)$$

For any cycle $c \in \mathcal{C}_0$, it is possible for the Markov chain $(X_n)_n$ to pass through all the edges of c in the direction and order induced by c . These passages can be counted given a realization $(X_n)_{1 \leq n \leq T}$ of length T of $(X_n)_n$. The sequence (X_1, X_2, \dots, X_T) of states visited by the realization forms a path in $G(P)$ which has a finite number of self-intersections or **recurrences**. Each recurrence corresponds to a cycle in $G(P)$ that has been passed through by $(X_n)_{1 \leq n \leq T}$. For each cycle $c \in \mathcal{C}_0$ we can count the number of times N_c^T it has been passed through by $(X_n)_{1 \leq n \leq T}$. To make this precise, we introduce the **derived chain** $(\eta_n)_n$ associated to $(X_n)_n$ which takes values in the space of finite ordered sequences of elements in \mathbb{X} . $(\eta_n)_n$ acts as a memory which allows us to detect when recurrences occur and which cycles are completed. We set $N_c^T := 0$ for all $c \in \mathcal{C}_0$ and $\eta_1 := X_1$. The other values of $(\eta_n)_n$ are constructed iteratively together with the passage counts N_c^T as follows:

- If $\eta_i = (y_1, \dots, y_s)$ and $X_{i+1} \neq y_l$ for all $1 \leq l \leq s$, then set $\eta_{i+1} := (y_1, \dots, y_s, y_{s+1})$ with $y_{s+1} := X_{i+1}$.

2.1. Cycle Decompositions

n	1	2	3	4	5	6	7	8
X_n	1	4	2	3	2	6	7	4
η_n	(1)	(1, 4)	(1, 4, 2)	(1, 4, 2, 3)	(1, 4, 2)	(1, 4, 2, 6)	(1, 4, 2, 6, 7)	(1, 4)
Cycles					[2, 3]			[4, 2, 6, 7]

Table 2.1 – The first few states of the Markov chain $(X_n)_n$ and the derived chain $(\eta_n)_n$, and the cycles passed through by $(X_n)_n$.

- If $\eta_i = (y_1, \dots, y_s)$ and $X_{i+1} = y_l$ for some $1 \leq l \leq s$, then the cycle $c := [y_l, \dots, y_s]$ has been completed. We set $N_c^T \rightarrow N_c^T + 1$ and $\eta_{i+1} := (y_1, \dots, y_l)$.

In table 2.1, an example of this construction is shown. Now we set $\omega_T(c) := N_c^T / T$ for all $c \in \mathcal{C}_0$, and $\mathcal{C}_T := \{c \in \mathcal{C}_0 : \omega_T(c) > 0\}$. Both ω_T and \mathcal{C}_T are \mathcal{F}_T -measurable random variables, and Qian et al. [JQQ04] have shown the following:

Lemma 2.4 (Cycle decomposition, stochastic version.) *The sequences of random variables $(\mathcal{C}_T)_{T \in \mathbb{N}}$ and $(\omega_T)_{T \in \mathbb{N}}$ converge almost surely. With $\mathcal{C}_\infty := \bigcup_{T=0}^\infty \mathcal{C}_T$ and $\omega_\infty := \lim_{T \rightarrow \infty} \omega_T$, we have*

$$\mu_x p_{xy} = \sum_{c \in \mathcal{C}_\infty} \omega_\infty(c) J_c(x, y) \quad \forall x, y \in \mathbb{X}, \quad (2.6)$$

$$\mu_x = \sum_{c \in \mathcal{C}_\infty} \omega_\infty(c) J_c(x) \quad \forall x \in \mathbb{X}. \quad (2.7)$$

Furthermore, for any $c = [x_1, \dots, x_s] \in \mathcal{C}_\infty$ the weight $\omega_\infty(c)$ is given by

$$\omega_\infty(c) = p_{x_1 x_2} p_{x_2 x_3} \cdots p_{x_{s-1} x_s} p_{x_s x_1} \frac{D(\{x_1, \dots, x_s\})}{\sum_{x \in \mathbb{X}} D(\{x\})} \quad (2.8)$$

where $D(\{x_1, \dots, x_s\})$ is the determinant of the matrix $I - P$ with rows and columns indexed by $\{x_1, \dots, x_s\}$ deleted (we set $D(\mathbb{X}) = 1$ by convention).

Example: In the graph shown in Figure 2.2, $\mathcal{C}_0 = \{c_1, c_2, c_3\}$ with $c_1 = \gamma_1$, $c_2 = \gamma_2$ and $c_3 = e_1 + e_2 + e_3 + e_6 + e_7$. The corresponding weights for the simple random walk are given by $\omega_\infty(c_1) = 2/15$, $\omega_\infty(c_2) = 1/15$ and $\omega_\infty(c_3) = 1/15$, and (2.6) for the flow $F(x, y) = \mu_x p_{xy}$ of the simple random walk reads $F = \frac{1}{15}(2c_1 + c_2 + c_3)$.

Note that (2.7) is an easy consequence of (2.6) and (2.5) and can be traced back to the fact that we only allow positive coefficients s_i in $c = \sum_i s_i e_i$ for $c \in \mathcal{C}_\infty$. The function ω_∞ is termed **circulation distribution** by Qian. $\omega_\infty(c)$ can be interpreted as the probability that $(X_n)_n$ completes the cycle c with its next step if it was initialized in equilibrium. Both ω_∞ and \mathcal{C}_∞ are unique, and (2.8) shows that $\omega_\infty(c)$ is the product of the transition probabilities along the edges of c times a normalization constant times the minor $D(\{x_1, \dots, x_s\})$ which accounts for the excursions $(X_n)_n$ is allowed to take while completing c . For example, while completing the cycle $[4, 2, 6, 7]$ the realization in table 2.1 performs the excursion $(2, 3, 2)$.

Computational aspects. The uniqueness and stochastic meaning of (2.6) are the advantages of (2.6) over (2.2). However, (2.6) is hard to compute: \mathcal{C}_∞ is much larger than the Betti basis Γ , and computing it directly amounts to finding all simple cycles in the directed graph G , which is at least an NP-hard problem. Here is the reason: If one has a list of all simple cycles in G , one may take the largest one(s) and check if it is a Hamiltonian cycle (a cycle visiting every node exactly once). But the problem of determining if a directed graph G has a Hamiltonian cycle is already NP-hard. In addition, for every $c \in \mathcal{C}_\infty$ a different minor of $I - P$ must be computed in order to obtain $\omega_\infty(c)$ according to (2.8). In practice, we will always work with a finite realization $(X_n)_{1 \leq n \leq T}$ and sample ω_T and \mathcal{C}_T according to the algorithm presented above.

2.1.4 Cycle Decompositions and Reversibility

We study the effects of time reversal on the cycle decompositions (2.2) and (2.6). Let μ be the unique invariant distribution of P and D_μ be the diagonal matrix with diagonal entries given by μ . We form the probability flow matrix $F = D_\mu P$ whose components are given by $F_{xy} = \mathbf{P}(X_n = x, X_{n+1} = y)$. Note that $\mu > 0$ by ergodicity, so that $F_{xy} = \mu_x p_{xy} > 0$ iff $p_{xy} > 0$, therefore $G(P) = G(F)$. Let P^- be the transition matrix of the time-reversed Markov chain $(X_n^-)_n$. In components, $p_{xy}^- = \frac{\mu_y}{\mu_x} p_{yx}$. The probability flow matrix $F^- = D_\mu P^-$ associated to P^- is just the transpose of F . The Coates graph $G(P^-)$ is $G(P)$ with the orientations of all edges reversed. We denote by e^- the edge e with reversed orientation, that is $s(e^-) = t(e)$ and $t(e^-) = s(e)$. Similarly, we denote by $c^- \in \mathcal{C}$ the cycle c with reversed orientation, that is if $c = \sum_i s_i e_i$ then $c^- = \sum_i s_i e_i^-$. The space of one-cycles of $G(P^-)$ is $\mathcal{C}^- = \{c^- : c \in \mathcal{C}\}$.

If T is a spanning tree of $G(P)$ with chords e_α , then T^- is a spanning tree of $G(P^-)$ with corresponding chords e_α^- . Correspondingly, if $\Gamma = \{\gamma_1, \dots, \gamma_b\}$ is a Betti basis of \mathcal{C} , then $\Gamma^- := \{\gamma_1^-, \dots, \gamma_b^-\}$ is a Betti basis of \mathcal{C}^- . We write the one-chain $\underline{F}^- \in \mathcal{C}^-$ in terms of Γ^- as

$$\underline{F}^- = \sum_{\alpha=1}^b \langle e_\alpha^-, F^- \rangle \gamma_\alpha^- = \sum_{\alpha=1}^b \langle e_\alpha, F \rangle \gamma_\alpha^- = \sum_{\alpha=1}^b \omega(\gamma_\alpha) \gamma_\alpha^-.$$

Thus $\underline{F} - \underline{F}^- = \sum_\alpha \omega(\gamma_\alpha) [\gamma_\alpha - \gamma_\alpha^-]$, or in components

$$F_{xy} - F_{xy}^- = \mu_x p_{xy} - \mu_y p_{yx} = \sum_{\alpha=1}^b \omega(\gamma_\alpha) \langle e_{xy}, \gamma_\alpha - \gamma_\alpha^- \rangle. \quad (2.9)$$

If P is reversible, then F is symmetric and the LHS of (2.9) is identically zero. Edges in $G(P)$ then come in pairs $\{e, e^-\}$, and for any Betti cycle $\gamma_\alpha \in \Gamma$, one of the following must be true:

- The chord e_α is such that $e_\alpha^- \in T$. Then γ_α is equal to the cycle $e_\alpha + e_\alpha^-$ so that $\gamma_\alpha = \gamma_\alpha^-$.
- The chord e_α is such that $e_\alpha^- \notin T$. Then e_α^- is itself a chord, and the corresponding Betti

cycle is γ_α^- , so that γ_α and γ_α^- are both in Γ . Furthermore,

$$0 = \langle e_\alpha, \underline{F} - \underline{F}^- \rangle = \sum_{\alpha'} \omega(\gamma_{\alpha'}) \langle e_\alpha, \gamma_{\alpha'} - \gamma_{\alpha'}^- \rangle = \omega(\gamma_\alpha) - \omega(\gamma_\alpha^-)$$

so that $\omega(\gamma_\alpha) = \omega(\gamma_\alpha^-)$.

As a consequence, we have $\Gamma = \Gamma^-$ and $\omega(\gamma_\alpha) = \omega(\gamma_\alpha^-)$ for all $\gamma_\alpha \in \Gamma$ if P is reversible. The converse follows directly from (2.9). On the other hand, from (2.6) we get

$$\mu_x p_{xy} - \mu_y p_{yx} = \sum_{c \in \mathcal{C}_\infty} \omega_\infty(c) [J_c(x, y) - J_{c^-}(y, x)].$$

If we extend the function ω_∞ to $\mathcal{C}_\infty \cup \mathcal{C}_\infty^-$ by putting $\omega(c^-) = 0$ for all $c^- \in \mathcal{C}_\infty^- \setminus \mathcal{C}_\infty$, then we can write the last equation as

$$\mu_x p_{xy} - \mu_y p_{yx} = \sum_{c \in \mathcal{C}_\infty} [\omega_\infty(c) - \omega_\infty(c^-)] J_c(x, y). \quad (2.10)$$

If P is reversible, we can verify that $\omega_\infty(c) = \omega_\infty(c^-)$ holds for all $c \in \mathcal{C}_\infty$ by inspecting (2.8): For $c = [x_1, \dots, x_s]$, we write $\mathbf{P}(c) = p_{x_1 x_2} p_{x_2 x_3} \dots p_{x_{s-1} x_s}$ for the product of the transition probabilities along the edges of c . Then

$$\frac{\omega_\infty(c)}{\omega_\infty(c^-)} = \frac{p_{x_1 x_2} p_{x_2 x_3} \dots p_{x_{s-1} x_s} p_{x_s x_1}}{p_{x_s x_{s-1}} p_{x_{s-1} x_{s-2}} \dots p_{x_2 x_1} p_{x_1 x_s}} = \frac{\mathbf{P}(c)}{\mathbf{P}(c^-)} \quad (2.11)$$

which equals one by detailed balance if P is reversible. On the other hand, if $\omega_\infty(c) = \omega_\infty(c^-) \forall c \in \mathcal{C}_\infty$, then P is reversible since the RHS of (2.10) is identically zero. In summary, the following conditions are all equivalent characterizations of the reversibility of P :

1. Detailed balance holds: $\mu_x p_{xy} = \mu_y p_{yx}$ for all $x, y \in \mathbb{X}$.
2. The matrix $F = D_\mu P$ is symmetric.
3. Iff Γ is a Betti basis of $G(P)$, then $\Gamma = \Gamma^-$ and $\omega(\gamma_\alpha) = \omega(\gamma_\alpha^-)$ for all $\gamma_\alpha \in \Gamma$.
4. The circulation distribution is symmetric under time reversal: $\omega_\infty(c) = \omega_\infty(c^-) \forall c \in \mathcal{C}_\infty$.

2.1.5 Cycle Decomposition and Entropy Production

Recall the definition (1.10) of the entropy production rate functional e_P in section 1.1.1. For Markov chains with transition matrix P and invariant distribution μ , we can write the entropy production rate as

$$e_P = \frac{1}{2} \sum_{x, y \in \mathbb{X}} (\mu_x p_{xy} - \mu_y p_{yx}) \log \left(\frac{\mu_x p_{xy}}{\mu_y p_{yx}} \right) \quad (2.12)$$

Chapter 2. Irreversible Markov Chains and Cycles

if the **weak reversibility condition** $\mu_x p_{xy} > 0 \Leftrightarrow \mu_y p_{yx} > 0 \forall x, y \in \mathbb{X}$ is satisfied (otherwise $e_P = +\infty$ by definition). Qian [JQQ04] has shown that e_P can also be expressed in terms of the circulation distribution ω_∞ as

$$e_P = \frac{1}{2} \sum_{c \in \mathcal{C}_\infty} (\omega_\infty(c) - \omega_\infty(c^-)) \log \left(\frac{\omega_\infty(c)}{\omega_\infty(c^-)} \right) \quad (2.13)$$

which shows that e_P measures the degree to which ω_∞ is antisymmetric under orientation reversal $c \mapsto c^-$. In terms of the Betti basis, we can express e_P as

$$e_P = \sum_{\alpha=1}^b \omega(\gamma_\alpha) \log \left(\frac{\mathbf{P}(\gamma_\alpha)}{\mathbf{P}(\gamma_\alpha^-)} \right). \quad (2.14)$$

Equation (2.14) shows that for $e_P = 0$ it is enough that $\omega_\infty(\gamma_\alpha) = \omega_\infty(\gamma_\alpha^-)$ for all Betti cycles γ_α , we do not need $\omega_\infty(c) = \omega_\infty(c^-) \forall c \in \mathcal{C}_\infty$. We prove equations (2.12), (2.13) and (2.14) in appendix A.

Example: Biased RW on a circle. As an example, we consider a biased random walk on a circle with n nodes. The transition probabilities of the RW are given by $p_{x,x\pm 1} = \frac{1}{2}(1 \pm a)$ with $0 \leq a \leq 1$. The Coates graph $G(P)$ is the graph with n nodes and $2n$ edges shown in Figure 2.3. The invariant distribution is $\mu = \frac{1}{n}(1, \dots, 1)^T$. By using (2.12) one readily checks that e_P only depends on a and is given by

$$e_P(a) = a \log \left(\frac{1+a}{1-a} \right) \quad (2.15)$$

with the extremal cases $e_P(0) = 0$ (the reversible case) and $e_P(1) = +\infty$ (where weak non-reversibility fails). We now check (2.14). If we pick the spanning tree T shown in Figure 2.3, then Γ consists of two cycles γ_r and $\gamma_l = \gamma_r^-$ of length n with corresponding chords e_r and e_l and $n-1$ cycles of length 2. The cycles of length 2 do not contribute to (2.14), therefore

$$e_P = \omega(\gamma_r) \log \frac{\mathbf{P}(\gamma_r)}{\mathbf{P}(\gamma_l)} + \omega(\gamma_l) \log \frac{\mathbf{P}(\gamma_l)}{\mathbf{P}(\gamma_r)}.$$

From Lemma 2.3 we see that the weights are given by $\omega(\gamma_r) = \frac{1}{2n}(1+a)$ and $\omega(\gamma_l) = \frac{1}{2n}(1-a)$. We furthermore have $\mathbf{P}(\gamma_r) = \frac{1}{2n}(1+a)^n$ and $\mathbf{P}(\gamma_l) = \frac{1}{2n}(1-a)^n$. Plugging everything in, we get

$$e_P = \frac{1}{2n}(1+a) \log \frac{(1+a)^n}{(1-a)^n} + \frac{1}{2n}(1-a) \log \frac{(1-a)^n}{(1+a)^n} = a \log \frac{1+a}{1-a}$$

which agrees with (2.15). Now we check (2.13). The class \mathcal{C}_∞ for this graph consists of γ_r, γ_l and n cycles of length 2 which do not contribute to (2.13). Hence

$$e_P = [\omega_\infty(\gamma_r) - \omega_\infty(\gamma_l)] \log \frac{\omega_\infty(\gamma_r)}{\omega_\infty(\gamma_l)}.$$

We now need to compute $\omega_\infty(\gamma_r)$ and $\omega_\infty(\gamma_l)$ according to (2.8). We have $D(\{x_1, \dots, x_n\}) = D(\mathbb{X}) = 1$ by definition, hence $\omega_\infty(\gamma_{lr}) = \mathbf{P}(\gamma_{lr})Z^{-1}$ with $Z = \sum_x D(\{x\})$. Because of symmetry, the minors $D(\{x\})$ are all equal, and because $D = I - P$ is a tridiagonal matrix they can be computed easily. One obtains

$$D(\{x\}) = \frac{(1+a)^n - (1-a)^n}{2^n a} = \frac{\mathbf{P}(\gamma_r) - \mathbf{P}(\gamma_l)}{a}$$

and putting everything together:

$$e_P = \frac{1}{nD(\{x\})} [\mathbf{P}(\gamma_r) - \mathbf{P}(\gamma_l)] \log \frac{\mathbf{P}(\gamma_r)}{\mathbf{P}(\gamma_l)} = \frac{a}{n} \log \frac{(1+a)^n}{(1-a)^n} = a \log \frac{1+a}{1-a}$$

which also agrees with (2.15).




Figure 2.3 – Coates graph of the biased random walk on a cycle with n states. The only nontrivial cycles are produced by the chords e_r and e_l .

2.2 Reversible Surrogates

When studying an irreversible process $(X_t)_t$, it is often natural to ask for some other reversible process $(X_t^s)_t$ that approximates $(X_t)_t$. Of course one wishes as many structural properties as possible to be inherited by $(X_t^s)_t$. In the context of Markov processes, we call $(X_t^s)_t$ a **reversible surrogate** of $(X_t)_t$ if $(X_t^s)_t$ is a reversible Markov process on the same state space and with the same invariant distribution as $(X_t)_t$. If $(X_t)_t$ is a Markov chain with transition matrix P , then the transition matrix P^s of $(X_t^s)_t$ has to satisfy $\mu_x p_{xy}^s = \mu_y p_{yx}^s$ for all $x, y \in \mathbb{X}$, where μ is the invariant distribution of P .

We can reformulate this in terms of the flow $F_{xy} = \mu_x p_{xy}$ as follows: A reversible surrogate of F is a flow F^s that

(a) is symmetric, $F^s = (F^s)^T$,

(b') has the same out-flow as F at every vertex x : $\sum_y F_{xy}^s = \mu_x = \sum_y F_{xy} \forall x \in \mathbb{X}$.

It is clear from section 2.1 that if F is decomposed into cycles according to (2.3) or (2.6), then the irreversibility of F is the degree to which either the Betti weights $\omega(\gamma_\alpha)$ or the circulation distribution $\omega_\infty(c)$ are not symmetric under orientation reversal of γ_α resp. c . If we have a strategy for curing this asymmetry for each cycle independently, then we have a general strategy to construct F^s . In the next section, we examine possible strategies to symmetrize flows cycle by cycle. This is used in section 2.2.2 to discuss reversible surrogates of general Markov chains.

2.2.1 Reversible Surrogates of Unit Cycle Flows

Let $c = [x_1, \dots, x_s]$ be a simple cycle with associated one-cycle \underline{c} , and let $F = \underline{c}$ be a flow³ of unit one along c . Note that F is invariant under the right-shift $R: [x_1, \dots, x_s] \mapsto [x_2, \dots, x_s, x_1]$. We now pose the following problem: Find a flow F^s on $\mathbb{X} = \{x_1, \dots, x_s\}$ that has the following properties:

- (a) F^s is symmetric, $F^s = (F^s)^T$.
- (b) $\sum_y F_{xy}^s = 1$ for all $x \in \mathbb{X}$.
- (c) F^s is invariant under R .

Here (a) and (b) correspond to (a') and (b'), and we require the additional property (c) because F also possesses (c) and we want F^s to inherit as many structural properties as possible. F^s also has a cycle decomposition, and the cycles making up F^s can be given by any sequence of the form $[y_1, \dots, y_l]$ with $l \leq s$ and pairwise different $y_i \in \{x_1, \dots, x_s\}$. Call such a sequence $[y_1, \dots, y_l]$ a **motif**. An example of a motif and the action of R on it are shown in Figure 2.4. We can use any motif $[y_1, \dots, y_l]$ to construct a flow F^s that satisfies (a)-(c) in the following way:

Choose $1 \leq l \leq s$ and a motif $m = [y_1, \dots, y_l]$. Average with respect to R :

$$\bar{m} = m + Rm + \dots + R^{s-1}m. \quad (2.16)$$

Then we let F^s to be a flow of unit $1/2l$ through \bar{m} and the reverse \bar{m}^- :

$$F^s[m] = \frac{1}{2l} [\bar{m} + \bar{m}^-]. \quad (2.17)$$

That F^s as constructed in (2.17) is symmetric and invariant under R is clear by construction. That leaves to check (b). Pick any $x \in \mathbb{X}$. As we rotate the motif m with R , every node y_i in m will coincide with x once. Therefore, l out of the s terms in (2.16) contribute to the out-flow of

³We use the identification of flows and one-cycles explained in section 2.1.2.

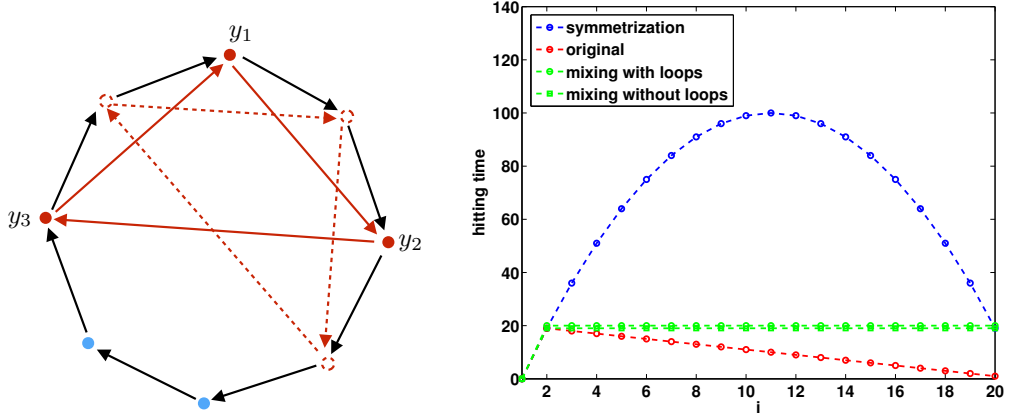


Figure 2.4 – Left: The cycle $c = [x_1, \dots, x_s]$ (black edges), the motif $m = [y_1, y_2, y_3]$ (red edges) and the rotated motif Rm (red dotted edges). Right: Hitting time signatures of different reversible surrogates.

x and we get $\sum_y \bar{m}_{xy} = l$. The same applies for \bar{m}^- , and we get

$$\sum_y F_{xy}^s[m] = \frac{1}{2l} [l + l] = 1$$

which is (b). Thus indeed, a reversible surrogate F^s can be constructed from any motif $m = [y_1, \dots, y_l]$ via (2.17). There are very many different motifs (though some might lead to the same $F^s[m]$): After averaging with R , there is one motif for $l = 1$, $(s - 1)/2$ motifs for $l = 2$ and $(s - 1)!$ motifs for $l = s$, one of which coincides with the original cycle $c = [x_1, \dots, x_s]$. Motifs can also be combined to give even more ways to construct F^s . However, some choices are clearly simpler than others. Table 2.2 shows the simplest motifs one might choose and the corresponding flow F^s .

Clearly the simplest choice is $m = [x_i]$, but the resulting flow $F^s = I$ is trivial and the Coates graph $G(F^s)$ is disconnected, so we have to discard that option. Another very simple choice is $m = [x_1, \dots, x_s]$. This amounts to $F^s = \frac{1}{2}(F + F^T)$, i.e. standard **symmetrization**. The next simplest choice is to use some motifs of length $l = 2$. There are $(s - 1)/2$ different motifs of length 2, corresponding to diagonals of different length in $G(F^s)$. The simplest choice is to use all of them, this leads to $F_{xy}^s = \frac{1}{s-1}(1 - \delta_{xy})$ and to $G(F^s)$ being the complete graph on s nodes without selfloops with each edge carrying a flow of $1/(s - 1)$ such that the total out-flow of any node is 1. We refer to this as **mixing without selfloops**. If one includes all motifs of length 1 and 2, then $F_{xy}^s = \frac{1}{s}$ and $G(F^s)$ is the complete graph on s nodes with selfloops with each edge carrying a flow of $1/s$. This is referred to as **mixing with selfloops**.

At first glance, it seems that mixing changes the original flow F much more dramatically than symmetrization, but this is deceiving. Since we are interested in dynamical properties of the original, non-reversible process $(X_t)_t$, we consider the hitting time $h(x) = \mathbf{E}_x[\tau_{x_1}]$ of

Chapter 2. Irreversible Markov Chains and Cycles

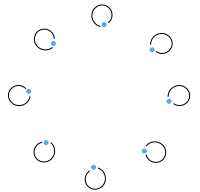
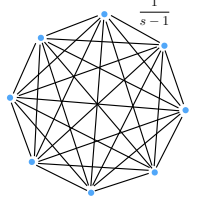
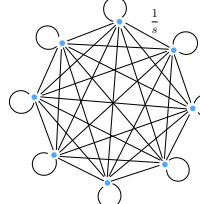
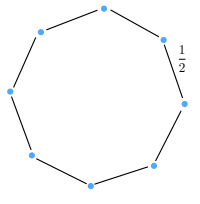
strategy	motifs taken	F^s	Coates graph of F^s
static	$[x_i]$	$F^s = I$	
mixing without selfloops	all $\frac{s-1}{2}$ different $[x_i, x_j]$	$F_{xy}^s = \frac{1}{s-1}(1 - \delta_{xy})$	
mixing with selfloops	$[x_i]$ and all $\frac{s-1}{2}$ different $[x_i, x_j]$	$F_{xy}^s = \frac{1}{s}$	
symmetrization	$[x_1, \dots, x_s]$	$F^s = \frac{1}{2}(F + F^T)$	

Table 2.2 – Some motifs and the corresponding flows F^s . For simplicity, edge directions are not shown in the Coates graph (all edges appear in pairs of opposite directions).

state x_1 as a dynamical signature. In Figure 2.4 on the right, $h(x)$ is shown for $s = 20$ for the original process and the reversible surrogates obtained by symmetrization, mixing with and mixing without self-loops. The hitting time of the original process is simply $h(x_i) = s - i + 1$ for $i \neq 1$. For mixing without self-loops, $h(x_i) = s - 1$ for $i \neq 1$ and for mixing with self-loops, $h(x_i) = s$ for $i \neq 1$. This is still very close to the hitting times of the original process, and more importantly it has the correct linear scaling with s . In contrast, the hitting time of the reversible surrogate obtained by symmetrization is completely off. It is given by the parabola $h(x_i) = \frac{1}{4}s^2 - (i - 1 - \frac{s}{2})^2$ for $i \neq 1$, which is much larger than the hitting time of the original process, and more importantly it scales quadratically with s , while the hitting time of the original process scales linearly with s .

What can one learn from this for general Markov processes? The hitting times are dynamical signatures of what we want to call the **mobility** of a stochastic process, i.e. its ability to reach other states quickly. Strong cyclic flows strongly increase mobility locally, i.e. between the states along the cycle. This is an inherent feature of non-reversible dynamics that we have to reproduce if we want to construct a reversible surrogate. We have seen that for one cycle, the

reversible surrogate obtained by symmetrization is far less mobile. This is in fact a general result that holds for arbitrary non-reversible Markov chains, we discuss this in section 2.2.3. The reversible surrogates constructed by the mixing with and without self-loops strategies emulate the strong mobility of the non-reversible process by adding many new edges in the interior of c . This changes the structure of $G(P)$ significantly and it will not be a faithful representation of all the dynamical properties of the original irreversible process (this is likely to be impossible), but both strategies do give faithful representations of the hitting times we are interested in. In section 2.2.2, we will discuss how reversible surrogates of general Markov chains can be constructed using cycle decompositions and the mixing with self-loops strategy for each cycle. In section 2.2.3, we discuss the relationship between certain hitting times of $(X_t)_t$ and $(X_t)_t^s$ for the symmetrization strategy in more detail.

2.2.2 Reversible Surrogates of General Markov Chains

In section 2.2.1 we discussed different symmetrization strategies for a unit flow $F = \underline{c}$ along a cycle c . Here, we will discuss the 'mixing with selfloops' strategy for general Markov processes. For a single loop, it consisted of replacing F with F^s given by $F^s(x, y) = \frac{1}{|c|} J_c(x) J_c(y)$ where $J_c(x)$ is the passage function introduced in (2.4) and $|c|$ denotes the **length** of the cycle c , i.e. the number of non-zero coefficients s_i in $c = \sum_i s_i e_i$. Now let $F: \mathbb{X} \times \mathbb{X} \rightarrow [0, \infty)$ be a general divergence-free flow. By (2.3) we can write $F = \sum_{\alpha=1}^b \omega(\gamma_\alpha) \gamma_\alpha$ as a linear combination of unit flows along Betti cycles γ_α . Applying mixing with selfloops to every one of them gives

$$F_b^s(x, y) = \sum_{\alpha=1}^b \frac{\omega(\gamma_\alpha)}{|\gamma_\alpha|} J_{\gamma_\alpha}(x) J_{\gamma_\alpha}(y) \quad \forall x, y \in \mathbb{X}, \quad (2.18)$$

the index b denotes that F_b^s is constructed by using the Betti basis (as such, F_b^s is basis-dependent). On the other hand, if $F(x, y) = \mu_x p_{xy}$ then we can use (2.6) to write F as a linear combination of unit flows along cycles $c \in \mathcal{C}_\infty$, and applying mixing with selfloops to every $c \in \mathcal{C}_\infty$ gives

$$F_\infty^s(x, y) = \sum_{c \in \mathcal{C}_\infty} \frac{\omega_\infty(c)}{|c|} J_c(x) J_c(y) \quad \forall x, y \in \mathbb{X}, \quad (2.19)$$

the index ∞ denotes that the circulation distribution ω_∞ is used to construct F_∞^s . $F_\infty^s(x, y)$ enjoys the additional property that, due to (2.7) and $\sum_y J_c(y) = |c|$,

$$\sum_y F_\infty^s(x, y) = \sum_{c \in \mathcal{C}_\infty} \omega_\infty(c) J_c(x) = \mu_x \quad \forall x \in \mathbb{X}. \quad (2.20)$$

This enables us to write $F_\infty^s(x, y) = \mu_x P_\infty^s(x, y)$ where the **loop transition matrix** P_∞^s is given by

$$P_\infty^s(x, y) = \frac{1}{\mu_x} \sum_{c \in \mathcal{C}_\infty} \frac{\omega_\infty(c)}{|c|} J_c(x) J_c(y) \quad \forall x, y \in \mathbb{X}. \quad (2.21)$$

Chapter 2. Irreversible Markov Chains and Cycles

By construction, P_∞^s is reversible with respect to the stationary distribution μ , thus P_∞^s is a reversible surrogate of the non-reversible transition matrix P . This is a special feature of the stochastic cycle decomposition (2.6): From (2.18) we obtain

$$\tilde{\mu}_x := \sum_y F_b^s(x, y) = \sum_{\alpha=1}^b \omega(\gamma_\alpha) J_{\gamma_\alpha}(x) \quad \forall x \in \mathbb{X} \quad (2.22)$$

similar to (2.20), enabling us to write $F_b^s(x, y) = \tilde{\mu}_x P_b^s(x, y)$ with a reversible transition matrix P_b^s . On the other hand, the decomposition of $F(x, y) = \mu_x p_{xy}$ according to (2.2) leads to

$$\mu_x = \sum_y \sum_{\alpha=1}^b \omega(\gamma_\alpha) \langle e_{xy}, \gamma_\alpha \rangle = \sum_{\alpha=1}^b \omega(\gamma_\alpha) J_{\gamma_\alpha}(x) s_{\gamma_\alpha}(x) \quad (2.23)$$

where $s_{\gamma_\alpha}(x) \in \{-1, 0, 1\}$, as follows: If γ_α passes through $x \in \mathbb{X}$ then there must be exactly two edges e_1 and e_2 with $\langle e_{1/2}, \gamma_\alpha \rangle \neq 0$ that are incident to x . The value of s_{γ_α} depends on the relative orientations of these edges and γ_α :

$$s_{\gamma_\alpha}(x) = \begin{cases} 1 & \text{if } \langle e_1, \gamma_\alpha \rangle = \langle e_2, \gamma_\alpha \rangle = 1, \\ -1 & \text{if } \langle e_1, \gamma_\alpha \rangle = \langle e_2, \gamma_\alpha \rangle = -1, \\ 0 & \text{if } \langle e_1, \gamma_\alpha \rangle = -\langle e_2, \gamma_\alpha \rangle = \pm 1. \end{cases} \quad (2.24)$$

In other words $s_{\gamma_\alpha}(x) = \pm 1$ if e_1 is incoming to x , e_2 is outgoing from x and their orientations agree (disagree) with γ_α , and $s_{\gamma_\alpha}(x) = 0$ if e_1 and e_2 are either both outgoing from or incoming to x . By comparing (2.22) with (2.23), we see that the invariant distribution $\tilde{\mu}$ of P_b^s is in general different from the invariant distribution μ of P . In other words, P_b^s is not a reversible surrogate in the sense that was discussed in section 2.2.1, but P_∞^s is. This is the disadvantage of the algebraic cycle decomposition (2.2) compared to the stochastic one (2.6).

2.2.3 Reversible Processes are Always Slower than Irreversible Processes

In this section we compare an irreversible ergodic Markov process $(X_t)_t$ with generator L and invariant distribution μ with the corresponding reversible process $(X_t^s)_t$ with generator $L^s = \frac{1}{2}(L + L^-)$, which also has the stationary distribution μ , and we assume that $(X_t^s)_t$ is also ergodic. There are several results in the literature indicating that $(X_t^s)_t$ is always 'slower' than $(X_t)_t$: For diffusions, it was shown in [RS14] that the asymptotic variance of MCMC estimators based on $(X_t)_t$ is strictly smaller when compared to those based on $(X_t^s)_t$ for a large class of observables f , provided a non-degeneracy condition on the added irreversible drift holds. In [HHMS05] it was shown that under similar non-degeneracy conditions, the absolute value of the second-largest eigenvalue of L is strictly larger than the absolute value of the second-largest eigenvalue of L^s , indicating a faster convergence rate towards μ . For Markov chains, similar results on the asymptotic variance were given in [SSG10], and the irreversible terms added to L^s were characterized in terms of additional flows along cycles in $G(L^s)$.

The result we give below complements these findings in terms of **commuting times**. First we will need some language from potential theory. We assume that \mathbb{X} is finite. Let $A, B \subset \mathbb{X}$ be disjoint and let

$$\tau_A = \inf\{t > 0 : X_t \in A\}, \quad \tau_B = \inf\{t > 0 : X_t \in B\}$$

be the first hitting times (after time $t = 0$) of A and B . We define the **equilibrium measure**

$$e_{AB}(x) = \begin{cases} \mathbf{P}_x[\tau_A > \tau_B] & \text{for } x \in A, \\ 0 & \text{otherwise} \end{cases} \quad (2.25)$$

which is the probability of escaping A from the point $x \in A$, and the **capacity** of A and B

$$\text{cap}(A, B) = \sum_{x \in A} \mu(x) e_{AB}(x) \quad (2.26)$$

which is the total rate of escape to B if starting from A . It is not hard to see [BEGK02] that $e_{AB}(x) = (Lq^+)(x)$ for $x \in A$ with the forward committor q^+ as in (1.47), and that $\text{cap}(A, B)$ equals the reaction rate k_{AB} that was defined in (1.55), thus potential theory uses the same objects as TPT. Following [Slo13], we establish a relation between first hitting times, capacities and the committor function via the **last exit biased distribution** on A , which we define as

$$\nu_{AB}(x) := \frac{\mu(x) e_{AB}^-(x)}{\text{cap}(A, B)}$$

with $e_{AB}^-(x) = \mathbf{P}_x^-[\tau_A > \tau_B] = (-L^- q^-)(x)$ being the equilibrium potential of the reversed process⁴. ν_{AB} is the distribution of the points on ∂A that $(X_t^-)_t$ visits last before escaping to B , or equivalently the distribution of points on ∂A that $(X_t)_t$ visits first when coming from B . Now we define the mean hitting time of B when starting in A and the **commuting time** $T(A, B)$ between A and B as

$$\mathbf{E}_A[\tau_B] = \sum_{x \in A} \nu_{AB}(x) \mathbf{E}_x[\tau_B], \quad T(A, B) = \mathbf{E}_A[\tau_B] + \mathbf{E}_B[\tau_A]. \quad (2.27)$$

The distribution ν_{AB} is concentrated on the boundary of A . Defining $\mathbf{E}_A[\tau_B]$ by averaging $\mathbf{E}_x[\tau_B]$ with respect to ν_{AB} leads to the following intuitive definition of $T(A, B)$: Given a realization $(X_t)_{0 < t \leq T}$, define a discrete process $(\hat{X}_t)_{0 < t \leq T}$ taking values in $\{a, b\}$ by the following set of rules: (i) $(\hat{X}_t)_t$ only changes when $(X_t)_t$ enters $A \cup B$, (ii) when $X_t \in A$ then $\hat{X}_t = a$, (iii) when $X_t \in B$ then $\hat{X}_t = b$. See also Figure 2.5. If we let N_{AB}^T be the number of times $(\hat{X}_t)_t$ switches from a to b in time T , then a.s.

$$T(A, B) = \lim_{T \rightarrow \infty} \frac{T}{N_{AB}^T}.$$

⁴The extra minus sign in $e_{AB}^-(x) = (-L^- q^-)(x)$ results from the fact that the boundary conditions for q^+ and q^- are flipped.




Figure 2.5 – Visualization of the discrete process $(\hat{X}_t)_t$: The trajectory of $(X_t)_t$ is green if $\hat{X}_t = a$ and red if $\hat{X}_t = b$. The support of the distributions ν_{AB} and ν_{BA} is illustrated in blue.

Also note that in the most interesting case where A is a metastable set, $\mathbf{E}_x[\tau_B]$ is almost constant on A and it does not matter how we average it over A . With our choice of ν_{AB} , we have the following relation between the hitting time $\mathbf{E}_A[\tau_B]$, the capacity $\text{cap}(A, B)$ and the backward committor q^- according to [Slo13]:

Lemma 2.5 *The mean hitting time $\mathbf{E}_A[\tau_B]$ of B when starting in A is given by*

$$\mathbf{E}_A[\tau_B] = \frac{1}{\text{cap}(A, B)} \sum_{x \notin B} \mu(x) q^-(x). \quad (2.28)$$

As a consequence,

$$T(A, B) = \frac{1}{\text{cap}(A, B)}. \quad (2.29)$$

We prove Lemma 2.5 in appendix A. The main result of this section is the following theorem:

Theorem 2.6 *Let L be the generator of an ergodic Markov process on \mathbb{X} and let $L^s = \frac{1}{2}(L + L^-)$. Let $A, B \subset \mathbb{X}$ be disjoint and let q be the forward committor of L^s with source A and target B . Denote by $\text{cap}^s(A, B)$ and $T^s(A, B)$ the capacity and commuting time of L^s , respectively. Then*

$$\text{cap}(A, B) \geq \text{cap}^s(A, B) \quad \text{and} \quad T(A, B) \leq T^s(A, B). \quad (2.30)$$

The inequalities are strict unless $Lq - L^s q = 0$.

