

**Mathematics of evolving ecologies in chemical reaction networks and
the origins of biochemical life**

by

Praful Gagrani

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

Doctor of Philosophy

(Physics)

at the

UNIVERSITY OF WISCONSIN–MADISON

2023

Date of final oral examination: 05/09/2023

The dissertation is approved by the following members of the Final Oral Committee:

Eric Smith, Honorary Associate, Physics

David Baum, Professor, Botany

Sridhara Dasu, Professor, Physics

Alex Levchenko, Professor, Physics

Gheorghe Craciun, Professor, Mathematics

© Copyright by Praful Gagrani 2023

All Rights Reserved

*Dedicated to philosophers, teachers, friends, family, and lovers for deepening my
understanding and appreciation of life.*

ACKNOWLEDGMENTS

Akin to life, which only comes from life, the explorations in this thesis are preceded by those of others. I am indebted to our ancestors, for their innovations in thinking have provided solid foundations, on which, this thesis has been articulated and brought to fruition. An incomplete list of philosophers whose insights have, in one way or another, shaped this research is: Vedic and Buddhist philosophers, Baruch Spinoza, Isaac Newton, Ludwig Boltzmann, Henri Bergson, Alfred Whitehead, Lev Landau, Gilles Deleuze, V. I. Arnold, Manuel De Landa, and Matt Segall.

I want to thank my family and relatives for both inspiring and helping nurture my childhood fascination with physics. I would be remiss in not mentioning the teachers, researchers, professors, students, and administrative staff at every academic institution I have been fortunate enough to attend or work with. In chronological order, Texmaco International School (Karawang, Indonesia), Trakya Koleji (Çorlu, Turkey), Delhi Public School (Kanpur and Pune, India), Hume McHenry Memorial High School (Pune, India), Cookeville High School (Cookeville, USA), Indian Institute of Technology-Delhi, University of California-Davis, University of Texas-Austin, Kansas State University (Manhattan, USA), and most of all, the University of Wisconsin-Madison.

A special thanks to Ahmet Alacaoglu and Stephen Wright for discussion and references on optimization, David Anderson, Gheorghe Craciun

and Tung Nguyen for explanations and references on mathematical theory of chemical reaction networks, Melvin Leok for references on symplectic integrators and information geometry, Alex Levchenko for clarifications and references regarding spectrum of the action functional, Emily Dolson, Philippe Nghe, and Anton Zadorin for discussions on ecology, and professors V Ravishankar, Dave Auckly, Larry Weaver, Lado Samushia, and Akikazu Hashimoto for mentorship.

I am indebted to Victor Blanco and Vladimir Sotirov for collaboration on two manuscripts. The sections written by them in the manuscripts, however, have been omitted in the thesis. Neither multidisciplinary work in this thesis nor graduate life in Madison would be possible without friends from the physics and other departments. I want to thank friends Vladimir Sotirov for \TeX support, Connor Simpson for help with combinatorics and programming, Diego Rojas, Gautam Neelakantham, Asvin G., and Vladimir Sotirov for mathematical discussions, Pavani Ganju, Alex Plum, Zhen Peng, and Tymofii Sokolskyi for discussions in ecology, chemistry, and biology, and Vivek Bagaria, Ashish Bora, Abhimanyu Dubey, Jeff Lazar, Neil Leonard, Merritt Losert, Tyler Mitchell, Soren Ormseth, Rahul Parhi, Aditya Sarma Phukon, Erika Pirnes, John Podczerwinski, Ibrahim Safa, Dylan Teague, and Michael Wolfe for occasional good advice and a lifetime of graduate school memories.

A gigantic thanks to my co-advisor David Baum for giving someone like me, who had no more than middle school education in biology, the

opportunity to research on the origins of life. David, thank you for patiently providing consistent boundary conditions through which I could organize my thoughts, and for teaching by demonstration how to navigate the busy academic life of teaching, mentoring, and research. Special thanks to my nominal advisor, Sridhara Dasu, for allowing me to pursue research outside the Physics department while being enrolled in it.

Finally, I would like to express my deepest gratitude to Eric Smith, who has played the role of my advisor, collaborator, critic, and friend these past few years. Eric, thank you for teaching me how to think, speak, and write clearly, through your long emails, weekly meetings, and detailed feedback on my manuscripts. While conversations with David provided me with the scientific questions that needed to be formalized, under your tutelage I learned how to research, formalize, develop, and present the relevant concepts. I am extremely thankful to both my advisors for respecting my creativity, never imposing anything, and for believing in me so that now I am able to stand on my own feet, albeit on their shoulders.

CONTENTS

Contents	v
List of Tables	x
List of Figures	xi
Abstract	xiv
1 Motivating question and the presentation of this thesis	1
1.1 <i>A question concerning the origins of biochemical life</i>	1
1.2 <i>Layout of the thesis</i>	5
1.2.1 Chapter 2	6
1.2.2 Chapter 3	7
1.2.3 Chapter 4	8
2 The geometry and combinatorics of an autocatalytic ecology in chemical and cluster chemical reaction networks	9
2.1 <i>Introduction</i>	9
2.2 <i>Chemical reaction networks and autocatalytic cycles</i>	12
2.2.1 Hypergraphs, stoichiometric matrix, conservation laws, and null flows	13
2.2.2 Formal, exclusive and stoichiometric autocatalysis, and autocatalytic cores	17

2.3	<i>Organizing an autocatalytic ecology</i>	23
2.3.1	Definitions	24
2.3.2	Mathematical results	32
2.3.3	Geometry and visualization	40
2.4	<i>Cluster chemical reaction networks (CCRN) framework</i>	46
2.4.1	Formalism	49
2.4.2	Complete CCRN	52
2.4.3	Rule generated CCRNs	63
2.5	<i>Discussion and future research</i>	66
3	Action Functional Gradient Descent (AFGD) algorithm for estimating escape paths in stochastic Chemical Reaction Networks	70
3.1	<i>Introduction</i>	70
3.2	<i>From master equation to Hamilton-Jacobi equation for CRN</i>	74
3.2.1	Hamilton-Jacobi equation and Non-Equilibrium Potential (NEP) for stochastic dynamics	75
3.2.2	Hamilton-Jacobi theory for stochastic Chemical Reaction Networks (CRN)	85
3.2.3	Switching dynamics for a multistable stochastic CRN	95
3.3	<i>Action Functional Gradient Descent (AFGD) algorithm</i>	97
3.3.1	Formulation as a MinMax problem	99
3.3.2	Algorithm and pseudocode	102
3.3.3	Performance costs and relation to other methods	111