According to Theorem 2.6, the commuting time between any two sets A and B is always smaller for the non-reversible process $(X_t)_t$ when compared to its reversible counterpart $(X_t^s)_t$. This can also be read the other way around: If one adds a non-reversible perturbation $\delta L = D_\mu^{-1} H$ to L^s such that $H^T = -H$ is an antisymmetric matrix⁵, then $T(A, B)$ always decreases and $\text{cap}(A, B)$ always increases. We prove Theorem 2.6 in appendix A using variational formulas

⁵This guarantees that $L = L^s + D_\mu^{-1} H$ satisfies $L^s = \frac{1}{2}(L + L^-)$.

for the capacities. For non-reversible processes, we have the following saddle point formula for $\text{cap}(A, B)$ [Slo13]:

$$\text{cap}(A, B) = \inf_{f \in \mathcal{H}_{AB}} \sup_{h \in \mathcal{G}_{AB}} \{-2\langle L^- f, h \rangle_\mu - \langle h, -L^s h \rangle_\mu\} \quad (2.31)$$

where $\mathcal{G}_{AB} = \{h : \mathbb{X} \rightarrow \mathbb{R} : h_A = \text{const}, h_B = \text{const}\}$ and $\mathcal{H}_{AB} = \{h : \mathbb{X} \rightarrow [0, 1] : h_A = 1, h_B = 0\}$. For reversible processes, we have the simpler variational formula [Slo13, BEGK02]

$$\text{cap}^s(A, B) = \inf_{f \in \mathcal{H}_{AB}} \langle f, -L^s f \rangle_\mu. \quad (2.32)$$

The infimum in (2.32) is attained for $f = q^+$. The infimum in (2.31) is attained for $f = \frac{1}{2}(q^+ + 1 - q^-)$.

2.3 Application I: Modules in Directed Networks

In the previous sections we studied graph-theoretical properties of the Coates graph $G(P)$ of a Markov transition matrix P . This can also be turned around: Starting from a graph G , one may look for a suitable transition matrix P or generator L such that $G(P)$ resp. $G(L)$ equals G . This defines a Markov process $(X_t)_t$ on the graph G ; in this context one often says that $(X_t)_t$ is a **random walk** on G . For example, if G is represented by an **adjacency matrix** A such that $A_{xy} = 1$ if the edge (x, y) is present in G and $A_{xy} = 0$ otherwise, then we may define the transition matrix P of the simple random walk on G by setting

$$p_{xy} = \frac{A_{xy}}{d_x} \quad \forall x, y \in \mathbb{X} \quad (2.33)$$

where $d_x = \sum_y A_{xy}$ is the out-degree of $x \in \mathbb{X}$. The main idea behind the study of random walks on graphs is that one may infer **structural** properties of G by studying **dynamical** properties of $(X_t)_t$. This connection is well established in the case of undirected networks, i.e. where A is symmetric. In this case P is reversible and one knows that the stationary distribution of (2.33) is given by the out-degree $\mu_x = Z^{-1} d_x$ with $Z = \sum_x d_x$. Further, one knows that if there is a subset $C \subset \mathbb{X}$ such that nodes in C have many more connections to other nodes in C than to nodes outside of C , then C is a metastable set of (2.33). In fact, one can check that

$$\mathbf{P}(X_1 \in C | X_0 \in C) = \frac{1}{\mathbf{P}(X_0 \in C)} \sum_{x, y \in C} \frac{1}{Z} A_{xy} = 1 - \frac{\sum_{x \in C, y \notin C} A_{xy}}{\sum_{x \in C, y \in \mathbb{X}} A_{xy}}. \quad (2.34)$$

If there are many more connections within C than between C and $\mathbb{X} \setminus C$, then this number will be close to one. Such densely connected subsets C are called **modules**, and their identification is of great interest in the network literature. An algorithm that finds an incomplete partition $\{C_1, \dots, C_m\}$ of network modules by detecting the m most metastable sets with MSM techniques was presented in [DBCS11, SCB⁺14]. The input parameter m is inferred from the spectrum of P . We refer to this algorithm as **MSM clustering**. MSM clustering and a few other popular

clustering methods are explained in more detail in appendix B.

In the case of directed networks where A is not symmetric, $(X_t)_t$ is irreversible and much less is known about properties and the detection of metastable sets of $(X_t)_t$. This problem can be circumvented by using a reversible surrogate $(X_t^s)_t$ of $(X_t)_t$, but the relationship between metastable sets of $(X_t^s)_t$ and metastable sets of $(X_t)_t$ and/or specific structures in G is often not clear. For example, if $(X_t^s)_t$ is obtained by symmetrizing $(X_t)_t$, i.e. $(X_t^s)_t$ has the transition matrix $P^s = \frac{1}{2}(P + P^-)$, then we know from section 2.2.3 that the commuting times between any two subsets $A, B \subset \mathbb{X}$ are slower for $(X_t^s)_t$ than for $(X_t)_t$, and we have already seen in section 2.2.1 that this effect can be severe. We therefore expect $(X_t^s)_t$ to have more metastable sets than $(X_t)_t$, and we expect clustering algorithms based on $(X_t^s)_t$ to show an **overpartitioning effect**. Two examples of overpartitioning are shown in Figure 2.6. On the left, the barbell graph is shown which consists of two loops of length n glued together. The loops themselves are metastable sets of $(X_t)_t$, but a typical clustering (here obtained by the infomap algorithm [DYB10]) obtained from the symmetric process $(X_t^s)_t$ partitions the loops further. These smaller subsets are metastable sets for $(X_t^s)_t$, but not for $(X_t)_t$. The effect grows as n gets larger, as we have seen in section 2.2.1. The graph in Figure 2.6 on the right has three subsets C_1, C_2 and C_3 with high link densities, and all three are metastable sets for $(X_t^s)_t$. However, C_3 is not a metastable set for the process $(X_t)_t$ which follows the link directions. In fact, $(X_t)_t$ can stay for at most 6 consecutive steps in C_3 before it is forced to go out.




Figure 2.6 – Two examples of overpartitioning. Left: Overpartitioning of large cycles in the barbell graph. Clustering produced by Infomap [DYB10]. Right: The set C_3 is a metastable set for the symmetric random walk $(X_t^s)_t$, but not for $(X_t)_t$. Clustering produced by MSM clustering of P^s .

In this section we propose to use the reversible process $(X_t^s)_t$ with transition matrix P_∞^s from (2.21) or P_b^s to detect modules in G . This work has been published in [CBS15, BC14]. We have seen in section 2.2.3 that the mixing along cycles which is performed in the construction of P_∞^s and P_b^s can help to overcome the slow mixing behaviour and hence the overpartitioning

effect of the symmetric process. Additionally, in [BC14] it has been argued that the quantity

$$\mu_x P_\infty^s(x, y) = \sum_{c \in \mathcal{C}_\infty} \frac{\omega_\infty(c)}{|c|} J_c(x) J_c(y) \quad (2.35)$$

is a measure of communication, similar to the commuting time, between x and y as seen by $(X_t)_t$. Indeed, (2.35) sums over all $c \in \mathcal{C}_\infty$ passing through x and y , which are ways to commute from x to y and then back to x . Each such commuting cycle c is weighted by $\omega_\infty(c)$ which measures how often $(X_t)_t$ completes c and thus how important c is to $(X_t)_t$ as a communication channel, and by the length $|c|$ so that short commuting cycles are preferred.

In the context of module detection however, the main question we have to answer is if there is a relation between metastable sets of $(X_t^s)_t$ and subsets of nodes in G with an interesting structure. The following lemma will help answering this question in the affirmative:

Lemma 2.7 *Let $(X_t)_t$ be a random walk on the graph G . Let ω_∞ be the circulation distribution and γ_α be the Betti cycles with weights $\omega(\gamma_\alpha)$ of $(X_t)_t$. Let $C \subset \mathbb{X}$. If $(X_t^s)_t$ is the reversible surrogate of $(X_t)_t$ with transition matrix P_∞^s (2.21), then*

$$\mathbf{P}(X_1^s \in C | X_0^s \in C) = 1 - \frac{\sum_{x \in C, y \notin C} \sum_{c \in \mathcal{C}_\infty} \frac{\omega_\infty(c)}{|c|} J_c(x) J_c(y)}{\sum_{x \in C, y \in \mathbb{X}} \sum_{c \in \mathcal{C}_\infty} \frac{\omega_\infty(c)}{|c|} J_c(x) J_c(y)}. \quad (2.36)$$

If $(X_t^s)_t$ has the transition matrix P_b^s , then

$$\mathbf{P}(X_1^s \in C | X_0^s \in C) = 1 - \frac{\sum_{x \in C, y \notin C} \sum_{\alpha=1}^b \frac{\omega(\gamma_\alpha)}{|\gamma_\alpha|} J_{\gamma_\alpha}(x) J_{\gamma_\alpha}(y)}{\sum_{x \in C, y \in \mathbb{X}} \sum_{\alpha=1}^b \frac{\omega(\gamma_\alpha)}{|\gamma_\alpha|} J_{\gamma_\alpha}(x) J_{\gamma_\alpha}(y)}. \quad (2.37)$$

We prove Lemma 2.7 in appendix A. The structural similarity between (2.34) and (2.36) resp. (2.37) is clear: The adjacency matrix A_{xy} has been replaced by $F_\infty^s(x, y) = \sum_{c \in \mathcal{C}_\infty} \frac{\omega_\infty(c)}{|c|} J_c(x) J_c(y)$ resp. $F_b^s(x, y) = \sum_{\alpha=1}^b \frac{\omega(\gamma_\alpha)}{|\gamma_\alpha|} J_{\gamma_\alpha}(x) J_{\gamma_\alpha}(y)$. However, (2.34) only holds if A is symmetric, while (2.36) and (2.37) hold for general networks G . Structurally, (2.34) compares the number of links between x and $\mathbb{X} \setminus C$ with the total number of links leaving x for any node $x \in C$, so C is metastable if it has many more internal than external links. On the other hand, (2.36) and (2.37) compare the number of cycles connecting x and $\mathbb{X} \setminus C$ with the total number of cycles at x for every $x \in C$, so a metastable set of $(X_t^s)_t$ is a set with very many internal cycles and few cycles connecting C and $\mathbb{X} \setminus C$. Cycles are weighted according to ω_∞ resp. $\omega(\gamma_\alpha)$. This becomes even more apparent with the following corollary of (2.36): Let $\mathcal{C}_\infty^{\mathbb{X} \setminus C} \subset \mathcal{C}_\infty$ be the subset of cycles $c \in \mathcal{C}_\infty$ that leave the set C . Then

$$\mathbf{P}(X_1^s \in C | X_0^s \in C) \geq 1 - \frac{\sum_{x \in C} \sum_{c \in \mathcal{C}_\infty^{\mathbb{X} \setminus C}} \omega_\infty(c) J_c(x)}{\sum_{x \in C} \sum_{c \in \mathcal{C}_\infty} \omega_\infty(c) J_c(x)}.$$

For example, the set C_3 in Figure 2.6 on the right does not have any internal cycles $c \in \mathcal{C}_\infty$.

Depending on the spanning tree chosen for the construction of the Betti basis, there are some Betti cycles γ_α which lie completely in C_3 , but their coefficients $\omega(\gamma_\alpha)$ should be small compared to the coefficients of the external cycles. We will see this in more detail in section 2.3.2.

So far we have established a relation between metastable sets of $(X_t^s)_t$, where $(X_t^s)_t$ is the Markov chain with transition matrix P_∞^s or P_b^s , and subsets of nodes of G with an interesting structure. We also have a way to detect these metastable sets with MSM clustering once P_∞^s resp. P_b^s is known. That leaves us with the task of computing P_∞^s resp. P_b^s . We discuss this in section 2.3.1 before moving to numerical examples in section 2.3.2.

2.3.1 Computational Aspects

Computing P_∞^s . A possible way to compute P_∞^s is via its definition (2.21). This requires knowledge of \mathcal{C}_∞ and ω_∞ . There is an immediate problem: As we discussed in section 2.1.3, a direct computation of \mathcal{C}_∞ and ω_∞ is NP-hard and thus numerically infeasible. A possible remedy is to approximate \mathcal{C}_∞ and ω_∞ by sampling \mathcal{C}_T and ω_T for some finite time $T > 0$; as we know from section 2.1.3 these approximations converge almost surely to \mathcal{C}_∞ resp. ω_∞ . An algorithm to sample \mathcal{C}_T and ω_T from a realization $(X_n)_{1 \leq n \leq T}$ has also been discussed in section 2.1.3. We state the algorithm again for convenience:

Algorithm 1 (sampling the cycle decomposition):

(i) Initialization: Set $\mathcal{C}_T = \emptyset$, $N_c^T = 0$ for all $c \in \mathcal{C}_0$ and $\eta_1 = X_1$.

(ii) for $i = 2$ to T :

(a) If $\eta_i = (y_1, \dots, y_s)$ and $X_{i+1} \neq y_l$ for all $1 \leq l \leq s$, then set $\eta_{i+1} := (y_1, \dots, y_s, y_{s+1})$ with $y_{s+1} := X_{i+1}$.

(b) If $\eta_i = (y_1, \dots, y_s)$ and $X_{i+1} = y_l$ for some $1 \leq l \leq s$, then the cycle $c := [y_l, \dots, y_s]$ has been completed. Set $\mathcal{C}_T \rightarrow \mathcal{C}_T \cup \{c\}$, $N_c^T \rightarrow N_c^T + 1$ and $\eta_{i+1} := (y_1, \dots, y_l)$.

Algorithm 1 requires to store a list of all the cycles that have been found, and in (b) one has to check if the newly found cycle is already in the list or not. That makes the algorithm inefficient both from the perspective of storage and computation time. It is more efficient to directly sample an approximation P_T^s of P_∞^s , which can be done by utilizing the following Lemma:

Lemma 2.8 (Sampling P_T^s) Let $P_T^s : \mathbb{R}^X \rightarrow \mathbb{R}^X$ be the matrix with components

$$P_T^s(x, y) = \frac{\tilde{N}_T(x, y)}{N_T(x)} \quad \forall x, y \in X, \quad (2.38)$$

where $\tilde{N}_T(x, y) = \sum_{c \in \mathcal{C}_T} \frac{N_c^T}{|c|} J_c(x) J_c(y)$ and $N_T(x)$ is the number of times $(X_n)_{1 \leq n \leq T}$ passes through x . Then P_T^s is a reversible transition matrix with stationary distribution $\mu_T(x) = N_T(x)/T$, and $P_T^s(x, y) \rightarrow P_\infty^s(x, y)$ a.s. as $T \rightarrow \infty$.

Proof: We first show the identity

$$\sum_y \tilde{N}_T(x, y) = \sum_{c \in \mathcal{C}_T} N_c^T J_c(x) = N_T(x), \quad (2.39)$$

which proves that P_T^s has row sum one. With $\mu_T(x) = N_T(x)/T$, (2.39) together with the fact that $\tilde{N}_T(x, y) = \tilde{N}_T(y, x) \forall x, y \in \mathbb{X}$ also shows that $\mu_T(x) P_T^s(x, y) = \mu_T(y) P_T^s(y, x) \forall x, y \in \mathbb{X}$. Thus P_T^s is reversible with stationary distribution μ_T . The first equality in (2.39) follows from $\sum_y J_c(y) = |c|$. The last equality is true for the following reason: Suppose $X_i = x$ and $X_j = x$ with $j > i$ are two successive visits of x . Then either η_{j-1} contains $X_i = x$, in which case adding $X_j = x$ triggers the completion of the cycle $c = [X_i, \dots, X_{j-1}]$ with $J_c(x) = 1$, or η_{j-1} does not contain X_i , in which case there must be a $i < j' < j$ such that $X_{j'}$ triggered the completion of a cycle c containing X_i . In both cases, two successive visits of x correspond to the completion of one cycle c with $J_c(x) = 1$.

Finally, to show that $P_T^s(x, y) \rightarrow P_\infty^s(x, y)$ as $T \rightarrow \infty$, write (2.38) out:

$$P_T^s(x, y) = \frac{T}{N_T(x)} \sum_{c \in \mathcal{C}_T} \frac{1}{T} \frac{N_c^T}{|c|} J_c(x) J_c(y) = \frac{1}{\mu_T(x)} \sum_{c \in \mathcal{C}_T} \frac{\omega_T(c)}{|c|} J_c(x) J_c(y).$$

Then $\omega_T \rightarrow \omega_\infty$ and $\mathcal{C}_T \uparrow \mathcal{C}_\infty$ a.s. by Lemma 2.4, and $\mu_T(x) \rightarrow \mu(x)$ a.s. by the ergodic theorem 1.5 (take $f = \mathbf{1}_x$). ■

Lemma 2.8 provides us with a way of sampling P_T^s , all we need to do is obtain the counts $\tilde{N}_T(x, y)$ for all $x, y \in \mathbb{X}$ from the realization $(X_n)_{1 \leq n \leq T}$. The algorithm 2 we present below does exactly that.

Algorithm 2 (sampling P_T^s):

(i) Initialization: Set $\eta_1 = X_1$ and $\tilde{N}_T(x, y) = 0$ for all $x, y \in \mathbb{X}$.

(ii) for $i = 2$ to T :

(a) If $\eta_i = (y_1, \dots, y_s)$ and $X_{i+1} \neq y_l$ for all $1 \leq l \leq s$, then set $\eta_{i+1} := (y_1, \dots, y_s, y_{s+1})$ with $y_{s+1} := X_{i+1}$.

(b) If $\eta_i = (y_1, \dots, y_s)$ and $X_{i+1} = y_l$ for some $1 \leq l \leq s$, then the cycle $c := [y_l, \dots, y_s]$ has been completed. We set $\tilde{N}_T(x, y) \rightarrow \tilde{N}_T(x, y) + \frac{1}{|c|}$ for all $x, y \in c$ and $\eta_{i+1} := (y_1, \dots, y_l)$.

Algorithm 2 is asymptotically optimal in terms of both data storage and computation time: For data storage, the only additional data besides the counts $\tilde{N}_T(x, y)$ themselves that needs to be stored is the current state of the derived chain η . Since η can contain at most $|\mathbb{X}|$ elements, the cost of storing η is $\mathcal{O}(|\mathbb{X}|)$ and is therefore negligible to the cost of storing the counts $\tilde{N}_T(x, y)$, which is $\mathcal{O}(|\mathbb{X}|^2)$ in the worst case. For the computation time let us examine the number $N(T)$ of operations performed by the algorithm as a function of T . For each $1 \leq i \leq T$, X_i has to be compared to all elements in η_i to determine if we are in (a) or (b), for a total of $\sum_{i=1}^T |\eta_i|$ operations. Every time a cycle c is completed, we are in (b) and need to perform $|c|^2$ additions. Finally, η_i needs to be updated to η_{i+1} for a total of T operations. In total,

$$\begin{aligned} N(T) &= \sum_{i=1}^T |\eta_i| + \sum_{c \in \mathcal{C}_T} N_c^T |c|^2 + T \\ &\leq |\mathbb{X}|T + |\mathbb{X}| \sum_{c \in \mathcal{C}_T} N_c^T |c| + T \\ &= (2|\mathbb{X}| + 1) T. \end{aligned}$$

The last line follows from $\sum_{c \in \mathcal{C}_T} N_c^T |c| = \sum_x \sum_{c \in \mathcal{C}_T} N_c^T J_c(x) = \sum_x N_T(x) = T$, see also (2.39). This shows that $N(T)$ is $\mathcal{O}(T)$, which is asymptotically optimal since every element of $(X_n)_{1 \leq n \leq T}$ has to be accessed at least once.

Convergence of $P_T^s \rightarrow P_\infty^s$. While Lemma 2.8 guarantees a.s. convergence $P_T^s \rightarrow P_\infty^s$ as $T \rightarrow \infty$ by Theorem 1.5, it does not give information about the rate of convergence. If the samples X_1, \dots, X_T were drawn independently, we could apply the law of large numbers to the random variables ω_T and μ_T , which would give a convergence with $1/T^{1/2}$. In practice, the correlations between the X_i induced by the Markov chain result in a convergence rate for both μ_T and ω_T that is dominated by the second-largest eigenvalue of P_∞^s [SNS10]. If P_∞^s has strong metastabilities, this convergence rate can be slow. If the stationary distribution μ is already known, one can use e.g. $\|\mu - \mu_T\|$ where $\|\cdot\|$ is any vector norm as an indicator for convergence of P_∞^s .

Time series perspective. Networks are often constructed from data. For example, let Y_1, \dots, Y_T be a time series of observations with $Y_i \in \mathbb{X}$. After a suitable discretization, we may assume the observational space \mathbb{X} to be discrete and define the counts $N_T(x) = \sum_{i=1}^T \delta(Y_i = x)$ and $N_T(x, y) = \sum_{i=1}^{T-1} \delta(Y_i = x, Y_{i+1} = y)$ for all $x, y \in \mathbb{X}$. If we assume that the time series $(Y_i)_i$ was generated by a Markov chain on \mathbb{X} , then the Maximum-Likelihood estimator of the transition matrix of that chain is given by

$$P(x, y) = \frac{N_T(x, y)}{N_T(x)} \quad \forall x, y \in \mathbb{X}.$$

Finally, we may construct the **transition network** of $(Y_i)_i$ [DZD⁺10], which is just the Coates graph $G(P)$. With this perspective, detecting modules in G really means detecting structure in

the data $(Y_i)_i$. Then that same data can and in fact should be used as input for computing P_T^s with algorithm 2, and the inherent finiteness of $(Y_i)_i$ removes the issue of convergence of P_T^s to some P_∞^s since P_∞^s is inaccessible and the observational window $[0, T]$ is dictated by the data. Since algorithm 2 is asymptotically optimal, computing P_T^s is as fast as computing P . We conclude that time series analysis is the most natural use case for module detection with P_T^s .

Computing P_b^s . We can compute P_b^s from its definition $P_b^s(x, y) = \tilde{\mu}_x^{-1} F_b^s(x, y)$ with $\tilde{\mu}$ given by (2.22) and $F_b^s(x, y)$ given by (2.18). This requires the computation of the Betti basis $\{\gamma_1, \dots, \gamma_b\}$ and the corresponding weights $\omega(\gamma_\alpha)$. We discussed in section 2.1.2 that this can be done in a computationally efficient way. The most demanding part is to find a spanning tree T in the graph G . We use the Prim algorithm [Pri57] for this, which is $\mathcal{O}(|E| \log |X|)$. The Prim algorithm is capable (at no additional cost) of finding a **minimal spanning tree**, i.e. a tree T that minimizes $R(T) := \sum_{e \in T} r(e)$ where $r : E \rightarrow [0, \infty)$ are custom edge weights. We utilize this to find a unique spanning tree T instead of having to chose one arbitrarily. More specifically, we let $r(e)$ be the **resistance** $r(e) = (\mu_{s(e)} p_{(s(e), t(e))})^{-1}$. The minimal spanning tree returned by the Prim algorithm has the property that the total resistance of the cochords is minimized, or equivalently that the total resistance of the chords is maximized.




Figure 2.7 – Left: The barbell graph with minimal spanning tree T , chords e_1 , e_2 and e_3 and corresponding cycles γ_1 , γ_2 and γ_3 . Right: $G(P_\infty^s)$, modules detected by MSM clustering (red and blue nodes) and two sets C_1 and C_2 whose metastability index we compute (red and blue shaded regions).

2.3.2 Numerical Examples

Barbell Graph. Our first example is the barbell graph G with $2n$ nodes and uniform edge weights. Figure 2.7 shows G together with a minimal spanning tree T on the left. We have three chords e_1 , e_2 and e_3 corresponding to the Betti cycles γ_1 , γ_2 and γ_3 . The weights are all equal and given by $\omega(\gamma_1) = \omega(\gamma_2) = \omega(\gamma_3) = \frac{1}{2(n+1)}$. One readily sees that $\mathcal{C}_\infty = \{\gamma_1, \gamma_2, \gamma_3\}$ and $\omega_\infty(\gamma_i) = \omega(\gamma_i)$. Hence in the particular case here, the algebraic and stochastic cycle decompositions are exactly equal.

In Figure 2.7 on the right, $G(P_\infty^s)$ is shown. The weights correspond to $F_\infty^s(x, y)$. Not surprisingly, MSM clustering detects two modules which correspond to the clusters of red and blue nodes shown. We can compute the metastability index of e.g. the red cluster C_1 according to (2.36):

$$p_1(C_1, C_1) = \mathbf{P}(X_1^s \in C_1 | X_0^s \in C_1) = 1 - \frac{\frac{\omega(\gamma_2)}{|\gamma_2|}}{\omega(\gamma_2) + \omega(\gamma_1) \cdot n} = 1 - \frac{1}{2(n+1)}.$$

In Figure 2.8, we show the metastability index $p_t(C, C)$ as defined in Definition 1.8 for different times t and $C = C_1$ (red) and $C = C_2$ (blue). $p_t(C, C)$ is computed for the process with transition matrix P_∞^s (square markers) and for the symmetric process with transition matrix $P^s = \frac{1}{2}(P + P^-)$ (solid lines). The full cycle C_1 is very metastable for both processes. On the other hand, for the half-cycle C_2 the metastability index $p_t(C_2, C_2)$ measured by P^s decays a lot slower than $p_t(C_2, C_2)$ measured by P_∞^s .




Figure 2.8 – Metastability index for $C = C_1$ (red) and $C = C_2$ (blue), computed with P_∞^s (square markers) and $P^s = \frac{1}{2}(P + P^-)$ (solid lines).

Two or Three Modules? Our next example is the graph that already showed up in Figure 2.6. It is shown again in Figure 2.9 together with the minimal spanning tree T and two possible clusterings: One with three modules C_1, C_2 and C_3 on the left, and one with only the two modules C_1 and C_2 and a larger transition region on the right. MSM clustering of the symmetric process with transition matrix P^s produces the three modules. MSM clustering of P_∞^s and P_b^s produces the clustering with two modules. C_3 is not a module for P_∞^s because there are no internal cycles $c \in \mathcal{C}_\infty$ in C_3 , cf. (2.36). In fact every cycle passing through C_3 also has to pass

through C_1 and C_2 . C_3 is not a module for P_b^s because examining the spanning tree T reveals that most chords in C_3 correspond to Betti cycles that also pass through C_1 and C_3 , and only a single chord corresponds to a Betti cycle that lies completely in C_3 .




Figure 2.9 – Directed graph G with minimal spanning tree T (black edges). Left: Clustering with P^s produces three modules C_1 , C_2 and C_3 . Right: Clustering with P_∞^s or P_b^s produces two modules C_1 and C_2 . Grey nodes are in the transition region.

In Figure 2.10 we examine this further by studying the metastability index $p_t(C, C)$ for different times t and $C = C_1$ (red), $C = C_2$ (blue) and $C = C_3$ (green). The metastability index was computed for the simple random walk with transition matrix P given by (2.33) (solid lines), the symmetrized random walk with transition matrix $P^s = \frac{1}{2}(P + P^-)$ (round markers) and finally the random walk with transition matrix P_T^s with $T = 10^6$ (square markers, results for P_b^s are similar). The sampling error induced by the finite sampling time T was monitored by the error $\|\mu - \mu_T\|_1$ and is $\|\mu - \mu_T\|_1 \approx 0.0036$.

For all three processes, C_1 and C_2 are metastable sets. The metastability indices of P and P_∞^s are in good agreement while the metastability index of P^s is significantly higher. This is consistent with the observations in section 2.2.3. For C_3 , the metastability index of P has a sharp drop and reaches zero for $t = 6$, reflecting the fact that the simple RW is forced out of C_3 after at most 6 steps. For P_∞^s , $p_t(C_3, C_3)$ is also very small and never larger than 0.5. In contrast, $p_t(C_3, C_3)$ decays much slower for the symmetric process given by P^s , e.g. $p_6(C_3, C_3) \approx 0.45$ is still an order of magnitude larger than the equilibrium probability $\mathbf{P}_\mu(X_t \in C_3) \approx 0.04$.

Langevin Dynamics. Our next example uses the cycle decomposition to analyse a Markov State Model (MSM) of the Langevin System (1.27)

$$\begin{aligned} dX_t &= m^{-1} P_t dt \\ dP_t &= -(\nabla V(X_t) + \gamma m^{-1} P_t) dt + \sqrt{2\gamma\epsilon} dW_t \end{aligned} \quad (2.40)$$




Figure 2.10 – Metastability index for $C = C_1$ (red), $C = C_2$ (blue) and $C = C_3$ (green) computed with P (solid lines) P_∞^s (square markers) and $P^s = \frac{1}{2}(P + P^-)$ (round markers).

with one-dimensional x and p coordinates and double well potential $V(x) = (x^2 - 1)^2$. So far MSMs have been considered mostly for reversible processes, Langevin dynamics however is non-reversible. Figure 2.11 shows the phase space of this system together with some periodic orbits of the associated Hamiltonian system. In this case the process has two metastable sets around the minima $x_\pm = \pm 1$ of the potential V and $p = 0$. We set $m = 1$, $\epsilon = 0.2$ and $\gamma = 0.2$, so that the Langevin dynamics is still 'close' to the associated Hamiltonian system

$$\dot{x} = p, \quad \dot{p} = -\nabla_x V(x), \quad (2.41)$$

i.e., if we start the Langevin process in (x_0, p_0) with energy $E_0 = H(x_0, p_0) < 0.9$ then the dynamics will approximately follow the periodic orbit $H(x, p) = E_0$ (with $x < 0$ if $x_0 < 0$ and $x > 0$ if $x_0 > 0$) of the associated Hamiltonian system for some time interval of order 1. Therefore the typical transition from the vicinity of one of the wells across the energy barrier at $x = 0$ towards the other well will look as follows: first the trajectory will orbit the initial well for some period of time before it crosses the barrier and starts to orbit the target well until it finally hits the close vicinity around the respective energy minimum.

For this Langevin process MSM building is done as follows [BPN14, SS13a]: We construct a uniform box covering of the essential state space $[-1.8, 1.8] \times [-1.8, 1.8]$ where the invariant pdf is larger than the square root of the machine precision. We took square boxes $B_i, i = 1, \dots, n$ of size $\Delta x = 0.2$ and $\Delta p = 0.2$. Next, $M = 1000$ trajectories of length $t = 0.25$ were started in every




Figure 2.11 – Some periodic orbits of the system with Hamiltonian $H = \frac{1}{2}p^2 + V(x)$. The coloring is according to the energy $E_0 = H(x_0, p_0)$.

box (distributed due to μ), and the transition matrix

$$P_{ij} = \frac{m_{ij}}{M},$$

was constructed, where m_{ij} is the number of trajectories starting in B_i and ending up in B_j . The length $t = 0.25$ of the trajectories is very small compared to the expected transition time (which is larger than 100 here) and still shorter than the period of the periodic orbits of the associated Hamiltonian system. MSM [SS13a] theory tells us that the leading eigenvalues of the transition matrix P are very close approximations of the leading eigenvalues of the Langevin transfer operator and thus allow for an approximation of the transition statistics between the metastable wells.

The metastable sets C_1 and C_2 found by MSM clustering of the loop transition matrix P_∞^s are shown in Figure 2.12 on the left. They respect the rotational symmetry $(x, p) \mapsto (-x, -p)$ that the Hamiltonian system (2.41) enjoys. In Figure 2.12 on the right, we show the committor function $q_1(x) = \mathbf{P}(X_\tau \in C_1 | X_0 = x)$ corresponding to the left set C_1 . The function q_1 determines the degree to which the nodes in the transition region $\mathbb{X} \setminus (\bigcup_{i=1}^m C_i)$ are affiliated to C_1 . We observe that q_1 picks out the separatrix, i.e. the curve $E_0 = 1$ which separates orbits that are restricted to one of the minima and orbits around both minima, in the following sense: Boxes inside the separatrix either have $q_1(x) \approx 0$ or $q_1(x) \approx 1$, while boxes outside have $q_1(x) \approx 0.5$. This is in perfect agreement with the intuition that the Langevin system is still close to the Hamiltonian system. Note that detailed balance is strongly violated here: In the $p < 0$ half-plane, transitions from right to left are vastly more probable than transitions from left to right. In the $p > 0$ half-plane, the situation is the other way around. This is also reflected by q_1 in (2.12b), which is closer to one in the $p < 0$ half-plane and closer to zero in the $p > 0$ half-plane.

A clustering based on $P^S = \frac{1}{2}(P + P^-)$ could not reflect this.

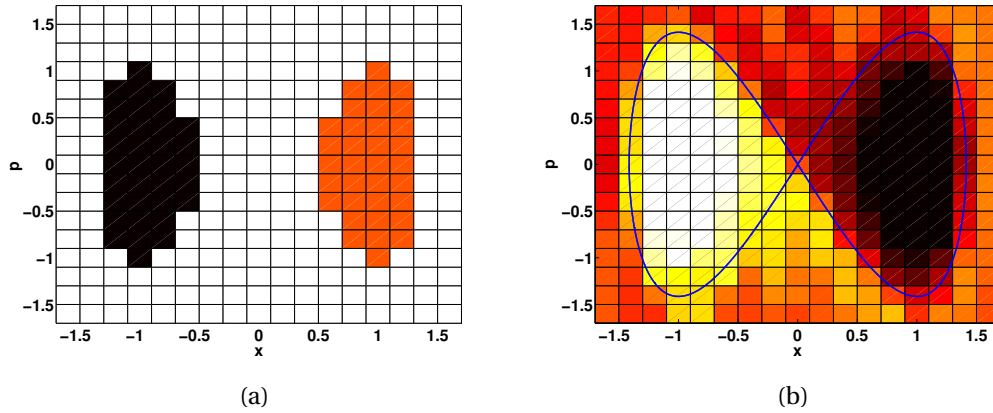


Figure 2.12 – Clustering results for the Langevin system. (a) MSM clustering for P_∞^S , cluster centers C_1 (left) and C_2 (right). (b) MSM clustering for P_∞^S , affiliation q_1 and separatrix $E_0 = 1$ (blue curve).

Timeseries example. Our final example is a time series $\{X_1, \dots, X_T\}$ of seismic events in California from 1952 to 2012, obtained from the SCEC⁶. See [AS06, AS04, DGP08, FDK13] for a discussion of several approaches to analyse seismic data using recurrence networks, including the approach based on a discretization of the observational space that we use here. Only events with magnitude larger than $m_c = 2.5$ are considered (these are 48669 events). The observational space \mathbb{X} is the rectangle from 32° to 37° in latitude and -122° to -114° in longitude, and we partition \mathbb{X} into 4000 quadratic boxes S_i of length $\Delta l = 0.1^\circ$. Finally, the boxes which don't see any events are discarded. The transition matrix thus constructed corresponds to a network with 2175 nodes and 28839 edges.

A Matlab implementation of algorithm 2 constructs P_T^S in 2.05 seconds on a laptop. MSM clustering takes 7.6 seconds. This clearly shows that performance is not an issue when algorithm 2 is used on time series data. The detected modules are shown in Figure 2.13, where a node x receives the color of module C_i if the corresponding committor function q_i satisfies $q_i(x) \geq 0.8$, and is colored grey if $q_i(x) < 0.8$ for all modules C_i . In fact the latter is the case for 80% of the nodes, but these correspond to only 25% of all events. This illustrates that the incomplete partition constructed here correctly reflects the uncertainty coming from limited data. A full partitioning, which is obtained by most of the available clustering algorithms, e.g. Markov Stability [SDYB12] or Infomap [DYB10], would have to cluster the grey nodes as well, even though not enough data is available to do so. 9 modules are found, all of which correspond to important faults or groups of faults, the largest one containing the San Andreas fault. This demonstrates that this method can successfully uncover structure in the dataset $\{X_1, \dots, X_T\}$ - in this case, the presence of geological faults that influence the earthquake pattern.

⁶Southern California Earthquake Center, www.scec.org

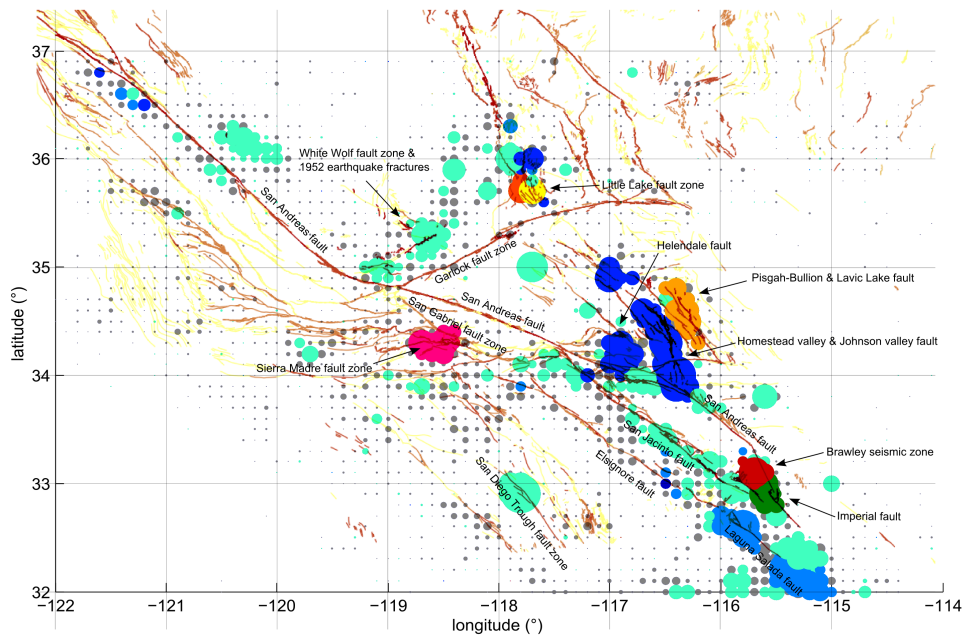


Figure 2.13 – Quaternary faults [USG] in Southern California and the clustering of the SCEC time series found by CSM. Node size is proportional to the number of events, color indicates the modules found.

2.4 Application II: Transition Pathways in Irreversible Markov Processes

In section 2.3, we studied how metastable sets of a non-reversible Markov process can be found. In this section, we are interested in the transition mechanism between sets. So let $A, B \subset \mathbb{X}$ be two disjoint sets. We think of A and B as dynamically separated metastable sets, but this is not necessary. The transition mechanism from A to B is completely described by the TPT current of reactive trajectories J_{AB} (or f_{AB} in the discrete case) that was introduced in section 1.3. We will see below that for reversible processes, J_{AB} is a gradient vector field with the forward committor function q^+ as potential. This greatly simplifies the study of the transition mechanism. For non-reversible processes, J_{AB} also has rotational parts and it makes sense to ask for a decomposition of J_{AB} into a part with a gradient structure and a purely rotational part. The gradient part may then be studied in isolation to understand the transition mechanism. We discuss two approaches to this problem: A *geometrical* approach via a projection analogous to a Hodge-Helmholtz decomposition in section 2.4.1 and a purely *dynamical* approach via a decomposition of the ensemble of reactive trajectories into an ensemble of directed pathways and an ensemble of cycles. The former approach is easier to compute but we shall see that it lacks a stochastic interpretation. The latter approach uses the cycle decompositions from section 2.1. It is computationally more demanding, but it has a clear stochastic interpretation and allows us to give correct stochastic weights to all possible

directed pathways from A to B .