3.4	<i>Application and results</i>	114
3.4.1	Selkov model	116
3.4.2	N-D Schlögl model	118
3.5	<i>Discussion</i>	127
4	Mathematics of evolving ecologies in stochastic chemical reaction networks and the origins of biochemical life	133
4.1	<i>A mathematical formalism for modeling evolving ecologies in chemical reaction networks</i>	133
4.1.1	The role of mathematics in our formalism and the physical assumptions leading to our basic set-up. . .	133
4.1.2	A topological definition of an ecology	138
4.1.3	A kinetic definition of evolution	143
4.2	<i>CRN models of evolving ecologies</i>	149
4.2.1	Procedure for ascertaining evolution	151
4.2.2	Open ended evolution in a one-dimensional CRN .	153
4.2.3	Directional evolution in polymer models in one monomer	160
4.2.4	Historically contingent evolution in polymer models in two monomers	168
4.3	<i>Discussion</i>	176
A	Appendix to Chapter 2	179
A.1	<i>Deficiency theory and multistability in CRNs</i>	179

A.2	<i>Kernel of the stoichiometric matrix for the complete 1-constituent CCRN</i>	181
A.3	<i>Deficiency-one algorithm for L=3 1-constituent CCRN</i>	183
B	Appendix to Chapter 3	186
B.1	<i>Experimental applications of AFGD</i>	186
B.2	<i>Rederiving previous results in the Hamilton-Jacobi formalism</i> . .	190
B.3	<i>Non-equilibrium action, action functional and its first and second variational derivatives</i>	196
B.3.1	Deriving stochastic and quantum dynamics from the non-equilibrium action (NEA)	196
B.3.2	Derivation of the path integral formula	198
B.3.3	Optimality condition: Hamilton's equations of motion and their relation to the Hamilton-Jacobi equation	202
B.3.4	Second variational derivative of action functional: Onsager-Machlup action and convexity of instanton	204
B.4	<i>Details of AFGD</i>	208
B.4.1	Notation	208
B.4.2	Initial condition	209
B.4.3	Lift curve to trajectory	212
B.4.4	Functional gradient - obtaining and filtering	215
B.4.5	Filtering and resampling routines	217
B.4.6	Pick step size	219

B.4.7	Updating cutoff frequency and increasing sample points during descent	220
B.4.8	Integrating Hamilton's equations of motion and convergence criteria	222
C	Appendix to Chapter 4	225
C.1	<i>Hamilton-Jacobi theory for stochastic population processes</i>	225
C.1.1	States, events, and probability	225
C.1.2	Event transition rates and probabilities	226
C.2	<i>Species interaction networks and their properties</i>	233
C.2.1	Chemical reaction network	233
C.2.2	Minimal null subnetwork	235
C.2.3	Siphon subnetwork	236
C.3	<i>Modeling a mechanism and obtaining quantitative predictions</i>	238
C.3.1	Stochastic chemical kinetics:	238
C.3.2	Stability, transitions, and cost of discrete attractors	240
	References	242

LIST OF TABLES

2.1	Data statistics for the 1-constituent CCRN.	54
2.2	Examples of sparse rule generated CCRNs.	64
3.1	Notation used for denoting paths in configuration and phase space.	98
4.1	Examples of chemical reaction networks in cosmology, astro- physics, and particle physics.	137
B.1	Correspondence between stochastic and quantum dynamics .	195

LIST OF FIGURES

1.1	The CRN of nuclear fusion in the Sun (Commons, 2015). . . .	2
1.2	Directional evolution in CRNs.	4
2.1	Ecologies in CRNs.	10
2.2	E-graph and flows on a CRN.	16
2.3	The list of MASs and their visualization for the $L = 4$ 1-constituent CCRN.	21
2.4	Food-waste-member-core and resource-member partitions. . .	26
2.5	An example of flow and species-productive cones.	28
2.6	The intersection data of partition-productive cones for the com- plete $L = 5$ 1-constituent CCRN	40
2.7	All MASs belonging to the $\{\{\bar{4}\}, \{\bar{1}\}, \{\{\bar{2}, \bar{3}, \bar{5}\}\}\}$ class in the com- plete $L = 5$ 1-constituent CCRN.	42
2.8	The pairwise intersection information for all the species-productive cones of the $L = 5$ 1-constituent CCRN.	44
2.9	Tree of polymers in two letters and its induced CCRN.	46
2.10	Number of MASs for the 1-constituent CCRN.	53
2.11	Partial CRN for the 1-constituent CCRN.	54
2.12	Stoichiometric subspace for the $L = 3$ complete 1-constituent CCRN.	57
2.13	Fixed point for $L = 3$ CCRN calculated using Feinberg's defi- ciency one algorithm.	59

3.1	Diagrammatic representations of the N-Schlögl models.	87
3.2	Diagrammatic proof of the Lyapunov property of the non-equilibrium potential with respect to mass action kinetics. . .	92
3.3	Escape paths and heteroclinic network for the 2-Schlögl model.	94
3.4	Comparing AFGD with Gillespie algorithm.	94
3.5	Lifting a curve to a trajectory.	102
3.6	Functional gradient descent without filtering.	102
3.7	The AFGD algorithm in a nutshell.	107
3.8	Selkov model diagram and proxy-heteroclinic network.	113
3.9	Comparing AFGD and the shooting method.	114
3.10	Descent progress for the Selkov model.	130
3.11	$H(p, q) = 0$ submanifold in the tangent and cotangent space for the 1-D Schlögl model.	131
3.12	Fixed points labelled by repelling directions (left) and proxy-heteroclinic network using relaxation trajectories (right) for the 2-Schlögl model.	131
3.13	Functional gradient and power spectrum for the 2-Schlögl model.	132
3.14	AFGD against integrated Hamilton's EoM.	132
4.1	1-species CRN admitting directed evolution.	153
4.2	Some polymerization reactions of polymers in a single monomer.	160
4.3	A polymer model in 1 monomer exhibiting directional evolution.	161

4.4	Polymer model in two monomers exhibiting historical contingency.	168
A.1	Representation of the complete 1-constituent CCRN with $L = 6$.	181
B.1	Evolving distribution by Δt	200
B.2	Initial conditions for the AFGD algorithm.	209
B.3	Phase space trajectory obtained by lifting the initial conditions.	213
B.4	Extracting functional gradient from phase space trajectory and smoothing it.	216
B.5	Descent progress and summary for the 2-Schlögl model for a trajectory with 500 points.	217
B.6	Descent progress and summary for the Selkov model for a trajectory with 2500 points.	221
B.7	Integrating Hamilton's equations of motion against AFGD algorithm output for the 6-D Schlögl model.	223
C.1	Extracting a large-deviations function from a probability distribution for a stochastic population process.	230

ABSTRACT

Modern physics and biology satisfactorily explain the passage from the Big Bang to the formation of Earth and the first cells to present-day life, respectively. However, the origins of biochemical life still remain an open question. Any answer to the question must explain how an evolving ecology of polymers of ever-increasing length could come about on a planet that otherwise consisted only of small molecules. In this work, we use chemical reaction network (CRN) theory to develop a formalism for defining a notion of ecology as well as different categories of evolution. We formalize an ecology as any collection of autocatalytic subnetworks (ASs). For a class of ASs, which we term *stoichiometric ASs*, we show that it is possible to classify them in equivalence classes and develop mathematical results about their behavior. We formalize evolution of a stochastic CRN (SCRN) as attractor-switching in a system with multiple attractors. To quantify evolution, we first derive the Hamilton-Jacobi theory underlying continuous-time Markov processes, and then use the construction to develop a variational algorithm for estimating switching dynamics for a generic SCRN that exhibits multiple attractors. We then summarize our mathematical formalism for identifying evolving ecologies in CRNs, and employ it to construct abstract polymer ecologies that can exhibit dynamics such that attractors in the polymer composition space with a higher average polymer length are also more probable. Furthermore,

we demonstrate a polymerization model in two monomers that exhibits historical contingency. This thesis is intended to provide a step forwards towards understanding the origins of biochemical life, and we conclude with a discussion of the limitations of our formalism and directions where future work is needed. While we employ theoretical techniques from mathematical physics, numerical optimization, polyhedral geometry, and CRN theory, we hope that this work finds practical applications with an inter-disciplinary audience including chemists, evolutionary biologists, optimal control theorists, and game theorists.

1 MOTIVATING QUESTION AND THE PRESENTATION OF THIS THESIS

1.1 A question concerning the origins of biochemical life

The mathematical foundations of modern physics provide a framework in which we can probe and comprehend the non-living mechanical entities of the universe (Aad et al., 2012; Gardner et al., 2006). Similarly, it is well established that mathematical, computational, and information-theoretic principles can have the same quantitative effectiveness in understanding and predicting the behavior of living entities (Eisenstein et al., 2021; Schneider, 2010; Dolson and Ofria, 2021). In particular, chemical reaction network (CRN) theory, the mathematical foundations of which were laid in the 1970s (Horn and Jackson, 1972; Feinberg, 2019; Yu and Craciun, 2018), has ubiquitous applications in understanding both abiotic systems, like nuclear fusion in stars (see Fig. 1.1), cellular biotic systems, like genetics and metabolism (Smith and Morowitz, 2016). In this work, we propose a rigorous characterization of *evolution* and *ecology* in CRN, and demonstrate its applications to the origins of biochemical life.

Our definition of evolution relies primarily on probability theory, which Jaynes termed ‘the logic of science’ (Jaynes, 2003). Broadly construed, given a mechanism of interactions and observational constraints, the role of




Figure 1.1: The CRN of nuclear fusion in the Sun (Commons, 2015).

a mathematical theory of probability in dynamical systems is to summarize the information of the dynamics in a probability distribution along the space of all possibilities. In particular, suppose that a stochastic system modeling a physical scenario is calculated to exhibit two steady states with unequal probabilities, and at an instance is observed to be in the less-probable steady state. Then it must be the case that, given sufficient time, as long as the two steady states are connected by a path, the stochastic system will discover the other more-probable steady state. Thus, using probability theory, we can predict a directionality to the dynamics of the simulation, as well as the physical system it simulates.

Probability theoretic considerations alone, however, are not sufficient to quantify evolution in living systems. For a start, as articulated by Oono in Ref. Oono (2012), a qualitative difference between *complex* living systems and their *simple* non-living counterparts is that complex systems have information in their internal state that is not reducible to their immediate boundary conditions, and thus relies on memory to be transmitted through time (for an information theoretic formulation, see Ref. Krakauer et al. (2020)). It is readily verified that the affordances of each biochemical entity, whether a protein, cell, or organism, are, at least partially, specified by its evolutionary history. This characterization is also cognizant with Louis Pasteur's maxim, '*omne vivum ex vivo*' or that 'all life is from life', which posits that life lacks spontaneous-organizing capability and only can only come from other complex systems.

In the last century, it was discovered that life has a genetic basis (Watson and Crick, 1953) and its memory is stored *biochemically* as an interacting ecology of autocatalytic and inhibitory polymers. A mathematical viewpoint of understanding the genetic basis of life as an ecology of autocatalytic polymers was formulated by Eigen and Schuster in their trilogy (Eigen and Schuster, 1977, 1978a,b). Since then, there has been a proliferation of mathematical models, such as (Kauffman and Levin, 1987; Nowak and Ohtsuki, 2008; Hordijk and Steel, 2018), that attempt to elucidate the mathematical structure underlying autocatalytic polymer dynamics and the transition from a prebiotic to a biotic world. However, it must be

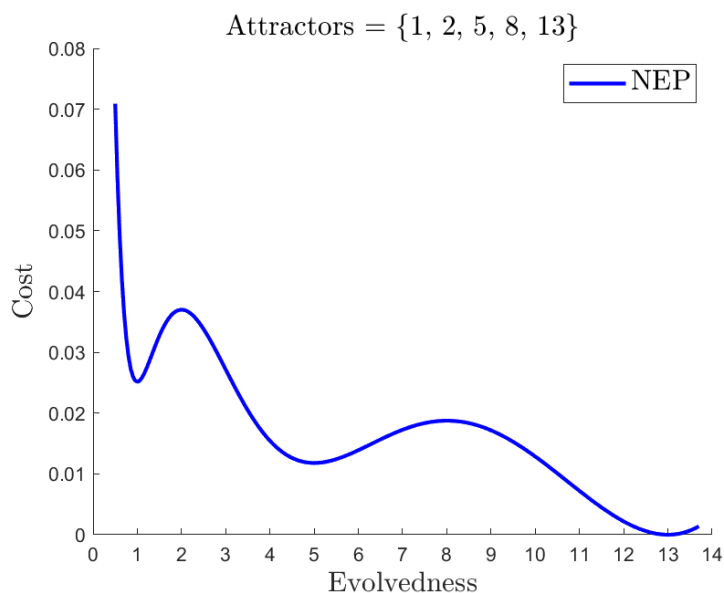


Figure 1.2: A 1-dimensional CRN whose attractors are such that their occupation is more probable as their distance increases from the origin. If the distance from the origin is seen as a measure of *evolvedness* and their log-improbability as a measure of *cost*, then the more evolved attractors are less costly. Such a CRN will be said to exhibit *directional evolution* from any but the most evolved attractor.

noted, while understanding the emergence of autocatalytic long polymers is necessary to understand the non-life to life transition, it is certainly not sufficient. To explain the sustenance of such an ecology, there have also been considerable efforts to understand the emergence of spatial compartments using small molecule chemistry, both phenomenologically (Ruiz-Mirazo et al., 2017; Higgs, 2021) and mathematically (Segré et al., 1998).