2.4.1 Hodge-Helmholtz Decomposition

Let $\mathbb{X} \subset \mathbb{R}^d$. Recall that the probability current of reactive trajectories J_{AB} introduced in section 1.3 is a vector field $J_{AB} : \mathbb{X} \rightarrow \mathbb{R}^d$ on \mathbb{X} . To understand the structure of the transition from A to B in more detail, it is helpful to decompose J_{AB} into simpler parts. The basic idea here is the concept of **Hodge-Helmholtz decompositions**: Let F be a vector field on some bounded domain $D \subset \mathbb{R}^3$. Then F can always be decomposed into a gradient and a rotation:

$$F = \nabla\Phi + \nabla \times R. \quad (2.42)$$

The gradient part $\nabla\Phi$ is rotation-free ($\nabla \times (\nabla\Phi) = 0$) while the rotational part $\nabla \times R$ is divergence-free ($\nabla \cdot (\nabla \times R) = 0$). The potential Φ can be computed by taking the divergence of (2.42), which yields the Poisson equation $\nabla \cdot F = \Delta\Phi$ together with appropriate boundary conditions on ∂D for Φ and R , e.g. $\vec{n} \cdot R = 0$ and $\vec{n} \cdot \nabla\Phi = \vec{n} \cdot F$, where \vec{n} is the normal vector of ∂D .

When looking for decompositions of the type (2.42) for J_{AB} , we have to make additional assumptions on the dynamics. We first consider the case of Markov Diffusions on $\mathbb{X} \subset \mathbb{R}^d$. Then we consider the case of MJPs on discrete state spaces, that is we look for decompositions of the type (2.42) for the effective probability current f_{AB}^+ introduced in (1.52).

Markov Diffusions

Let $(X_t)_t$ be a diffusion on $\mathbb{X} \subset \mathbb{R}^d$ with generator L given by (1.18). Let A and B be two disjoint closed subsets of \mathbb{X} with smooth boundaries. The probability current J_{AB} is given by [Met07]

$$J_{AB}(x) = q^-(x)q^+(x)J(x) + q^-(x)\mu(x)a(x)\nabla q^+(x) - q^+(x)\mu(x)a(x)\nabla q^-(x) \quad (2.43)$$

with the equilibrium current $J(x) = \mu(x)b(x) - \nabla \cdot (a(x)\mu(x))$. We assume that the diffusion matrix a is either symmetric positive definite (denoted by $a > 0$) or symmetric positive semidefinite (denoted by $a \geq 0$). We show in appendix A that if $(X_t)_t$ is reversible, J_{AB} reduces to

$$J_{AB}(x) = \mu(x)a(x)\nabla q^+(x), \quad (2.44)$$

that is J_{AB} is $\mu(x)a(x)$ times a gradient vector field with potential q^+ . This motivates the following definition:

Definition 2.9 *A Hodge-Helmholtz decomposition of the probability current of reactive trajectories $J_{AB} : \mathbb{X} \subset \mathbb{R}^d$ is given by a potential $\Phi : \mathbb{X} \subset \mathbb{R}$ and another vector field $R : \mathbb{X} \subset \mathbb{R}^d$ such that*

$$J_{AB}(x) = \mu(x)a(x)\nabla\Phi(x) + R(x) \quad \forall x \in \mathbb{X} \quad (2.45)$$

2.4. Application II: Transition Pathways in Irreversible Markov Processes

where $\operatorname{div} R = 0$ on \mathbb{X} and $\vec{n}_A \cdot R = \vec{n}_B \cdot R = 0$ with \vec{n}_A and \vec{n}_B being the unit normal vector fields on ∂A and ∂B . We write $J_\Phi(x) := \mu(x) a(x) \nabla \Phi(x)$.

We collect some properties:

Theorem 2.10 *The decomposition (2.45) has the following properties:*

- (i) *The Potential Φ solves the following Dirichlet problem with Neumann boundary conditions:*

$$\begin{aligned} L^s \Phi &= 0 && \text{on } \mathbb{X} \setminus (A \cup B), \\ \vec{n}_A \cdot (a \nabla \Phi) &= \vec{n}_A \cdot (a \nabla q^+) && \text{on } \partial A, \\ \vec{n}_B \cdot (a \nabla \Phi) &= -\vec{n}_B \cdot (a \nabla q^-) && \text{on } \partial B. \end{aligned} \quad (2.46)$$

If $a > 0$ then there is a (up to an overall additive constant) unique weak solution Φ to (2.46) in the Sobolev space $H^1(\mathbb{X} \setminus (A \cup B))$, and Φ achieves its minimum on ∂A and its maximum on ∂B .

- (ii) *If $(X_t)_t$ is reversible, then $\Phi = q^s$ and $R = 0$ are solutions for (2.45).*

- (iii) *Let $\gamma: [0, T] \rightarrow \mathbb{X}$ be a path so that for every $t \in [0, T]$, the tangent vector $\dot{\gamma}(t)$ is parallel to $J_\Phi(\gamma(t))$. If $a \geq 0$ then Φ is nondecreasing along γ . If $a > 0$ then Φ is strictly increasing along γ .*

Proof:

- (i) To compute the potential Φ , we take the divergence of (2.45). For $x \in \mathbb{X} \setminus (A \cup B)$ we know that $\operatorname{div} J_{AB}(x) = 0$. On the other hand, the divergence of the RHS of (2.45) is given by

$$\begin{aligned} \operatorname{div}(\mu(x) a(x) \nabla \Phi(x)) &= \sum_i \nabla^i \left(\mu(x) \sum_j a_{ij}(x) \nabla^j \Phi(x) \right) \\ &= \sum_j \left(\nabla^j \Phi(x) \right) \sum_i \nabla^i (\mu(x) a_{ij}(x)) + \mu(x) \sum_{i,j} a_{ij}(x) \frac{\partial^2 \Phi(x)}{\partial x_i \partial x_j} \\ &= \mu(x) b^s(x) \cdot \nabla \Phi + \mu(x) a(x) : \nabla^2 \Phi(x) \end{aligned}$$

with, in view of (1.20),

$$b_j^s(x) := \frac{1}{\mu(x)} \sum_i \nabla^i (\mu(x) a_{ij}(x)) = \frac{1}{2} (b_j(x) + b_j^-(x)).$$

Putting this together, we obtain that

$$\operatorname{div}(\mu(x) a(x) \nabla \Phi(x)) = \mu(x) L^s \Phi(x) \quad \forall x \in \mathbb{X} \setminus (A \cup B)$$

with the symmetrized generator $L^s = \frac{1}{2}(L + L^-)$. Together with the boundary conditions obtained by multiplying (2.43) with \vec{n}_A resp. \vec{n}_B , this gives (2.46). If a is positive definite then L^s is elliptic and (2.46) has a unique weak solution (up to an additive constant) in the Sobolev space $H^1(\mathbb{X} \setminus (A \cup B))$ [Bra07, Bra12]. Furthermore, by the maximum principle for elliptic operators, Φ achieves its maximum and minimum on $\partial A \cup \partial B$. Let x_M be the maximum of Φ and suppose that $x_M \in \partial A$. Then by the Hopf Lemma [EL98], either Φ is constant (But that would imply $\nabla\Phi = 0$ and thus $\text{div} J_{AB} = 0$ on \mathbb{X} by (2.45). Hence unless $k_{AB} = 0$, Φ is not constant.) or $\vec{n}_A(x_M) \cdot \nabla\Phi(x_M) < 0$. Since a is positive definite, this implies that $\vec{n}_A(x_M) \cdot (a(x_M)\nabla\Phi(x_M)) < 0$. On the other hand, since ∂A is the surface where $q^+ = 0$, either $\nabla q^+(x_M) = 0$ or $\nabla q^+(x_M)$ is parallel to $\vec{n}_A(x_M)$. Since a is positive definite, this implies $\vec{n}_A(x_M) \cdot (a(x_M)\nabla q^+(x_M)) \geq 0$, in contradiction to (2.46).

(ii) If $(X_t)_t$ is reversible, then $\Phi = q^s$ solves (2.46) since then $q^+ = q^s$ and $q^- = 1 - q^s$. Furthermore, by (2.44) and $\Phi = q^s$, we have $J_{AB} = J_\Phi$ and by (2.45), we have $J_{AB} = J_\Phi + R$. Thus $R = 0$.

(iii) Since $\dot{\gamma}(t)$ is parallel to $J_\Phi(\gamma(t))$ for every $t \in [0, T]$, there is a scalar function $\alpha : [0, T] \rightarrow (0, \infty)$ so that $\dot{\gamma}(t) = \alpha(t)J_\Phi(\gamma(t))$. Let $\dot{\Phi}_t = \dot{\Phi}(\gamma(t))$ and $\xi(x) = \nabla\Phi(x)$. Then

$$\dot{\Phi}_t = \nabla\Phi(\gamma(t)) \cdot \dot{\gamma}(t) = \alpha(t)\nabla\Phi(\gamma(t)) \cdot J_\Phi(\gamma(t)) = \alpha(t)\mu(\gamma(t)) \sum_{ij} a_{ij}(\gamma(t)) \xi^i(\gamma(t)) \xi^j(\gamma(t)).$$

Since $\alpha > 0$ and $\mu > 0$, $a \geq 0$ implies $\dot{\Phi}_t \geq 0$ so that Φ is non-decreasing, and $a > 0$ implies $\dot{\Phi}_t > 0$ so that Φ is strictly increasing. ■

Theorem (2.10) gives a way to compute Φ by solving (2.46) and guarantees weak uniqueness of the solution in the case where $a > 0$. It also allows for an interpretation of the decomposition (2.45). Firstly, in the reversible case we obtain $\Phi = q^s$ and $R = 0$. Thus the appearance of the rotation R is due to nonreversibility. R is divergence-free and thus the surface integral of R over any dividing surface between A and B is zero, so R does not describe how reactive trajectories are transported from A to B . The gradient part J_Φ shows the opposite behaviour: The potential Φ is a harmonic function of the elliptic operator L^s , and differs from q^s only in the boundary conditions in (2.46). If $a > 0$, then by Theorem (2.10) Φ is strictly increasing along flowlines of J_Φ . Since Φ achieves its minimum at ∂A and its maximum at ∂B , this means that flowlines of J_Φ go from ∂A to ∂B , they do not form cycles and they cross every equipotential surface of Φ exactly once. So J_Φ does describe the transport mechanism of the reactive trajectories from A to B .

These properties show that J_Φ is very similar to J_{q^s} , the reactive current of the symmetrized process with generator L^s . But note that by Theorem 2.6 $k_{AB} \geq k_{AB}^s$, so J_Φ will generically transport more flow from A to B than J_{q^s} . This is achieved by stretching the potential Φ relative to q^s , allowing for a comparatively larger gradient. That this really happens becomes clear in Theorem 2.14 which discusses the discrete case and shows that $\Phi = (k_{AB}/k_{AB}^s) q^s$ when $A = \{a\}$ and $B = \{b\}$ are single points. Note the similarity to the construction of reversible surrogates in

2.4. Application II: Transition Pathways in Irreversible Markov Processes

section 2.2.1: In order to emulate the stronger mobility of a non-reversible process, we had to add additional edges. Here, we have to allow for more flow to pass along the existing edges in order to emulate the stronger reaction rate of the non-reversible process.

Markov Jump Processes

Let $(X_t)_t$ be a MJP on $\mathbb{X} = \{1, \dots, n\}$, and let A and B be two disjoint subsets of \mathbb{X} . The discrete analogon of the vector field J_{AB} is the effective probability current $f_{AB}^+ : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$ introduced in (1.52). In order to talk about Hodge-Helmholtz decompositions of f_{AB}^+ , we need discrete gradient and divergence operators. The divergence was introduced in (1.53). We define a discrete gradient operator:

Definition 2.11 (Discrete gradient.) *Let L be the generator of a MJP on $\mathbb{X} = \{1, \dots, n\}$. For any $(x, y) \in \mathbb{X} \times \mathbb{X}$ let*

$$c^s(x, y) = \frac{1}{2} (\mu_x l_{xy} + \mu_y l_{yx})$$

*be the **capacity** of the edge (xy) . For a function $f : \mathbb{X} \rightarrow \mathbb{R}$, let the **gradient** of f be the function $\text{grad } f : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$ given by*

$$\text{grad } f(x, y) = \frac{1}{2} c^s(x, y) (f(y) - f(x)).$$

This definition is motivated by the following observation: If $(X_t)_t$ is reversible, then f_{AB}^+ reduces to [Met07]

$$f_{AB}^+(x, y) = \frac{1}{2} \mu_x l_{xy} (q^+(y) - q^+(x)) = \text{grad } q^+(x, y), \quad (2.47)$$

thus f_{AB}^+ is a gradient flow with potential q^+ , in perfect agreement with the diffusion situation. Note that $\text{grad } f$ is an antisymmetric function on $\mathbb{X} \times \mathbb{X}$, and we have the following Lemma:

Lemma 2.12 *For any $f : \mathbb{X} \rightarrow \mathbb{R}$, we have*

$$\text{div grad } f = D_\mu L^s f$$

where $L^s = \frac{1}{2} (L + L^-)$.

Chapter 2. Irreversible Markov Chains and Cycles

Proof: Let $x \in \mathbb{X}$. We compute

$$\begin{aligned} \operatorname{div} \operatorname{grad} f(x) &= \sum_y \operatorname{grad} f(x, y) - \sum_y \operatorname{grad} f(y, x) = 2 \sum_y \operatorname{grad} f(x, y) \\ &= \sum_y c^s(x, y) [f(y) - f(x)] \\ &= \sum_y \mu_x L_{xy}^s f(y) \end{aligned}$$

using (1.53), the antisymmetry of $\operatorname{grad} f$ and definition 2.11. ■

We now make the following definition, in analogy with definition 2.9:

Definition 2.13 A **Hodge-Helmholtz decomposition** of the effective current $f_{AB}^+ : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$ is given by a **potential** $\Phi : \mathbb{X} \subset \mathbb{R}$ and a **flow** $R : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$ such that

$$f_{AB}^+(x, y) = \operatorname{grad} \Phi(x, y) + R(x, y) \quad \forall x, y \in \mathbb{X} \quad (2.48)$$

and $\operatorname{div} R = 0$ on \mathbb{X} .

Note that there are no additional boundary conditions for R in (2.48). The reason for this is that in contrast to the diffusion case, (2.48) together with $\operatorname{div} R = 0$ already leads to an equation for Φ which has a unique solution. This also means that in general $R(x, y) \neq 0$ for $x \in A$ and $y \notin A$ or $x \notin B$ and $y \in B$. Again we collect properties of this decomposition:

Theorem 2.14 The decomposition (2.48) has the following properties:

- (i) If $(X_t)_t$ is ergodic then Φ and R are unique up to an overall additive constant in Φ . The potential Φ solves the linear system

$$L^s \Phi = b, \quad b(x) = \begin{cases} 0 & x \in \mathbb{X} \setminus (A \cup B), \\ (Lq^+)(x) & x \in A, \\ -(L^- q^-)(x) & x \in B. \end{cases} \quad (2.49)$$

In addition, Φ achieves its minimum on A and its maximum on B .

- (ii) If $(X_t)_t$ is reversible, then $\Phi = q^s$ and $R = 0$.
 (iii) If $A = \{a\}$ and $B = \{b\}$ are singletons, then

$$\Phi = \frac{k_{AB}}{k_{AB}^s} q^s \quad (2.50)$$

where $k_{AB} = \sum_y \mu_a l_{ay} q^+(y)$ and $k_{AB}^s = \sum_y \mu_a L_{ay}^s q^s(y)$.

- (iv) If (x_1, \dots, x_T) is a path in $G(L^s)$ such that $\operatorname{grad} \Phi(x_i, x_{i+1}) > 0$ for all $i = 1 \dots T-1$, then Φ is strictly increasing along (x_1, \dots, x_T) .

Proof:

- (i) In complete analogy with the diffusion case, we obtain the equation that Φ must solve by taking the divergence of (2.48). From (1.51) and (1.54), we obtain for the divergence of the LHS

$$\operatorname{div} f_{AB}^+(x) = \begin{cases} 0 & x \in \mathbb{X} \setminus (A \cup B), \\ \sum_y \mu_x l_{xy} q^+(y) & x \in A, \\ -\sum_y q^-(y) \mu_y l_{yx} & x \in B. \end{cases} \quad (2.51)$$

The divergence of the RHS of (2.48) gives $D_\mu L^s \Phi$ in view of Lemma (2.12). If we divide both sides by D_μ and use the definition of L^- , we obtain (2.49). Since $(X_t)_t$ is ergodic, the kernel of L^s is given by the constant functions on \mathbb{X} . Thus the solution to (2.49) is unique up to a constant. By (2.49) and the maximum principle for the operator $L^s \leq 0$, the potential Φ must achieve its maximum and minimum on $A \cup B$. Suppose $x \in A$ is the maximum of Φ . Then

$$L^s \Phi(x) = \sum_y L_{xy}^s \Phi(y) = \sum_{y \neq x} L_{xy}^s (\Phi(y) - \Phi(x)) < 0.$$

On the other hand,

$$b(x) = Lq^+(x) = \sum_{y \neq x} l_{xy} (q^+(y) - q^+(x)) \geq 0$$

since $q^+(x) = 0$, in contradiction to (2.49).

- (ii) If $(X_t)_t$ is reversible then $q^+ = q^s$ and $q^- = 1 - q^s$, so $\Phi = q^s$ solves (2.49). On the other hand, $f_{AB}^+ = \operatorname{grad} q^+$ by (2.47), thus $R = 0$.

- (iii) If $A = \{a\}$ and $B = \{b\}$, then by (i) and (1.55), Φ solves $L^s \Phi = b$ with

$$\mu_x b(x) = \begin{cases} 0 & x \in \mathbb{X} \setminus (A \cup B), \\ k_{AB} & x = a, \\ -k_{AB} & x = b. \end{cases}$$

On the other hand, the committor function q^s solves (1.48) with L replaced by L^s , which can be written as $L^s q^s = b^s$ with

$$\mu_x b^s(x) = \begin{cases} 0 & x \in \mathbb{X} \setminus (A \cup B), \\ k_{AB}^s & x = a, \\ -k_{AB}^s & x = b. \end{cases}$$

This shows (2.50).

(iv) Clear from definition (2.11). ■

Theorem (2.14) guarantees that the decomposition (2.48) always has a solution up to an overall additive constant in Φ as long as $(X_t)_t$ is ergodic. It also shows that the interpretation of the diffusion case (2.45) carries over to (2.48). Again $R = 0$ in the reversible case, so that the appearance of R can be attributed to the non-reversibility of $(X_t)_t$. Since $\operatorname{div} R = 0$ on \mathbb{X} , the total flow transported by R across any dividing surface between A and B is zero, so R does not describe the transport mechanism of reactive trajectories from A to B , but additional rotations instead. The potential Φ again differs from q^s only due to the different boundary conditions in (2.49). If A and B are singletons, then (2.50) shows that Φ and q^s are equal up to a multiplicative constant given by the fraction of the reaction rates k_{AB} and k_{AB}^s . This multiplicative constant is always larger than one and stretches Φ compared to q^s so that $\operatorname{grad} \Phi$ can transport the rate $k_{AB} \geq k_{AB}^s$. The situation $A = \{a\}$ and $B = \{b\}$ can always be reached by **lumping** the states in A to a single state $\{a\}$ and the states in B to a single state $\{b\}$: Let $\mathbb{T} = \mathbb{X} \setminus (A \cup B)$. We set $\mu_a = \sum_{x \in A} \mu_x$, $\mu_b = \sum_{x \in B} \mu_x$ and define the lumped generator \bar{L} on $\mathbb{T} \cup \{a\} \cup \{b\}$ by setting

$$\begin{aligned} \bar{l}_{xy} &= l_{xy} & \forall x, y \in \mathbb{T}, \\ \bar{l}_{ay} &= \frac{1}{\mu_a} \sum_{x \in A} \mu_x l_{xy} & \forall y \in \mathbb{T}, \\ \bar{l}_{xb} &= \sum_{y \in B} l_{xy} & \forall x \in \mathbb{T}. \end{aligned}$$

It can be checked easily that lumping preserves $f_{AB}^+(x, y)$ for all $x, y \in \mathbb{T}$ and that $f_{AB}^+(a, y) = \sum_{x \in A} f_{AB}^+(x, y)$ for all $y \in \mathbb{T}$ as well as $f_{AB}^+(x, b) = \sum_{y \in B} f_{AB}^+(x, y)$ for all $x \in \mathbb{T}$.

The following Lemma is a corollary of Theorem 2.14.

Lemma 2.15 *Let Φ be the potential of f_{AB}^+ as in (2.48), and write $\operatorname{grad} \Phi = \frac{1}{2} (\operatorname{grad} \Phi_+ - \operatorname{grad} \Phi_-)$ where $\operatorname{grad} \Phi_+ \geq 0$ and $\operatorname{grad} \Phi_- \geq 0$ are the positive and negative parts of $\operatorname{grad} \Phi$. Then the Coates graph $G(\operatorname{grad} \Phi_+)$ is a directed tree with all of its roots in A and all of its leaves in B , and Φ is strictly increasing along paths in $G(\operatorname{grad} \Phi_+)$.*

Proof: By (iv) in Theorem (2.14), $G(\operatorname{grad} \Phi_+)$ is cycle-free, and Φ is strictly increasing along paths in $G(\operatorname{grad} \Phi_+)$. If x is a root in $G(\operatorname{grad} \Phi_+)$, then $\operatorname{grad} \Phi(y, x) \leq 0$ for all $y \in \mathbb{X}$, and there must be an $x' \in \mathbb{X}$ with $\operatorname{grad} \Phi(x, x') > 0$. Thus $\operatorname{div} \operatorname{grad} \Phi(x) > 0$. Likewise, if x is a leaf in $G(\operatorname{grad} \Phi_+)$ then $\operatorname{div} \operatorname{grad} \Phi(x) < 0$. By (2.48) and (2.51), $\operatorname{div} \operatorname{grad} \Phi(x) = \operatorname{div} f_{AB}^+(x) = 0$ if $x \in \mathbb{X} \setminus (A \cup B)$, $\operatorname{div} \operatorname{grad} \Phi(x) > 0$ if $x \in A$ and $\operatorname{div} \operatorname{grad} \Phi(x) < 0$ if $x \in B$. Thus all the roots are in A and all the leaves are in B . ■

Examples

Barbell Graph. As our first example we consider the Barbell graph G that was shown in Figure (2.7). We define a random walk on G that goes with probability $p = 0.8$ along edge directions and with probability $(1 - p)$ against edge directions. We select the source $A = \{a\}$ and target $B = \{b\}$ as shown in Figure 2.14. The resulting effective current f_{AB}^+ is also shown in Figure 2.14. The dotted lines transport almost no current, so that f_{AB}^+ is concentrated on a single pathway from a to b . A typical reactive trajectory uses this pathway. In Figure 2.14 on the right, the potential Φ and the splitting of f_{AB}^+ into $\text{grad}\Phi$ and R are shown. We observe that $\text{grad}\Phi$ uses the loops in both directions with equal preference. The fact that one of the directions is actually preferred is encoded in R . In other words: The flowlines of $\text{grad}\Phi$ do not coincide with typical pathways of reactive trajectories.




Figure 2.14 – Left: The effective current f_{AB}^+ on the barbell graph with source $A = \{a\}$ and target $B = \{b\}$. The dotted edges carry almost no current. Right: Gradient flow $\text{grad}\Phi$ (green edges) and rotation R (red edges). Thickness denotes magnitude of flow. The colorscale of the nodes indicates the value of Φ .




Figure 2.15 – Transition tube graph. Left: Effective current. Edge thickness is proportional to the magnitude of the current. Right: Node coloring indicates Φ . Edge directions according to $\text{grad}\Phi$. Edge thickness is proportional to $\text{grad}\Phi$. Edge color is proportional to R .

Transition tube with vortices. The next example is a graph G which represents a short transition tube with 4 additional vortices. We construct this by starting with the simple RW (2.33) on the undirected graph G and then adding a strong preference to travel along the

vortices in the direction shown in Figure 2.15, which results in the current f_{AB}^+ shown on the left. On the right we show the splitting of f_{AB}^+ into $\text{grad}\Phi$ and R . The color of the edges is proportional to R while the thickness is proportional to $\text{grad}\Phi$. The rotation R picks the vortices out. We observe that $\text{grad}\Phi$ is larger for the outer pathways than it is for the central pathway and, as in the barbell graph example, $\text{grad}\Phi$ passes through the vortices in both directions - in the direction given by f_{AB}^+ and against it. The potential Φ itself is also shown in Figure 2.15, and it is monotonically increasing from the left to the right.

Random block model. Our next example is a simple random walk on a random graph G generated by a so-called **random block model**. Given parameters m , n_b , p_i and p_o , the random block model first groups the $n = mn_b$ nodes into m blocks C_1, \dots, C_m of n_b vertices each. Then the model decides independently for each directed edge if it is present or not. In the simplest version of the model, edges that connect nodes in the same block (i.e. interior edges) are present with probability p_i , while edges between blocks (i.e. exterior edges) are present with probability p_o . We make the additional modification of only allowing exterior edges between a specified central block C_{m_0} and any other block C_i . This arranges the blocks in a star-shaped fashion, see Figure 2.16.

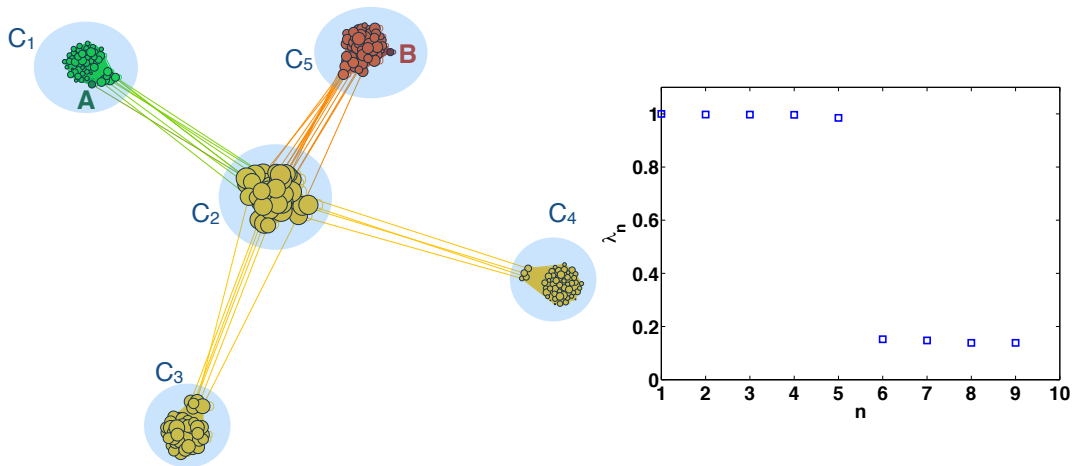


Figure 2.16 – Left: A realization of the random block model with $m = 5$, $n_b = 50$, $p_i = 0.5$ and $p_o = 0.0016$, with a random source node A in block C_1 and a random target node B in block C_3 selected. Node color is according to q^+ , with red denoting $q^+(x) = 1$ and green denoting $q^+(x) = 0$. Node size proportional to $\mu_R(x)$ (log scale). Right: Dominant eigenvalues of the transition matrix P (2.33) for this model.

For $p_i \gg p_o$, the block model produces a directed graph G with a pronounced modular structure. The simple random walk (2.33) on G is expected to display strong metastability with metastable sets C_1, \dots, C_m . As an example, we consider the parameters $m = 5$, $n_b = 50$, $p_i = 0.5$ and $p_o = 0.0016$. This produces 5 well connected blocks C_1, \dots, C_5 with on average 4 connections per direction between the central block C_2 and the other blocks⁷. Figure 2.16

⁷There is a small probability that the random graph G created by the block model is not strongly connected, so

2.4. Application II: Transition Pathways in Irreversible Markov Processes

shows an example realization. Figure 2.16 also shows the dominant spectrum of the transition matrix P of the simple RW on G given by (2.33). As expected, there are 5 eigenvalues close to one which shows that P displays very strong metastability.

Now we study transitions from C_1 to C_5 by selecting a random node in C_1 as the source A and a random node in C_5 as the target B . Figure 2.16 shows the forward committor q^+ (colour of the nodes) and the probability of reactive trajectories μ_R given by (1.50) (size of the nodes). We see that q^+ is almost constant on C_1 , C_5 and $C_2 \cup C_3 \cup C_4$. Most of the reactive trajectories are in C_2 , followed by C_3 . This further confirms the presence of strong metastability. A typical reactive trajectory first equilibrates in C_1 , then transitions to C_2 , equilibrates in $C_2 \cup C_3 \cup C_4$ and finally transitions to C_5 where it goes to B . Since it cannot go from C_3 or C_4 to C_5 , transitions $C_2 - C_3 - C_5$ or $C_2 - C_4 - C_5$ are detours that we would like to separate from the actual transition $C_1 - C_2 - C_5$. We may conclude that $\text{grad}\Phi$ describes the transition from A to B , but $\text{grad}\Phi$ is positive on edges where the reactive trajectories themselves cannot go.




Figure 2.17 – The decomposition (2.48) of f_{AB}^+ (top) into $\text{grad}\Phi$ (bottom left) and R (bottom right) for the random block model. Colour according to q^+ with red denoting $q^+(x) = 1$ and green denoting $q^+(x) = 0$. Node size proportional to the total amount of flow passing through the node (log scale). Edge color proportional to the total amount of flow passing through the edge (log scale).

that the simple RW on G would not be ergodic. If this happens, we discard G and redraw.

In Figure 2.17, the effective current f_{AB}^+ and its splitting into $\text{grad}\Phi$ and R according to (2.48) is visualized. For all the plots in Figure 2.17, the color of the edges is proportional to the logarithm of the amount of flow the edge transports, from light grey (small) to black (large). The color of the nodes corresponds to q^+ . The potential Φ and q^+ are indistinguishable. The size of the nodes corresponds to the total traffic, i.e. the total incoming flow, the node sees, again on a logarithmic scale. We see that the effective current f_{AB}^+ is large on all exterior edges between C_2, C_3 and C_4 , indicating a lot of traffic between C_2, C_3 and C_4 before the transition to C_5 is made. The edges from C_5 to C_2 also carry a lot of current. In contrast, $\text{grad}\Phi$ only goes from C_1 to C_2 and then to C_5 . All exterior edges between $C_1 - C_2$ and $C_2 - C_5$ are now oriented in the direction of the reaction, and there is almost no traffic between C_2, C_3 and C_5 . The rotation R captures back and forth transitions between C_2 and C_1, C_3, C_4 , and C_5 .

Langevin system. The last example is the Langevin system (2.40) from section 2.3.2. Recall that with the parameters from section 2.3.2, the Langevin system is close to the corresponding Hamiltonian system. For the Hamiltonian system, the energy $H(x, p) = \frac{1}{2m}p^2 + V(x)$ is conserved. For the Langevin system, $E_\mu[H_t]$ is conserved and $E_\mu[|H_t - H_0|]$ varies slowly: The system will sample constant energy surfaces quickly, but take a long time to reach energies which are significantly higher or lower. We study reactions from the right basin $A = \{0.9 \leq x \leq 1.1, -0.3 \leq p \leq 0.3\}$ to the left basin $B = \{-1.1 \leq x \leq -0.9, -0.3 \leq p \leq 0.3\}$. To go from the right basin to the left basin, the system must first increase its energy from $H_0 = 0$ to $H_{1/2} \geq 1$ to reach the separatrix, then it must travel to the left and finally decrease its energy from $H_{1/2} \geq 1$ to $H_1 = 0$. Therefore we expect the potential Φ to coincide with H in the right well and with $-H$ in the left well. Figure 2.18 shows Φ and the contour lines of H , and the result indeed confirms our expectations exactly.

Figure 2.18 also shows the splitting (2.48) of f_{AB}^+ into $\text{grad}\Phi$ (green arrows) and R (red arrows). The rotation R essentially coincides with the flow of the Hamiltonian system itself. That is, R describes the fast dynamics within regions of constant energy. The gradient flow $\text{grad}\Phi$ is approximately symmetric with respect to the momentum flip $p \mapsto -p$. In particular, the total flow crossing the $x = 0$ hypersurface in the $p > 0$ region is very similar to the total flow crossing $x = 0$ in the $p < 0$ region. In contrast, a typical reactive trajectory would cross the $x = 0$ hypersurface many times, but it would almost always cross from right to left in the $p < 0$ region and from left to right in the $p > 0$ region, and it would almost always make its last crossing in the $p < 0$ region. This shows that we cannot interpret $\text{grad}\Phi$ as i.e. the flowlines of reactive trajectories after their detours have been removed.

2.4.2 Transition Pathways and Cycles

Parts of the work presented in this section are published in [BCS15]. The numerical results of section 2.4.1 all indicate one drawback if we wish to interpret the decomposition (2.48): The potential part $\text{grad}\Phi$ is **insensitive** to the edge orientations of f_{AB}^+ . This renders a stochastic interpretation of the splitting (2.48) impossible. By this we have the following in mind: Consider

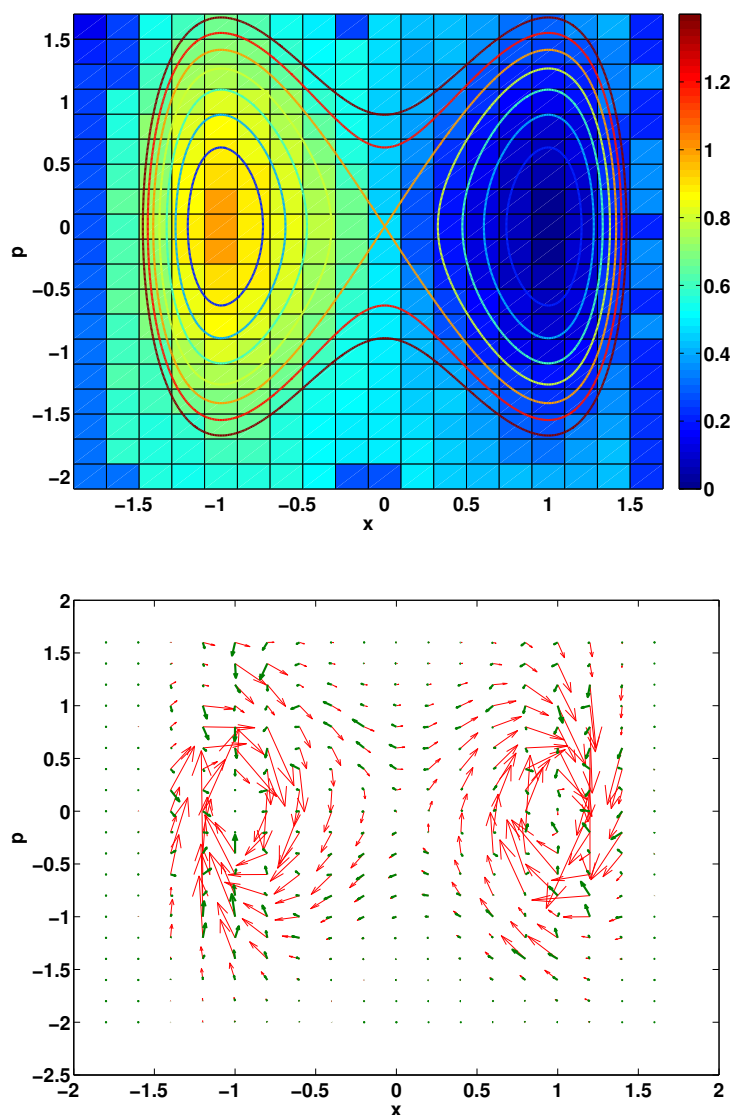


Figure 2.18 – Langevin system. Top: Potential Φ and contour lines of H . For the boxes at the boundary, Φ could not be computed reliably due to lack of data. Bottom: Gradient flow $\text{grad } \Phi$ (green arrows) and rotation R (red arrows). The length of the arrows is proportional to the magnitude of flow transported. The size of the arrows representing $\text{grad } \Phi$ has been enlarged by a factor of 2 relative to R for visibility.

some way to interpret $\text{grad } \Phi$ as the probability flow generated by an ensemble of 'pruned' reactive trajectories. By the properties of $\text{grad } \Phi$, the pruned reactive trajectories will have no cycles and will produce the same reaction rate k_{AB} . However, other statistical properties of the ensemble of pruned trajectories might be completely different when compared to the original ensemble of reactive trajectories. For example, pruned trajectories might perform transitions frequently that are extremely unlikely or even outright forbidden for the reactive trajectories.

The numerical results in 2.4.1 indicate that this behaviour of $\text{grad } \Phi$ is generic. As a way out, we will now discuss an alternative to the Hodge-Helmholtz-decomposition presented in section 2.4.1 that constructs the ensemble of pruned reactive trajectories explicitly and therefore has a direct stochastic interpretation. The cycle decompositions of section 2.1 provide the tool for achieving this.

To set things up, we make the following definition, following [Met07]:

Definition 2.16 A reaction pathway is a nonintersecting directed path in $G(f_{AB}^+)$ that starts in A and ends in B .

The idea here is the following: A single reactive trajectory (A, x_1, \dots, x_n, B) can always be decomposed into one uniquely defined reaction pathway and several cycles. This can be done via the algorithm presented in section 2.1.3. The only difference is that we have two different kinds of 'cycles': Reaction pathways and cycles in the transition region $\mathbb{T} = \mathbb{X} \setminus (A \cup B)$. The corresponding decomposition of f_{AB}^+ into a current generated by reaction pathways and a current generated by cycles will be given by Lemma 2.4.

In order to use Lemma 2.4, we need to construct a Markov chain $(X_n^R)_n$ with transition matrix P^R and invariant distribution μ^R such that $\mu_x^R p_{xy}^R$ gives the current f_{AB}^+ . However, f_{AB}^+ is not divergence-free in A and B , but $\mu_x^R p_{xy}^R$ has to be divergence-free everywhere in order to use Lemma 2.4. As a first step we therefore map all of $A \cup B$ to a single state s and define a lumped flow F^R on $\mathbb{T} \cup \{s\}$, as shown in Figure 2.19. The next Lemma tells how this is done.




Figure 2.19 – Left: Transition region \mathbb{T} (grey), source A (green) and target B (red). Right: The lumping procedure consists of keeping \mathbb{T} and mapping all of $A \cup B$ to a single node s .

Lemma 2.17 Let the lumped flow F^R on $\mathbb{T} \cup \{s\}$ be defined by

$$F^R(x, y) = \begin{cases} \max \{f_{AB}(x, y) - f_{AB}(y, x), 0\} & x, y \in \mathbb{T} \\ \sum_{a \in A} f_{AB}(a, y) & x = s, y \in \mathbb{T} \\ \sum_{b \in B} f_{AB}(x, b) & x \in \mathbb{T}, y = s. \end{cases} \quad (2.52)$$

2.4. Application II: Transition Pathways in Irreversible Markov Processes

Then $\frac{1}{2}(F^R(x, y) - F^R(y, x)) = f_{AB}^+(x, y)$ for all $x, y \in \mathbb{T}$, and F^R is divergence-free on $\mathbb{T} \cup \{s\}$. In addition, the total outflow $\sum_y F^R(s, y)$ of s equals k_{AB} .

Proof: $\frac{1}{2}(F^R(x, y) - F^R(y, x)) = f_{AB}^+(x, y)$ for all $x, y \in \mathbb{T}$ is clear. The outflow of F^R from s is given by

$$\sum_y F^R(s, y) = \sum_{a \in A, y \in \mathbb{X}} f_{AB}(a, y) = k_{AB}$$

and equals the inflow of F^R into s :

$$\sum_x F^R(x, s) = \sum_{x \in \mathbb{X}, b \in B} f_{AB}(x, b) = k_{AB}.$$

Thus $\text{div} F^R(s) = 0$. It remains to show that $\text{div} F^R(x) = 0$ for $x \in \mathbb{X}$:

$$\begin{aligned} \text{div} F^R(x) &= \sum_{y \in \mathbb{T} \cup \{s\}} (F^R(x, y) - F^R(y, x)) \\ &= \sum_{y \in \mathbb{T}} (f_{AB}(x, y) - f_{AB}(y, x)) + F^R(x, s) - F^R(s, x) \\ &= \sum_{y \in \mathbb{T}} (f_{AB}(x, y) - f_{AB}(y, x)) + \sum_{y \in B} f_{AB}(x, y) - \sum_{y \in A} f_{AB}(y, x) \\ &= \sum_{y \in \mathbb{X}} (f_{AB}(x, y) - f_{AB}(y, x)) = 0 \end{aligned}$$

since f_{AB} is divergence-free on $\mathbb{X} \setminus (A \cup B)$. ■

Thus F^R describes the transport of reactive trajectories from A to B in the same way f_{AB}^+ does, but it transports them from s back to s instead. Every reaction pathway $(x_A, x_1, \dots, x_n, x_B)$ corresponds to a simple loop (s, x_1, \dots, x_n, s) in $G(F^R)$ which passes through s . The next step is to define a transition matrix P^R on $\mathbb{T} \cup \{s\}$ and a corresponding stationary distribution μ^R such that $\mu^R(x)P^R(x, y) = F^R(x, y)$ for all $x \neq y \in \mathbb{T} \cup \{s\}$. This can indeed always be done, as the next lemma shows.