The question concerning the origins of biochemical life that we raise in this work is: How can an ecology of autocatalytic polymers emerge in

an environment with otherwise small molecules and directionally evolve towards more stable states that include a greater abundance of long polymers? Are there ecologies that have the potential to exhibit open-ended evolution, or demonstrate some form of historical contingency? While we do not argue for the physical existence of such systems, mathematically, we demonstrate several classes of CRNs that can satisfy some or all of the above properties. A simple example 1-dimensional CRN that exhibits directional evolution, explained further in Sec. 4.2, can be seen in Fig. 1.2.

1.2 Layout of the thesis

This thesis is a compilation of three manuscripts:

- The geometry and combinatorics of an autocatalytic ecology in chemical and cluster chemical reaction networks (*preprint* available at the time of thesis submission, (Gagrani et al., 2023a)).
- Action Functional Gradient Descent algorithm for estimating escape paths in stochastic chemical reaction networks (*published* in Physical Review E, (Gagrani and Smith, 2023)).
- The mathematics of evolving ecologies in chemical reaction networks and the origins of biochemical life (*under preparation*, (Gagrani et al., 2023b)).

Two of the three manuscripts were prepared in collaboration with co-authors. However, I have omitted all the parts contributed by the collaborators and, where necessary, have referred to the manuscript in the main text.

1.2.1 Chapter 2

In Sec. 2.2, we briefly review chemical reaction networks (CRNs) and a few different notions of autocatalysis. In Sec. 2.3, we develop a formalism for analyzing minimal autocatalytic subnetworks (MASs) that are detectable from the stoichiometric matrix. In the manuscript Ref. Gagrani et al. (2023a), we provide a linear-programming algorithm to exhaustively enumerate them. We also propose a visualization scheme in which to represent MAS combinatorics and geometry. In Sec. 2.4, we present the cluster chemical reaction network (CCRN) framework and argue that it provides a natural coarse-graining of realistic chemical and biochemical reaction networks. We then explore the combinatorics of MASs in CCRN framework and provide a worked-out example to demonstrate the geometry of autocatalysis in the stoichiometric subspace of species concentration. Multiple examples from the CCRN framework are also spread throughout this work. We conclude, in Sec. 2.5, with an overview of our contributions and possible avenues for future investigation.

1.2.2 Chapter 3

In Section 3.2, we derive the Hamilton-Jacobi theory for stochastic processes starting from a master equation and apply it to stochastic CRN. We show that the Hamilton-Jacobi theory of the non-equilibrium potential (NEP) arises naturally as a result of a variational principle applied to the *action functional*. In Appendix B.1, we investigate the relevance our work might have to a stochastic modelling practitioner by posing a general practical problem, giving an overall picture of the solution and explaining where our algorithm fits in. In Appendix B.2, we employ the Hamilton-Jacobi formalism to recover the ACK theorem and NEP for complex-balanced systems. In Appendix B.3, we define a ‘non-equilibrium action’, of which both the master equation and Schrödinger equation can be seen as a variational solution, derive the path integral formula and action functional for stochastic population processes, and calculate the first and second variational derivatives of the action functional. In Section 3.3, we propose an *action functional gradient descent* (AFGD) algorithm that finds the variational solution in the space of paths constrained at end points for a given Hamiltonian value constraint. We also explain how the algorithm can be used to find least-improbable escape paths out of a stable attractor and assign a value to the NEP along them. The details of the implementation are provided in Appendix B.4, and a MATLAB implementation is made available at (Gagrani, 2022). In Section 3.4, we demonstrate the applications of the AFGD algorithm on several CRNs. We first consider the

Selkov model, and compare the result of the algorithm against the escape trajectory found by the ‘shooting-method’ (Figure 3.9). We then define a class of high dimensional birth-death models, namely ‘N-Schlögl model’, and compare the results of the algorithm against a stochastic simulation for the 2-Schlögl model (Figure 3.4). We also use the algorithm on the six dimensional 6-Schlögl model, and compare the result against the integration of Hamilton’s equations of motion (Figure B.7). Finally, in Section 3.5, we conclude with a discussion of our contribution and potential avenues of future research.

1.2.3 Chapter 4

In Sec. 4.1, we explain the mathematical formalism for modeling evolving ecologies in CRNs. We rigorously define the concepts species composition, ecology, attractor, evolvedness, and cost using CRN theory and Hamilton-Jacobi theory for stochastic processes developed in the previous chapters. The mathematical preliminaries have been presented in appendices C.1, C.2, and C.3, where the passage from probability theory to the dynamics of stochastic CRNs is detailed. In Sec. 4.2, we employ our formalism on a one-dimensional CRN, and polymer models in one and two monomers, and demonstrate that they can evolve. Finally, in Sec. 4.3, we provide a summary of our contributions, acknowledge the limitations of our formalism in understanding the origins of biochemical life, and discuss directions for future research.

2 THE GEOMETRY AND COMBINATORICS OF AN AUTOCATALYTIC ECOLOGY IN CHEMICAL AND CLUSTER CHEMICAL REACTION NETWORKS

2.1 Introduction

Chemical reaction network (CRN) theory offers a versatile mathematical framework in which to model complex systems, ranging from biochemistry and game theory to the origins of life (Veloz et al., 2014; Smith and Morowitz, 2016). The usefulness of CRNs in modelling these phenomena stems from its ability to exhibit a wide range of nonlinear dynamics and it is widely recognized that *autocatalysis* can be seen as a basis for many of them (Epstein and Pojman, 1998; Schuster, 2019). Broadly, autocatalysis is framed as the ability of a given chemical species to make more copies of itself or otherwise promote its own production. Thus, from the perspective of kinetics, one would expect autocatalysts to be able to show super-linear growth (Hordijk and Steel, 2004; Andersen et al.). For some examples of nonlinear kinetics obtained by composing two autocatalytic cycles and their relevance to ecological dynamics, see Fig. 2.1.