Lemma 2.18 *Define*

$$\mu^R(x) = \begin{cases} \sum_y F^R(x, y) & x \in \mathbb{T}, \\ 1 - \sum_{x \in \mathbb{T}} \mu^R(x) & x = s \end{cases} \quad (2.53)$$

and

$$P^R(x, y) = \begin{cases} \frac{F^R(x, y)}{\mu^R(x)} & x \neq y \in \mathbb{T} \cup \{s\}, \\ 0 & x = y \in \mathbb{T}, \\ 1 - \sum_{y \in \mathbb{T}} P^R(s, y) & x = y = s. \end{cases} \quad (2.54)$$

Then P^R is a transition matrix with stationary distribution μ^R and $\mu^R(x)P^R(x, y) = F^R(x, y)$

Chapter 2. Irreversible Markov Chains and Cycles

holds for all $x \neq y \in \mathbb{T} \cup \{s\}$.

Proof: That $\sum_y P^R(x, y) = 1$ and $\sum_{x \in \mathbb{T} \cup \{s\}} \mu^R(x) = 1$ is clear from the definition. $\mu^R(x) p^R(x, y) = F^R(x, y)$ for all $x \neq y \in \mathbb{T} \cup \{s\}$ follows directly from (2.54). We show that μ^R is the stationary distribution of P^R : For $y \in \mathbb{T} \cup \{s\}$,

$$\begin{aligned} \sum_x \mu^R(x) p^R(x, y) &= \mu^R(y) p^R(y, y) + \sum_{x \neq y} F^R(x, y) \\ &= \mu^R(y) p^R(y, y) + \sum_{x \neq y} F^R(y, x) \\ &= \mu^R(y) p^R(y, y) + \sum_{x \neq y} \mu^R(y) p^R(y, x) = \mu^R(y) \end{aligned}$$

using the results above and the fact that F^R is divergence-free. This finishes the proof. ■

Note that we needed to include a self-loop at the state s , cf. Figure 2.19. This ensures that the process $(X_t^R)_t$ spends the appropriate amount of time in the non-reactive state s and that the overall transition rate k_{AB} is correctly given by the outflow of s . There is no need to include self-loops for any other state $x \in \mathbb{T}$.

The transition matrix P^R constructed in Lemma 2.18 is the generalization of the 'no detour transition path process' to non-reversible processes that was constructed in [CVE14] for reversible ones. In fact, the name 'no detour process' is only justified in the reversible case, where the current F^R of P^R equals $\text{grad } q^+$ and thus the trajectories produced by P^R are directed pathways from A to B along which q^+ always increases. In the non-reversible case, paths generated by P^R indeed perform detours in the form of loops, which we proceed to remove in what follows.

Now we are ready to use the cycle decomposition (2.6) for $\mu^R(x) p^R(x, y)$. This will decompose the flow $F^R(x, y)$ into simple cycles in the graph $G(F^R)$. We can write the set \mathcal{C}_∞ of simple cycles in $G(F^R)$ as $\mathcal{C}_\infty = \{[s]\} \cup \mathcal{C}_\infty^s \cup \mathcal{C}_\infty^\mathbb{T}$ where \mathcal{C}_∞^s is the set of cycles that contain the node s and at least one node in \mathbb{T} and therefore correspond to reaction pathways, and $\mathcal{C}_\infty^\mathbb{T}$ is the set of cycles that do not contain s and therefore correspond to detours of the reactive trajectories.

Theorem 2.19 *The decomposition formula*

$$\begin{aligned} F^R(x, y) = \mu^R(x) p^R(x, y) &= \sum_{c \in \mathcal{C}_\infty^s} \omega_\infty(c) J_c(x, y) + \sum_{c \in \mathcal{C}_\infty^\mathbb{T}} \omega_\infty(c) J_c(x, y) \\ &= F^P(x, y) + R(x, y) \end{aligned} \quad (2.55)$$

holds for all $x \neq y \in \mathbb{T} \cup \{s\}$, and has the following properties:

- (i) $\sum_{c \in \mathcal{C}_\infty^s} \omega_\infty(c) = k_{AB}$
- (ii) If $(X_t)_t$ is reversible, then $\mathcal{C}_\infty^\mathbb{T} = \emptyset$ and thus $R = 0$.

2.4. Application II: Transition Pathways in Irreversible Markov Processes

Proof: The decomposition (2.55) is a consequence of (2.6) and $\mathcal{C}_\infty = \{[s]\} \cup \mathcal{C}_\infty^s \cup \mathcal{C}_\infty^\mathbb{T}$. We show (i) and (ii):

- (i) We use Lemma 2.17, equation (2.55) and the fact that $J_c(s, y) = 0$ for all $c \in \mathcal{C}_\infty^\mathbb{T}$ and all $y \in \mathbb{T}$ to write

$$k_{AB} = \sum_{y \in \mathbb{T}} F^R(s, y) = \sum_{y \in \mathbb{T}} \sum_{c \in \mathcal{C}_\infty^s} \omega_\infty(c) J_c(s, y) = \sum_{c \in \mathcal{C}_\infty^s} \omega_\infty(c) J_c(s).$$

But by definition $J_c(s) = 1$ for all $c \in \mathcal{C}_\infty^s$, which shows (i).

- (ii) If $(X_t)_t$ is reversible, then we know from section 2.4.1 that $f_{AB}^+ = \text{grad } q^+$ and $G(f_{AB}^+)$ contains no cycles. Then the only cycles $G(F^R)$ can have must pass through s . ■

The weights $\omega_\infty(c)$ for the reaction pathways $c \in \mathcal{C}_\infty^s$ have the following interpretation: Suppose $X_0^R = s$ and let $\tau = \inf\{t > 0 : X_t^R = s\}$ be the next time the state s is visited. By definition, this completes a cycle in \mathcal{C}_∞^s , and the probability that the completed cycle is c is given by $\omega_\infty(c)/k_{AB}$. Thus $\omega_\infty(c)$ endows \mathcal{C}_∞^s , and thus the space of reaction pathways, with a probability distribution. If we generate a long realization of the process $(X_t^R)_t$ that gives many transitions from s back to s , remove all detours (i.e. all loops in \mathbb{T}) and then compute a histogram of the different reaction pathways that have been used by $(X_t^R)_t$, then this histogram converges to $\omega_\infty(c)/k_{AB}$ as the length of the realization goes to infinity. This follows from the properties of the circulation distribution discussed in section 2.1.3. Note that the ensemble of trajectories generated by $(X_t^R)_t$ is very different from the ensemble of reactive trajectories: They have the same effective current, but the trajectories generated by $(X_t^R)_t$ are typically much shorter since they cannot go back and forth along the same edge (which reactive trajectories can do).

Comparison with the Hodge-Helmholtz decomposition. With (2.48) and (2.55) we have two different decompositions of the effective current f_{AB}^+ on our hands⁸. Both decompositions are unique and split up f_{AB}^+ into a part responsible for the transition from A to B - in the case of (2.48) this is $\text{grad } \Phi$ while in the case of (2.55) this is the flow F^P carried by the reactive pathways \mathcal{C}_∞^s - and an additional rotational part R . Both decompositions share the property that $R = 0$ iff $(X_t)_t$ is reversible. The difference is that the decomposition (2.48) is based on a projection and thus neither R nor $\text{grad } \Phi$ have an immediate stochastic interpretation. Indeed $\text{grad } \Phi > 0$ on edges where no reactive trajectory can go, and thus $\text{grad } \Phi$ cannot be interpreted as a flow of pruned reactive trajectories. In contrast, this is exactly what F^P is by construction: If all the reactive trajectories are pruned so that only the directed pathways from A to B remain, then their flow is given by F^P . Consequently, $F^P(x, y) > 0$ only on edges (x, y) where $f_{AB}^+(x, y) > 0$. On the downside, the decomposition (2.55) is much harder to compute than (2.48) which only requires to solve a linear system, and F^P cannot in general be written

⁸Recall that by lemma 2.17, F^R antisymmetrized essentially equals f_{AB}^+ .

Chapter 2. Irreversible Markov Chains and Cycles

as the gradient of a potential Φ , which is desirable because it gives us a reaction coordinate to monitor the transition from A to B (namely Φ). If one wishes to obtain such a reaction coordinate, then Φ from (2.48) is our best answer. If one wishes to know which reaction pathways are actually used during that transition, then (2.55) gives the answer. We will now look at examples to highlight this.

Examples

We start with a tutorial example. Consider a system with source $A = \{1\}$, target $B = \{5\}$ and transition region $\mathbb{T} = \{2, 3, 4\}$. The effective probability current is shown in Figure 2.20. The reaction rate is given by $k_{AB} = k$ and equals the flow from 1 to 2. At node 3, a fraction of δ of the incoming current is rerouted to 2 via 4, the rest is pushed onward to 5. The lumped flow F^R is shown on the right in Figure 2.20. It is given by




Figure 2.20 – Left: Effective current f_{AB}^+ of a tutorial example network. Right: The lumped current F^R of the same network.

$$F^R = k \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & \frac{1}{1-\delta} & 0 \\ 1 & 0 & 0 & \frac{\delta}{1-\delta} \\ 0 & \frac{\delta}{1-\delta} & 0 & 0 \end{pmatrix}.$$

According to Lemma 2.18, we define $\mu^R = \left(\alpha, \frac{1}{1-\delta}k, \frac{1}{1-\delta}k, \frac{\delta}{1-\delta}k \right)^T$ with $\alpha = 1 - \frac{2+\delta}{1-\delta}k$ so that $\sum_x \mu^R(x) = 1$, and

$$P^R = \begin{pmatrix} 1 - \frac{k}{\alpha} & \frac{k}{\alpha} & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 1 - \delta & 0 & 0 & \delta \\ 0 & 1 & 0 & 0 \end{pmatrix}.$$

Thus according to P^R , from 3 one goes next to 4 (and thus uses the loop) with probability δ , and one goes to 5 and arrives at the target with probability $1 - \delta$. The decomposition (2.55) is given

2.4. Application II: Transition Pathways in Irreversible Markov Processes

in terms of $\mathcal{C}_\infty^s = \{[1, 2, 3]\}$ with $\omega_\infty([1, 2, 3]) = k$ and $\mathcal{C}_\infty^\mathbb{T} = \{[2, 3, 4]\}$ with $\omega_\infty([2, 3, 4]) = \frac{\delta}{1-\delta}k$. We can make out two regimes: If $\delta > 1/2$ then $\omega_\infty([2, 3, 4]) > \omega_\infty([1, 2, 3])$, and rotation in \mathbb{T} dominates. If $\delta < 1/2$ then $\omega_\infty([2, 3, 4]) < \omega_\infty([1, 2, 3])$, and the reaction from A to B dominates.

The analysis so far assumed that k is independent of δ . But actually k depends on δ because α , which is the probability of being in the non-reactive state s , stays fixed as δ is varied. Therefore $k = (1 - \alpha) \frac{1-\delta}{2+\delta}$, and accordingly

$$\begin{aligned}\omega_\infty([1, 2, 3]) &= (1 - \alpha) \frac{1 - \delta}{2 + \delta}, \\ \omega_\infty([2, 3, 4]) &= (1 - \alpha) \frac{\delta}{2 + \delta}.\end{aligned}$$

Not surprisingly, $k = \omega_\infty([1, 2, 3])$ is maximized for $\delta = 0$ and becomes vanishingly small as δ approaches one. On the other hand, $\omega_\infty([2, 3, 4]) = 0$ for $\delta = 0$. $\omega_\infty([2, 3, 4])$ grows monotonically with δ and reaches its limit value of $\frac{1}{3}(1 - \alpha)$ for $\delta = 1$, where \mathbb{T} and s disconnect.

Now we revisit the examples from section 2.4.1. In all of them, the lumped flow F^R is first constructed from f_{AB}^+ using the procedure described above, then the splitting of F^R into F^P and R according to (2.55) was studied by sampling a random walk of length $N = 10^6$ from P^R , which is long enough to obtain convergence for the quantities of interest (i.e. the probabilistic weights $\omega_\infty(c)$ of the most important pathways) in our examples.

Barbell graph. We revisit the barbell graph from Figure 2.7. As in section 2.4.1, we define a random walk on the barbell graph that goes with probability $p = 0.8$ along edge directions and with probability $(1 - p)$ against edge directions. The resulting effective current f_{AB}^+ was shown in Figure 2.14 and is shown again in Figure 2.21. One observes that f_{AB}^+ has no cycles, hence $\mathcal{C}_\infty^\mathbb{T} = \emptyset$ and consequently $R = 0$ and $F^P = F^R$. The set \mathcal{C}_∞^s has four elements that we may denote as $\{(tt), (tb), (bt), (bb)\}$ where (tt) is the path from a to b that uses the top arc in the left loop and in the right loop, (tb) is the path that uses the top arc in the left and the bottom arc in the right loop, and so on. Almost all of the current is transported by (bt) .




Figure 2.21 – The effective current f_{AB}^+ on the barbell graph with source $A = \{a\}$ and target $B = \{b\}$. The dotted edges carry almost no current.




Figure 2.22 – Transition tube graph. Node coloring indicates q^+ . Edge directions according to F^P . Edge thickness is proportional to the magnitude of F^P . Edge color is proportional to the magnitude of R .

Transition tube with vortices. Now we revisit the transition tube graph with additional vortices from Figure 2.15. In Figure 2.22, the splitting of f_{AB}^+ into the potential part F^P and the rotation R is shown. The edge thickness is proportional to F^P , the edge color is proportional to the magnitude of R with large values in red and small values in green. The only difference to the splitting shown in Figure 2.15 is the way the red edges, i.e. the edges in the vortices, are used by F^P resp. $\text{grad}\Phi$. The red edges which allow transport from A to B are used by F^P , the others are not used. This is because F^P by construction has to respect the sign structure of f_{AB}^+ . In contrast, $\text{grad}\Phi$ reverses the orientation of the red edges which do not allow transport from A to B .

As expected, the set \mathcal{C}_∞^\top has four members with large weights which correspond to the four vortices shown in red in Figure 2.22. The other members in \mathcal{C}_∞^\top are 3 orders of magnitude less likely, and correspond to slightly longer versions of the vortices that also use a few of the green edges. The most important pathways in \mathcal{C}_∞^s are the central, top and bottom pathway which are visible in Figure 2.22. They are all equally likely and together they carry more than 82% of the reactive flow, that is, more than 82% of the reactive trajectories used one of these three pathways.

Random block model. We revisit the random block model from Figure 2.16. On the same realization that was studied in Figures 2.16 and 2.17, the current F^P carried by directed pathways from A to B was sampled. The result is shown in Figure 2.23. The result looks similar to $\text{grad}\Phi$ shown in Figure 2.17: There is almost no flow in C_3 and C_4 , and no flow back from C_5 to C_2 or from C_2 to C_1 . This is to be expected: If a trajectory is e.g. in C_3 , it must come from and go back to C_2 . Because of the strong metastability, it is very likely to loop back on itself in C_2 before going to C_5 . Then all the states it has visited in C_3 become part of a loop which gets




Figure 2.23 – Current $F^P(x, y)$ carried by directed pathways in (2.55). Colour according to q^+ with red denoting $q^+(x) = 1$ and green denoting $q^+(x) = 0$. Node size proportional to the total amount of flow passing through the node (log scale). Edge color proportional to the total amount of flow passing through the edge (log scale).

erased. The biggest difference between F^P and $\text{grad } \Phi$ is that F^P uses only the exterior edges which are present in G and point in the direction of the reaction, while not using edges that point in the direction opposite to the reaction. In contrast, $\text{grad } \Phi$ changes the orientations of exterior edges so that they all point in the direction of the reaction, and then it uses all of them.

Langevin system. We revisit the Langevin system from section 2.3.2. We use the same parameters as for the numerical study in sections 2.3.2 and 2.4.1, so that the discussions therein apply. In particular, the Langevin system is close to the Hamiltonian system, the energy $H(x, p) = \frac{1}{2m}p^2 + V(x)$ is a slow variable, and a typical reactive trajectory consists of many orbits around A followed by a transition from the $x > 0$ to the $x \leq 0$ half-plane followed by many orbits around B . Because of the dynamics (2.40), all reaction pathways must cross from the $x > 0$ to the $x \leq 0$ half-plane in the $p < 0$ region.

Figure 2.24 shows the splitting of f_{AB}^+ into F^P (blue arrows) and R (red arrows). And indeed, the only qualitative difference of this splitting to the one shown in Figure 2.18 is that F^P crosses the $x = 0$ plane only in the $p < 0$ region whereas $\text{grad } \Phi$ crossed the $x = 0$ plane with equal probability in the $p < 0$ and the $p > 0$ region. In this respect, F^P captures the behaviour of the




Figure 2.24 – Langevin system. Productive flow F^P (blue arrows) and rotation R (red arrows). The length of the arrows is proportional to the magnitude of flow transported. The size of the arrows representing F^P has been enlarged by a factor of 2 relative to R for visibility.

reactive trajectories and may tell which reaction pathways the reactive trajectories are using, whereas $\text{grad}\Phi$ does not.

To obtain further insight into the structure of $\mathcal{C}_\infty^\mathbb{T}$, we show the empirical distribution of cycle lengths $d(n) = \sum_{c \in \mathcal{C}_\infty^\mathbb{T}} \omega_\infty(c) \delta(|c| = n)$ in Figure 2.25. The number $Nd(n)$ is the number of times a cycle which has length n and is in $\mathcal{C}_\infty^\mathbb{T}$ was passed through by the realization $(X_n)_{1 \leq n \leq N}$ of length N that we use to sample. The distribution $d(n)$ has a pronounced global maximum at $n \approx 5$ and another local maximum at $n \approx 10$. These correspond to orbits with winding number one and two respectively: The Hamiltonian system completes an orbit with winding number w in time $t = w$. Since the discretization time step is $t = 0.25$ and the Langevin system is close to the Hamiltonian system, we expect orbits of winding number one to take 4 – 5 steps and orbits of winding number two to take 9 – 10 steps. To confirm this intuition, we order the members of $\mathcal{C}_\infty^\mathbb{T}$ according to $\omega_\infty(c)$ in descending order. Figure 2.24 shows the cycles c_1, c_4, c_{13} and c_{15} on the right and confirms that these are orbits around A or B with winding number one. All other cycles c_i with $1 \leq i \leq 20$ also correspond to similar orbits, but they are not shown since they would overlap with the other cycles.

Common observations. When comparing the numerical results for the decomposition (2.48) with those for the decomposition (2.55), we can make the following common observation for all the examples studied here:

- (i) The support of R is essentially the same for both decompositions, but the magnitude of

2.4. Application II: Transition Pathways in Irreversible Markov Processes




Figure 2.25 – Langevin system. Left: Distribution $d(n)$ of lengths of cycles in $\mathcal{C}_\infty^\mathbb{T}$. Right: The sets A and B (light yellow) and the cycles c_1 , c_4 , c_{13} and c_{15} .

R is different,

- (ii) edge directions in $G(\text{grad } \Phi)$ may be different from edge directions in $G(f_{AB}^+)$ while edge directions in $G(F^P)$ and $G(f_{AB}^+)$ are always the same,
- (iii) while F^P is by construction a current generated by pruned reactive trajectories, $\text{grad } \Phi$ cannot be interpreted as such.

These observations lead to the conclusion that the splitting (2.48) is more *geometrical*, i.e. dictated by the graph topology and the metric given by the capacities $c^s(x, y)$, while the splitting (2.55) is more *dynamical*, i.e. dictated by properties of actual realizations of $(X_n)_n$.

3 Optimal Control Theory in Molecular Dynamics

In chapter 2, we have seen that numerical estimators based on reversible processes often display strong metastabilities which results in poor convergence properties. In particular, section 2.2.3 showed that these metastabilities can always be reduced by adding a non-reversible force. In this chapter, we study problems where the additional force term can be adjusted by an external controller that wishes to minimize a certain objective function. These **optimal control problems** have already been introduced in section 1.4. Now we restrict our attention to a special class of control problems, so-called linear quadratic (LQ) control problems, which are dual to sampling problems in e.g. Molecular Dynamics applications. Solving the optimal control problem exactly solves the sampling problem in the sense that it gives rise to a zero-variance estimator. The duality between control and sampling is explained in detail in section 3.1.

The remainder of the chapter is devoted to the approximate solution of LQ control problems. We have to perform a discretization in order to be able to do numerics. We choose a Galerkin method in conjunction with a logarithmic transformation. After the discretization, a control problem for a discrete Markov jump process is obtained. The method is explained in section 3.2. We prove error bounds in section 3.2.2 and study examples in section 3.2.4.

Section 3.3 is concerned with the method of averaging [PS08]. In large and complex systems with metastabilities, one usually has few slow and many fast variables. The method of averaging allows to reduce the problem size by averaging over the fast variables. We demonstrate in section 3.3 that the method also applies to LQ control problems. The averaging method is one way to arrive at discretizations of the type studied in 3.2 which are numerically tractable and still capture the important slow degrees of freedom, even for larger systems. As an example, we study alanine dipeptide in section 3.2.4.

3.1 Optimal Control as a Means to Accelerate MD Sampling

In this chapter, we are concerned with the following problem: Given a (possibly large) state space \mathbb{X} with a dynamics governed by a Markov process $(X_t)_t$ on \mathbb{X} and a **target region** $A \subset \mathbb{X}$ which we assume to be a closed set with smooth boundary ∂A , we are interested in the hitting time $\tau_A = \inf\{t > 0 : X_t \in A\}$. In particular, we are interested in the statistics of τ_A (i.e. mean, variance, higher moments) conditioned on $X_0 \in B$ in a situation where the starting region B is strongly dynamically separated from A so that transitions from B to A are rare events. This is a typical situation in Molecular Dynamics. Here the space $\mathbb{X} = \mathbb{R}^{3N}$ is the configuration space of a large molecule with N atoms, $(X_t)_t$ describes the dynamics of the molecule in its environment (e.g. a water bath), and the regions A and B correspond to two different **molecular conformations**, i.e. different geometric configurations of the molecule which influence its chemical behaviour. For example, A might be an unfolded and B a folded structure. Transitions between A and B are rare events since they happen on a time scale which is several orders of magnitude larger than the simulation time scale one is forced to use.

One way to obtain statistics of a bounded random variable Z is to compute its **cumulant generating function** or **free energy**

$$\gamma_s[Z] = -\log \mathbf{E} \left[e^{-sZ} \right] \quad (3.1)$$

which, by using the Taylor formula around $s = 0$, admits a moment expansion of the form

$$\gamma_s[Z] = s\mathbf{E}[Z] - \frac{s^2}{2}\mathbf{E}[(Z - \mathbf{E}[Z])^2] + \dots \quad (3.2)$$

so that we can compute the n th moment $M_n[Z] = \mathbf{E}[(Z - \mathbf{E}[Z])^n]$ of Z by taking the n th derivative of $\gamma_s[Z]$ around $s = 0$:

$$M_n[Z] = (-1)^{n-1} \left. \frac{d^n}{ds^n} \gamma_s[Z] \right|_{s=0}. \quad (3.3)$$

If we therefore know $\gamma_s[Z]$ for all values s in a small neighbourhood around $s = 0$, then at least theoretically we can compute all moments of Z from (3.3), which gives us complete information about the statistics of Z . In the case we are interested in, $Z = \tau_A$ and $\mathbf{E}[\dots]$ is expectation conditional on $X_0 \in B$. In order to be slightly more general, we will seek the function

$$V(x) := -\varepsilon \log \mathbf{E}_x \left[\exp \left(-\varepsilon^{-1} \int_0^{\tau_A} f(X_s) ds \right) \right] \quad (3.4)$$

for $x \in B$, a temperature parameter $\varepsilon > 0$ and some measurable and nonnegative function $f : \mathbb{X} \rightarrow [0, \infty)$. The cumulant generating function (3.1) of $Z = \tau_A$ is obtained from (3.4) by taking $f = s\mathbf{1}_{\mathbb{X}}$ and averaging $V(x)$ over $x \in B$. In the situations we are interested in, $V(x)$ will always be approximately constant for all $x \in B$ so that the averaging step isn't necessary. Conditioning on a single $x \in B$ has advantages that will become clear shortly.

3.1. Optimal Control as a Means to Accelerate MD Sampling

One way to compute $V(x)$ would be through sampling: Suppose we have obtained N realizations $(X_t^{(n)})_t$ of the dynamics $(X_t)_t$ with $X_0^{(n)} = x$, e.g. via numerical integration. Let $\tau_A^{(n)}$ be the value of τ_A for the n th realization. Then a naive estimator of $V(x)$ would be

$$\hat{V}(x) = -\varepsilon \log \frac{1}{N} \sum_{n=1}^N \exp \left(-\varepsilon^{-1} \int_0^{\tau_A^{(n)}} f(X_s^{(n)}) ds \right). \quad (3.5)$$

However there are three problems with the estimator (3.5): (i) By our assumption, transitions from B to A are rare and thus difficult to observe in a numerical simulation, (ii) because of the nonlinearity of the logarithm, the estimator (3.5) is biased, (iii) because of the exponential inside the expectation value in (3.4) which gives large weights to realizations where the value of the integral inside the exponential is small, the estimator (3.5) has a large variance.

We will instead pursue an entirely different way of computing $V(x)$ in this chapter: It will be shown below that $V(x)$ can be interpreted as the value function of an optimal control problem. The rest of the chapter is then devoted to the development of techniques for solving this optimal control problem approximately. For concreteness, we will assume that the dynamics $(X_t)_t$ is a Markov diffusion (1.17) on $\mathbb{X} = \mathbb{R}^d$ with a generator L given by (1.18), i.e. $(X_t)_t$ evolves according to the SDE

$$dX_t = b(X_t)dt + \sigma(X_t)dW_t, \quad X_0 = x \quad (3.6)$$

with W_t being Brownian motion in \mathbb{R}^m , drift $b: \mathbb{R}^d \rightarrow \mathbb{R}^d$ and diffusion coefficient $\sigma: \mathbb{R}^d \rightarrow \mathbb{R}^{d \times m}$. We further assume L to be uniformly elliptic, see (1.19). As a consequence, the bilinear form

$$(\xi, \eta)_{a(x)} := \sum_{i,j} a_{ij}(x) \xi_i \eta_j \quad \forall \xi, \eta \in \mathbb{R}^d \quad (3.7)$$

is an inner product on \mathbb{R}^d for every $x \in \mathbb{X}$. Consider the following optimal control problem:

LQ control problem. Minimize the functional

$$J^u(x) = \mathbf{E}_x \left[\int_0^{\tau_A} \left(f(X_s^u) + \frac{1}{4} |u_s|^2 \right) ds \right] \quad (3.8)$$

over all $u \in \mathcal{U}_0$ with the space of admissible control strategies \mathcal{U}_0 as in Theorem 1.14 and control space $\mathbb{U} = \mathbb{R}^m$, subject to the dynamics of the controlled process $(X_t^u)_t$ satisfying the SDE

$$dX_t^u = \left(b(X_t^u) + \frac{1}{\sqrt{2\varepsilon}} \sigma^T(X_t^u) u_t \right) dt + \sigma(X_t^u) dW_t, \quad X_0^u = x. \quad (3.9)$$

Thus the control enters linearly in the dynamics (3.9) and quadratically in the cost (3.8). The factor $1/4$ in (3.8) is merely convention. The quadratic penalization term in (3.8) equals the relative entropy $H(Q, P)$ between the reference measure P of the uncontrolled diffusion (3.6)

Chapter 3. Optimal Control Theory in Molecular Dynamics

and the probability measure Q of the controlled process (3.9), as can be shown using Girsanov's theorem [Øks03, Thm. 8.6.8]. In other words, it holds that (cf. [HS12, HBS⁺14]):

$$\mathbf{E}_Q \left[\frac{1}{4} \int_0^{\tau_A} |u_s|^2 ds \right] = \varepsilon \int \log \frac{dQ}{dP} dQ = \varepsilon H(Q, P),$$

cf. the definition (1.9) of the relative entropy.

The rigorous statement of the duality result that now follows requires several strong assumptions, in particular the existence of a classical solution to a linear elliptic PDE. We first state and prove the result and then comment on the assumptions.

Theorem 3.1 (Duality between control and sampling problem.) *Let $(X_t)_t$ be a Markov diffusion on $\mathbb{X} = \mathbb{R}^d$ evolving according to the SDE (3.6) with uniformly elliptic generator L . Let $A \subset \mathbb{X}$ be closed with smooth boundary and such that the hitting time $\tau_A = \inf\{t > 0 : X_t \in A\}$ is a.s. finite and $\kappa := \sup_{x \in \mathbb{X} \setminus A} \mathbf{E}_x [\tau_A] < \infty$. Assume further that L , A and $K \geq f \geq 0$ with $K > 0$ are such that the PDE*

$$\begin{aligned} L\phi(x) - \varepsilon^{-1} f(x)\phi(x) &= 0 && \text{on } \mathbb{X} \setminus A, \\ \phi(y) &= 1 && y \in \partial A. \end{aligned} \tag{3.10}$$

has a bounded solution $\phi \in C^2(\mathbb{X} \setminus A)$. Let $V : \mathbb{X} \setminus A \rightarrow \mathbb{R}$ be given by

$$V(x) = -\varepsilon \log \mathbf{E}_x \left[\exp \left(-\varepsilon^{-1} \int_0^{\tau_A} f(X_s) ds \right) \right].$$

Then $V \in C^2(\mathbb{X} \setminus A)$, and $V = \inf_{u \in \mathcal{U}_0} J^u$ is the value function of the LQ control problem (3.8), (3.9). The minimizer $u^* = \operatorname{argmin} J^u$ is unique and given by the feedback law

$$u_t^* = \alpha^*(X_t^u), \quad \alpha^*(x) = -\sqrt{2}\varepsilon^{-1/2} \sigma^T \nabla V(x).$$

The strategy u^* is in \mathcal{U}_0 if $\mathbf{E}_x [\tau_A] < \infty$ under $(X_t^{u^*})_t$ and $u_t^* \leq K'$ for all $t \leq \tau_A$ and some $K' > 0$. If this holds then u^* is an optimal Markov control strategy, i.e. we have $V(x) = J^{u^*}(x)$ and $V(x) \leq J^u(x)$ for all $u \in \mathcal{U}_0$.

Proof. By the Feynman-Kac formula (1.23) in Theorem 1.4, the solution ϕ to (3.10) is given by

$$\phi(x) = \mathbf{E}_x \left[\exp \left(-\varepsilon^{-1} \int_0^{\tau_A} f(X_s) ds \right) \right], \tag{3.11}$$

and it follows that

$$V(x) = -\varepsilon \log \phi(x). \tag{3.12}$$

Next we show that ϕ is uniformly bounded away from zero: From $f \leq K$, Jensen's inequality

3.1. Optimal Control as a Means to Accelerate MD Sampling

and the assumption $\kappa := \sup_{x \in \mathbb{X} \setminus A} \mathbf{E}_x[\tau_A] < \infty$, we have

$$\phi(x) \geq \mathbf{E}_x[\exp(-\varepsilon^{-1}K\tau_A)] \geq \exp(-\varepsilon^{-1}K\mathbf{E}_x[\tau_A]) \geq \exp(-\varepsilon^{-1}K\kappa) > 0,$$

thus ϕ is uniformly bounded away from zero, and since $\phi \in C^2(\mathbb{X} \setminus A)$, we can take derivatives in (3.12) and conclude that $V \in C^2(\mathbb{X} \setminus A)$. From (3.12), we obtain by the chain rule and (1.18)

$$\varepsilon \frac{L\phi}{\phi} = -LV + \varepsilon^{-1}(\nabla V, \nabla V)_{a(x)}. \quad (3.13)$$

Now consider the relation

$$-\varepsilon^{-1}(\nabla V, \nabla V)_{a(x)} = \min_{\alpha \in \mathbb{R}^m} \left\{ \frac{1}{\sqrt{2\varepsilon}} \alpha \cdot \sigma^T \nabla V + \frac{1}{4} |\alpha|^2 \right\}, \quad (3.14)$$

which holds in view of $a = \frac{1}{2}\sigma\sigma^T$. If we combine (3.10) with (3.13) and (3.14), we see that V solves the boundary value problem

$$\begin{aligned} \min_{\alpha \in \mathbb{R}^m} \left\{ LV(x) + \frac{1}{\sqrt{2\varepsilon}} \alpha \cdot \sigma^T \nabla V + f(x) + \frac{1}{4} |\alpha|^2 \right\} &= 0 && \text{on } \mathbb{X} \setminus A, \\ V(y) &= 0 && y \in \partial A. \end{aligned} \quad (3.15)$$

Now the statements of the Theorem all follow from the verification Theorem (1.14). The only thing left to show is that the minimum $\alpha^*(x)$ exists and the corresponding Markov control strategy $u_t^* = \alpha^*(X_t^u)$ is in \mathcal{U}_0 . That the minimum exists and is equal to $\alpha^*(x) = -\sqrt{2\varepsilon}^{-1/2} \sigma^T \nabla V(x)$ is clear from (3.14). To show that $u_t^* \in \mathcal{U}_0$, we need that $\tau_A < \infty$ a.s. under $(X_t^{u^*})_t$, which holds by assumption, and $\mathbf{E}_x[\int_0^{\tau_A} \nabla V(X_s^u) \cdot \sigma(X_s^u) dW_s] = 0$. The latter follows from the assumption $u_t^* \leq K'$ for all $t \leq \tau_A$ [vH07]. ■

Remark: On the assumptions in Theorem 3.1. The question of existence of a classical solution $\phi \in C^2(\mathbb{X} \setminus A)$ to the PDE (3.10) is difficult in general. We are guaranteed existence of a classical solution ϕ if one of the following sets of assumptions holds:

- (i) a, b and f are in $C^\infty(\mathbb{X} \setminus A)$, the boundary ∂A is in C^∞ , and A is such that $\mathbb{X} \setminus A$ is bounded [EL98]. The latter may be achieved by setting $A := A' \cup E$ where E is the exterior of the bounded region of state space we are interested in, and A' is the compact target region we want to reach. We have different options of specifying boundary conditions on $\partial A' = \partial A \cup \partial E$. We may use Neumann boundary conditions on E , i.e. replace (3.10) with

$$\begin{aligned} L\phi(x) - \varepsilon^{-1}f(x)\phi(x) &= 0 && \text{on } \mathbb{X} \setminus A', \\ \phi(y) &= 1 && y \in \partial A, \\ \frac{\partial \phi}{\partial n}(y) &= 0 && y \in \partial E. \end{aligned} \quad (3.16)$$

This translates to Neumann boundary conditions $\frac{\partial V}{\partial n} = 0$ on E in (3.15). Alternatively we may put Dirichlet boundary conditions $\phi(y) = \exp(-\varepsilon^{-1}C)$ for $y \in \partial E$ with some $C > 0$, which leads to Dirichlet boundary conditions $V(y) = C$ for $y \in \partial E$ in (3.15). If E is sufficiently far away and such that the probability of reaching E before A is negligible, then these modifications to (3.10) do not change the solution ϕ much in the region of interest.

- (ii) Instead of (3.10), we may consider the following parabolic PDE on $Q = O \times [t_0, t_1]$ with bounded $O \subset \mathbb{X}$:

$$\begin{aligned} L\phi - \varepsilon^{-1}f\phi &= \frac{\partial \phi}{\partial t} && \text{on } O \times [t_0, t_1), \\ \phi(y, t) &= \exp(-\varepsilon^{-1}g(y)) && y \in \partial O, t \in [t_0, t_1), \\ \phi(y, t_1) &= \exp(-\varepsilon^{-1}C) && y \in O. \end{aligned}$$

Here, $C > 0$ is some (possibly large) constant. Then a solution $\phi \in C^{1,2}(Q) \cap \phi(\bar{Q})$ exists [FS06] which is given by a Feynman-Kac formula similar to (1.22). The functional J^u to be minimized in the control problem changes to

$$J^u(x) = \mathbf{E}_x \left[\int_0^\tau \left(f(X_s^u) + \frac{1}{4}|u_s|^2 \right) ds + z(X_\tau, \tau) \right]$$

where τ is the first exit time of Q , $z(y, t) = g(y)$ for $y \in \partial O$ and $z(y, t_1) = C$ for $y \in O$. As before, we may separate $\partial O = \partial A \cup \partial E$ where ∂A is the boundary of the target region and ∂E is the boundary of the region of interest in the state space \mathbb{X} , and we may set $g(y) = 0$ for $y \in \partial A$ and $g(y) = C$ for $y \in \partial E$. Formally, we may take the limit $t_1 \rightarrow \infty$ to recover (3.16). For very large t_1 exiting Q is the same as exiting O , and ϕ becomes time-independent.

Concerning the additional assumptions needed to ensure that $u_t^* \in \mathcal{U}_0$, we have that since $\alpha^* = -\sqrt{2\varepsilon}\sigma^T \nabla V$ is in $C^1(\mathbb{X} \setminus A)$, the condition $u_t^* \leq K'$ for all $t \leq \tau$ automatically holds if we restrict $\mathbb{X} \setminus A$ to a bounded region by one of the procedures outlined above. The condition that $\mathbf{E}_x[\tau_A] < \infty$ under $(X_t^{u^*})_t$ also automatically holds in this case, this follows from the uniform ellipticity of a [vH07].

Special case: Overdamped Langevin dynamics. An interesting special case is when the dynamics of the uncontrolled process is given by the overdamped Langevin equation

$$dX_t = -\nabla v(X_t)dt + \sqrt{2\varepsilon}dW_t \tag{3.17}$$

that was introduced in (1.25). We denote the potential by v in order to avoid confusion with the value function V . The control space \mathbb{R}^m equals the state space \mathbb{R}^d here. The generator L acts on functions $u \in C^2(\mathbb{X})$ as $Lu = \varepsilon \Delta u - (\nabla v) \cdot \nabla u$, and is uniformly elliptic with $\sigma = \sqrt{2\varepsilon} \mathbf{1}$

3.1. Optimal Control as a Means to Accelerate MD Sampling

and $a = \varepsilon \mathbf{1} > 0$. The controlled equation (3.9) in Theorem 3.1 becomes

$$dX_t^u = (-\nabla v(X_t^u) + u_t) dt + \sqrt{2\varepsilon} dW_t.$$

The HJB equation (3.15) becomes

$$\begin{aligned} \min_{\alpha \in \mathbb{R}^d} \left\{ L^\alpha V(x) + f(x) + \frac{1}{4} |\alpha|^2 \right\} &= 0 && \text{on } \mathbb{X} \setminus A, \\ V(y) &= 0 && y \in \partial A. \end{aligned}$$

with $L^\alpha = L + \alpha \cdot \nabla$. The optimal control u^* is given by $u_t^* = \alpha^*(X_t^u)$ with $\alpha^*(x) = -2\nabla V(x)$, and thus the optimally controlled process $(X_t^*)_t$ evolves according to the SDE

$$dX_t^* = -\nabla U(X_t^*) dt + \sqrt{2\varepsilon} dW_t \quad (3.18)$$

with the **effective potential** $U(x) = v(x) + 2V(x)$.

3.1.1 Duality Between Control and Path Sampling for Jump Processes

The duality result of Theorem 3.1 between a path sampling problem and an optimal control problem via a logarithmic transformation also holds for Markov jump processes. This construction goes back to [She85], and we repeat it here in condensed form, see also [FS06, Sec. VI.9]. In the following, let $(\hat{X}_t)_{t \geq 0}$ be an ergodic MJP on the discrete state space $\hat{\mathbb{X}} = \{1, \dots, n\}$ with infinitesimal generator $G \in \mathbb{R}^{n \times n}$, and define the stopping time $\tau_A = \inf\{t > 0: \hat{X}_t \in A\}$ to be the first hitting time of a subset $A \subset \hat{\mathbb{X}}$. Consider the following optimal control problem:

LQ control problem for MJPs. Minimize the functional

$$\hat{J}^v(i) = \mathbf{E}_i \left[\int_0^{\tau_A} \{ \hat{f}(\hat{X}_s^v) + k(\hat{X}_s^v, v_s) \} ds \right] \quad (3.19)$$

over all controls $v \in \mathcal{U}_0$ with \mathcal{U}_0 as in section 1.4.3. The controls v take values in the control space $\mathbb{U} = \{ \alpha : \hat{\mathbb{X}} \rightarrow (0, \infty) \}$. The function $\hat{f} : \hat{\mathbb{X}} \rightarrow \mathbb{R}$ is assumed to be nonnegative. If $v_s = \alpha$ for $s \in [t, t')$, then the generator of the controlled process $(\hat{X}_s^v)_s$ for times $s \in [t, t')$ is given by

$$G^\alpha = (G_{ij}^\alpha)_{i,j \in \hat{\mathbb{X}}}, \quad G_{ij}^\alpha = \frac{G_{ij} \alpha(j)}{\alpha(i)}. \quad (3.20)$$

The control-dependent costs $k(i, \alpha)$ in (3.19) are given by

$$k(i, \alpha) = \varepsilon (G^\alpha (\log \alpha))(i) - \varepsilon \frac{(G\alpha)(i)}{\alpha(i)}. \quad (3.21)$$

Chapter 3. Optimal Control Theory in Molecular Dynamics

The next lemma records some important properties of the control problem (3.19), (3.20). This will also make the connection to the continuous case clear.

Lemma 3.2 *Let the LQ control problem (3.19), (3.20) be defined as above.*

- (i) *Let G be reversible with unique stationary distribution μ . Then $\mu^\alpha(i) = Z_\alpha^{-1} \alpha^2(i) \mu(i)$, with Z_α an appropriate normalization constant, is the unique probability distribution such that G^α is reversible with stationary distribution μ^α .*
- (ii) *Let \hat{P} denote the probability measure on the space of trajectories generated by \hat{X}_t with initial condition $\hat{X}_0 = i$, and let \hat{Q} be the corresponding probability measure generated by \hat{X}_t^v with the same initial condition $\hat{X}_0^v = i$. Then \hat{Q} is absolutely continuous with respect to \hat{P} and the expected value of the running cost k is the relative entropy (1.9) between \hat{Q} and \hat{P} , i.e.,*

$$\mathbf{E}_{\hat{Q}} \left[\int_0^{\tau_A} k(\hat{X}_s^v, v_s) ds \right] = \int \log \frac{d\hat{Q}}{d\hat{P}} d\hat{Q} = \varepsilon H(\hat{Q}, \hat{P})$$

where $\mathbf{E}_{\hat{Q}}[\dots]$ is the expectation over all realizations of \hat{X}_t^v starting at $\hat{X}_0^v = i$.

Proof. We first show (i). By assumption we have $\mu(i)G_{ij} = \mu(j)G_{ji}$. Now, let μ^α be such that $\mu^\alpha(i)G_{ij}^\alpha = \mu^\alpha(j)G_{ji}^\alpha$. We will show that μ^α has the proposed form:

$$\mu^\alpha(i)G_{ij}^\alpha = \frac{\alpha(j)}{\alpha(i)} \frac{\mu^\alpha(i)}{\mu(i)} \mu(i)G_{ij} = \frac{\alpha(j)}{\alpha(i)} \frac{\mu^\alpha(i)}{\mu(i)} \mu(j)G_{ji} = \frac{\alpha^2(j)}{\alpha^2(i)} \frac{\mu(j)}{\mu(i)} \frac{\mu^\alpha(i)}{\mu^\alpha(j)} \mu^\alpha(j)G_{ji}^\alpha$$

But since $\mu^\alpha(i)G_{ij}^\alpha = \mu^\alpha(j)G_{ji}^\alpha$, we must have

$$\frac{\mu^\alpha(j)}{\mu(j)\alpha^2(j)} = \frac{\mu^\alpha(i)}{\mu(i)\alpha^2(i)} \quad \forall i \neq j.$$

This can only be true if the quantity $Z_\alpha^{-1} := \frac{\mu^\alpha(i)}{\mu(i)\alpha^2(i)}$ is independent of i . This gives $\mu^\alpha(j) = Z_\alpha^{-1} \alpha^2(j) \mu(j)$ as desired. The constant Z_α is uniquely determined by the requirement that μ^α be normalized. Finally, from reversibility it follows directly that μ^α is also a stationary distribution of G^α .

To show (ii), note that the running cost $k(i, \alpha)$ can be written as

$$k(i, \alpha) = \varepsilon \sum_{j \neq i} G_{ij} \left\{ \frac{\alpha(j)}{\alpha(i)} \left[\log \frac{\alpha(j)}{\alpha(i)} - 1 \right] + 1 \right\}, \quad (3.22)$$

the integral of which is the relative entropy between \hat{Q} and \hat{P} (see [DPMR96, Sec. 3.1.4]). The absolute continuity between \hat{Q} and \hat{P} simply follows from the fact that α in the definition of G^α was required to be component-wise strictly positive. ■

3.1. Optimal Control as a Means to Accelerate MD Sampling

Lemma 3.2 explains the rather unusual form (3.21) of the control-dependent cost $k(i, \alpha)$: As for the continuous problem (3.8), (3.9), the expectation value of k is just the relative entropy between the probability measures \hat{Q} of the controlled process $(\hat{X}_t^v)_t$ and \hat{P} of the uncontrolled process $(X_t)_t$. In the continuous case the relative entropy had the simple form of a quadratic penalty term. Here it is more complicated, but structurally still the same.

Lemma 3.2 also shows that the generator G^α is still reversible for all controls $\alpha \in \mathbb{U}$. In particular, note the following: Suppose μ is of the form $\mu(i) = Z^{-1} \exp(-\varepsilon^{-1} \hat{v}(i))$ of a Boltzmann distribution for some potential $\hat{v} : \hat{\mathbb{X}} \rightarrow \mathbb{R}$. Then the stationary distribution of the optimally controlled process is given by

$$\mu^{\alpha^*}(i) \propto \exp(-\varepsilon^{-1}(\hat{v}(i) + 2\hat{V}(i))),$$

i.e. μ^{α^*} is a Boltzmann distribution in the effective potential $\hat{U} = \hat{v} + 2\hat{V}$. This is in complete analogy to the continuous case, cf. (3.18). We have the following duality result:

Theorem 3.3 (Duality for MJPs.) *Let $(\hat{X}_t)_t$ be an ergodic MJP on $\hat{X} = \{1, \dots, n\}$ with generator G , $\tau_A = \inf\{t > 0 : \hat{X}_t \in A\}$ for some $A \subset \hat{X}$ and $f : \hat{X} \rightarrow [0, \infty)$. Let $\varepsilon > 0$. Define*

$$\hat{V}(i) = -\varepsilon \log \mathbf{E}_i \left[\exp \left(-\frac{1}{\varepsilon} \int_0^{\tau_A} \hat{f}(\hat{X}_s) ds \right) \right].$$

Then \hat{V} is the value function of the LQ control problem (3.19), (3.20). The optimal control $v^ = \operatorname{argmin} \hat{J}^v$ exists and is unique, and is given by*

$$v_t^* = \alpha^*(X_t^v), \quad \alpha^*(i) = \exp(-\varepsilon^{-1} \hat{V}(i)).$$

Proof. Consider the linear boundary value problem

$$\begin{aligned} \sum_{j \in \hat{\mathbb{X}}} G_{ij} \hat{\phi}(j) - \varepsilon^{-1} \hat{f}(i) \hat{\phi}(i) &= 0, \quad i \in \hat{\mathbb{X}} \setminus A \\ \hat{\phi}(i) &= 1, \quad i \in A. \end{aligned} \tag{3.23}$$

By Theorem 1.7, the unique solution to (3.23) is given by

$$\hat{\phi}(i) = \mathbf{E}_i \left[\exp \left(-\frac{1}{\varepsilon} \int_0^{\tau_A} \hat{f}(\hat{X}_s) ds \right) \right]$$

and thus $\hat{V} = -\varepsilon \log \hat{\phi}$, and by multiplying (3.23) with $\hat{\phi}^{-1}(i)$, we see that \hat{V} satisfies the equation

$$\begin{aligned} \exp(\hat{V}/\varepsilon) G \exp(-\hat{V}/\varepsilon) - \varepsilon^{-1} \hat{f} &= 0, \quad i \in \hat{\mathbb{X}} \setminus A \\ \hat{V}(i) &= 0, \quad i \in A. \end{aligned}$$

The exponential term in the above equation for \hat{V} can be recast as

$$\exp(\hat{V}/\varepsilon)G\exp(-\hat{V}/\varepsilon) = \varepsilon^{-1} \min_{\alpha>0} \{-(G^\alpha \hat{V})(i) + k(i, \alpha)\} \quad (3.24)$$

using the identity $\min_{y \in \mathbb{R}} \{e^{-y} + ay\} = a - a \log a$ for $a > 0$ and the definition (3.21) of $k(i, \alpha)$. As a consequence, (3.23) is equivalent to

$$\begin{aligned} \min_{\alpha \in \mathbb{U}} \{ (G^\alpha \hat{V})(i) + k(i, \alpha) + \hat{f}(i) \} &= 0, \quad i \in \hat{\mathbb{X}} \setminus A \\ \hat{V}(i) &= 0, \quad i \in A. \end{aligned} \quad (3.25)$$

which is the dynamic programming equation of a Markov decision problem. It readily follows from (3.24) that the minimizer in (3.25) exists and is given by $\alpha^* = \exp(-\varepsilon^{-1}\hat{V})$. The rest follows from the verification theorem 1.16. ■

3.2 MSM Discretizations of LQ Type Control Problems

3.2.1 Galerkin Projection Point of View

In this section we will develop a discretization for the LQ control problem (3.8), (3.9) discussed in Section 3.1. The discretization will approximate the continuous control problem with a control problem for a Markov jump process on finite state space. We will focus on overdamped dynamics (3.17). Overdamped dynamics is reversible, we thus study reversible dynamics with external control. This work has been published in [SBHS13, BH14].

Because of the nonlinearity of the problem, a general theory for discretizing continuous optimal control problems is unavailable. However, we saw in Section 3.1 that for the control problems we are interested in, a logarithmic transform to a linear PDE (3.10) is available. For linear PDEs, discretization theory in terms of Galerkin projections onto finite-dimensional subspaces of the PDE solution space exists. Our strategy will therefore be the one indicated in Figure 3.1, where the term 'linear PDE' refers to the PDE (3.10), and the term 'constrained linear system' refers to the linear system (3.23).

In the first part of this section, we will develop the Galerkin projection for general subspaces and obtain some control of the discretization error. To refine this control, we specify the subspace D we project onto. As the state space is unbounded and possibly high-dimensional, a grid-based discretization is prohibitive. Here we suggest a meshless discretization based on an incomplete partition of state space into so called *core sets*, that are the metastable regions of the uncontrolled dynamics. We will prove an error bound which gives us detailed control over the discretization error, even if very few basis functions are used. We should mention that clearly other choices are possible, such as radial basis functions or moving least-squares, but for metastable and high-dimensional systems like in molecular dynamics or chemical reaction kinetics using core sets and the associated basis of committor functions is beneficial.

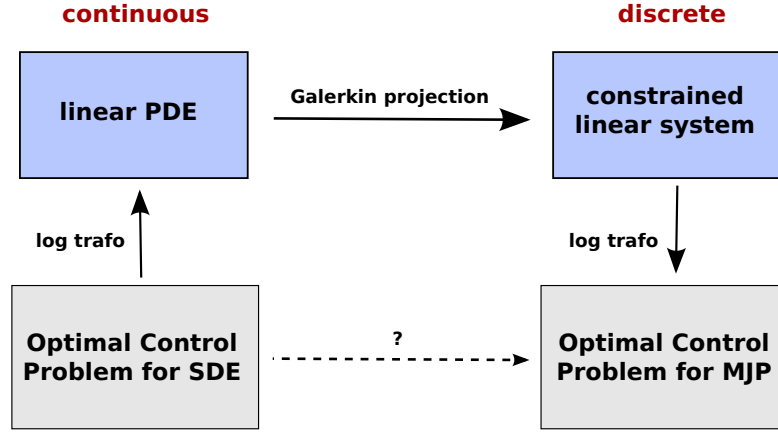


Figure 3.1 – Discretization of continuous control problems via a logarithmic transform.

In the second part of this section, we will develop the stochastic interpretation of the resulting matrix equation as the backward Kolmogorov equation of a MJP. We will study the resulting discrete control problem and make a connection to Transition Path Theory [VE06] and core set MSMs [SNL⁺11].

3.2.2 Galerkin Projection of the Dirichlet Problem

We consider the boundary value problem

$$\begin{aligned} (L - \varepsilon^{-1}f)\phi(x) &= 0, & x \in \mathbb{X} \setminus A \\ \phi(x) &= 1, & x \in \partial A. \end{aligned} \tag{3.26}$$

with $L = \varepsilon\Delta - \nabla v \cdot \nabla$ being the uniformly elliptic generator of (3.17) and $f : \mathbb{X} \rightarrow [0, K]$ bounded and continuous. We declare that $\phi|_A = 1$, so that the domain of ϕ is \mathbb{X} . In order to be able to apply theorem 3.1 and to have a problem which is amenable to numerical discretization, we think of \mathbb{X} as being bounded and impose homogeneous Neumann boundary conditions $\frac{\partial \phi}{\partial n} = 0$ on ∂X , see Figure 3.2.

Following standard references (e.g. [Bra07]) we construct a Galerkin projection of (3.26). For this purpose, we introduce the L^2 -based Sobolev space H^1 with norm $\|\phi\|_{H^1} = \|\nabla u\|_{\mu}^2 + \|u\|_{\mu}^2$ and the Hilbert spaces $V = \{\psi \in L^2(\mathbb{X}, \mu), \|\psi\|_{H^1} < \infty\}$ and $V_0 = \{\psi \in V, \psi|_{\partial A} = 0\}$. We further define the symmetric and positive bilinear form

$$\mathcal{B} : V \times V \rightarrow \mathbb{R}, \quad \mathcal{B}(\phi, \psi) = \varepsilon^{-1} \langle f\phi, \psi \rangle_{\mu} + \varepsilon \langle \nabla \phi, \nabla \psi \rangle_{\mu}.$$




Figure 3.2 – The domain \mathbb{X} and target set A .

Now if ϕ is a solution of (3.26), then it also solves the weak problem¹

$$\mathcal{B}(\phi, \psi) = 0 \quad \forall \psi \in V_0. \quad (3.27)$$

A Galerkin solution $\hat{\phi}$ is now any function satisfying

$$\mathcal{B}(\hat{\phi}, \hat{\psi}) = 0 \quad \forall \hat{\psi} \in D_0, \quad (3.28)$$

with D_0 being a suitable finite dimensional subspace V_0 . Specifically, we choose basis functions $\chi_1, \dots, \chi_{n+1}$ with the following properties:

- (S1) The functions $\chi_i: \mathbb{X} \rightarrow \mathbb{R}$ are in V .
- (S2) The χ_i form a partition of unity, that is $\sum_{i=1}^{n+1} \chi_i = \mathbf{1}$.
- (S3) The χ_i satisfy $\chi_{n+1}|_A = 1$ and $\chi_i|_A = 0$ for $i \in \{1, \dots, n\}$.

All elements of $D_0 := \text{lin}\{\chi_1, \dots, \chi_n\}$ will satisfy homogeneous Dirichlet boundary conditions in (3.26), and we will sometimes write $D := \chi_{n+1} \oplus D_0$ and think of the Galerkin solution $\hat{\phi}$ as an element in D . Now define the matrices

$$F_{ij} = \frac{\langle \chi_i, f \chi_j \rangle_\mu}{\langle \chi_i, \mathbf{1} \rangle_\mu}, \quad K_{ij} = -\varepsilon \frac{\langle \nabla \chi_i, \nabla \chi_j \rangle_\mu}{\langle \chi_i, \mathbf{1} \rangle_\mu}.$$

¹Neumann boundary conditions of the form $\frac{\partial \phi}{\partial n} = g_N$ on $\partial \mathbb{X}$ only contribute to the RHS of (3.27) with a term $\langle g_N, \psi \rangle_{L^2(\partial \mathbb{X})}$ [Bra07]. For homogeneous Neumann boundary conditions $g_N = 0$ and therefore this contribution is identically zero.

3.2. MSM Discretizations of LQ Type Control Problems

Setting $\hat{\phi} = \sum_i \hat{\phi}_i \chi_i$, the weak form (3.28) becomes a matrix equation for the unknown coefficients $\hat{\phi}_i$:

$$\begin{aligned} \sum_{j=1}^{n+1} (K_{ij} - \varepsilon^{-1} F_{ij}) \hat{\phi}_j &= 0, \quad i \in \{1, \dots, n\} \\ \hat{\phi}_{n+1} &= 1, \end{aligned} \tag{3.29}$$

which is the discretization of (3.26).

In order to control the discretization error of the Galerkin method, we choose a norm $\|\cdot\|$ on V and introduce the two error measures:

1. The *Galerkin error* $\varepsilon = \|\phi - \hat{\phi}\|$, i.e. the difference between original and Galerkin solution measured in $\|\cdot\|$.
2. The *best approximation error* $\varepsilon_0 = \inf_{\hat{\psi} \in D} \|\phi - \hat{\psi}\|$, i.e. the minimal difference between the solution ϕ and any element $\hat{\psi} \in D$.

In order to obtain full control over the discretization error, we need bounds on ε , and we will get them by first obtaining a bound on the performance $\mathbf{p} := \varepsilon/\varepsilon_0$ and then a bound on ε_0 . The latter will depend on the choice of subspace D . For the former, standard estimates assume the following $\|\cdot\|$ -dependent properties of A :

- (i) Boundedness: $\mathcal{B}(\phi, \psi) \leq \alpha_1 \|\phi\| \|\psi\|$ for some $\alpha_1 > 0$
- (ii) Ellipticity: for all $\phi \in V$ holds $\mathcal{B}(\phi, \phi) \geq \alpha_2 \|\phi\|^2$ for some $\alpha_2 > 0$.

If both (i) and (ii) hold, Céa's lemma states that $\mathbf{p} \leq \frac{\alpha_1}{\alpha_2}$, see e.g. [Bra07]. For the *energy norm* $\|\phi\|_{\mathcal{B}}^2 := \mathcal{B}(\phi, \phi)$ we have $\alpha_1 = \alpha_2 = 1$ and therefore $\mathbf{p} = 1$, thus the Galerkin solution $\hat{\phi}$ is the best-approximation to ϕ in the energy norm.

Performance bound

The next two statements give a bound on \mathbf{p} if errors are measured in the L^2 -norm. In this case, $\mathcal{B}(\cdot, \cdot)$ is still elliptic but possibly unbounded. Later in this section, we will specify the bound on ε_0 for a specific Galerkin basis.

Theorem 3.4 *Let \mathcal{B} be elliptic. Further let*

$$Q: L^2(\mathbb{X}, \mu) \rightarrow D_0 \subset L^2(\mathbb{X}, \mu), \quad Qw = \sum_{i=1}^n \langle \chi_i, w \rangle \chi_i$$

be the orthogonal projection onto D_0 . Then

$$p^2 = \left(\frac{\boldsymbol{\varepsilon}}{\boldsymbol{\varepsilon}_0} \right)^2 \leq 1 + \frac{1}{\alpha_2^2} \sup_{v \in V} \frac{\|QB(\mathbf{1} - Q)v\|_\mu^2}{\|v\|_\mu^2}$$

where $B = \varepsilon^{-1}f - L$ is the linear operator associated with $\phi \mapsto \mathcal{B}(\cdot, \phi)$.

Proof. In Appendix C. ■

Remark 3.5 Note that $\|QB(\mathbf{1} - Q)v\|_\mu \leq \|QBv\|_\mu$ is always finite even though B is possibly unbounded since $v \in V \subset L^2(\mathbb{X}, \mu)$ and Q is the projection onto a finite-dimensional subspace of $L^2(\mathbb{X}, \mu)$.

The bottom line of Theorem 3.4 is that if B leaves the subspace D almost invariant, then $\hat{\phi}$ is almost the best-approximation of ϕ in $\|\cdot\|_\mu$. The following lemma gives a more detailed description. In the following, we will write $\|\cdot\| = \|\cdot\|_\mu$ for convenience.

Lemma 3.6 Let $Q^\perp = \mathbf{1} - Q$ and define

$$\delta_L := \max_k \|Q^\perp L \chi_k\|, \quad \delta_f := \max_k \|Q^\perp \varepsilon^{-1} f \chi_k\|$$

to be the maximal projection error of the images of the χ_k 's under L and f . Then

$$\|QBQ^\perp\| = \|Q^\perp BQ\| \leq (\delta_L + \delta_f) \sqrt{\frac{n}{m}}$$

where m is the smallest eigenvalue of \hat{M} .

Proof. The first statement is true since A is essentially self-adjoint. For the second statement, first of all

$$\|Q^\perp BQ\| = \|Q^\perp (\varepsilon^{-1}f - L)Q\| \leq \|Q^\perp \varepsilon^{-1}fQ\| + \|Q^\perp LQ\|$$

holds from the triangle inequality. We now bound the term involving L . Notice that for $\hat{\phi} = \sum_i \hat{\phi}_i \chi_i \in D$:

$$\|Q^\perp L\hat{\phi}\| = \left\| \sum_i \hat{\phi}_i Q^\perp L \chi_i \right\| \leq \delta_L \sum_i |\hat{\phi}_i| = \delta_L \|\hat{\phi}\|_1.$$

Then, with $\hat{M} := (\langle \chi_i, \chi_j \rangle_\mu)_{ij}$:

$$\|Q^\perp LQ\| = \sup_{\phi = \phi_{\parallel} + \phi_{\perp} \in V} \frac{\|Q^\perp L\phi_{\parallel}\|}{\|\phi\|} \leq \sup_{\phi_{\parallel} \in D} \frac{\|Q^\perp L\phi_{\parallel}\|}{\|\phi_{\parallel}\|} \leq \delta_L \sup_{\hat{\phi} \in \mathbb{R}^n} \frac{\|\hat{\phi}\|_1}{\sqrt{\langle \hat{\phi}, \hat{\phi} \rangle_M}}$$

A similar result holds for the term involving f . The statement now follows from a standard equivalence between finite-dimensional norms, $\|\hat{\phi}\|_1 \leq \sqrt{n} \|\hat{\phi}\|_2$, and the fact that \hat{M} is symmetric, which implies that $\langle \hat{\phi}, \hat{\phi} \rangle_M = \hat{\phi}^T \hat{M} \hat{\phi} \geq m \hat{\phi}^T \hat{\phi} = m \|\hat{\phi}\|_2^2$. ■

3.2. MSM Discretizations of LQ Type Control Problems

To summarize, Theorem 3.4 and Lemma 3.6 give us a formula for the projection performance p which states that

$$p^2 \leq 1 + \frac{n}{m} \frac{(\delta_L + \delta_f)^2}{\alpha_2^2}.$$

How large or small δ_f is will depend on the behaviour of f , e.g., if $f = \text{const}$ then $\delta_f = 0$. Both δ_f and δ_L are always finite even though L is possibly unbounded.

Choice of basis functions and bound on ϵ_0

Case A: Linear basis functions. We let the χ_i be piecewise linear basis functions. Specifically, in $d = 1$ we use the **Lagrange basis** on a regular grid with grid points x_i and grid spacing h , given by

$$\chi_i(x) = \begin{cases} 1 + \frac{1}{h}(x - x_i) & x \in [x_{i-1}, x_i) \\ 1 - \frac{1}{h}(x - x_i) & x \in [x_i, x_{i+1}), \\ 0 & \text{else.} \end{cases}$$

See also Figure 3.3. The derivatives of these basis functions are particularly simple. A simple computation shows that

$$K_{i,i\pm 1} = -\varepsilon \frac{\langle \nabla \chi_i, \nabla \chi_j \rangle_\mu}{\langle \chi_i, \mathbf{1} \rangle_\mu} = \frac{\varepsilon}{h^2} \left[\exp(-\varepsilon^{-1}(v(x_{i\pm 1/2}) - v(x_i))) + \mathcal{O}(h^2) \right] \quad (3.30)$$

where $x_{i\pm 1/2} = x_i \pm h/2$, and

$$F_{i,i\pm 1} = \frac{\langle \chi_i, \chi_j \rangle_\mu}{\langle \chi_i, \mathbf{1} \rangle_\mu} = \frac{1}{6} f(x_{i\pm 1/2}) \exp(-\varepsilon^{-1}(v(x_{i\pm 1/2}) - v(x_i))) + \mathcal{O}(h^2) \quad (3.31)$$

with $F_{ii} = 1 - F_{i,i+1} - F_{i,i-1}$. It follows from standard finite element theory [Bra07] that the L^2 best-approximation error ϵ_0 is given by

$$\epsilon_0 = Ch^2 \quad (3.32)$$

with a constant $C > 0$.

Case B: Committor basis functions. We now generalize results [DSS12] on the approximation quality of MSMs for reversible equilibrium diffusions and estimate the best-approximation error ϵ_0 for the case that the subspace D is spanned by committor functions associated with the metastable sets of the dynamics. To this end suppose that the potential $V(x)$ has $n + 1$ deep minima x_1, \dots, x_{n+1} . Let C_1, \dots, C_{n+1} be convex *core* sets around x_1, \dots, x_{n+1} and such that $A = C_{n+1}$. We write $C = \cup_{i=1}^{n+1} C_i$ and $\mathbb{T} = \mathbb{X} \setminus C$ and introduce $\tau_C = \inf\{t \geq 0 : X_t \in C\}$. We




Figure 3.3 – Left: Lagrange basis functions χ_i and χ_{i+1} . Right: Committor basis associated to the core sets C_1, C_2, C_3 .

take χ_i to be the committor function associated with the set C_i , that is

$$\chi_i(x) = \mathbf{P}(X_{\tau_C} \in C_i | X_0 = x). \quad (3.33)$$

See also Figure 3.3. These χ_i satisfy the assumptions (S2)–(S3) and (S1) except on the core set boundaries, which is a set of measure zero. Since we do not have a grid parameter, by which the approximation error can be controlled, standard PDE techniques for bounding ϵ_0 fail. Indeed, typically we will have very few basis functions compared to a grid-like discretization. The following theorem gives a bound on ϵ_0 .

Theorem 3.7 *Let Q be the orthogonal projection onto the subspace D spanned by the committor functions (3.33), and let ϕ be the solution of (3.26). Then we have*

$$\epsilon_0 = \|Q^\perp \phi\|_\mu \leq \|P^\perp \phi\|_\mu + \mu(\mathbb{T})^{1/2} [\kappa \|f\|_\infty + 2\|P^\perp \phi\|_\infty]$$

where $\|\cdot\| = \|\cdot\|_\mu$, $\kappa = \sup_{x \in \mathbb{T}} \mathbf{E}_x[\tau_C]$, and P is the orthogonal projection onto $V_c = \{v \in L^2(\mathbb{X}, \mu), v|_{C_i} = \text{const on every } C_i\} \subset L^2(\mathbb{X}, \mu)$, with $P^\perp = \mathbf{1} - P$.

Proof. In Appendix C.2. ■

In Theorem 3.7, κ is the maximum expected time of hitting the metastable set from outside (which is short). Note further that $P^\perp \phi = 0$ on T . The errors $\|P^\perp \phi\|_\mu$ and $\|P^\perp \phi\|_\infty$ measure how constant the solution ϕ is on the core sets. Theorem 3.7 suggest the following strategy to minimize ϵ_0 : (i) Place a core set C_i in every metastable region where ϕ is expected to be almost constant, (ii) place core sets in regions with high invariant density μ in order to minimize $\mu(\mathbb{T})$. This strategy requires knowledge of the invariant density μ . Identifying the metastable regions requires additional dynamical information. If this is not available, then a good guess is usually to use the deepest wells of μ .

Theorem 3.7 together with Theorem 3.4 gives us full control over the discretization error ϵ . These error bounds are along the lines of MSM projection error bounds [SNS10], [DSS12].

Remark 3.8 *It would be nice to have an error estimate also for the value function. In general such an estimate is difficult to get, because of the nonlinear logarithmic transformation $V = -\varepsilon \log \phi$ involved. However we know that ϕ and its discrete approximation are both uniformly bounded and bounded away from zero. Hence the logarithmic transformation is uniformly Lipschitz continuous on its domain, which implies that the L^2 error bounds holds for the value function with an additional prefactor given by the Lipschitz constant squared; for a related argument see [HLPZ14]*

3.2.3 Interpretation in Terms of a Markov Decision Problem

We derive an interpretation of the discretized equation (3.29) in terms of a MDP. We introduce the diagonal matrix Λ with entries $\Lambda_{ii} = \sum_j F_{ij}$ (zero otherwise) and the full matrix $G = K - \varepsilon^{-1}(F - \Lambda)$, and rearrange (3.29) as follows:

$$\begin{aligned} \sum_{j=1}^{n+1} (G_{ij} - \varepsilon^{-1} \Lambda_{ij}) \hat{\phi}_j &= 0, \quad i \in \{1, \dots, n\} \\ \hat{\phi}_{n+1} &= 1, \end{aligned} \tag{3.34}$$

This equation can be given a stochastic interpretation. To this end let us introduce the vector $\pi \in \mathbb{R}^{n+1}$ with nonnegative entries $\pi_i = \langle \chi_i, \mathbf{1} \rangle$ and notice that $\sum_i \pi_i = 1$ follows immediately from the fact that the basis functions χ_i form a partition of unity, i.e. $\sum_i \chi_i = \mathbf{1}$. This implies that π is a probability distribution on the discrete state space $\hat{\mathbb{X}} = \{1, \dots, n+1\}$. We summarize properties of the matrices K , F and G :

Lemma 3.9 *Let K , G , F and π be as above.*

(i) *K is a generator matrix (i.e. K is a real-valued square matrix with row sum zero and positive off-diagonal entries) with stationary distribution π that satisfies detailed balance*

$$\pi_i K_{ij} = \pi_j K_{ji}, \quad i, j \in \hat{\mathbb{X}}$$

(ii) *$F \geq 0$ (entry-wise) with $\pi_i F_{ij} = \pi_j F_{ji}$ for all $i, j \in \hat{\mathbb{X}}$.*

(iii) *G has row sum zero and satisfies $\pi^T G = 0$ and $\pi_i G_{ij} = \pi_j G_{ji}$ for all $i, j \in \hat{\mathbb{X}}$.*

(iv) *If $\|f\|_\infty \leq -\varepsilon^2 \frac{\langle \nabla \chi_i, \nabla \chi_j \rangle_\mu}{\langle \chi_i, \chi_j \rangle_\mu}$ for all $i \neq j$, then $G_{ij} \geq 0$ for all $i \neq j$. In this case equation (3.34) admits a unique and strictly positive solution $\hat{\phi} > 0$.*

Proof. (i) follows from $\sum_i \chi_i(x) = \mathbf{1}$ and reversibility of L : We have $\sum_i \pi(i) K_{ij} = \sum_i \langle \chi_i, L \chi_j \rangle_\mu = \langle L \mathbf{1}, \chi_j \rangle_\mu = 0$ and $\pi(i) K_{ij} = \langle \chi_i, L \chi_j \rangle_\mu = \langle L \chi_i, \chi_j \rangle_\mu = \pi(j) K_{ji}$. (ii) follows from $f(x)$ being real and positive for all x . As for (iii), G has row sum zero by (i) and the definition of Λ . $\pi(i) G_{ij} = \pi(j) G_{ji}$ follows from (i), (ii) and the fact that Λ is diagonal, and $\pi^T G = 0$ follows

directly. For (iv), first note that by the assumption on $\|f\|_\infty$, we have

$$F_{ij} = \varepsilon^{-1} \langle \chi_i, f \chi_j \rangle_\mu \leq \|f\|_\infty \langle \chi_i, \chi_j \rangle_\mu \leq -\varepsilon \langle \nabla \chi_i, \nabla \chi_j \rangle_\mu = K_{ij},$$

so that $G_{ij} \geq 0$ for all $i \neq j$. Existence, uniqueness and strict positivity of the solution $\hat{\phi}$ to (3.34) then all follow from Theorem 1.7. ■

It follows that if the running costs f are such that (iv) in Lemma 3.9 holds, then G is a generator matrix of a MJP that we shall denote by $(\hat{X}_t)_{t \geq 0}$, and by Theorem 1.7, (3.34) has a unique and positive solution of the form

$$\hat{\phi}(i) = \mathbf{E} \left[\exp \left(-\varepsilon^{-1} \int_0^{\tau_A} \hat{f}(\hat{X}_s) ds \right) \middle| \hat{X}_0 = i \right]$$

with $\hat{f}(i) = \Lambda_{ii}$ and $\tau_A = \inf\{t \geq 0 \mid \hat{X}_t = i + 1\}$. We can then apply the Duality Theorem 3.3 to interpret $\hat{V} = -\varepsilon \log \hat{\phi}$ as the value function of the Markov decision problem (3.19), (3.20). This completes the construction of the discrete control problem. Note that in general² $G \neq K$, but both K and G are reversible with stationary distribution π .

Elber's milestoning process

The discretized equation admits a useful stochastic representation, by which its coefficients can be computed without knowing the committor functions. Define the *forward milestoning process* \tilde{X}_t^+ to be in state $\tilde{X}_t^+ = i$ if X_t visits core set C_i next, and the *backward milestoning process* \tilde{X}_t^- to be in state $\tilde{X}_t^- = i$ if X_t came from C_i last. Then the discrete costs can be written as

$$\hat{f}(i) = \frac{1}{\pi_i} \langle \chi_i, f \sum_j \chi_j \rangle = \int v_i(x) f(x) dx = \mathbf{E}_\mu [f(X_t) \mid \tilde{X}_t^- = i] \quad (3.35)$$

where $v_i(x) = \pi_i^{-1} \chi_i(x) \mu(x) = \mathbf{P}(X_t = x \mid \tilde{X}_t^- = i)$ is the probability density of finding the system in state x given that it came last from i . Hence $\hat{f}(i)$ is the average costs conditioned on the information $\tilde{X}_t^- = i$, i.e. X_t came last from A_i , which is the natural extension to the full partition case where $\hat{f}(i)$ was the average costs conditioned on the information that $X_t \in A_i$.

The matrix $K_{ij} = \pi_i^{-1} \langle \chi_i, L \chi_j \rangle$ is reversible with stationary distribution

$$\pi_i = \langle \chi_i, \mathbf{1} \rangle = \mathbf{P}_\mu(\tilde{X}_t^- = i)$$

and is related to so called *core MSMs*. To see this, define the core MSM transition matrix P^τ with components $P_{ij}^\tau = \mathbf{P}(\tilde{X}_{t+\tau}^+ = j \mid \tilde{X}_t^- = i)$, and the mass matrix M with components $M_{ij} = \mathbf{P}(\tilde{X}_t^+ = j \mid \tilde{X}_t^- = i)$. Then, it is not hard to show that for reversible processes we have

²In case of a full partition, $\cup_i C_i = \mathbb{X}$, the χ_i become stepfunctions and $K = G$ is the generator of a full partition MSM. Our method then becomes a finite volume method. Stepfunctions are not regular enough to be in V however.

$P_{ij}^\tau = \pi_i^{-1} \langle \chi_i, T^\tau \chi_j \rangle_\mu$ and $M_{ij} = \pi_i^{-1} \langle \chi_i, \chi_j \rangle_\mu$ so that

$$K = \frac{1}{\pi_i} \langle \chi_i, L \chi_j \rangle_\mu = \lim_{\tau \rightarrow 0} \frac{1}{\tau} (P^\tau - M).$$

Thus K is formally³ the generator of the P^τ . If the core sets are chosen as the metastable states of the system, K can be sampled directly from \tilde{X}_t^\pm . See [SNS10, SNL⁺11] for more details on the construction and sampling of core MSMs. F can also be sampled using

$$F_{ij} = \mathbf{E}_\mu \left[f(X_t) \chi_{\{\tilde{X}_t^+ = j\}} \middle| \tilde{X}_t^- = i \right] \quad (3.36)$$

Therefore, as in the construction of core MSMs, we do not need to compute committor functions explicitly.

3.2.4 Numerical Results

We will present two examples to illustrate the approximation of LQ-type stochastic control problems based on a sparse Galerkin approximation using MSMs.

1D triple well potential

To begin with we study diffusion in the triple well potential which is presented in Figure 3.5a. This potential has three minima at $x_{0/1} = \mp 3.4$ and $x_2 = 0$. We take $\mathbb{X} = [-4, 4]$ and choose $A = [x_0 - \delta, x_0 + \delta]$ with $\delta = 0.2$ as the target set and the running cost $f = f_0 = \text{const}$, such that the control goal is to steer the particle into C_0 in minimum time. In Figure 3.5a the potential v and effective potential U are shown for $\varepsilon = 0.5$ and $f_0 = 0.08$ (solid lines), cf. equation (3.18). One can observe that the optimal control lifts the second and third well up such that the system is driven into C_0 quickly.

First we validate our method with a convergence test using linear finite elements as basis functions χ_i . To do so, we compute a reference solution $\hat{\phi}$ of (3.29) using linear finite elements on a uniform grid with spacing $h_r = 10^{-4}$. The resulting interpolation $\phi_I = \sum_i \chi_i \hat{\phi}(i)$ is very close to the true solution ϕ of (3.26). We also compute a reference solution for the value function $V_I = \sum_i \chi_i \hat{V}(i)$ with $\hat{V} = -\varepsilon \log \hat{\phi}$ and for the optimal control $u_I = -2\nabla V_I$. Then we compute coarser solutions $\phi_{I,h}$ using various grid spacings $1 \geq h \geq 10^{-3}$ and compute the L^2 error $\|\phi_{I,h} - \phi_I\|_\mu$, and L^2 errors for V_I and u_I similarly. The result is shown in Figure 3.4. The L^2 error of $\phi_{I,h}$ is quadratic in h , as expected from the theory. Additionally, the L^2 error of $V_{I,h}$ is also quadratic in h which, given that the transformation between ϕ and V is nonlinear, is surprising. The error of $u_{I,h}$ is only linear in h ; as expected one order of convergence is lost due to the fact that u_I is the gradient of V_I .

Next we use a committor basis. In accordance with the strategy to minimize minimize ε_0 in

³The P^τ do not form a semigroup since $M \neq \mathbf{1}$, thus K cannot be interpreted as i.e. the generator of \tilde{X}_t^- . However, the entries of K are the transition rates between the core sets as defined in transition path theory [VE06].




Figure 3.4 – L^2 error of $\phi_{l,h}$, $V_{l,h}$ and $u_{l,h}$ using linear finite elements (dashed lines) and using the committor basis (circles).

Theorem 3.7, we placed core sets $C_i = [x_i - \delta, x_i + \delta]$ in each of the three wells of the potential shown in Figure 3.5a, resulting in the set of three basis functions shown in Figure 3.5b. The L^2 errors achieved by solving (3.29) in this basis are shown as circles in Figure 3.4. We observe that the 3 committor functions achieve the same performance as linear finite elements with grid spacing $h \approx 0.2$, which corresponds to ≈ 40 basis functions. Theorem 3.7 gives $\epsilon_0 \leq 0.08$, while the actual error is one order of magnitude smaller. The dashed line in Figure 3.5a gives the approximation to the effective potential U calculated in the committor basis, which is in good agreement to the reference solution. In Figure 3.5c the optimal control u (solid line) and its approximation u_l (dashed line) are shown. The core sets are shown in blue. The jumps in u_l at the left boundaries of the core sets are due to the fact that the committor functions are only piecewise C^1 .