A comprehensive understanding of autocatalysis in CRNs has great practical values, but its full mathematical treatment remains to be developed. In (Andersen et al.), it is conjectured that it is impossible for a CRN to show a temporary speed-up of the reaction before settling down to reach

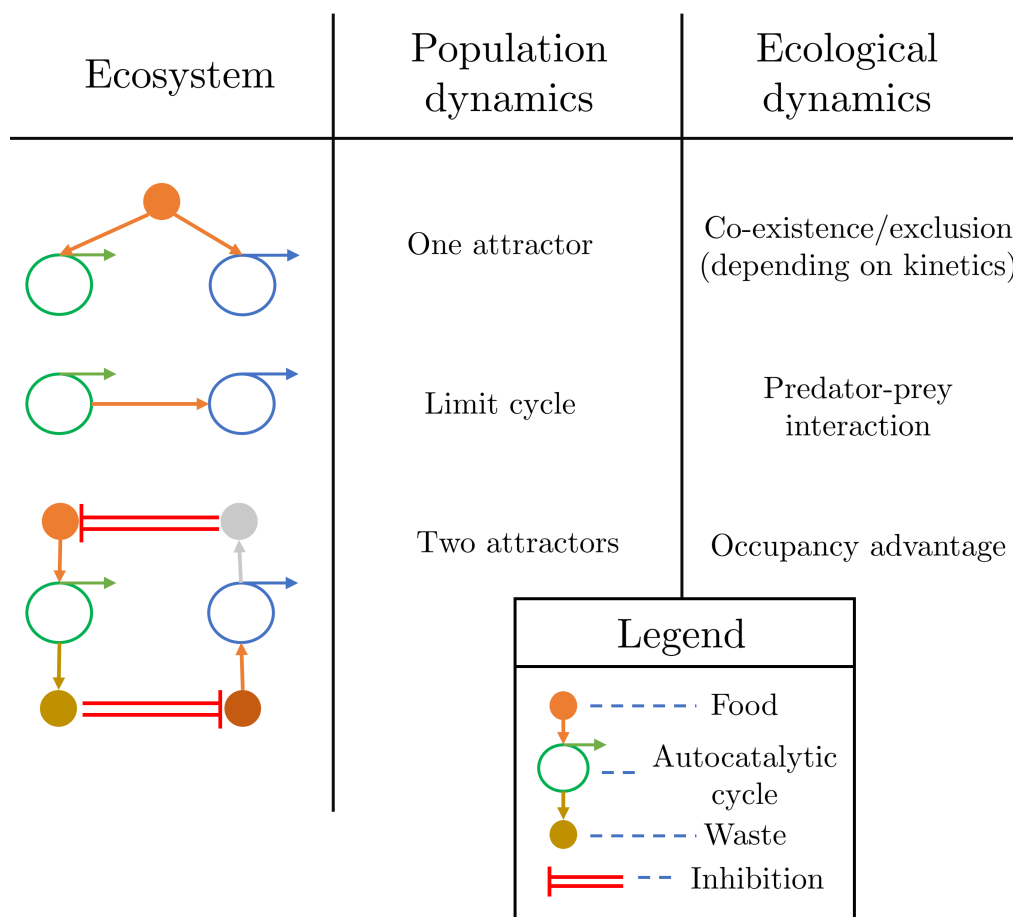


Figure 2.1: Ecologies obtained by composing autocatalytic cycles. The member species of each autocatalytic cycle are shown to consume food and produce waste resources. Two cycles can be composed in several ways such as by sharing food (line one), the member species of one can be the food of another (line two), or the waste of one can deplete the food of another (line three). For a detailed account of such interactions and their ecological counterpart, see Peng et al. (2020).

an equilibrium if, at least, it is not *formally* autocatalytic. Also, it is suggested in (Deshpande and Gopalkrishnan, 2014) that understanding the interactions between autocatalytic and drainable subnetworks might provide a basis for obtaining a proof for the long-standing *persistence conjecture* (Feinberg, 1989).

In recent work (Blokhuys et al., 2020), Blokhuys et al. defined a notion of autocatalysis strictly in terms of the stoichiometric matrix without referring to the underlying CRN. While their notion is more restrictive than the notion of autocatalysis discussed in (Andersen et al.; Deshpande and Gopalkrishnan, 2014), it permits richer mathematical analysis. Using their stoichiometric criterion for autocatalysis, Blokhuys et al. show that there are only five types of minimal autocatalytic motifs. They also found that autocatalysis is more abundant than previously thought, and simple reaction networks with a few reactions can contain very many minimal autocatalytic motifs in them.

In this chapter, we extend the definition of Blokhuys et al. to subnetworks, and provide tools for exhaustively enumerating autocatalytic subnetworks and understanding their combinatorics. The layout of the chapter is as follows. In Sec. 2.2, we briefly review chemical reaction networks (CRNs) and a few different notions of autocatalysis. In Sec. 2.3, we develop a formalism for analyzing minimal autocatalytic subnetworks (MASs) that are detectable from the stoichiometric matrix. In the manuscript Ref. Gagrani et al. (2023a), we provide a linear-programming algorithm to

exhaustively enumerate them. We also propose a visualization scheme in which to represent MAS combinatorics and geometry. In Sec. 2.4, we present the cluster chemical reaction network (CCRN) framework and argue that it provides a natural coarse-graining of realistic chemical and biochemical reaction networks. We then explore the combinatorics of MASs in CCRN framework and provide a worked-out example to demonstrate the geometry of autocatalysis in the stoichiometric subspace of species concentration. Multiple examples from the CCRN framework are also spread throughout this work. We conclude, in Sec. 2.5, with an overview of our contributions and possible avenues for future investigation.

2.2 Chemical reaction networks and autocatalytic cycles

The mathematical theory of chemical reaction networks (CRNs), pioneered by Horn, Jackson, and Feinberg (Horn and Jackson, 1972; Feinberg, 2019; Yu and Craciun, 2018), has ubiquitous applications in understanding nature. In particular, for biological and ecological applications, the formalism allows for a notion of self-replication or *autocatalysis* where certain species increase the population of these same species. Consequently, there is value in rigorously understanding the organization of a collection (*ecology*) of autocatalytic cycles.

In (Feinberg, 2019), Feinberg defines a CRN as a triple of the set of species, complexes, and reactions $(\mathcal{S}, \mathcal{C}, \mathcal{R})$, which we review in Appendix A.1. It has been argued in (Deshpande and Gopalkrishnan, 2014), however, that the data of the set \mathcal{C} of complexes is redundant as it can be recovered from the reaction set \mathcal{R} and, thus, a CRN can be defined more economically as $(\mathcal{S}, \mathcal{R})$. Since the two representations are equivalent, we will pick either, dictated by convenience. The concept of autocatalysis, dating back to the late 1800s (Ostwald, 1890; Peng et al., 2022b), has several distinct formulations in CRN theory, for a review of which we refer the readers to Andersen et al. (Andersen et al.). In our work, we closely follow the formulation of Blokhuis et al. (Blokhuis et al., 2020), while acknowledging its relationship to other formulations of autocatalysis, as shown in Sec. 2.2.2.