The computations so far require explicit knowledge of the basis functions χ_i to compute the matrices K and F . For high-dimensional systems the committor basis is usually not explicitly known. To mimic this situation, we construct a core MSM to sample the matrices K and F . 100 trajectories of length $T = 20,000$ were used to build the MSM. In Figure 3.5d, the optimal cost starting from the rightmost well $V(x_1)$ and its estimate using the core MSM are shown for $\epsilon = 0.5$ and different values of f_0 . Each of the 100 trajectories has seen about four transitions. For comparison, a direct sampling estimate of $V(x_1)$ using the same data is shown (green). The direct sampling estimate suffers from a large bias and variance. In contrast, the MSM estimator for $V(x_1)$ performs well for all considered values of σ . The constant C which ensures

$\hat{\phi} > 0$ when $f_0 \leq C$ is approximately 0.2 in this case. This seems restrictive but still allows to capture all interesting information about ϕ and $V(x_1)$.

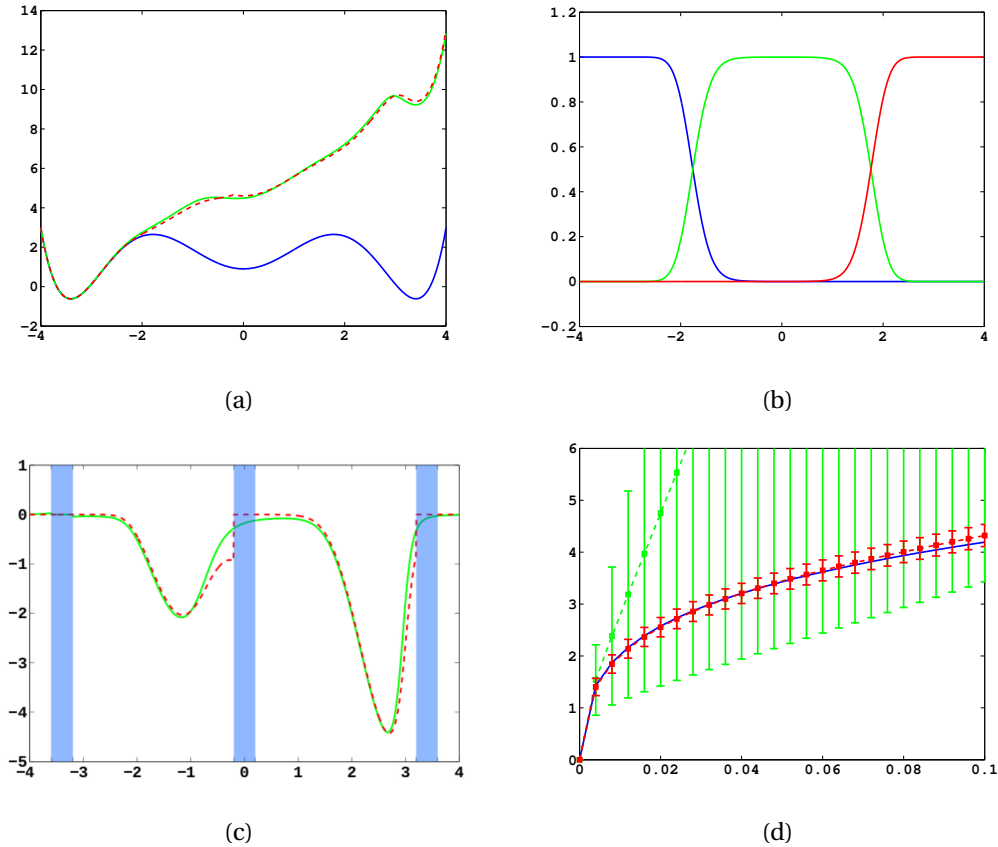


Figure 3.5 – Three well potential example for $\varepsilon = 0.5$ and $f_0 = 0.08$. (a) Potential $v(x)$ (blue), effective potential $U = v + 2V$ (green) and approximation of U with committors (dashed red). (b) The three committors. (c) The optimal control $u(x)$ (solid line) and its approximation (dashed line). Core sets are shown in blue. (d) Optimal cost $V(x_1)$ for $\varepsilon = 0.5$ as a function of f_0 . Blue: Exact solution. Red: Core MSM estimate. Green: Direct sampling estimate.

Rugged 1D triple well potential

We now make the potential in Figure 3.5a rugged by adding a sinusoidal function. This is an important test case since in many applications, in particular in MD, potential energy landscapes are not smooth and instead have many local minima. We choose the same control target, the same parameters and the same core sets C_0 , C_1 and C_2 as in the previous section. As before, we compute a reference solution on a uniform grid with grid spacing $h_r = 10^{-4}$. Then we compute coarser solutions with linear finite elements, and we compare with the committor basis based on the three core sets C_0 , C_1 and C_2 . The rugged potential $v(x)$, the effective potential $U(x) = v(x) + 2V(x)$ and the approximation computed with the committor basis are shown in Figure 3.7a. Qualitatively, the solution computed from the committor basis seems

to perform quite well and is able to track both the large-scale behaviour and the small-scale oscillations of $U(x)$.




Figure 3.6 – L^2 error of $\phi_{l,h}$, $V_{l,h}$ and $u_{l,h}$ using linear finite elements (dashed lines) and using the committor basis (circles) for the rugged potential. The right set of circles corresponds to the small committor basis with 3 basis functions. The left set of circles corresponds to the larger committor basis with 15 additional basis functions.

Finally, we want to investigate if it is possible to improve upon the performance of the committor basis by enlarging it. To do so, we put 13 further core sets C_0^2, \dots, C_{12}^2 (the upper index should indicate that these are 'level 2' core sets) equidistantly in the region $\mathbb{X} \setminus (C_0 \cup C_1 \cup C_2)$. No information about the position of the local minima is used when placing the additional core sets, since it is unrealistic to assume that this is possible. The level 1 and level 2 core sets are also shown in Figure 3.7a. Then, a larger committor basis is computed based on the level 1 core sets C_0, C_1, C_2 together with the level 2 core sets C_0^2, \dots, C_{12}^2 .

The numerical errors for the finite element method and both the small and large committor basis are all summarized in Figure 3.6. We observe that the linear finite elements perform a lot worse when compared to the results for the smooth potential in Figure 3.4: The quadratic (for ϕ and v) resp. linear (for u) convergence only happens for $h < 10^{-1}$, and we lose two orders of magnitude in accuracy for the same value of h when compared to Figure 3.4. This is because the small oscillations in the potential happen on the length scale $h = 10^{-1}$, and in order to achieve good accuracy, the linear basis functions $\chi_{i,h}$ have to resolve them. In contrast, the results computed with the small committor basis do not lose much accuracy when compared to the smooth potential. In terms of the error in ϕ , the small committor basis

3.2. MSM Discretizations of LQ Type Control Problems

achieves the same performance as linear finite elements with $h \approx 0.2$, as before. In terms of u , the performance is as good as linear finite elements with $h \approx 0.1$. In terms of V , the performance is as good as linear finite elements with $h \approx 2 \cdot 10^{-2}$, corresponding to ≈ 400 basis functions. The rugged potential is an example of a dynamics with slow and fast variables. We study these systems in more detail in section 3.3.

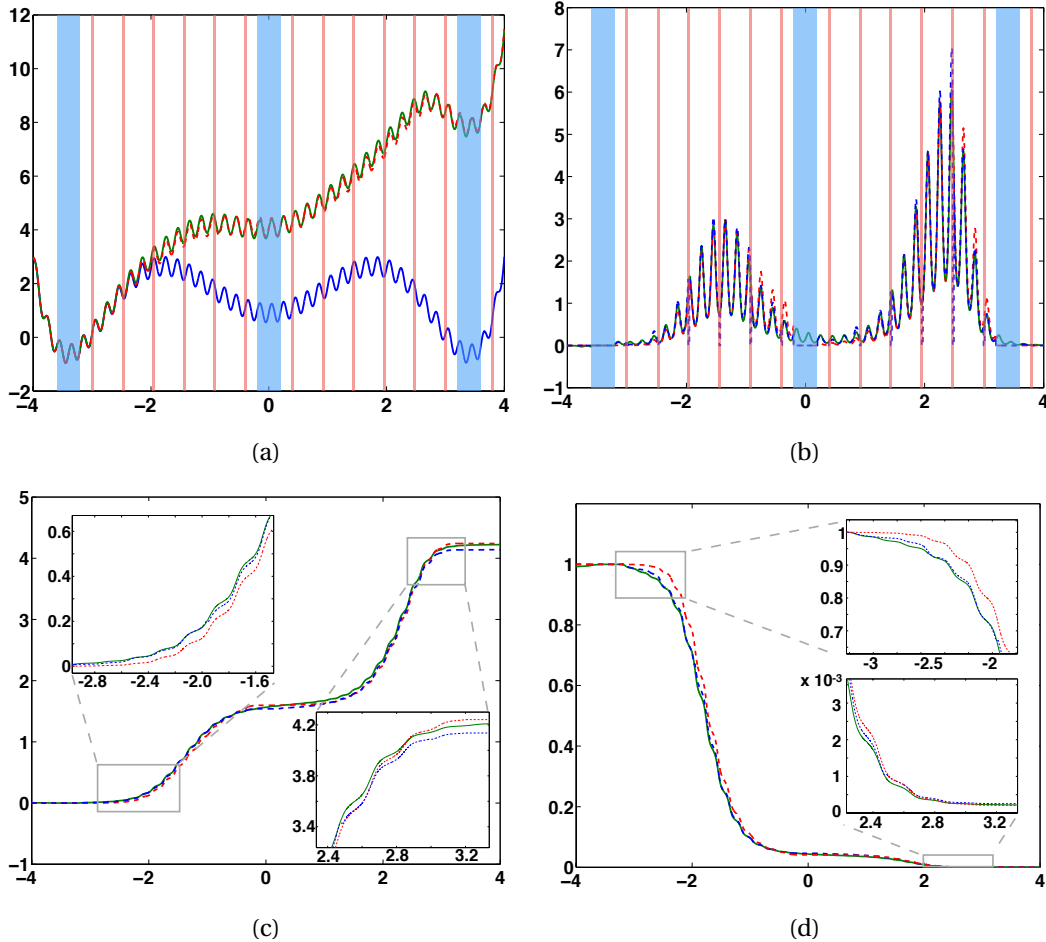


Figure 3.7 – Rugged three well potential example for $\varepsilon = 0.5$ and $f_0 = 0.08$. (a) Potential $v(x)$ (blue), effective potential $U = v + 2V$ (green) and approximation of U with small committor basis (dashed red). First level core sets C_0, C_1, C_2 are shown in blue, second level core sets C_2^0, \dots, C_{12}^2 are shown in red. (b)-(d) Comparison of reference solution (green line) with small (red dashed line) and large committor basis approximation (blue dashed line). (b) Optimal control, (c) value function, (d) solution ϕ of the boundary value problem.

Note however that the bound given by Theorem 3.7 does decrease by adding more core sets: Theorem 3.7 gives $\varepsilon_0 \leq 0.1$ for the small committor basis, and $\varepsilon_0 \leq 0.012$ for the large committor basis. In summary, the small committor basis still performs well for the rugged potential. However, while adding more core sets and thus basis functions decreases both the bound on ε_0 given by Theorem 3.7 and the actual error in ϕ , we cannot guarantee that the error in V and/or u also decreases.

When using the large committor basis, the error in ϕ decreases, but the error in u and V increase. To see why this happens, we compare the reference solution and the approximations given by the small and large committor basis for u , V and ϕ in Figures (3.7b) - (3.7d). We can see that the control oscillates very strongly, and while the approximation from the large committor basis can track the oscillations a bit better outside the core sets, it also forced to zero more often, i.e. on all the additional core sets. The value function improves near C_0 if one goes to the large committor basis, but deteriorates around C_2 . This is understood if one considers the error in ϕ , which is the error that is controlled by Theorem 3.7: The increase in error near C_2 is of the order 10^{-4} and is barely visible. In contrast, the decrease in error near C_0 is of the order 10^{-1} and vastly outweighs the increase near C_2 . The problem for the value function is the large **relative** error near C_2 since ϕ becomes very small, while the Galerkin discretization only controls the **absolute** error.

Alanine dipeptide

As a second, non-trivial example we study conformational transitions in Alanine dipeptide (ADP), a well-studied test system in molecular dynamics. We performed an all-atom simulation of ADP in explicit water (TIP3P) with the Amber FF99SB force field [HAO⁺06] using the GROMACS 4.5.5 simulation package [VDSLH⁺05]. The simulations were performed in the NVT ensemble, where the temperature was restrained to 300 K using the V-Rescale thermostat [BDP07]. 20 trajectories of 200ns with 100ps equilibration runs were simulated. Covalent bonds to hydrogen atoms were constrained using the LINCS algorithm11 [HBBF97] (lincs iter = 1, lincs order = 4), allowing for an integration timestep of 2 fs. The leap-frog integrator was used. Lennard-Jones interactions were cut off at 1 nm. Electrostatic interactions were treated by the Particle-Mesh Ewald (PME) algorithm12 [DYP93] with a real space cut-off of 1 nm, a grid spacing of 0.15 nm, and an interpolation order of 4. Periodic boundary conditions were applied in the x, y, and z-direction.

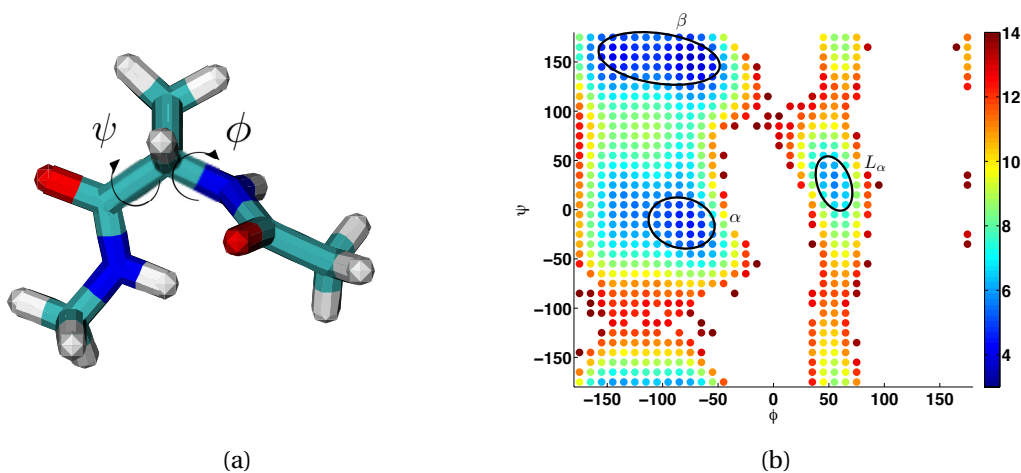


Figure 3.8 – (a) Alanine dipeptide. (b) Free energy $g_i = -\log \pi_i$.

3.2. MSM Discretizations of LQ Type Control Problems

In Figure 3.8a, a cartoon of the molecule is shown. The full system including the water molecules has about 4000 degrees of freedom. However, it is well known that the conformational dynamics, which are the slowest dynamical processes in the system, can be monitored via the backbone dihedral angles ϕ and ψ . The dynamics along the other degrees of freedom happens on much faster time scales. We show in section 3.3 that this time scale separation implies that the value function is essentially a function of ϕ and ψ and constant in all other directions. We will use this to build a Markov State Model which partitions the $\phi - \psi$ -plane.

Validation of the MSM approximation

We construct a full partition MSM using a uniform clustering into 36×36 boxes A_i of size $10^\circ \times 10^\circ$ in the $\phi - \psi$ -plane, and we use characteristic basis functions $\chi_i(x) = \mathbf{1}_{A_i}$ for the discretization⁴. Figure 3.8b shows the free energy $g_i = -\log \pi_i = -\log \mathbf{P}[X_t \in A_i]$ together with the three largest molecular conformations α , β and L_α . The missing boxes have not seen any data. The slowest dynamical process is the switching between the left-handed L_α structure and the right-handed α and β sheet structures. As is customary in MSM theory [SNS10], we estimate the slowest *implied timescale* (ITS) as follows: For different lagtimes τ , we construct the MSM transfer operator $P_{ij}^\tau = \pi_i^{-1} \langle \chi_i, T^\tau \chi_j \rangle_\mu$ and compute

$$t_1(\tau) = -\frac{\tau}{\log \lambda_1(\tau)}$$

where $\lambda_1(\tau)$ is the 2nd largest eigenvalue of P^τ . The result is shown in Figure 3.9a. We observe a plateau for $6\text{ps} \leq \tau \leq 30\text{ps}$ which indicates that the time-discrete snapshots $\hat{X}_n^\tau := \hat{X}_{n\tau}$ are well described by a Markov chain with transition matrix P^τ for $\tau \geq 6\text{ps}$. The plateau is used to compute $t_1 = 1560\text{ps} \pm 6\text{ps}$. For smaller values of τ , \hat{X}_n^τ is not Markovian due to recrossing effects. For this reason we also cannot sample K directly and have to work with the finite-time transfer operator P^τ instead. Before proceeding to the optimal control problem, we study the effect of this time discretization on the mean first passage time (MFPT) $t(x) = \mathbf{E}_x[\tau_{\alpha \cup \beta}]$ where $\tau_{\alpha \cup \beta}$ is the first hitting time of $\alpha \cup \beta$. Since the α , β and L_α conformations are very metastable, $t(x)$ is almost constant on L_α and $t_{L_\alpha} = \mathbf{E}[t(x)|x \in L_\alpha]$ can be computed from t_1 via $t_{L_\alpha} = t_1 / \pi_{\alpha \cup \beta}$ where $\pi_{\alpha \cup \beta} = 0.96$ is the invariant measure of the α and β conformation combined [SNL⁺11]. This gives $t_{L_\alpha} = 1626\text{ps} \pm 6\text{ps}$.

On the other hand, let $N_{\alpha \cup \beta}^\tau = \inf\{n > 0 : \hat{X}_n^\tau \in \alpha \cup \beta\}$. If the chain $(\hat{X}_n^\tau)_n$ is Markovian, then the time-discrete MFPT $\hat{t}^\tau(x) = \mathbf{E}_x[\tau N_{\alpha \cup \beta}^\tau]$ satisfies the matrix equation

$$(P^\tau - I)\hat{t}^\tau = -\tau \text{ outside } \alpha \cup \beta, \quad \hat{t}^\tau = 0 \text{ in } \alpha \cup \beta. \quad (3.37)$$

⁴Let $C_i \subset A_i$ be a core set in cell A_i . The characteristic basis functions can be obtained from the committors by expanding C_i to fill out all of A_i . This basis is used to construct full partition MSMs and produces particularly simple sampling formulas. The characteristic functions are not in V , but the sampling results are similar to what one would obtain using committors with cores C_i which fill out most of A_i .

Additionally, since $\tau_{\alpha\cup\beta} \in (\tau(N_{\alpha\cup\beta}^T - 1), \tau N_{\alpha\cup\beta}^T]$, we should expect that

$$\hat{t}_{L_\alpha}^\tau := \mathbf{E}[\hat{t}^\tau(x)|x \in L_\alpha] = t_{L_\alpha} + c\tau, \quad c \leq 1. \quad (3.38)$$

In Figure 3.9b, $\hat{t}_{L_\alpha}^\tau$ obtained by solving (3.37) is shown as a function of τ together with t_{L_α} . A linear interpolation using the values for $6\text{ps} \leq \tau \leq 30\text{ps}$ where the Markov assumption holds gives $\hat{t}^\tau = \hat{t}^0 + 0.8\tau$ with $\hat{t}^0 = 1628\text{ps}$, which is consistent with (3.38). This shows that the time discretization introduces only small, controllable errors for $\tau \geq 6\text{ps}$. In the following, we will work with $\tau = 10\text{ps}$. Notice that the time discretization amounts to the replacement $K \rightarrow \tau^{-1}(P^\tau - I)$ in (3.29).




Figure 3.9 – (a) Slowest implied timescale $t_1(\tau)$ (red) and average using the values for $6\text{ps} \leq \tau \leq 30\text{ps}$ (green). (b) Time-discrete MFPT $\hat{t}_{L_\alpha}^\tau$ as a function of τ (red), linear interpolation of $\hat{t}_{L_\alpha}^\tau$ (blue) and the reference value t_{L_α} (green). The confidence interval of t_{L_α} is shown as dashed lines.

Controlled transition to the α -helical structure

Next we consider an optimal control problem for steering the molecule into the α -structure. We choose as the target region $A = \alpha$ and define running costs in the (ϕ, ψ) variables as $f(\phi, \psi) = f_0 + f_1 \|\psi - \psi_\alpha\|^2$ where $\|\cdot\|$ is a simple metric on the torus, and we choose $f_0 = 0.01$ and $f_1 = 0.001$ representing a mild penalty for being away from the target region in the ψ -direction. We discretize this control problem using the same partition and time discretization as for the MSM construction in section 3.2.4 and sample P^τ and F from the MSM data. The resulting value function $\hat{V} = \log \hat{\phi}$ is shown in Figure 3.10a. Since the basis functions χ_i are not differentiable and some data is missing in \hat{V} , we have to construct an interpolation $V_I(\phi, \psi)$ from the point data \hat{V} to obtain an estimate for the optimal control force $u(\phi, \psi) = -2\nabla V_I(\phi, \psi)$. An interpolation based on a Delauney triangulation which is C^1 everywhere except at the data points is shown in Figure 3.10b.

To demonstrate that adding the control force $u(\phi, \psi)$ has the effect of speeding up the transi-

tion from L_α to α , we would have to implement it in the MD simulation software. We leave that for future work. We can make a prediction of the anticipated effect within the MSM framework: In accordance with (3.20), we compute the transition matrix $P_{\alpha^*}^\tau$ of the optimally controlled process by $P_{\alpha^*}^\tau(i, j) = P^\tau(i, j) \frac{\alpha^*(j)}{\alpha^*(i)}$ with $\alpha^* = \hat{\phi}$ for $i \neq j$. The discretized MFPT vector \hat{t}^* of the optimally controlled process can be computed from the Matrix equation

$$(P_{\alpha^*}^\tau - I) \hat{t}^* = -\tau \text{ outside } \alpha, \quad \hat{t}^* = 0 \text{ in } \alpha.$$

The result is shown in Figure 3.10c and gives a speed up compared to \hat{t}_0 of one order of magnitude. A larger speed up could easily be achieved by increasing f . In Figure 3.10d we show the free energy of the controlled process in log scale, which according to Lemma 3.2 is given by $g^{\alpha^*} = -\log \pi^{\alpha^*} = \log Z_{\alpha^*} - 2\hat{V} - \log \pi$. Observe that the L_α and β conformations are now much less populated compared to the equilibrium distribution in Figure 3.8b: As in the 1D example, the control mainly has the effect of lifting the wells which are not in the target region up such that they become less metastable.

3.3 Multiscale Problems: Averaging

In this section, we discuss the averaging approach for MJPs⁵ as it was described in [PS08] and apply it to the optimal control problems discussed earlier. This approach is useful when systems with many degrees of freedom and a strong time scale separation into a 'slow' and a 'fast' dynamics are considered. One can then average over the fast dynamics in order to obtain a reduced problem in terms of a MJP that only acts on the slow dynamics. This is the justification for working in the (ϕ, ψ) -space in the numerical example discussed in section 3.2.4: The Markov Model of the Molecular dynamics simulation of ADP does exhibit a strong time scale separation with the slow variables being the dihedral angles ϕ and ψ . We follow the reference textbook [PS08]. See also [PW14, HP14] for further details.

Following [PS08], we consider a MJP on finite state space \mathbb{X} with a generator G that admits a separation of the form

$$G = \frac{1}{\varepsilon} G_0 + G_1$$

with a small parameter $0 < \varepsilon \ll 1$ that controls the time scale separation. We will let \mathbb{X} be indexed by two variables x and y with x standing for the **slow** and y for the **fast** variables. The leading order contribution G_0 corresponds to fast ergodic dynamics in y with x frozen.

The precise situation is as follows: The state space is $\mathbb{X} = \mathbb{X}_x \times \mathbb{X}_y$. We assume \mathbb{X} to be finite and denote by $\|\cdot\| = \|\cdot\|_\infty$ the sup-norm on \mathbb{X} . We let $g((x, x'), (y, y'))$ be the element of G associated with a transition from $(x, y) \in \mathbb{X}_x \times \mathbb{X}_y$ to $(x', y') \in \mathbb{X}_x \times \mathbb{X}_y$, and similarly for G_0 and

⁵Everything discussed in this section works for Markov chains as well, if one replaces the generator by $P - I$.

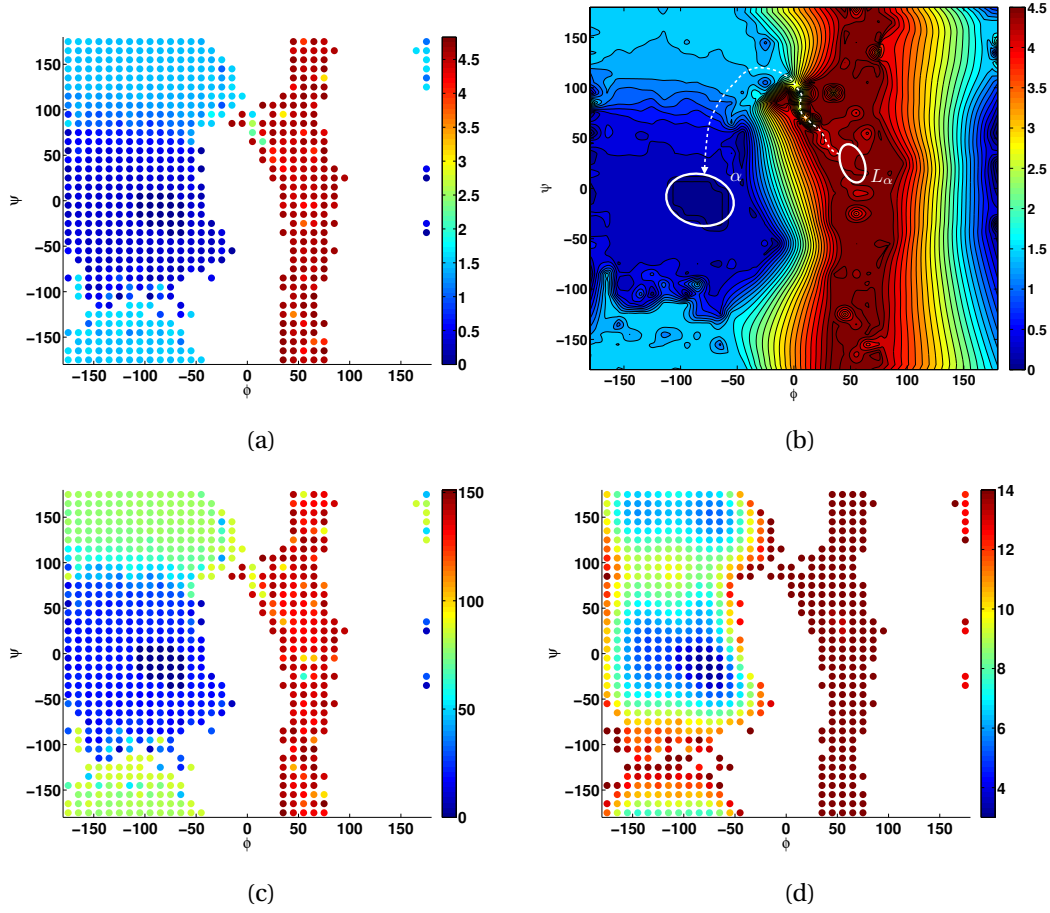


Figure 3.10 – (a) Optimal control \hat{V} for steering into the α -structure. (b) Interpolation $V_I(\phi, \psi)$ obtained from \hat{V} via a Delaunay triangulation, and a steepest descent path from L_α to α . (c) MFPT to the α conformation for the optimally controlled process. (d) Free energy $g^{\alpha^*} = -\log \pi^{\alpha^*}$ of the optimally controlled process.

G_1 . We now assume that the elements of G_0 and G_1 are of the form

$$g_0((x, y), (x', y')) = a_0(y, y'; x) \delta_{xx'} \quad (3.39)$$

$$g_1((x, y), (x', y')) = a_1(x, x'; y) \delta_{yy'} \quad (3.40)$$

where $a_0(y, y'; x)$ denotes the entry corresponding to a transition from y to y' of a generator $A_0(x)$ that acts only on \mathbb{X}_y , for fixed $x \in \mathbb{X}_x$. Similarly, $a_1(x, x'; y)$ denotes the entry of a transition from x to x' of a generator $A_1(y)$ that acts only on \mathbb{X}_x , for fixed $y \in \mathbb{X}_y$. We assume the family of generators $A_0(x)$ to be ergodic for every $x \in \mathbb{X}_x$, such that $A_0(x)$ has a one-

dimensional left and right nullspace for every $x \in \mathbb{X}_x$, characterized by

$$\sum_{y'} a_0(y, y'; x) \mathbf{1}(y') = 0, \quad (3.41)$$

$$\sum_y \mu^\infty(y; x) a_0(y, y'; x) = 0. \quad (3.42)$$

Next we introduce the **averaged** generator \bar{G}_1 on \mathbb{X}_x by:

$$\bar{g}_1(x, x') = \sum_y \mu^\infty(y; x) a_1(x, x'; y). \quad (3.43)$$

In [PS08], the following averaging result is proved:

Theorem 3.10 (Averaging for MJPs) *Let $(X_t, Y_t)_t$ be a MJP on $\mathbb{X} = \mathbb{X}_x \times \mathbb{X}_y$ with generator $G = \frac{1}{\varepsilon} G_0 + G_1$ with G_0 and G_1 given by (3.39) and (3.40). Then as $\varepsilon \rightarrow 0$, $(X_t)_t$ converges in distribution⁶ to a Markov chain $(\bar{X}_t)_t$ on \mathbb{X}_x with generator \bar{G} given by (3.43).*

One should note that $(X_t)_t$ is not a Markov process, only $(X_t, Y_t)_t$ is. Theorem 3.10 now states that $(X_t)_t$ can be approximated by a Markov process $(\bar{X}_t)_t$ on \mathbb{X}_x . The proof is done by applying the method of averaging to the backward Kolmogorov equation (see (1.14))

$$\frac{\partial u}{\partial t} = Lu \quad \text{on } (0, \infty) \times \mathbb{X}, \quad u(0, x, y) = f(x, y). \quad (3.44)$$

for $u(t, x, y) = \mathbf{E}_{(x,y)} [f(X_t, Y_t)]$. One shows that u is approximated by $u_0(x, y) = u_0(x) \mathbf{1}(y)$ which is governed by the averaged equation

$$\frac{\partial u_0}{\partial t} = \bar{G} u_0.$$

In applications, we cannot compute \bar{G} from (3.43) since we do not have access to the conditional distributions $\mu^\infty(y; x)$. However, from a realization of $(X_t, Y_t)_t$ we can construct an estimator for \bar{G} using the X_t component only. This is because for $\varepsilon \ll 1$ the fast variables sample $\mu^\infty(y; x)$ on the fast time scale where $X_t \approx x$ stays approximately constant. The MSM transfer operator that was constructed for the ADP example in section 3.2.4 is effectively $P^\tau = \exp(\tau \bar{G})$, with the slow variable space being the space of dihedral angles (ϕ, ψ) .

In the following, we will show a similar averaging result for discrete control problems. The idea is to use the duality theorem 3.3 and the Feynman Kac formula (1.39) to express the value function of the control problem in terms of the solution ϕ to a boundary value problem of the form (3.29). For boundary value problems, we can establish the following averaging result:

⁶I.e. the finite-dimensional distributions (1.1) of $(X_t)_t$ converge weakly to the finite-dimensional distributions of $(\bar{X}_t)_t$.

Theorem 3.11 (Averaging for BVPs). Let $A = A_x \times \mathbf{1}_y \subset \mathbb{X}_x \times \mathbb{X}_y$ and let ϕ be the solution to the boundary value problem

$$\begin{aligned} (G - f)\phi &= 0 & \text{on } \mathbb{X} \setminus A \\ \phi &= g & \text{on } A \end{aligned} \quad (3.45)$$

with functions $f : \mathbb{X} \rightarrow [0, \infty)$ and $g : A \rightarrow [0, \infty)$ such that $g(x, y) = g_0(x)\mathbf{1}(y)$, and generator $G = \frac{1}{\varepsilon}G_0 + G_1$ with G_0 and G_1 given by (3.39) and (3.40). Let τ_A^ε be the first hitting time of A under $G = \frac{1}{\varepsilon}G_0 + G_1$ and suppose that there is $C > 0$ and $\varepsilon_0 > 0$ such that $\sup_{x,y} \mathbf{E}_{x,y}[\tau_A^\varepsilon] \leq C$ holds for all $0 < \varepsilon \leq \varepsilon_0$. Then for $\varepsilon \leq \varepsilon_0$, $\|\phi - \phi_0\| \leq C'\varepsilon$ with C' independent of ε , where $\phi_0(x, y) = \psi_0(x)\mathbf{1}(y)$ and ψ_0 is the solution to the boundary value problem

$$\begin{aligned} (\bar{G} - \bar{f})\psi_0 &= 0 & \text{on } \mathbb{X}_x \setminus A_x \\ \psi_0 &= g_0 & \text{on } A_x \end{aligned} \quad (3.46)$$

with \bar{G} given by (3.43) and

$$\bar{f}(x) = \sum_y \mu^\infty(y; x) f(x, y). \quad (3.47)$$

Proof. First we derive (3.46). We insert the ansatz $\phi = \phi_0 + \varepsilon\phi_1$ into (3.45) and compare coefficients. This gives

$$G_0\phi_0 = 0 \quad \text{on } \mathbb{X} \setminus A, \quad \phi_0 = g \quad \text{on } A \quad (3.48)$$

$$G_0\phi_1 + G_1\phi_0 - f\phi_0 = 0 \quad \text{on } \mathbb{X} \setminus A, \quad \phi_1 = 0 \quad \text{on } A. \quad (3.49)$$

We introduce the orthogonal projection P (w.r.t. the euclidean scalar product $\langle u, v \rangle = u^T v$) onto the space $\{v \in \mathbb{R}^\mathbb{X} : v = 0 \text{ on } A\}$ and use the notation

$$G = \begin{pmatrix} H & M \\ N & O \end{pmatrix}$$

for a matrix $G : \mathbb{R}^\mathbb{X} \rightarrow \mathbb{R}^\mathbb{X}$ where $H = PGP$, $M = PGP^\perp$ and so on. In this notation, the matrix G_0 is of the form

$$G_0 = \begin{pmatrix} H & \mathbf{0} \\ \mathbf{0} & O \end{pmatrix}$$

since in order to go from $A = A_x \times \mathbf{1}_y$ to $\mathbb{X} \setminus A$ one needs to change the x -variable, but by (3.39) the x -variable is left unchanged by G_0 . We can therefore rewrite equation (3.48) as

$$\begin{pmatrix} H & \mathbf{0} \\ \mathbf{0} & \mathbb{1} \end{pmatrix} \begin{pmatrix} P\phi_0 \\ P^\perp\phi_0 \end{pmatrix} = \begin{pmatrix} \mathbf{0} \\ g \end{pmatrix}$$

which has the solution $\phi_0(x, y) = \psi_0(x)\mathbf{1}(y)$ with boundary condition $\psi_0(x) = g_0(x)$ for $x \in A_x$.

We rewrite (3.49) in the same manner:

$$\begin{pmatrix} H & \mathbf{0} \\ \mathbf{0} & \mathbb{1} \end{pmatrix} \begin{pmatrix} P\phi_1 \\ P^\perp\phi_1 \end{pmatrix} = \begin{pmatrix} P(L_1\phi_0 - f\phi_0) \\ 0 \end{pmatrix}$$

The second equation has the unique solution $P^\perp\phi_1 = 0$. In order for the first equation $H\psi_1 = P(L_1\phi_0 - f\phi_0)$ to have a solution $\psi_1 := P\phi_1$, the Fredholm alternative [PS08] implies the solvability condition

$$P(L_1\phi_0 - f\phi_0) \perp \ker(H^T).$$

By our assumptions, the nullspace $\ker(H^T)$ is characterized by

$$\sum_{x \in \mathbb{X}_x \setminus A_x, y \in \mathbb{X}_y} \mu^\infty(y; x) c(x) g_0((x, y), (x', y')) = 0$$

for any vector $c = \{c(x)\}$ on $\mathbb{X}_x \setminus A_x$. Thus imposing the solvability condition gives

$$0 = \sum_{x \in \mathbb{X}_x \setminus A_x, y \in \mathbb{X}_y} \mu^\infty(y; x) c(x) [(G_1\phi_0 - f\phi_0)(x, y)]. \quad (3.50)$$

From (3.40) and $\phi(x, y) = \psi_0(x)\mathbf{1}(y)$, we have

$$(G_1\phi_0)(x, y) = \sum_{x', y'} g_1((x, y), (x', y')) \phi_0(x', y') = \sum_{x'} a_1(x, x'; y) \psi_0(x').$$

Therefore (3.50) is equivalent to

$$\begin{aligned} 0 &= \sum_{x \in \mathbb{X}_x \setminus A_x, y \in \mathbb{X}_y} \mu^\infty(y; x) c(x) \left[\sum_{x'} a_1(x, x'; y) \psi_0(x') - f(x, y) \psi_0(x) \right] \\ &= \sum_{x \in \mathbb{X}_x \setminus A_x} c(x) \left[\sum_{x'} \bar{g}_1(x, x') \psi_0(x') - \bar{f}(x) \psi_0(x) \right] \end{aligned}$$

where the last line follows from the definitions of the averaged generator (3.43) and the averaged costs (3.47). Since c is an arbitrary vector in $\mathbb{R}^{\mathbb{X}_x \setminus A_x}$, the sum in the last line has to vanish componentwise. This together with the boundary condition $\psi_0(x) = g_0(x)$ implies (3.46).

Now we have to show that $\|\phi - \phi_0\| \leq C'\varepsilon$. To do this, we insert $\phi = \phi_0 + \varepsilon\phi_1 + r$ into (3.45), use (3.48) and (3.49) and arrive at

$$(G - f)r + \varepsilon(G_1\phi_1 - f\phi_1) = 0 \quad \text{on } \mathbb{X} \setminus A, \quad r = 0 \quad \text{on } A.$$

With the shorthand $b := G_1\phi_1 - f\phi_1$, this becomes

$$(-G + f)r = \varepsilon b \quad \text{on } \mathbb{X} \setminus A, \quad r = 0 \quad \text{on } A.$$

Taking norms on both sides gives $\lambda_m^f \|r\| \leq \varepsilon \|b\|$ where λ_m^f is the smallest eigenvalue of the operator $P(-G + f)P$ where P is orthogonal projection onto the space $\{v : \mathbb{X} \rightarrow \mathbb{R} : v|_A = 0\}$. Let $\mathcal{P}(\mathbb{X})$ be the space of probability measures on \mathbb{X} . By the Perron-Frobenius theorem, λ_m^f has a positive eigenvector that we can characterize by a variational formula due to Donsker and Varadhan [DV75]:

$$\begin{aligned} \lambda_m^f &= \inf_{v \in \mathcal{P}(\mathbb{X} \setminus A)} \sup_{u \geq 0, u|_A = 0} \left(\sum_{x \in \mathbb{X} \setminus A} \left(\frac{-Gu}{u} \right) (x) v(x) + \mathbf{E}_v[f] \right) \\ &\geq \inf_{v \in \mathcal{P}(\mathbb{X} \setminus A)} \sup_{u \geq 0, u|_A = 0} \sum_{x \in \mathbb{X} \setminus A} \left(\frac{-Gu}{u} \right) (x) v(x) =: \lambda_m^0. \end{aligned}$$

Here, the inequality holds because $f \geq 0$. λ_m^0 is the smallest eigenvalue of $P(-G)P$ and can be bounded by $\lambda_m^0 \geq 1 / \sup_{x,y} \mathbf{E}_{x,y}[\tau_A^\varepsilon] \geq C^{-1}$, see [Bov09]. This can also be seen by inserting $u = \mathbf{E}_x[\tau_A^\varepsilon]$ into the variational formula above. Putting everything together gives

$$\|r\| \leq \frac{\varepsilon}{\lambda_m^f} \|b\| \leq \frac{\varepsilon}{\lambda_m^0} \|b\| \leq \varepsilon C \|b\|$$

and thus $\|\phi - \phi_0\| = \|\varepsilon \phi_1 + r\| \leq \varepsilon \|\phi_1\| + \|r\| \leq \varepsilon (\|\phi_1\| + C \|b\|)$ which finishes the proof. ■

The conditions in theorem 3.11 demand that transitions from $\mathbb{X} \setminus A$ to A happen on the slow time scale, and that boundary conditions encoded by g do not vary along the fast time scale. This is satisfied e.g. if $A = A_1 \cup \dots \cup A_n$ is a union of metastable sets A_i and $g = g_1 \mathbf{1}_{A_1} + \dots + g_n \mathbf{1}_{A_n}$ is constant on the A_i . The control problems we considered in section 3.2.2 are a special case of this situation. With the help of theorem 3.11, we get the following averaging result for control problems:

Theorem 3.12 (Averaging for control problems.) *Let $V = \inf_{v \in \mathcal{U}} J^v$ be the value function of the LQ control problem⁷ (3.19)-(3.20) for a MJP $(X_t, Y_t)_t$ on $\mathbb{X} = \mathbb{X}_x \times \mathbb{X}_y$ with cost function*

$$J^v(x, y) = \mathbf{E}_{xy} \left[\int_0^{\tau_A} \{f(X_s^v, Y_s^v) + k(X_s^v, Y_s^v, v_s)\} ds + z(X_{\tau_A}) \right],$$

target region $A = A_x \times \mathbf{1}_y \subset \mathbb{X}$, control-dependent costs k as in (3.21) and terminal costs $z(X_{\tau_A})$. Let $G = \frac{1}{\varepsilon} G_0 + G_1$ with G_0 and G_1 as in (3.39)-(3.40). Then for $\varepsilon \ll 1$, $\|V - V_0\| \leq C\varepsilon$ where $V_0(x, y) = \bar{V}_0(x) \mathbf{1}(y)$ and \bar{V}_0 is the value function of an LQ control problem for a MJP $(\bar{X}_t)_t$ on \mathbb{X}_x with averaged cost functional

$$\bar{J}^v(x) = \mathbf{E}_x \left[\int_0^{\tau_{A_x}} \{\bar{f}(\bar{X}_s^v) + \bar{k}(\bar{X}_s^v, v_s)\} ds + z(\bar{X}_{\tau_{A_x}}) \right] \quad (3.51)$$

with \bar{f} given by (3.47), control space $\mathbb{U} = \{\alpha : \mathbb{X}_x \rightarrow (0, \infty)\}$, modified generator \bar{G}^α given by $\bar{G}^\alpha(x, x') = \bar{G}(x, x') \alpha(x') / \alpha(x)$ with \bar{G} given by (3.43), and control-dependent costs $\bar{k}(x, \alpha)$ given

⁷In order to not confuse the multiscale parameter ε with the temperature parameter in (3.19)-3.20), we set temperature to one here.

by

$$\bar{k}(x, \alpha) = (\bar{G}^\alpha(\log \alpha))(x) - \frac{(\bar{G}\alpha)(x)}{\alpha(x)}.$$

Proof. By the duality theorem 3.3 and theorem 1.7, $V = -\log \phi$ and ϕ is the unique solution of the boundary value problem

$$\begin{aligned} (G - f)\phi &= 0 && \text{on } \mathbb{X} \setminus A \\ \phi &= g && \text{on } A \end{aligned}$$

with $g(x, y) = g_0(x)\mathbf{1}(y)$ and $g_0(x) = -\log z(x)$. By theorem 3.11, $\phi = \phi_0 + \mathcal{O}(\varepsilon)$ with $\phi_0(x, y) = \psi_0(x)\mathbf{1}(y)$, and ψ_0 solves the averaged boundary value problem

$$\begin{aligned} (\bar{G} - \bar{f})\psi_0 &= 0 && \text{on } \mathbb{X}_x \setminus A_x \\ \psi_0 &= g_0 && \text{on } A_x. \end{aligned}$$

By applying theorem 3.3 again, we see that $\bar{V}_0 := -\log \psi_0$ is the value function of the averaged control problem (3.51). Set $V_0(x, y) = \bar{V}_0(x)\mathbf{1}(y)$. Since ϕ_0 is bounded away from zero, we have for $\varepsilon \ll 1$

$$\|V - V_0\| = \left\| \log \phi_0 - \log(\phi_0 + \varepsilon \phi_1 + \mathcal{O}(\varepsilon^2)) \right\| = \left\| \varepsilon \frac{\phi_1}{\phi_0} + \mathcal{O}(\varepsilon^2) \right\| = C\varepsilon + \mathcal{O}(\varepsilon^2).$$

This finishes the proof. ■

Theorem (3.12) tells that under the same assumptions as for theorem (3.11), the value function V can be approximated to lowest order in ε by a function V_0 which is constant along \mathbb{X}_y , and V_0 itself is the value function of an averaged control problem. Recall that the optimal strategy for the problem in theorem 3.12 was $\alpha^* = \exp(-V) = \phi$, hence $\alpha^* = \alpha_0^* + \mathcal{O}(\varepsilon)$ with α_0^* being constant along \mathbb{X}_y as well. Therefore, the modified generator $G_{xy}^{\alpha^*} = G_{xy}\alpha^*(y)/\alpha^*(x)$ only modifies G_1 and leaves G_0 invariant. In other words: Only transitions along the slow variable are affected by the control α^* , while transitions along the fast variable are unchanged. Note also that we make no assumptions on f , in particular f is allowed to vary along \mathbb{X}_y .

Example. As an example, we revisit the **random block model** from section 2.4.1. We choose the parameters $n_b = 5$, $m = 100$, $p_i = 0.4$ and $p_o = 10^{-3}$. We arrange the 5 blocks A_1, \dots, A_5 on a ring and only allow external edges between neighbouring blocks. We also disallow edges from A_1 to A_5 , such that a transition from A_1 to A_5 has to happen via A_2, A_3, A_4 . The resulting graph is shown in Figure 3.11.

We construct a generator $G = \frac{1}{\varepsilon}G_0 + G_1$ on this graph, where G_0 describes transitions within the blocks and G_1 describes transitions along edges between blocks. We can describe the position of the MJP on the graph by two coordinates (x, y) , where the slow coordinate $x \in \{1, \dots, n_b\}$ is

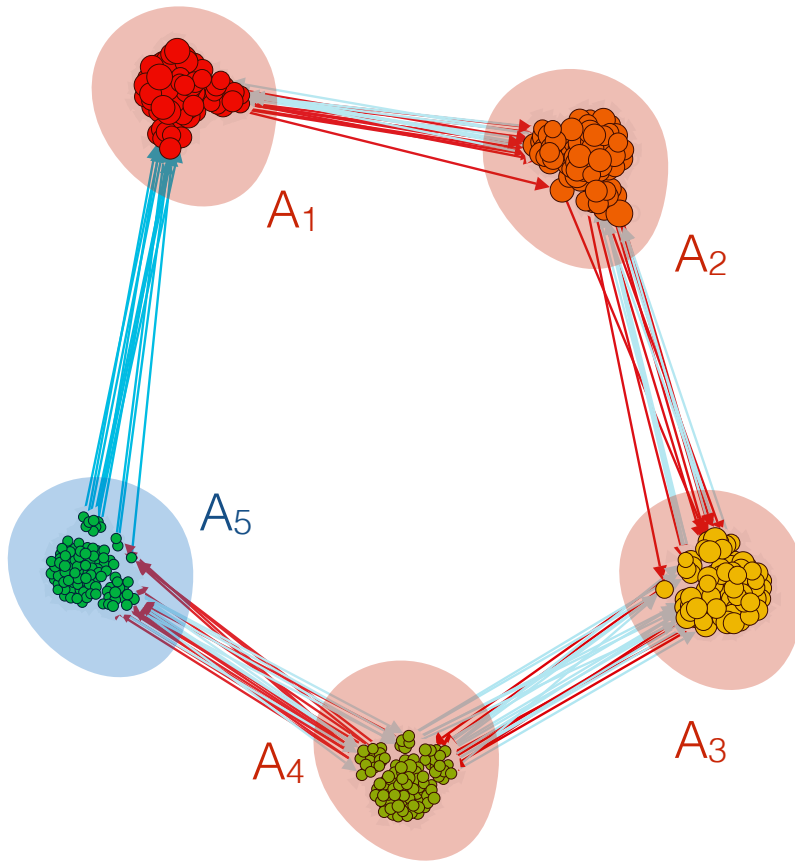


Figure 3.11 – Blockmodel with 5 metastable sets A_1 - A_5 . The control target is $A = A_5$. Nodes are colored according to the value function. Node size is proportional to μ . Edges are colored according to how their transition rates are modified by α^* : For red edges $G_{xy}^{\alpha^*} > G_{xy}$. For blue edges $G_{xy}^{\alpha^*} < G_{xy}$.

the index of the block and the fast coordinate $y \in \{1, \dots, m\}$ is the position within the block. We set $\varepsilon = 0.01$. To demonstrate averaging, we study the LQ control problem of theorem 3.12 with target region $A = A_5$, no terminal costs and normally distributed running costs $f(x) = N(f_0, f_0/2)$ with $f_0 = 10$. In Figure 3.12, the resulting value function V and the averaged value function V_0 are shown. Indeed V is almost constant within blocks, with fluctuations of order ε . V_0 , which is computable from a 5×5 -system, approximates V almost perfectly.

As a consequence, the generator G^{α^*} of the resulting optimal strategy α^* only modifies jump rates along edges between blocks. In Figure 3.11, the coloring of the edges indicates how their jump rates are modified by α^* . Red edges have increased jump rates, blue edges have decreased jump rates. As expected, jump rates on clockwise edges are increased while jump rates on counterclockwise edges and jump rates from A_5 back to A_1 are decreased.




Figure 3.12 – The value function V for the full problem (blue line) and the averaged value function V_0 (red dotted line).

4 Conclusion and Outlook

In the introduction, two questions were asked:

- (1) How do the long-term dynamics of a reversible Markov process change when non-reversibility is added?
- (2) Can we utilize non-reversibility to accelerate the long-term dynamics?

We will now interpret the results of the thesis in order to give partial answers.

It is clear that we cannot expect a general answer to question (1): Without further constraints on the type of non-reversibility added, anything could happen. The most general answer to question (1) given in the thesis is in Theorem 2.6 in section 2.2.3. It compares a reversible and ergodic Markov process $(X_t^s)_t$ which has stationary distribution μ and generator L^s with a non-reversible Markov process $(X_t)_t$ which is obtained from $(X_t^s)_t$ by adding a term δL to the generator. The generator $L = L^s + \delta L$ of $(X_t)_t$ has to satisfy $L^s = \frac{1}{2}(L + L^-)$. For finite state spaces, this is satisfied if $\delta L = D_\mu^{-1} H$ with $H = -H^T$ an antisymmetric matrix. Theorem 2.6 now states that $(X_t^s)_t$ is always slower than $(X_t)_t$ in the sense that $T^s(A, B) \geq T(A, B)$ for any $A, B \subset \mathbb{X}$. In other words, adding the non-reversible term $\delta L = D_\mu^{-1} H$ will always decrease the commuting time for any pair of regions A, B . If A and B are metastable sets then the commuting time $T(A, B)$ is a way to quantify their degree of metastability. Theorem 2.6 therefore tells us that $(X_t)_t$ will **always be less metastable** than $(X_t^s)_t$.

Theorem 2.6 was stated and proven for finite \mathbb{X} , but the tools used to prove it - potential theory and the variational formulas (2.31) and (2.32) - also exist for diffusions on continuous \mathbb{X} . A possible generalisation of Theorem 2.6 to this case can therefore be expected.

Note that we cannot use any antisymmetric matrix H to construct $L = L^s + D_\mu^{-1} H$. L must still be a generator matrix, in particular all the off-diagonal entries of L have to be nonnegative. But if H is a suitable antisymmetric matrix, then we automatically have $\text{div } H = 0$ and this guarantees that we can decompose $G(H)$ into cycles. Since cycles in $G(H)$ must be cycles in $G(L^s)$ (otherwise $L = L^s + D_\mu^{-1} H$ would have negative off-diagonal entries), a way to construct

Chapter 4. Conclusion and Outlook

a suitable H is to look for cycles in $G(L^s)$. This highlights another important feature of non-reversible processes: The appearance of non-trivial cycles in the probability current. Section 2.1 compared two possible ways for finding these cycles: An algebraic construction in Lemma 2.3 and a probabilistic one in Lemma 2.4. Throughout the thesis, cycles appeared in many places that hint at their importance:

- Expressions (2.13) and (2.14) for the entropy production rate (1.10) in terms of cycles.
- If one wants to find a suitable H so that $L = L^s + D_\mu^{-1}H$ satisfies the conditions of Theorem 2.6, one needs to find cycles in $G(L^s)$.
- Module detection for irreversible processes: With the help of Lemma 2.7, we defined modules as subsets with many internal and few external cycles and identified them with metastable sets of the loop transition matrix P_∞^s (2.21) in section 2.3.
- The TPT current of reactive trajectories f_{AB}^+ has cycles if the dynamics is irreversible. These cycles cloud the interpretation of f_{AB}^+ as describing a transport mechanism, and we would like to remove them. This was addressed in section 2.4.

In summary, we may answer question (1) in the following way: The long-term dynamics is almost always accelerated if non-reversibility is added. If one wants to understand the precise mechanisms, then cycles are the key.

In order to give an answer to question (2), we need to specify what exactly we want to accelerate. We shall be interested in two kinds of observables: *Equilibrium observables*, which are of the form $\phi(f) = \mathbf{E}_\mu[f(X_t)]$ with $X_t \sim \mu$ and a measurable and bounded function $f : \mathbb{X} \rightarrow \mathbb{R}$, and *path observables* of the form $\psi(F) = \mathbf{E}[F]$ where the function $F : \Omega \rightarrow \mathbb{R}$ depends on the whole realization of $(X_t)_t$. What we want to accelerate is the time a numerical method needs to compute $\phi(f)$ or $\psi(F)$ to a certain accuracy.

Since sampling from an equilibrium distribution μ converges with a rate that is given by the second-largest eigenvalue which in turn corresponds to a pair of regions A and A^c such that $T(A, A^c)$ is maximal, Theorem 2.6 gives an answer to question (2) for equilibrium observables. The answer is that, within the confinements of Theorem 2.6, we can add any non-reversible term to a reversible process, it will always accelerate the sampling. This is consistent with results in [SSG10]. The question of how e.g. $\delta L = D_\mu^{-1}H$ should be chosen **best** is not answered. It might be possible to formulate an optimization problem to answer this question, but we leave that for future work (see below).

For path observables, Theorem 2.6 does not help since the dynamics and therefore $\psi(F)$ change under the addition of δL . In order to compute the unmodified observable $\psi(F)$ under the modified dynamics, a reweighting procedure is needed. In section 3.1, it was explained that for path observables of the form

$$\psi(x) = \mathbf{E}_x \left[\exp \left(\int_0^\tau f(X_s) ds \right) \right]$$

such a reweighting procedure is given in terms of the solution to an optimal control problem. Here, τ was the first hitting time of a region $A \subset \mathbb{X}$, and we are interested in situations where the initial state $x \in B$ and A and B are metastable sets. With the optimal control formulation, we can compute the **optimal way** to accelerate transitions from B to A by adding a control u_t . Under the optimal control u_t^* , a zero variance estimator for ψ is obtained [HS12].

In section 3.2, a numerical method to solve the resulting LQ control problems numerically was constructed based on Galerkin projections. An L^2 error bound on the performance $\mathbf{p} = \boldsymbol{\varepsilon}/\boldsymbol{\varepsilon}_0$ was given in Theorem 3.4 that complements C ea's lemma [Bra07]. Bounds on $\boldsymbol{\varepsilon}_0$ for standard choices of basis functions, e.g. piecewise polynomial or radial basis functions, can be found in the literature [Bra07, Wen99]. We gave special attention to the committor basis and bounded $\boldsymbol{\varepsilon}_0$ for this choice in Theorem 3.7. We have shown that the committor basis is extremely sparse and still very accurate, provided the core sets are placed in the metastable wells. Doing this however requires prior knowledge about the location of these metastable regions.

In MSM theory, the elements of the MSM transition matrix

$$P_{ij} = \frac{\langle \chi_i, P^\tau \chi_j \rangle_\mu}{\langle \chi_i, \mathbf{1} \rangle_\mu}$$

cannot be computed numerically due to the high-dimensional integrals involved and must be sampled instead. The same is true for the matrices K and F that we construct with our method in (3.29) if a committor basis or a basis of indicator functions of a full partition of state space is used. The benefits of using such a discretization for our method are the same as for MSM theory: Instead of having to sample ψ using few long trajectories, which would be extremely difficult, we can sample K and F using many short trajectories.

Outlook

We give a short account on questions left unanswered and possible future directions, without an effort of completeness:

- (1) The duality Theorem 3.1 holds for general Markov diffusions of the form

$$dX_t = b(X_t)dt + \sigma(X_t)dW_t.$$

However, we construct the Galerkin discretization method in section 3.2 for the case of X_t being reversible with $b(x) = -\nabla V(x)$ only. This is in accordance with the focus of the thesis on studying how **reversible** processes change if non-reversibility (in this case control) is added. The technical reason for this limitation is that then the generator L of the process is self-adjoint in $L^2(\mu)$ with scalar product $\langle \cdot, \cdot \rangle_\mu$, which leads to better properties for the Galerkin projection. However, it should be possible to extend the Galerkin projection formalism to the case where L is not self-adjoint, which enables the

study of control applied to irreversible processes.

- (2) In section 2.4, we studied transition pathways between two metastable regions A and B for irreversible processes with the help of TPT. It would be interesting to see how this system would react if a control is introduced, e.g. to accelerate the transition from A to B . In principle, it should be possible with the methods presented in section 2.4 to compute for every reaction pathway how it reacts to the control, e.g. it might shift slightly or the proportion of reactive flow it carries might change.
- (3) Theorem 2.6 does not tell how $\delta L = D_\mu^{-1}H$ should be chosen in order to optimally accelerate the commuting time $T(A, B)$. This question can be formulated as an optimal control problem in the following way: Consider the extended state space $\mathbb{X}^e = \mathbb{X} \otimes \{0, 1\}$. We denote a state in \mathbb{X}^e by (x, s) and think of $s = 0$ as signifying ' $(X_t)_t$ has not yet reached B ' and of $s = 1$ as signifying ' $(X_t)_t$ has reached B '. In other words, we define a process $(X_t^e)_t$ on \mathbb{X}^e by $X_t^e = (X_t, s_t)$ where the dynamics of X_t is Markovian with generator L and $s_t = \mathbf{1}_{\{t \geq \tau_B\}}$ where τ_B is the first hitting time of B . Now let $\tau^e = \inf\{t > 0 : X_t^e \in A \otimes \{1\}\}$ be the first hitting time of the state $A \otimes \{1\} \subset \mathbb{X}^e$ that denotes 'return to A after B has been visited'. With appropriate running costs $k(x, u)$, we may now think of minimizing the functional

$$J^u(x) = \mathbf{E}_x \left[\int_0^{\tau^e} (f(X_s^e) + k(X_s^e, u_s)) ds \right]$$

under the constraint that $L^u = L + D_\mu^{-1}H$ with $H = -H^T$. For the choice $f = \mathbf{1}$ and initial point $x \in A$, the functional $J^u(x)$ is essentially $T(A, B)$ plus an additional control penalty term. This is not an LQ control problem however, so new solution methods must be sought.

- (4) If we want to go beyond LQ control problems, then obtaining a linear PDE by logarithmic transformation will no longer be possible in general, and we are forced to work with the nonlinear control problem directly. A possible avenue for doing this is **approximate policy iteration**, see [Ber11] and the references therein for a description of this method in the context of Markov chains. We now outline how this method would look like in the context of diffusions. The starting point is ordinary policy iteration [Kus01], which begins with an initial Markov control policy u_0 and then iterates the following two steps:

- *Policy evaluation:* Given a Markov control policy u_n , determine the cost $J_n = J^{u_n}$. If J^u has the form (1.57b) for the indefinite horizon problem, then this can be done by solving the linear PDE

$$\begin{aligned} 0 &= L^{u_n} J_n(x) + f(x, u_n(x)), & x \in A, \\ J_n(x) &= z(x), & x \in \partial A. \end{aligned} \tag{4.1}$$

- *Policy improvement:* Compute a better policy u_{n+1} by finding the minimizer in the Bellman equation (1.66) with the value function V replaced by the cost J_n of the

current policy u_n . For the indefinite horizon problem with cost function (1.57b), policy improvement takes the form

$$u_{n+1}(x) = \operatorname{argmin}_{\alpha \in \mathbb{U}} \{L^\alpha J_n(x) + f(x, \alpha)\}, \quad x \in A. \quad (4.2)$$

Policy evaluation is usually more difficult than policy improvement (solving a PDE vs. pointwise minimization). The appeal of policy iteration is that it replaces the nonlinear Bellman equation (1.66) with an iterative scheme where a linear PDE (4.1) must be solved at each iteration step. The idea of approximate policy iteration is now to replace the difficult step (4.1) by an approximation obtained by projecting (4.1) onto a low-dimensional space of ansatz functions. This is precisely the Galerkin projection formalism that was developed in section 3.2. With the ansatz space $D = \{\chi_1, \dots, \chi_n\}$ and Q being the projection onto D , approximate policy iteration begins with an initial Markov control policy u_0 and iterates the following two steps:

- *Approximate policy evaluation:* Given a Markov control policy u_n , determine the approximate cost $\hat{J}_n = QJ_n$ by solving

$$0 = \hat{L}^{u_n} \hat{J}_n + r^{u_n}, \quad \hat{L}_{ij}^{u_n} = \frac{\langle \chi_i, L^{u_n} \chi_j \rangle_\mu}{\langle \chi_i, \mathbf{1} \rangle_\mu}, \quad r_i^{u_n} = \frac{\langle \chi_i, f(\cdot, u_n) \rangle_\mu}{\langle \chi_i, \mathbf{1} \rangle_\mu}. \quad (4.1')$$

- *Policy improvement:* Compute a better policy u_{n+1} by finding the minimizer in the Bellman equation (1.66) with the value function V replaced by the **approximate** cost \hat{J}_n of the current policy u_n . For the indefinite horizon problem with cost function (1.57b), policy improvement takes the form

$$u_{n+1}(x) = \operatorname{argmin}_{\alpha \in \mathbb{U}} \{L^\alpha \hat{J}_n(x) + f(x, \alpha)\}, \quad x \in A. \quad (4.2')$$

The approximate policy evaluation step (4.1') requires solving a linear system as well as sampling of L^{u_n} and r^{u_n} . This sampling can be performed with the techniques discussed in section 3.2.3. Approximate policy iteration is not restricted to LQ control problems, but for LQ control problems the policy improvement step (4.2') becomes particularly simple, namely $u_{n+1}(x) = -\sqrt{2}\varepsilon^{-1/2}\sigma^T \nabla J_n(x)$, cf. Theorem 3.1.

Figure 4.1 shows some preliminary results for the triple well potential with 20 evenly spaced gaussian basis functions as ansatz functions. The 'do nothing' policy $u_0 \equiv 0$ was chosen as starting policy. In Figure 4.1 on the right, the error of the iterates $\|J^* - \hat{J}_k\|_\mu$ is compared with the best approximation error $\|J^* - QJ^*\|_\mu$. One can observe that the error decreases very fast initially, but then it starts oscillating instead of converging. This is due to the projection in (4.1'). Oscillations around a limit cycle are typical for approximate policy iteration [Ber11].




Figure 4.1 – Approximate policy iteration, preliminary results for the triple well potential. Left: First six iterates \hat{J}_k (blue), value function $V = J^*$ (green) and best approximation QJ^* (red). Right: Error of the iterates \hat{J}_k compared to the best approximation error.

A Appendix A

A.1 Proof of theorem 1.1

We present the proof of theorem 1.1. Let $g(x) = u(t, x)$ for $t \in (0, \infty)$. Since f is bounded, $f \in L^\infty(\mathbb{X}, \mu)$ and $u(t, x) = (T_t f)(x)$. Since $t \mapsto u(t, x)$ is differentiable, we have

$$\begin{aligned} \frac{\partial u}{\partial t} &= \lim_{h \rightarrow 0} \frac{u(t+h, x) - u(t, x)}{h} = \lim_{h \rightarrow 0} \frac{(T_{t+h} f)(x) - (T_t f)(x)}{h} \\ &= \lim_{h \rightarrow 0} \frac{(T_h(T_t f))(x) - (T_t f)(x)}{h} = \lim_{h \rightarrow 0} \frac{(T_h g)(x) - g(x)}{h}. \end{aligned}$$

Since $g = u(t, \cdot) \in \mathcal{D}$, the last limit on the RHS exists and equals $Lg = Lu(t, \cdot)$. This proves the assertion. ■

A.2 Proof of theorem 1.7

We present the proof of theorem 1.7. Existence and uniqueness of the solution to (1.39) holds as follows: Note that we can write (1.39) as

$$(L_D - f_D)h_D = b \tag{A.1}$$

where L_D , f_D and h_D are L , f and h restricted to D respectively (i.e. columns and rows corresponding to $\mathbb{X} \setminus D$ have been erased) and b is the vector one obtains by plugging the boundary conditions $h(x) = g(x)$ into equation (1.39), i.e. $b(x) = \sum_{y \in \mathbb{X} \setminus D} l_{xy} g(y)$. Because L is a generator matrix, $L \leq 0$ and $L_D < 0$, and since $f \geq 0$ the matrix $L_D - f_D < 0$ is negative definite and therefore invertible, so that (A.1) has a unique solution.

Now we show that (1.40) is a solution to (1.39). To do so, we use the characterization of the MJP $\{X_t\}_t$ by the embedded chain $(X_n)_n$ and the sequence of waiting times $(\tau_n)_n$. Let $X_0 = x \in D$ (if $x \in \mathbb{X} \setminus D$, we immediately have $h(x) = g(x)$). We use the law of total expectation to rewrite $h(x)$

by conditioning on $\tau_0 = \tau$ and $X_1 = y \neq x$, this implies $X_t = x$ for $0 \leq t < \tau$ and $X_\tau = y$. Then,

$$\begin{aligned} h(x) &= \int_0^\infty \mathbf{P}[\tau_0 = \tau | X_0 = x] \sum_{y \neq x} \mathbf{P}[X_1 = y | X_0 = x] \\ &\quad \times \mathbf{E}_x \left[\exp \left(- \int_0^{\tau_D} f(X_s) ds \right) g(X_{\tau_D}) \middle| \tau_0 = \tau, X_1 = y \right] d\tau \end{aligned}$$

Since the τ_i are exponentially distributed, $\mathbf{P}[\tau_0 = \tau | X_0 = x] = \lambda e^{-\lambda\tau}$ with $\lambda = -l_{xx}$. Furthermore, $\mathbf{P}[X_1 = y | X_0 = x] = q_{xy} = l_{xy}/(-l_{xx})$ by (1.36). The $\exp(-\int_0^\tau \dots)$ part of the integral in the conditional expectation is \mathcal{F}_τ -measurable and, since we condition on \mathcal{F}_τ , can be pulled outside the conditional expectation. This results in

$$\begin{aligned} h(x) &= \int_0^\infty e^{-\lambda\tau} \sum_{y \neq x} l_{xy} \exp \left(- \int_0^\tau f(x) ds \right) \\ &\quad \times \mathbf{E}_x \left[\exp \left(- \int_\tau^{\tau_D} f(X_s) ds \right) g(X_{\tau_D}) \middle| \tau_0 = \tau, X_1 = y \right] d\tau \end{aligned}$$

Since $\{X_t\}_t$ is a homogeneous Markov process, we can restart it at $t = \tau$ with initial condition $X_\tau = y$, and then the conditional expectation above equals $h(y)$. We end up with

$$h(x) = \sum_{y \neq x} l_{xy} h(y) \int_0^\infty e^{-(\lambda+f(x))\tau} d\tau = \sum_{y \neq x} l_{xy} h(y) \frac{1}{\lambda + f(x)}.$$

This can be rearranged to give, in view of $\lambda = -l_{xx}$,

$$f(x)h(x) = \sum_{y \neq x} l_{xy} h(y) + l_{xx} h(x) = (Lh)(x)$$

which is (1.39). ■

A.3 Proof of theorem 1.16

Let $u \in \mathcal{U}_0$. Since V solves the Bellman equation (1.71) and $u_s \in \mathcal{U}$ for all $0 \leq s < \tau$, we have

$$L^{u_s} V(X_s^u) + f(X_s^u, u_s) \geq 0, \quad 0 \leq s < \tau. \tag{A.2}$$

Now use the Dynkin formula (1.68):

$$\mathbf{E}_x [V(X_\tau^u)] = V(x) + \mathbf{E}_x \left[\int_0^\tau L^{u_s} V(X_s^u) ds \right] \geq V(x) - \mathbf{E}_x \left[\int_0^\tau f(X_s^u, u_s) ds \right].$$

Because of the boundary condition in (1.71) we have $\mathbf{E}_x [V(X_T^u)] = \mathbf{E}_x [z(X_T^u)]$, hence

$$V(x) \leq \mathbf{E}_x \left[\int_0^T f(X_s^u, u_s) ds + z(X_T^u) \right] = J^u(x).$$

This proves the first part of the assertion. For the second part, let $u_t^* = \alpha^*(X_t^u)$ with $\alpha^*(x)$ given by (1.72). Then the inequality (A.2) becomes an equality, and therefore $V(x) = J^{u^*}(x)$. ■

A.4 Expressions for the entropy production rate

We first show (2.12). Recall that $e_P = \lim_{T \rightarrow \infty} \frac{1}{T} H(\mathbf{P}_{[0,T]}, \mathbf{P}_{[0,T]}^-)$. Let $v_x(t) = \mathbf{P}(X_t = x)$. We will now show that

$$H(\mathbf{P}_{[0,T]}, \mathbf{P}_{[0,T]}^-) = \sum_{t=0}^{T-1} e_P(t) \tag{A.3}$$

with the transient entropy production rate at time t given by

$$e_P(t) = \frac{1}{2} \sum_{x,y \in \mathbb{X}} [v_x(t)p_{xy} - v_y(t)p_{yx}] \log \left(\frac{\mu_x p_{xy}}{\mu_y p_{yx}} \right). \tag{A.4}$$

Equation (2.12) follows from (A.3) and (A.4) since $\lim_{t \rightarrow \infty} \|v(t) - \mu\| = 0$ by the second statement in the ergodic theorem 1.5 and by the fact that if a sequence of real numbers converges as $\lim_{n \rightarrow \infty} |f_n - f_\infty| = 0$, then also $\lim_{n \rightarrow \infty} |F_n - f_\infty| = 0$ for $F_n := \frac{1}{n} \sum_{i=1}^n f_i$. To show (A.3), we note that the probability to observe the trajectory (x_0, x_1, \dots, x_T) under the forward and backward dynamics is given by

$$\begin{aligned} \mathbf{P}_{[0,T]}(x_0, \dots, x_T) &= v_{x_0}(0) p_{x_0 x_1} \cdots p_{x_{T-1} x_T}, \\ \mathbf{P}_{[0,T]}^-(x_0, \dots, x_T) &= v_{x_0}(0) p_{x_0 x_1}^- \cdots p_{x_{T-1} x_T}^- \\ &= p_{x_T x_{T-1}} \cdots p_{x_1 x_0} v_{x_0}(0) \frac{\mu_{x_T}}{\mu_{x_0}}. \end{aligned}$$

Therefore,

$$\begin{aligned}
 H(\mathbf{P}_{[0,T]}, \mathbf{P}_{[0,T]}^-) &= \mathbf{E}_{\mathbf{P}_{[0,T]}} \left[\log \frac{\mu_{x_0} p_{x_0 x_1} \cdots p_{x_{T-1} x_T}}{\mu_{x_T} p_{x_T x_{T-1}} \cdots p_{x_1 x_0}} \right] \\
 &= \sum_{t=0}^{T-1} \mathbf{E}_{\mathbf{P}_{[0,T]}} \left[\log \frac{\mu_{x_t} p_{x_t x_{t+1}}}{\mu_{x_{t+1}} p_{x_{t+1} x_t}} \right] \\
 &= \sum_{t=0}^{T-1} \sum_{x, y \in \mathbb{X}} \mathbf{P}(x_t = x, x_{t+1} = y) \left(\log \frac{\mu_x p_{xy}}{\mu_y p_{yx}} \right) \\
 &= \sum_{t=0}^{T-1} \sum_{x, y \in \mathbb{X}} v_x(t) p_{xy} \left(\log \frac{\mu_x p_{xy}}{\mu_y p_{yx}} \right) \\
 &= \sum_{t=0}^{T-1} e_P(t)
 \end{aligned}$$

which is (A.3).

Next we show that (2.13) follows from (2.12). By inserting (2.10) into (2.12), we have

$$e_P = \frac{1}{2} \sum_{x, y \in \mathbb{X}} \sum_{c \in \mathcal{C}_\infty} [\omega_\infty(c) - \omega_\infty(c^-)] J_c(x, y) \log \left(\frac{\mu_x p_{xy}}{\mu_y p_{yx}} \right).$$

This can be rearranged as

$$e_P = \frac{1}{2} \sum_{c \in \mathcal{C}_\infty} [\omega_\infty(c) - \omega_\infty(c^-)] \log \left(\frac{\mathbf{P}(c)}{\mathbf{P}(c^-)} \right).$$

Now we use (2.11) to arrive at (2.13). Finally, we show that (2.14) follows from (2.12). An equivalent way of writing (2.12) is

$$e_P = \sum_{x, y \in \mathbb{X}} \mu_x p_{xy} \log \left(\frac{\mu_x p_{xy}}{\mu_y p_{yx}} \right).$$

We can express this in the edge basis for the one-cycle $\underline{F} = \sum_{x, y} \mu_x p_{xy} e_{xy}$ and insert (2.3):

$$e_P = \sum_{e \in E} \langle e, \underline{F} \rangle \log \left(\frac{\langle e, \underline{F} \rangle}{\langle e^-, \underline{F} \rangle} \right) = \sum_{e \in E} \sum_{\alpha=1}^b \omega(\gamma_\alpha) \langle e, \gamma_\alpha \rangle \log \left(\frac{\langle e, \underline{F} \rangle}{\langle e^-, \underline{F} \rangle} \right).$$

Let us write $e \uparrow \gamma_\alpha$ if $\langle e, \gamma_\alpha \rangle = 1$ and $e \downarrow \gamma_\alpha$ if $\langle e, \gamma_\alpha \rangle = -1$. Note that $\langle e, \underline{F} \rangle = F(s(e), t(e))$ where $s(e)$ is the source and $t(e)$ the target of e . Then the last equation becomes

$$e_P = \sum_{\alpha=1}^b \omega(\gamma_\alpha) \left[\sum_{e \uparrow \gamma_\alpha} \log \left(\frac{F(s(e), t(e))}{F(t(e), s(e))} \right) + \sum_{e \downarrow \gamma_\alpha} \log \left(\frac{F(t(e), s(e))}{F(s(e), t(e))} \right) \right].$$

This can be rearranged to give

$$e_P = \sum_{\alpha=1}^b \omega(\gamma_\alpha) \log \left(\frac{\mathbf{P}(\gamma_\alpha)}{\mathbf{P}(\gamma_\alpha^-)} \right).$$

which is (2.14). ■

A.5 Proof of Lemma 2.5

In view of a theorem similar to theorem 1.7 (see for example [BEGK02, Øks03]), the function $h(x) = \mathbf{E}_x[\tau_B]$ solves the boundary value problem

$$\begin{aligned} Lh(x) &= -1, & x \notin B, \\ h(x) &= 0, & x \in B. \end{aligned} \tag{A.5}$$

Then, using that $e_{AB}^-(x) = (-L^- q^-)(x)$ holds for all $x \in A$, we get

$$\sum_{x \in A} \mu(x) e_{AB}^-(x) h(x) = - \sum_{x \in A} \mu(x) (L^- q^-)(x) h(x) = - \langle L^- q^-, h \rangle_\mu$$

where the last step is true because $(L^- q^-)(x) = 0$ for $x \in \mathbb{X} \setminus (A \cup B)$ and $h(x) = 0$ for $x \in B$. We proceed by using (A.5):

$$\langle L^- q^-, h \rangle_\mu = \langle q^-, Lh \rangle_\mu = - \sum_{x \notin B} \mu(x) q^-(x) + \sum_{x \in B} \mu(x) q^-(x) (Lh)(x).$$

The last term is zero since $q^-(x) = 0$ for $x \in B$. Putting everything together, we arrive at

$$\sum_{x \in A} \mu(x) e_{AB}^-(x) h(x) = \sum_{x \notin B} \mu(x) q^-(x).$$

Dividing both sides by $\text{cap}(A, B)$ gives (2.28). Equation (2.29) follows by adding up (2.28) and (2.28) with A and B interchanged, which replaces q^- with $1 - q^-$ on the RHS of (2.28). ■

A.6 Proof of Theorem 2.6

Recall the variational formula (2.31) for $\text{cap}(A, B)$. For any $f \in \mathcal{H}_{AB}$, we have

$$\sup_{h \in \mathcal{G}_{AB}} \{-2 \langle L^- f, h \rangle_\mu - \langle h, -L^s h \rangle_\mu\} \geq -2 \langle L^- f, f \rangle_\mu - \langle f, -L^s f \rangle_\mu$$

by choosing $h = f \in \mathcal{G}_{AB}$. But $\langle L^- f, f \rangle_\mu = \langle f, Lf \rangle_\mu = \langle f, L^s f \rangle_\mu$. Hence

$$\sup_{h \in \mathcal{G}_{AB}} \{-2 \langle L^- f, h \rangle_\mu - \langle h, -L^s h \rangle_\mu\} \geq \langle f, -L^s f \rangle_\mu \quad \forall f \in \mathcal{H}_{AB}.$$

Taking the infimum over all $f \in \mathcal{H}_{AB}$ gives in view of (2.31) and (2.32)

$$\text{cap}(A, B) \geq \text{cap}^s(A, B)$$

Appendix A. Appendix A

which is (2.30). To show strictness, we write $q^+ = q + \delta q$ with $\delta q(x) = 0$ for all $x \in A \cup B$ and then we write $\text{cap}(A, B)$ as

$$\text{cap}(A, B) = \langle q^+, -Lq^+ \rangle_\mu = \langle q^+, -L^s q^+ \rangle_\mu = \langle q + \delta q, -L^s(q + \delta q) \rangle_\mu.$$

Using the fact that $\langle \delta q, -L^s q \rangle_\mu = 0$ since $(L^s q)(x) = 0$ for $x \notin (A \cup B)$ and $\delta q(x) = 0$ for $x \in (A \cup B)$, we obtain

$$\text{cap}(A, B) = \text{cap}^s(A, B) + \langle \delta q, -L^s \delta q \rangle_\mu. \quad (\text{A.6})$$

Since L^s is ergodic, $-L^s \geq 0$ with $\ker L^s$ being the space of constant functions on \mathbb{X} . Since $\delta q = 0$ on $(A \cup B)$, we have that δq is not constant and therefore $\langle \delta q, -L^s \delta q \rangle_\mu > 0$ as long as $\delta q(x) \neq 0$ for some $x \in \mathbb{X}$. We demonstrate that this is indeed the case: q^+ solves the linear system

$$0 = \sum_{y \in \mathbb{X}} L_{xy} (q(y) + \delta q(y)) \quad \forall x \in \mathbb{X} \setminus (A \cup B).$$

If we take into account that $L^s q(x) = 0$ for $x \in \mathbb{X} \setminus (A \cup B)$ and rearrange terms, we get

$$\sum_{y \in \mathbb{X}} L_{xy} \delta q(y) + \sum_{y \in \mathbb{X}} (L_{xy} - L_{xy}^s) q(y) = 0 \quad \forall x \in \mathbb{X} \setminus (A \cup B).$$

Since $\delta q = 0$ on $A \cup B$, this can be rearranged to

$$\sum_{y \in \mathbb{X} \setminus (A \cup B)} L_{xy} \delta q(y) = b(x) \quad \forall x \in \mathbb{X} \setminus (A \cup B)$$

with $b(x) = -\sum_{y \in \mathbb{X}} (L_{xy} - L_{xy}^s) q(y)$, or shortly $\hat{L} \delta q = b$ where \hat{L} is L restricted to $\mathbb{X} \setminus (A \cup B)$. The matrix $-\hat{L}$ is a non-singular M-Matrix, which follows from $-\hat{L}e > 0$ where e is the vector whose components are all one [PB74]. By the properties of non-singular M-matrices, \hat{L} is invertible and $\ker \hat{L} = \{0\}$ [PB74]. Therefore $\delta q = 0$ if and only if $b = 0$, which is equivalent to $Lq - L^s q = 0$. ■

A.7 Proof of Lemma 2.7

We prove (2.36), (2.37) follows analogously. From the definition of conditional expectation,

$$\begin{aligned} \mathbf{P}(X_1^s \in C | X_0^s \in C) &= \frac{1}{\mathbf{P}(X_0^s \in C)} \sum_{x, y \in C} \mu_x P_\infty^s(x, y) = 1 - \frac{1}{\mathbf{P}(X_0^s \in C)} \sum_{x \in C, y \notin C} \mu_x P_\infty^s(x, y) \\ &= 1 - \frac{\sum_{x \in C, y \notin C} \mu_x P_\infty^s(x, y)}{\sum_{x \in C, y \in \mathbb{X}} \mu_x P_\infty^s(x, y)}. \end{aligned}$$

Now insert (2.19) to obtain (2.36). ■

A.8 The Current of Reactive Trajectories J_{AB}

We show that (2.44) follows from (2.43). First of all, in view of (1.20),

$$\frac{1}{2}\mu(x)(b(x) + b^-(x)) = \nabla \cdot (a(x)\mu(x))$$

so that the equilibrium current $J(x)$ becomes

$$J(x) = \mu(x)b(x) - \nabla \cdot (a(x)\mu(x)) = \frac{1}{2}\mu(x)(b(x) - b^-(x)).$$

If $(X_t)_t$ is reversible then $b = b^-$ and thus $J(x) = 0$ for all $x \in \mathbb{X}$. Furthermore, $q^- = 1 - q^+$ in this case. Thus (2.43) simplifies to

$$\begin{aligned} J_{AB}(x) &= (1 - q^+(x))\mu(x)a(x)\nabla q^+(x) + q^+(x)\mu(x)a(x)\nabla q^+(x) \\ &= \mu(x)a(x)\nabla q^+(x) \end{aligned}$$

which is (2.44). ■

B Appendix B: Module Detection Methods

B.1 MSM Clustering

Recall Definition 1.10 of an incomplete metastable partition. The goal of MSM clustering is to find core sets C_1, \dots, C_m such that the metastability index

$$\rho = \max_i \rho_i, \quad \rho_i = \frac{\sup_{x \in \mathbb{X} \setminus C_i} \mathbf{E}_x[\tau_{C_i}]}{\inf_{x \in C_i} \mathbf{E}_x[\tau_{C_i \setminus C_i}]} \quad (\text{B.1})$$

is minimized. This is an NP-hard problem which can only be solved with a heuristic which we now explain, for more details see [DBCS11, SCB⁺14]. The heuristic behind MSM clustering is to separate the problem into two subproblems:

- (1) Separate the state space \mathbb{X} into the union of the cores $C = \bigcup_i C_i$ and the transition region $\mathbb{T} = \mathbb{X} \setminus C$,
- (2) Find a full partition of C into m core sets C_1, \dots, C_m .