2.2.1 Hypergraphs, stoichiometric matrix, conservation laws, and null flows

A *chemical reaction network* (CRN) is defined by the pair $(\mathcal{S}, \mathcal{R})$ where \mathcal{S} and \mathcal{R} denote the finite sets of species and reactions, respectively.

$$\mathcal{S} = \{s_1, \dots, s_S\}$$

$$\mathcal{R} = \{\mathbf{y} \longrightarrow \mathbf{y}'\}$$

The *reaction hypergraph* of the CRN $(\mathcal{S}, \mathcal{R})$ is the directed graph with *complex* set $\mathcal{C} := \{\mathbf{y}, \mathbf{y}' | \mathbf{y} \rightarrow \mathbf{y}' \in \mathcal{R}\}$ and edge set \mathcal{R} . The complexes, are multisets of *species* (vertices), and are denoted by a column vector \mathbf{y} representing the stoichiometry or multiplicity of each species in the complex. Note that this is purely topological data, and no reference is made to the kinetics or dynamics of a system specified by the CRN.

Let (hyper)graph $\mathcal{G} = (\mathcal{S}, \mathcal{R})$ be a CRN. Then the *stoichiometric subspace* \mathcal{S} is defined as the linear span of the reaction vectors $\{\mathbf{y}' - \mathbf{y} | \mathbf{y} \rightarrow \mathbf{y}' \in \mathcal{R}\}$. The *stoichiometric matrix* \mathbb{S} is defined as the matrix whose columns are the reaction vectors,

$$\text{cols}(\mathbb{S}) = \{\mathbf{y}' - \mathbf{y} | \mathbf{y} \rightarrow \mathbf{y}' \in \mathcal{R}\}. \quad (2.1)$$

We denote the restriction of the rows of the stoichiometric matrix \mathbb{S} to a species set $\mathcal{S}' \subset \mathcal{S}$ by $\mathbb{S}_{\mathcal{S}'}$, and the restriction of the columns of \mathbb{S} to a reaction set $\mathcal{R}' \subset \mathcal{R}$ by $\mathbb{S}^{\mathcal{R}'}$. In general, for a subgraph $\mathcal{G}' = (\mathcal{S}', \mathcal{R}') \subset \mathcal{G}$, we denote the restriction of \mathbb{S} to \mathcal{G}' by $\mathbb{S}|_{\mathcal{G}'} := \mathbb{S}_{\mathcal{S}'}^{\mathcal{R}'}$.

We define a **subnetwork** of a graph \mathcal{G} to be the subgraph $(\mathcal{S}|_{\mathcal{R}'}, \mathcal{R}')$ where $\mathcal{R}' \subset \mathcal{R}$ and there is no omission of species in the subnetwork that participates in the set of reactions \mathcal{R} of the original graph \mathcal{G} .


Let vector $\mathbf{n} = [n_1, \dots, n_s]^T \in \mathbb{Z}_{\geq 0}^{\mathcal{S}}$ ($\in \mathbb{R}_{\geq 0}^{\mathcal{S}}$) denote the population (concentration) of each species in the system, and $\mathbf{v} = [v_1, \dots, v_{\mathcal{R}}]^T$ denote a (hyper)flow on \mathcal{G} . Then the change of species population (concentration)

$\Delta \mathbf{n}$ resulting from the flow $\mathbf{v} \in \mathbb{Z}_{\geq 0}^{\mathcal{R}}$ ($\mathbf{v} \in \mathbb{R}_{\geq 0}^{\mathcal{R}}$) on the graph \mathcal{G} is given by

$$\Delta \mathbf{n} = \mathbb{S} \cdot \mathbf{v}.$$

The span of the columns of the stoichiometric matrix is termed as the *stoichiometric subspace*, and is the subspace of the change in species concentration under an arbitrary flow on \mathcal{G} .

A flow on a graph can also be used to create a linear combination of the reactions, also called a composite reaction or \mathcal{G} -dilution in the terminology of (Andersen et al.) or (Deshpande and Gopalkrishnan, 2014; Cardelli, 2011), respectively. For example, a flow \mathbf{v} on graph \mathcal{G} would result in the composite reaction



The left null space of \mathbb{S} , i.e. vectors \mathbf{x} such that $\mathbf{x} \cdot \mathbb{S} = 0$, correspond to *conservation laws*

$$\mathbf{x} \cdot \Delta \mathbf{n} = 0.$$

Equivalently, the inner product of the conservation law with the population $\mathbf{x} \cdot \mathbf{n}$ is an invariant for the network under an arbitrary flow. The vector \mathbf{x} will be called a *positive conservation law* if all its coefficients are nonnegative.

Similarly, the right null space of \mathbb{S} , i.e. vectors \mathbf{w} such that $\mathbb{S} \cdot \mathbf{w} = 0$,






Figure 2.2: In the left panel, the Euclidean embedded graph (E-graph, see Craciun (2019)) of the CRN in Example 2.1 is shown. The E-graph is obtained by representing the complexes as lattice points in a Euclidean lattice and representing the reactions as edges between them. The flows on the reactions are labelled by κ and the change in concentration they induce, $\Delta \mathbf{n} = \mathbb{S}\kappa$, is schematically represented in the right panel.

correspond to *null flows* which leave the population (concentration) of the species unchanged ($\Delta \mathbf{n} = 0$). A careful accounting of the right null space of \mathbb{S} provides a topological classification of a CRN, which we review in Appendix A.1.

Example 2.1. Consider the CRN $\mathcal{G} = (\mathcal{S}, \mathcal{R})$ given as



For a visual representation of this example, see Fig. 2.2. Then the species set $\mathcal{S} = \{A, B\}$ and the reaction set $\mathcal{R} = \{A + B \rightarrow 2A, 2A \rightarrow A + B, A + B \rightarrow 2B\}$. As

can be read from \mathcal{R} , the complex set $\mathcal{C} = \{A + B, 2A, 2B\}$, and their corresponding vectors are

$$y_{A+B} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, y_{2A} = \begin{bmatrix} 2 \\ 0 \end{bmatrix}, y_{2B} = \begin{bmatrix} 0 \\ 2 \end{bmatrix}.$$

The stoichiometric matrix \mathbb{S} for \mathcal{G} is


$$\mathbb{S} = \begin{bmatrix} 1 & -1 & -1 \\ -1 & 1 & 1 \end{bmatrix},$$

with left null space (conservation laws) spanned by $x = [1, 1]$ and right null space (null flows) spanned by the basis $\{[1, 1, 0]^T, [1, 0, 1]^T\}$. Notice that the stoichiometric subspace is one-dimensional and perpendicular to the conservation law.

2.2.2 Formal, exclusive and stoichiometric autocatalysis, and autocatalytic cores


Formal autocatalysis: Following the formulation in (King, 1978), Andersen et al. in (Andersen et al.), define a CRN \mathcal{G} to be *formally autocatalytic* in the set \mathcal{M} if:

1. there exists a flow on \mathcal{G} such that the resulting composite reaction is of the form



where $\mathbf{0} \prec \mathbf{m} \prec \mathbf{n}^1$. Here \mathbf{m} and \mathbf{n} are the stoichiometries of the set \mathcal{M} in the input and output of the composite reaction, respectively, and $\mathbf{o}\mathcal{M} = \sum_i \mathbf{o}_i \mathcal{M}_i$.

Example 2.2. Consider the reaction network $\mathcal{G}_r = (\{A, B\}, \{A \rightarrow B, A + B \rightarrow B\})$ under the flow $\mathbf{v} = [1, 1]^T$. The resulting composite reaction is



Thus the reaction network is formally autocatalytic for the set $\mathcal{M} = \{B\}$.

Exclusive autocatalysis: In (Andersen et al.), Andersen et al. define a CRN \mathcal{G} to be *exclusively autocatalytic* in the set \mathcal{M} if:

1. there exists a flow on \mathcal{G} such that the population of each species in the set \mathcal{M} strictly increases.
2. for every reaction $y \rightarrow y' \in \mathcal{G}$, the input complex y and the output complex y' contain at least one species from the set \mathcal{M} . This ensures that the flow is inadmissible, or there is no flow, if the population of any species in the set \mathcal{M} is zero.

While the first condition is identical, it is the second condition that makes exclusive autocatalysis a more restrictive notion than formal autocatalysis.

¹ $A \prec B$, read as *A precedes B*, if all entries of $A - B$ are strictly negative. Analogously, $A \succ B$, read as *A succeeds B*, if all entries of $A - B$ are strictly positive.

Example 2.3. The CRN \mathcal{G}_f in Example 2.2 is not exclusively autocatalytic since the flow proceeds even if the population of B is 0. That is because the reaction $A \rightarrow B$ does not require the presence of B to proceed.

However, in Example 2.1, \mathcal{G} is exclusively autocatalytic in the set $\mathcal{M} = \{A\}$ since:

1. the flow $\mathbf{v} = [1, 0, 0]^T$ increases the population of A, n_A , by 1 if $n_A > 0$.
2. the reaction $A + B \rightarrow 2A$ requires the presence of at least one A and cannot occur if $n_A = 0$.

Similarly, one can show that \mathcal{G} is also exclusively autocatalytic in the set $\mathcal{M} = \{B\}$.

Note that this definition is consistent with an intuitive notion of autocatalysis and is cognate with those of prior works (Deshpande and Gopalkrishnan, 2014; Gopalkrishnan, 2011; Barenholz et al., 2017).

Stoichiometric autocatalysis: In (Blokhuys et al., 2020), Blokhuys et al. generalize the notion of autocatalysis for subgraphs rather than complete networks. They define a *motif* to be any subgraph \mathcal{G}' of a CRN, or equivalently any possible submatrix $\mathbb{S}|_{\mathcal{G}'}$ of the stoichiometric matrix \mathbb{S} . Then, an *autocatalytic motif* is defined to be a subgraph $\mathcal{A} = (\mathcal{A}_S, \mathcal{A}_R)$ such that:

1. There exists a nonnegative flow \mathbf{g} on \mathcal{A} such that the change in concentration of each species in the set \mathcal{A}_S is strictly positive or $\Delta \mathbf{n}_{\mathcal{A}_S} = \bar{\mathbb{S}} \cdot \mathbf{g} \succ 0$.

2. For every reaction $y \rightarrow y' \in \mathcal{A}_R$, the complexes y and y' contain at least one **distinct** species from the set \mathcal{A}_S . Equivalently, each column in $\bar{\mathbb{S}} := \mathbb{S}|_{\mathcal{A}}$ has at least one positive and one negative coefficient.

Notice that the first condition is identical to that of *exclusive autocatalysis*. However, the second condition, also termed *autonomy* of submatrix $\bar{\mathbb{S}}$, ensures that the production of any species *consumes* at least one other species of the motif, disallowing unconditional growth. The condition that there needs to be at least one autocatalytic species in the reactants distinct from the products makes this definition more restrictive than exclusive autocatalysis, due to which we term their criterion as **stoichiometric autocatalysis**. A useful feature of stoichiometric autocatalysis is that it can be inferred from the stoichiometric matrix without referring to the underlying CRN. In the rest of our work we will adhere to this criterion and will use *autocatalysis* to mean stoichiometric autocatalysis, unless specified otherwise.

Following the terminology in (Deshpande and Gopalkrishnan, 2014), if there exists a nonnegative flow such that $\bar{\mathbb{S}} \cdot \mathbf{g} \prec 0$, we would term the subgraph as a *drainable motif*, where the change in concentration of all the species in the subset is strictly negative. While we focus our attention on autocatalytic motifs, the results in our chapter are easily extended to their drainable counterparts. In (Deshpande and Gopalkrishnan, 2014), the authors prove that a network is persistent if it has no drainable motifs and also remark on the importance of autocatalytic and drainable motifs for

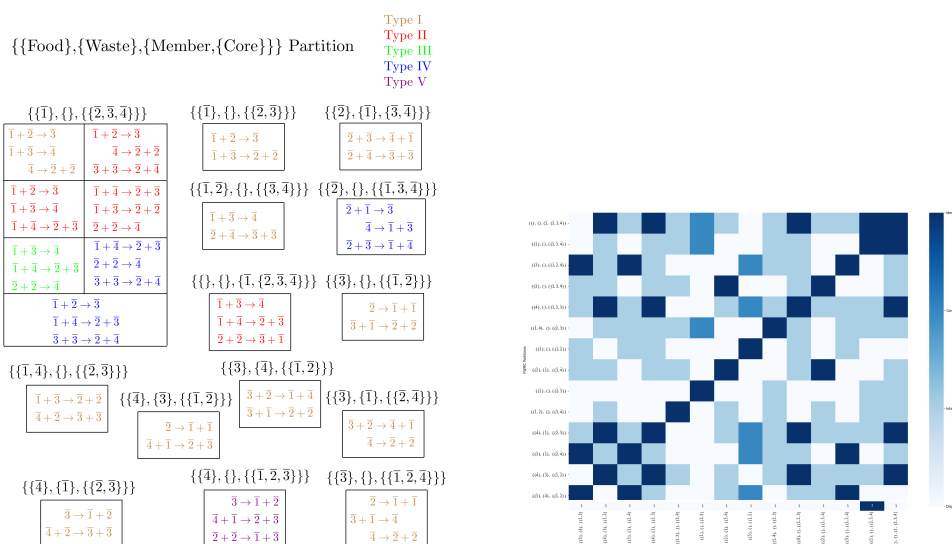


Figure 2.3: In the left panel, all minimal autocatalytic subnetworks (MASs) for the complete 1-constituent CCRN with $L = 4$ are shown. The core types are also color coded using the typology of Blokhuis et al. in (Blokhuis et al., 2020). In a CCRN, each species is denoted by a vector of positive integers with an overbar, and each reaction satisfies certain conservation laws. For a definition of a CCRN, see Sec. 2.4. The list of reactions can be found in Ref. Gagrani et al. (2023a). In the right panel, the list of the FWMC classes with the information of their intersection is shown.

understanding the dynamics of a network in ecological terms.