By examining (B.1), we see that we should choose C in step (1) in such a way that $t_R = \sup_{x \in \mathbb{X} \setminus C} \mathbf{E}_x[\tau_C]$ is small, i.e. C is always reached quickly from \mathbb{T} . To select such a C , one starts with a uniform distribution μ_0 on \mathbb{X} , propagates μ_0 for some time with the dynamics and then identifies regions where probability has accumulated. More algorithmically, we select a lagtime parameter $\alpha > 0$ and set

$$C_\alpha = \{x \in \mathbb{X} : P_\alpha^T \mu_0(x) > \mu_0(x)\}$$

where P_α is the transition matrix of the RW process $(X_t)_t$ with lagtime α . This condition is more restrictive the larger α is. We comment on the choice of parameters below. For step (2), we consider the RW process $(X_t)_t$ on C_α only, that is we construct the transition matrix

$$\hat{P}_\alpha(x, y) = \sum_{z \in \mathbb{X}} P(x, z) q_y(z), \quad x, y \in C_\alpha$$

where $q_y(z)$ is the probability that z is the next node in C_α that is hit by $(X_t)_t$ conditioned on $X_0 = y$. \hat{P}_α describes the dynamics of $(X_t)_t$ restricted to C_α . To compute P_α , we must compute $k = |C_\alpha|$ committor functions q_y . With \hat{P}_α at hand, we need to specify the number m of core sets we want to partition C_α into. Some hard clustering method, e.g. [DHFS00], can then be used to partition C_α into the core sets C_1, \dots, C_m . As a final step, MSM clustering computes m committor functions $q_i(x) = \mathbf{P}(X_\tau = C_i | X_0 = x)$, where $\tau = \inf\{t > 0 : X_t \in C_\alpha\}$, as fuzzy **affiliation functions** that tell how much any node $x \in \mathbb{X}$ is affiliated with each of the cores C_i . All committor functions are computed by solving linear systems of the form (1.48).

MSM clustering requires the specification of two input parameters m and α . The larger α is, the more modules one will find, so different values of α can be used in order to find a hierarchy of different clusterings. One way to obtain a good value for m once α is selected is the following: For some not too large k , compute the k largest eigenvalues $\lambda_1 = 1 > \dots > \lambda_k$ of \hat{P}_α . Select m in such a way that the gap $\Delta_m = \lambda_m - \lambda_{m+1}$ is maximal.

B.2 Alternative Module Detection Methods

Module detection in networks is a huge area of research, see e.g. [For10, SLB12, SDYB12, DYB10, EL09, New06, NBW06]. A comprehensive review of the different approaches that exist in the literature would be well beyond the scope of this thesis. A nice review can be found in [For10, vL07]. We will focus on methods based on random walks and describe a few of the most important ones briefly.

Modularity Optimization. This is one of the most popular methods, originally developed by Newman [New06]. Given a full partition of \mathbb{X} into m sets A_1, \dots, A_m , we define the **modularity** of the partition to be

$$Q = \sum_{i=1}^m [\mathbf{P}_\mu(X_1 \in A_i, X_0 \in A_i) - \mathbf{P}_\mu(X_1 \in A_i)\mathbf{P}_\mu(X_0 \in A_i)]. \quad (\text{B.2})$$

To find the best partition, Q is optimized over all full partitions A_1, \dots, A_m of \mathbb{X} . The optimization problem is NP-hard, but fast greedy heuristics exist [NBW06]. Modularity optimization is a one-step method, and if we apply it to an irreversible process with transition matrix P , it effectively performs a symmetrisation $P^s = \frac{1}{2}(P + P^-)$.

Markov Stability. Markov stability (MS) [LDB09] is a state-of-the-art approach for community detection based on revealing communities at different scales by looking at how the probability flow spreads out over time. At the heart of MS is stability function

$$r(t) = \sum_{i=1}^m [\mathbf{P}_\mu(X_t \in A_i, X_0 \in A_i) - \mathbf{P}_\mu(X_t \in A_i)\mathbf{P}_\mu(X_0 \in A_i)], \quad (\text{B.3})$$

which is optimized for every fixed $t > 0$ over all full partitions A_1, \dots, A_m in the same way as the modularity function Q (in fact $Q = r(1)$). For any fixed $t > 0$, $r(t)$ encodes information about paths of length t , and the method uses t as a resolution parameter: The optimization of $r(t)$ is carried out for all values of t in the desired range, and one searches for communities which persist for a range of values of t .

Infomap. Infomap [DYB10] is a popular method for detecting communities, relying on the idea that community structure can be used to describe the position of a random walker on the network compactly by reusing codewords in different communities. Given a full partition A_1, \dots, A_m , Infomap computes the average description length l that a two-level code would need to describe a realization $(X_t)_{1 \leq t \leq T}$ of length T of the RW. The code uses an alphabet to describe the position of the random walker within each module A_i (and compression is achieved because this alphabet is smaller than the alphabet we would need to describe a position in all of \mathbb{X}), and it uses special codewords for jumps between modules. Compression is maximal if these jumps are rare, so the modules found by Infomap are typically metastable sets.

Infomap can deal with directed networks and performed very well in a recent benchmark [LF09] with clique-like communities, but despite its information-theoretic origin it is inherently a one-step method and as such it can fail at detecting non-clique-like communities [SDYB12] by displaying an overpartitioning effect.

C Appendix C

C.1 Proof of Theorem 3.4

Here we give the proof of Theorem 3.4 from Section 3.2.2. For ease of notation, let $\|\cdot\| = \|\cdot\|_\mu$.

Proof. Let ϕ be the solution to (3.27), and write $\phi = Q\phi + \phi_\perp = \phi_{||} + \phi_\perp$ with $\phi_\perp \in D^\perp$. The first step is to show that $\|\phi - \phi_{||}\| = \inf_{\psi \in D} \|\phi - \psi\|$, i.e. the infimum in the definition of ϵ_0 is attained at $\phi_{||}$. But this is clear since for any $\psi \in D$, by orthogonality we have

$$\|\phi - \psi\|^2 = \|\phi_{||} - \psi + \phi_\perp\|^2 = \|\phi_{||} - \psi\|^2 + \|\phi_\perp\|^2$$

which attains its minimum of $\epsilon_0^2 = \|\phi_\perp\|^2$ for $\psi = \phi_{||}$. By (3.27), $\phi_{||}$ solves the equation

$$\mathcal{B}(\phi, \psi) = \mathcal{B}(\phi_{||}, \psi) + \mathcal{B}(\phi_\perp, \psi) = 0 \quad \forall \psi \in D,$$

and if we write $\phi_{||} = \sum_{i=1}^n \hat{\phi}_i^* \chi_i + 1\chi_{n+1}$ with n unknown coefficients $\hat{\phi}_i^*$ (note that a general element of D is of this form), this takes the matrix form

$$\hat{B}\hat{\phi}^* - c = F,$$

where in components we have $\hat{B}_{ij} = \mathcal{B}(\chi_i, \chi_j)$, $c_i = -\mathcal{B}(\phi_\perp, \chi_i) = -\langle \phi_\perp, B\chi_i \rangle_\mu$ and $F_i = -\langle \chi_i, B\chi_{n+1} \rangle_\mu$. On the other hand, the Galerkin solution $\hat{\phi} = \sum_i \hat{\phi}_i \chi_i$ satisfies $\hat{B}\hat{\phi} = F$ by 3.28, hence we obtain

$$\hat{B}(\hat{\phi}^* - \hat{\phi}) = c. \tag{C.1}$$

Now we can write

$$\begin{aligned}
 \boldsymbol{\varepsilon}^2 &= \|\phi_{\parallel} + \phi_{\perp} - \hat{\phi}\|^2 = \|\phi_{\parallel} - \hat{\phi}\|^2 + \|\phi_{\perp}\|^2 \\
 &= \left\langle \sum_i (\hat{\phi}_i^* - \hat{\phi}_i) \chi_i, \sum_j (\hat{\phi}_j^* - \hat{\phi}_j) \chi_j \right\rangle_{\mu} + \boldsymbol{\varepsilon}_0^2 \\
 &= (\hat{\phi}^* - \hat{\phi})^T \hat{M} (\hat{\phi}^* - \hat{\phi}) + \boldsymbol{\varepsilon}_0^2
 \end{aligned}$$

where $\hat{M}_{ij} = \langle \chi_i, \chi_j \rangle_{\mu}$. The scalar product $\langle \cdot, \cdot \rangle_{\mu}$ on $D_0 \subset V$ induces a natural scalar product on \mathbb{R}^n by the isomorphism $\hat{\phi} \mapsto \sum_i \phi_i \hat{\chi}_i$:

$$\left\langle \sum_i \hat{\phi}_i \chi_i, \sum_j \hat{\phi}'_j \chi_j \right\rangle_{\mu} = \hat{\phi}^T \hat{M} \hat{\phi}' =: \langle \hat{\phi}, \hat{\phi}' \rangle_M$$

The error $\boldsymbol{\varepsilon}^2$ is exactly $\boldsymbol{\varepsilon}_0^2$ plus the distance between Galerkin solution and best approximation measured in this scalar product. There is also a natural bilinear form inherited from \mathcal{B} on \mathbb{R}^n :

$$\mathcal{B} \left(\sum_i \hat{\phi}_i \chi_i, \sum_j \hat{\phi}'_j \chi_j \right) = \hat{\phi}^T \hat{B} \hat{\phi}' = \langle \hat{\phi}, \hat{M}^{-1} \hat{B} \hat{\phi}' \rangle_M$$

The Matrix $\hat{M}^{-1} \hat{B}$ is symmetric since $\mathcal{B}(\cdot, \cdot)$ is symmetric. Moreover, since $\mathcal{B}(\cdot, \cdot)$ is elliptic,

$$\langle \hat{\phi}, \hat{M}^{-1} \hat{B} \hat{\phi} \rangle_M = A \left(\sum_i \hat{\phi}_i \chi_i, \sum_j \hat{\phi}_j \chi_j \right) \geq \alpha_2 \left\langle \sum_i \hat{\phi}_i \chi_i, \sum_j \hat{\phi}_j \chi_j \right\rangle_{\mu} = \alpha_2 \langle \hat{\phi}, \hat{\phi} \rangle_M \quad (\text{C.2})$$

In particular, $\hat{M}^{-1} \hat{B}$ is positive, hence it has a positive and symmetric square root $\hat{S}^2 = \hat{M}^{-1} \hat{B}$. Now, for any $\hat{\phi} \in \mathbb{R}^n$ it holds by virtue of (C.2),

$$\begin{aligned}
 \langle \hat{\phi}, \hat{\phi} \rangle_M &\leq \frac{1}{\alpha_2} \langle \hat{\phi}, \hat{M}^{-1} \hat{B} \hat{\phi} \rangle_M = \frac{1}{\alpha_2} \langle \hat{S} \hat{\phi}, \hat{S} \hat{\phi} \rangle_M \\
 &\leq \frac{1}{\alpha_2^2} \langle \hat{S} \hat{\phi}, \hat{M}^{-1} \hat{B} \hat{S} \hat{\phi} \rangle_M = \frac{1}{\alpha_2^2} \langle \hat{M}^{-1} \hat{B} \hat{\phi}, \hat{M}^{-1} \hat{B} \hat{\phi} \rangle_M.
 \end{aligned} \quad (\text{C.3})$$

Now we apply the inequality (C.3) to $\hat{\phi}^* - \hat{\phi}$ and use (C.1):

$$\boldsymbol{\varepsilon}^2 \leq \boldsymbol{\varepsilon}_0^2 + \frac{1}{\alpha_2^2} \langle \hat{M}^{-1} c, \hat{M}^{-1} c \rangle_M. \quad (\text{C.4})$$

Now for some final simplifications, note that the orthogonal projection Q onto D_0 can be written as

$$Q\psi = \sum_{i,j=1}^n \hat{M}_{ij}^{-1} \langle \chi_j, \psi \rangle_{\mu} \chi_i.$$

Using this we can write

$$\begin{aligned}
 \langle \hat{M}^{-1}c, \hat{M}^{-1}c \rangle_M &= \sum_{ij} c_i \hat{M}_{ij}^{-1} c_j = \sum_{ij} \langle \chi_i, B\phi_\perp \rangle_\mu M_{ij}^{-1} \langle \chi_j, B\phi_\perp \rangle_\mu \\
 &= \left\langle \sum_{ij} M_{ij}^{-1} \langle \chi_j, B\phi_\perp \rangle_\mu \chi_i, B\phi_\perp \right\rangle_\mu = \langle QB\phi_\perp, B\phi_\perp \rangle_\mu \\
 &= \langle QB\phi_\perp, QB\phi_\perp \rangle_\mu
 \end{aligned}$$

To arrive at the final result, notice that

$$\begin{aligned}
 \langle QB\phi_\perp, QB\phi_\perp \rangle_\mu &\leq \left(\sup_{\phi'_\perp \in D^\perp} \frac{\langle QB\phi'_\perp, QB\phi'_\perp \rangle_\mu}{\langle \phi'_\perp, \phi'_\perp \rangle_\mu} \right) \cdot \langle \phi_\perp, \phi_\perp \rangle_\mu \\
 &= \left(\sup_{\phi'_\perp \in D^\perp} \frac{\langle QBQ^\perp \phi'_\perp, QBQ^\perp \phi'_\perp \rangle_\mu}{\langle \phi'_\perp, \phi'_\perp \rangle_\mu} \right) \cdot \langle \phi_\perp, \phi_\perp \rangle_\mu \\
 &\leq \left(\sup_{\phi' \in V} \frac{\langle QBQ^\perp \phi', QBQ^\perp \phi' \rangle_\mu}{\langle \phi', \phi' \rangle_\mu} \right) \cdot \langle \phi_\perp, \phi_\perp \rangle_\mu \\
 &= \|QBQ^\perp\|^2 \langle \phi_\perp, \phi_\perp \rangle_\mu
 \end{aligned}$$

Plugging these inequalities into (C.4) and dividing by ϵ_0^2 completes the proof. ■

C.2 Best-approximation error bound

In this appendix, we prove lemma 3.7:

$$\epsilon_0 = \|Q^\perp \phi\|_\mu \leq \|P^\perp \phi\|_\mu + \mu(\mathbb{T})^{1/2} [\kappa \|f\|_\infty + 2\|P^\perp \phi\|_\infty].$$

Recall that $\kappa = \sup_{x \in \mathbb{T}} \mathbf{E}_x[\tau_C]$ and P is the orthogonal projection onto the subspace $V_c = \{v \in L^2(\mathbb{X}, \mu), v = \text{const on every } C_i\} \subset L^2(\mathbb{X}, \mu)$. Note that $P^\perp \phi = 0$ on C . The errors $\|P^\perp \phi\|$ and $\|P^\perp \phi\|_\infty$ measure how constant the solution ϕ is on the core sets. We write $\|\cdot\| = \|\cdot\|_\mu$ throughout the proof for convenience.

Proof. The proof closely follows the proof of theorem (12) in [Sar11]. The first step of the proof is to realize that the committor subspace D where Q projects onto can be written as $D = \{v \in L^2(\mathbb{X}, \mu), v = \text{const on every } C_i, Lv = 0 \text{ on } C\}$. To see this, note that the values v takes on the C_i can be used as boundary values for the Dirichlet problem $Lv = 0$ on \mathbb{T} . A linear combination of committor functions is obviously a solution to this problem. But the solution to the Dirichlet problem must be unique, otherwise one can construct a contradiction to the uniqueness of the invariant distribution, see [Sar11].

By definition we have $\|Q^\perp \phi\| \leq \|\phi - I\phi\|$ for every interpolation $I\phi \in D$ of ϕ . With the definition

Appendix C. Appendix C

of P from above, we will take $q = I\phi$ such that

$$Lq = 0 \text{ on } \mathbb{T}, \quad q = P\phi \text{ on } \mathbb{X} \setminus \mathbb{T}. \quad (\text{C.5})$$

Now $D \subset V$, therefore $q \in V_c$ and $Pq = q$. Therefore (C.5) is equivalent to

$$PLPq = 0 \text{ on } \mathbb{T}, \quad q = P\phi \text{ on } \mathbb{X} \setminus \mathbb{T}. \quad (\text{C.6})$$

Now define $e := P\phi - q$. Then we have

$$PLPe = PLP(P\phi - q) = PLP\phi - PLPq = PL\phi - PLP^\perp\phi - PLPq$$

and by (C.6) and since $L\phi = f\phi$ on $\mathbb{X} \setminus A \supset \mathbb{T}$, we have

$$PLPe = Pf\phi - PLP^\perp\phi \text{ on } \mathbb{T}, \quad e = 0 \text{ on } \mathbb{X} \setminus \mathbb{T}. \quad (\text{C.7})$$

Therefore, $e \in E_\Theta = \{v \in L^2(\mathbb{X}, \mu), v = 0 \text{ on } \mathbb{X} \setminus \mathbb{T}\}$ and with Θ being the orthogonal projection onto E_Θ , e has to fulfil

$$\Theta PLP\Theta e = \Theta Pf\phi - \Theta PLP^\perp\phi.$$

Since $\Theta P = P\Theta = \Theta$, this can be written as

$$Re := \Theta L\Theta e = \Theta f\phi - \Theta LP^\perp\phi.$$

The operator $R = \Theta L\Theta$ is invertible on E_Θ : If this wasn't the case, there would be a nontrivial solution v to

$$Lv = 0 \text{ on } \mathbb{T}, \quad v = 0 \text{ on } \mathbb{X} \setminus \mathbb{T}.$$

But the solution to this boundary value problem is again unique, and hence there is only the trivial solution. This gives

$$e = R^{-1}\Theta f\phi - R^{-1}\Theta LP^\perp\phi, \quad (\text{C.8})$$

and $\|R^{-1}\| = \frac{1}{|\lambda_0|}$ where λ_0 is the principal eigenvalue of R . Due to an estimate by Varadhan we have

$$\frac{1}{|\lambda_0|} \leq \sup_{x \in \mathbb{T}} \mathbf{E}_x[\tau_{\mathbb{X} \setminus \mathbb{T}}] =: \kappa,$$

see e.g. [Bov09]. To complete the derivation we need to focus on the second term in (C.8). Since R^{-1} is an operator on E_Θ , we can write it as $R^{-1}\Theta LP^\perp\phi =: \Theta g$, where the function Θg solves

$$\Theta L\Theta g = R\Theta g = \Theta LP^\perp\phi \Leftrightarrow \Theta L[\Theta g - P^\perp\phi] = 0$$

by the definition of R and Θg . Therefore $w := \Theta g - P^\perp \phi$ solves the boundary value problem

$$Lw = 0 \text{ on } \mathbb{T}, \quad w = -P^\perp \phi \text{ on } \mathbb{X} \setminus \mathbb{T} \quad (\text{C.9})$$

which implies that $\|w\|_\infty \leq \|P^\perp \phi\|_\infty$, this follows from Dynkin's formula or Lemma 3 in [Sar11]. Finally,

$$\|\Theta g\| \leq \mu(\mathbb{T})^{1/2} \|\Theta g\|_\infty \leq \mu(\mathbb{T})^{1/2} (\|P^\perp \phi\|_\infty + \|w\|_\infty) \leq 2\mu(\mathbb{T})^{1/2} \|P^\perp \phi\|_\infty$$

holds by the triangle inequality and the above considerations. Now focus on the first term in (C.8). Note that by the maximum principle, ϕ achieves its maximum of 1 on the boundary of $\mathbb{X} \setminus A \supset \mathbb{T}$, therefore $\max_{x \in \mathbb{T}} |\phi(x)| \leq 1$. Then we have

$$\|\Theta f \phi\| \leq \mu(\mathbb{T})^{1/2} \|f\|_\infty \max_{x \in \mathbb{T}} |\phi(x)| \leq \mu(\mathbb{T})^{1/2} \|f\|_\infty.$$

Now putting everything together, we arrive at

$$\begin{aligned} \|e\| &\leq \|R^{-1}\| \|\Theta f \phi\| + \|R^{-1} \Theta L P^\perp \phi\| \\ &\leq \kappa \|\Theta f \phi\| + \|\Theta g\| \\ &\leq \mu(\mathbb{T})^{1/2} [\kappa \|f\|_\infty + 2 \|P^\perp \phi\|_\infty]. \end{aligned}$$

Finally, note that by the triangle inequality

$$\|Q^\perp \phi\| \leq \|\phi - q\| \leq \|\phi - P\phi\| + \|P\phi - q\| = \|P^\perp \phi\| + \|e\|$$

which completes the proof. ■

Bibliography

- [AS04] S. Abe and N. Suzuki. Scale-free network of earthquakes. *EPL (Europhysics Letters)*, 65(4):581, 2004. 64
- [AS06] S. Abe and N. Suzuki. Complex-network description of seismicity. *Nonlinear Processes in Geophysics*, 13(2):145–150, 2006. 64
- [BBI12] Alessandra Bianchi, Anton Bovier, and Dmitry Ioffe. Pointwise estimates and exponential laws in metastable systems via coupling methods. *The Annals of Probability*, 40(1):339–371, 01 2012. 21
- [BC14] Ralf Banisch and Nataša Djurdjevac Conrad. Cycle-flow based module detection in directed recurrence networks. *EPL (Europhysics Letters)*, 108(6):68008, 2014. 4, 54, 55
- [BCS15] R. Banisch, N.Djurdjevac Conrad, and Ch. Schütte. Reactive flows and unproductive cycles for random walks on complex networks. *The European Physical Journal Special Topics*, pages 1–19, 2015. 4, 76
- [BDP07] G. Bussi, D. Donadio, and M. Parrinello. Canonical sampling through velocity rescaling. *J. Chem. Phys.*, 126:014101, 2007. 112
- [BEGK02] A. Bovier, M. Eckhoff, V. Gayrard, and M. Klein. Metastability and low lying spectra in reversible markov chains. *Comm. Math. Phys.*, 228:219–255, 2002. 20, 22, 23, 51, 53, 135
- [Bel56] Richard Bellman. Dynamic programming and lagrange multipliers. *Proceedings of the National Academy of Sciences of the United States of America*, 42(10):767, 1956. 3
- [Ber11] Dimitri P. Bertsekas. Approximate policy iteration: a survey and some new methods. *Journal of Control Theory and Applications*, 9(3):310–335, 2011. 128, 129
- [BGK05] A. Bovier, V. Gayrard, and M. Klein. Metastability in reversible diffusion processes ii. precise asymptotics for small eigenvalues. *J. Eur. Math. Soc.*, 7:69–99, 2005. 20, 23

Bibliography

- [BH14] Ralf Banisch and Carsten Hartmann. Meshless galerkin discretization of lq-type stochastic control problems. *submitted to SICON*, 2014. 4, 98
- [Bov09] A. Bovier. *Methods of Contemporary Statistical Mechanics*. Metastability. Springer, 2009. 120, 146
- [BPN14] G. R. Bowman, V. S. Pande, and F. Noé, editors. *An Introduction to Markov State Models and Their Application to Long Timescale Molecular Simulation*, volume 797 of *Advances in Experimental Medicine and Biology*. Springer, 2014. 62
- [Bra07] D. Braess. *Finite Elements: Theory, Fast Solvers, and Applications in Solid Mechanics*. Cambridge University Press, 2007. 7, 68, 99, 100, 101, 103, 127
- [Bra12] M. Braack. *Finite Elemente*. 2012. 68
- [Bré99] P. Brémaud. *Markov chains, Gibbs Fields, Monte Carlo simulation, and Queues*. Springer, 1999. 19
- [CBS15] N. Djurdjevac Conrad, R. Banisch, and Ch. Schütte. Modularity of directed networks: cycle decomposition approach. *Accepted in JCD*, <http://arxiv.org/abs/1407.8039>, 2, 2015. 4, 54
- [CMZ11] T Chou, K Mallick, and RKP Zia. Non-equilibrium statistical mechanics: from a paradigmatic model to biological transport. *Reports on progress in physics*, 74(11):116601, 2011. 2
- [CVE14] Maria Cameron and Eric Vanden-Eijnden. Flows in complex networks: Theory, algorithms, and application to lennard–jones cluster rearrangement. *Journal of Statistical Physics*, 156(3):427–454, 2014. 80
- [DBCS11] N. Djurdjevac, S. Bruckner, T. O. F. Conrad, and Ch. Schütte. Random walks on complex modular networks. *Journal of Numerical Analysis, Industrial and Applied Mathematics*, 6:29–50, 2011. 5, 53, 139
- [DGP08] Jörn Davidsen, Peter Grassberger, and Maya Paczuski. Networks of recurrent events, a theory of records, and an application to finding causal signatures in seismicity. *Phys. Rev. E*, 77:066104, 2008. 64
- [DHFS00] P. Deuffhard, W. Huisinga, A. Fischer, and Ch. Schütte. Identification of almost invariant aggregates in reversible nearly uncoupled markov chains. *Linear Algebra and its Applications*, 315(13):39 – 59, 2000. 1, 140
- [DPMR96] P. Dai Pra, L. Meneghini, and W. Runggaldier. Connections between stochastic control and dynamic games. *Math. Control Signals Systems*, 9:303–326, 1996. 3, 96

- [DSS12] N. Djurdjevac, M. Sarich, and Ch Schütte. Estimating the eigenvalue error of markov state models. *Multiscale Modeling & Simulation*, 10:61–81, 2012. 20, 22, 103, 104
- [DV75] Monroe D Donsker and S R S Varadhan. On a variational formula for the principal eigenvalue for operators with maximum principle. *Proceedings of the National Academy of Sciences of the United States of America*, 72(3):780–783, 03 1975. 120
- [DYB10] J.-C. Delvenne, S. N. Yaliraki, and M. Barahona. Stability of graph communities across time scales. *Proceedings of the National Academy of Sciences*, 107(29):12755–12760, 2010. 5, 54, 64, 140, 141
- [DYP93] T. Darden, D. York, and L. Pedersen. Particle mesh Ewald: An $N \cdot \log(N)$ method for Ewald sums in large systems. *J. Chem. Phys.*, 98:10089–10092, 1993. 112
- [DZD⁺10] Reik V Donner, Yong Zou, Jonathan F Donges, Norbert Marwan, and Jürgen Kurths. Recurrence networks—a novel paradigm for nonlinear time series analysis. *New Journal of Physics*, 12(3):033025, 2010. 58
- [EL98] Evans and Lorence. *Partial Differential Equations*. American Mathematical Society, 1998. 68, 93
- [EL09] T. S. Evans and R. Lambiotte. Line graphs, link partitions, and overlapping communities. *Phys. Rev. E*, 80:016105, 2009. 5, 140
- [FDK13] Jonathan F Donges, Reik V. Donner, and Jürgen Kurths. Testing time series irreversibility using complex network methods. *Europhysics Letters*, 102(1), 2013. 64
- [For10] Santo Fortunato. Community detection in graphs. *Physics Reports*, 486(3–5):75 – 174, 2010. 5, 140
- [FS06] W.H. Fleming and H.M. Soner. *Controlled Markov Processes and Viscosity Solutions*. Springer Verlag, 2006. 16, 27, 28, 29, 31, 32, 33, 94, 95
- [HAO⁺06] V. Hornak, R. Abel, . Okur, B. Strockbine, A. Roitberg, and C. Simmerling. Comparison of multiple amber force fields and development of improved protein backbone parameters. *Proteins*, 65:712–725, 2006. 112
- [HBBF97] B. Hess, H. Bekker, H.J.C. Berendsen, and J. Fraaije. Lincs: a linear constraint solver for molecular simulations. *J. Comp. Chem.*, 18:1463–1472, 1997. 112
- [HBS⁺14] C. Hartmann, R. Banisch, M. Sarich, T. Badowski, and Ch. Schütte. Characterization of rare events in molecular dynamics. *Entropy*, 16:350–376, 2014. 4, 92

Bibliography

- [HHMS05] Chii-Ruey Hwang, Shu-Yin Hwang-Ma, and Shuenn-Jyi Sheu. Accelerating diffusions. *Ann. Appl. Probab.*, 15(2):1433–1444, 05 2005. 3, 4, 50
- [HLPZ14] C. Hartmann, J.C. Latorre, G.A. Pavliotis, and W. Zhang. Optimal control of multiscale systems using reduced-order models. *J. Comput. Dynamics*, 1:279–306, 2014. 105
- [HP86] U. G. Haussmann and E. Pardoux. Time reversal of diffusions. *Ann. Probab.*, 14(4):1188–1205, 10 1986. 15
- [HP14] CP Ho and P Parpas. Singularly perturbed markov decision processes: A multiresolution algorithm. *SIAM JOURNAL ON CONTROL AND OPTIMIZATION*, 52:3854–3886, 2014. 115
- [HS06] W. Huisinga and B. Schmidt. Metastability and dominant eigenvalues of transfer operators. *Lecture Notes in Computational Science and Engineering*, 49:167–182, 2006. 1, 20, 23
- [HS12] C. Hartmann and C. Schütte. Efficient rare event simulation by optimal nonequilibrium forcing. *J. Stat. Mech. Theor. Exp.*, 11:4, November 2012. 3, 92, 127
- [Hui01] W. Huisinga. Metastability of markovian systems: A transfer operator approach in application to molecular dynamics. *PhD Thesis at Freie Universität Berlin*, 2001. 23
- [JQQ04] D. Jiang, M. Qian, and M.-P. Quian. *Mathematical theory of nonequilibrium steady states: on the frontier of probability and dynamical systems*. Springer, 2004. 2, 5, 40, 41, 44
- [Kai12] K. Kaiser. Optimalsteuerung von diffusionsprozessen mit indefinitem zeithorizont. *Bachelor Thesis at Freie Universität Berlin*, 2012. 16
- [Kal06] Sophia L. Kalpazidou. *Cycle Representations of Markov Processes*. Springer, 2006. 2, 5, 36, 38
- [Kus01] H. Kushner. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer Verlag, 2001. 3, 7, 128
- [LDB09] R. Lambiotte, J. C. Delvenne, and M. Barahona. Laplacian dynamics and multi-scale modular structure in networks. *ArXiv*, 2009. 140
- [LF09] Andrea Lancichinetti and Santo Fortunato. Benchmarks for testing community detection algorithms on directed and weighted graphs with overlapping communities. *Phys. Rev. E*, 80:016118, 2009. 141
- [LM94] A. Lasota and M. Mackey. *Chaos, Fractals and Noise*. Springer Verlag, 1994. 14

- [Met07] P. Metzner. Transition path theory for markov processes. *PhD thesis*, 2007. 15, 24, 26, 27, 66, 69, 78
- [NBW06] MEJ Newman, AL Barabasi, and DJ Watts. *The Structure and Dynamics of Networks*. Princeton Univ Press, Princeton, NJ, 2006. 5, 140
- [New06] M. E. J. Newman. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, 103(23):8577–8582, 2006. 5, 140
- [Øks03] B. Øksendal. *Stochastic Differential Equations*. Springer Verlag, 2003. 13, 14, 15, 16, 28, 92, 135
- [PB74] George Poole and Thomas Boullion. A survey on m-matrices. *SIAM Review*, 16(4):pp. 419–427, 1974. 136
- [Pol15] M. Poletti. Cycle/Cocycle Oblique Projections on Oriented Graphs. *Letters in Mathematical Physics*, 105:89–107, January 2015. 2, 36, 37
- [Pri57] Robert Clay Prim. Shortest connection networks and some generalizations. *Bell system technical journal*, 36(6):1389–1401, 1957. 39, 59
- [PS08] G. A. Pavliotis and A.M. Stuart. *Multiscale Methods - Averaging and Homogenization*. Springer Verlag, 2008. 89, 115, 117, 119
- [PW14] Panos Pappas and Mort Webster. A stochastic multiscale model for electricity generation capacity expansion. *European Journal of Operational Research*, 232(2):359 – 374, 2014. 115
- [RS14] L. Rey-Bellet and K. Spiliopoulos. Irreversible Langevin samplers and variance reduction: a large deviation approach. *ArXiv e-prints*, March 2014. 3, 4, 50
- [Sar11] M. Sarich. Projected transfer operators. discretization of markov processes in high-dimensional state spaces. *PhD Thesis at Freie Universität Berlin*, 2011. 7, 21, 22, 145, 147
- [SBHS13] Marco Sarich, Ralf Banisch, Carsten Hartmann, and Christof Schütte. Markov state models for rare events in molecular dynamics. *Entropy*, 16(1):258, 2013. 4, 98
- [SCB⁺14] M. Sarich, N. Djurdjevac Conrad, S. Bruckner, T. O. F. Conrad, and Ch. Schütte. Modularity revisited: A novel dynamics-based concept for decomposing complex networks. *Journal of Computational Dynamics*, 1(1):191–212, 2014. 5, 53, 139
- [Sch76] J. Schnakenberg. Network theory of microscopic and macroscopic behavior of master equation systems. *Rev. Mod. Phys.*, 48:571–585, Oct 1976. 2, 37

Bibliography

- [SDYB12] Michael T. Schaub, Jean-Charles Delvenne, Sophia N. Yaliraki, and Mauricio Barahona. Markov dynamics as a zooming lens for multiscale community detection: Non clique-like communities and the field-of-view limit. *PLoS ONE*, 7(2):e32210, 02 2012. 5, 64, 140, 141
- [SFHD99] Ch Schütte, A Fischer, W Huisinga, and P Deuffhard. A direct approach to conformational dynamics based on hybrid monte carlo. *Journal of Computational Physics*, 151(1):146 – 168, 1999. 1
- [She85] S. Sheu. Stochastic control and exit probabilities of jump processes. *J. Control Optim.*, 23:306–328, 1985. 95
- [SLB12] Michael T. Schaub, Renaud Lambiotte, and Mauricio Barahona. Encoding dynamics for multiscale community detection: Markov time sweeping for the map equation. *Phys. Rev. E*, 86:026112, 2012. 5, 140
- [Slo13] M Slowik. A note on variational representations of capacities for reversible and non-reversible markov chains. *Preprint*, 2013. 51, 52, 53
- [SNL⁺11] Ch. Schütte, F. Noe, J. Lu, M. Sarich, and E. Vanden-Eijnden. Markov state models based on milestoning. *J. Chem. Phys.*, 134:204105, 2011. 1, 99, 107, 113
- [SNS10] M. Sarich, F. Noé, and Ch. Schütte. On the Approximation Quality of Markov State Models. *Multiscale Modeling and Simulation*, 8(4):1154–1177, 2010. 1, 58, 104, 107, 113
- [SS13a] Ch. Schütte and M. Sarich. *Metastability and Markov State Models in Molecular Dynamics: Modeling, Analysis, Algorithmic Approaches*, volume 24 of *Courant Lecture Notes*. American Mathematical Society, December 2013. 62, 63
- [SS13b] Ch. Schütte and M. Sarich. Metastability and markov state models in molecular dynamics: Modeling, analysis, algorithmic approaches. *Courant Lecture Notes*, 24, 2013. 9, 20
- [SS14] Marco Sarich and Christof Schütte. Utilizing hitting times for finding metastable sets in non-reversible markov chains. Technical Report 14-32, ZIB, Takustr.7, 14195 Berlin, 2014. 23
- [SSG10] Yi Sun, Jürgen Schmidhuber, and Faustino J Gomez. Improving the asymptotic performance of markov chain monte-carlo by inserting vortices. In *Advances in Neural Information Processing Systems*, pages 2235–2243, 2010. 3, 4, 50, 126
- [USG] *US Geological Survey*. <http://www.earthquake.usgs.gov>. 65
- [VDSLH⁺05] D. Van Der Spoel, E. Lindahl, B. Hess, G. Groenhof, A.E. Mark, and H.J. C. Berendsen. Gromacs: Fast, flexible, and free. *J. Comp. Chem.*, 26:1701–1718, 2005. 112

- [VE06] E. Vanden-Eijnden. Transition path theory. *Lect. Notes Phys.*, 703:439–478, 2006. 6, 24, 99, 107
- [vH07] Ramon van Handel. Stochastic calculus, filtering, and stochastic control. *Lecture Notes*, 2007. 31, 32, 93, 94
- [vL07] Ulrike von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007. 140
- [Wal82] P. Walters. *An Introduction to Ergodic Theory*, volume 79 of *Graduate Texts in Mathematics*. Springer, 1982. 10
- [Wen99] H. Wendland. Meshless galerkin methods using radial basis functions. *Math. Comput.*, 68:1521–1531, 1999. 7, 127

Colophon. This thesis was typeset with \LaTeX using the EPFL template available at <http://phd.epfl.ch/thesistemplates>. It uses the Utopia font created by C. Dellagiacoma, which is a combination of OpenSource Utopia by Adobe and Greek characters provided by P. Pichareau within the OpenSource MathDesign package for \LaTeX .

Declaration

Ich, Ralf Banisch, erkläre hiermit, dass ich die vorliegende Arbeit selbständig verfasst habe und dass die Arbeit zu keinem Zeitpunkt in einem früheren Promotionsverfahren eingereicht oder in Auszügen verwendet worden ist. Alle verwandten Hilfsmittel und Publikationen wurden im Text zitiert und sind im Anhang aufgelistet.

Berlin, Juli 2015

Ralf Banisch