Autocatalytic core: An *autocatalytic core* is a minimal autocatalytic motif that does not contain a smaller autocatalytic motif. In Ref. Blokhuis et al. (2020), it is shown that if $\mathcal{A} = (\mathcal{A}_C, \mathcal{A}_R)$ is an autocatalytic core, it has the following properties:

1. $\bar{\mathcal{S}}$ is a square and invertible matrix, and the number of reactions $|\mathcal{A}_R|$ must be equal to the number of autocatalytic species in the core $|\mathcal{A}_C|$.

2. Every forward reaction in \mathcal{A}_R contains only one species from the core set \mathcal{A}_C (possibly with stoichiometry higher than one).

Furthermore, the authors provide a geometric characterization of the autocatalytic cores and show that they can only be of five types. We label the autocatalytic subnetworks in Fig. 2.3 using their taxonomy, and for details we refer the readers to (Blokhuys et al., 2020).

Example 2.4. Recall from Example 2.1, the stoichiometric matrix was given by

$$\mathbb{S} = \begin{bmatrix} 1 & -1 & -1 \\ -1 & 1 & 1 \end{bmatrix}.$$

Notice that there is no autocatalytic core in this network, and thus a naive characterization by the Blokhuys et al. criteria would regard the network to be non-autocatalytic. On the other hand, from Example 2.3, we know that the network is exclusively autocatalytic. To remedy this, we modify the network by replacing each reaction whose reactant and product complex shares a species with two new reactions with a fictitious distinct intermediate species. This then yields a new network \mathcal{G}^ given by*



where C, D are the newly added fictitious species. The stoichiometric matrix for

\mathcal{G}^* is given by

$$\mathbb{S}^* = \begin{bmatrix} 2 & -2 & -1 & 1 & -1 & 0 \\ 0 & 0 & -1 & 1 & -1 & 2 \\ -1 & 1 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & -1 \end{bmatrix}.$$

Notice that the restriction of the stoichiometric matrix to the species set $\mathcal{S}' = \{A, C\}$ and the reactions set $\mathcal{R}' = \{A + B \rightarrow C, C \rightarrow 2A\}$, denoted by $\mathbb{S}_{\mathcal{S}'}^{*\mathcal{R}'}$, indeed satisfies the properties for an autocatalytic core, and thus \mathcal{G}^* is indeed autocatalytic with core species $\{A, C\}$. A similar construction shows that the network is also autocatalytic in the set $\{B, D\}$. (In the typology of Blokhuis et al. both of these cores are of type I.)

2.3 Organizing an autocatalytic ecology

Even mildly complicated CRNs can exhibit an abundance of autocatalytic cores (Personal conversations with Philippe Nghe; Peng et al., 2022a). Recall from Sec. 2.3.1 that we define a *subnetwork* of a graph $\mathcal{G} = (\mathcal{S}, \mathcal{R})$ to be a subgraph $(\mathcal{S}|_{\mathcal{R}'}, \mathcal{R}')$ where $\mathcal{R}' \subset \mathcal{R}$ and there is no omission of species in the subnetwork that participate in \mathcal{R} in the original graph \mathcal{G} . Then, we define a *minimal autocatalytic subnetwork* (MAS) to be a subnetwork that contains exactly one autocatalytic core. In Sec. 2.3.1, we demonstrate first, these subnetworks organize within equivalence classes, and, second, that the subnetworks within an equivalence class can also be organized by

analyzing, what we call their *species-productive cones*. We prove some results about autocatalytic subnetworks in Sec. 2.3.2, and provide algorithms to exhaustively enumerate the MASs and to identify their species-productive cones in Ref. Gagrani et al. (2023a). Finally, we discuss the polyhedral geometry of autocatalytic ecologies and develop a visualization scheme in Sec. 2.3.3. It must be noted that, to simplify construction, we have assumed that the graph \mathcal{G} is reversible throughout this section, i.e. for every reaction $y_1 \rightarrow y_2 \in \mathcal{G}$, the reverse reaction $y_2 \rightarrow y_1$ is also in the reaction set of \mathcal{G} .

2.3.1 Definitions

Let $\mathcal{G} = (\mathcal{S}, \mathcal{R})$ be a CRN. Let $A = (\mathcal{S}_A, \mathcal{A}_R)$ be a minimal autocatalytic subnetwork (MAS) containing exactly one core, where \mathcal{S}_A is the set of species with non-zero stoichiometry in the reaction set $\mathcal{A}_R \subseteq \mathcal{R}$. From property 1 of an autocatalytic core, there must be a subset $\mathcal{A}_C \subseteq \mathcal{S}_A$ containing the core species that has the same size as \mathcal{A}_R . In general, however, there will also be disjoint subsets of \mathcal{S}_A of species that only occur as co-reactants, co-products, and both co-reactants and co-products in \mathcal{A}_R . Correspondingly, if we consider the restriction of the stoichiometric matrix \mathbb{S} to the reaction set \mathcal{A}_R , denoted by $\bar{\mathbb{S}} := \mathbb{S}|_{\mathcal{A}_R}$, for the MAS A we define:

1. the species set corresponding to all rows of $\bar{\mathbb{S}}$ with nonpositive coefficients as the *food set* and denote it by A_F .
2. the species set corresponding to all rows of $\bar{\mathbb{S}}$ with nonnegative

coefficients as the *waste set* and denote it by A_W .

3. the species set corresponding to all rows of \bar{S} with both positive and negative coefficients as the *member set* and denote it by A_M .
4. the species set corresponding to all rows of autocatalytic core as the *core member set* and denote it by A_C .

We denote the exclusion of a subset K from the species set by $A_{/K}$. For example, all species in the species set of an autocatalytic subnetwork A except the core member species will be denoted by $A_{/C}$. Note that the core member set, or simply *core set*, is a subset of the member set, or $A_C \subseteq A_M$. We refer to the species in the member set that are not in the core set as *non-core member species* and denote their set as $A_{M/C}$. For an example of the partitioning process on a reaction network where each set is non-empty, see Fig. 2.4.

Food-waste-member-core partition of an autocatalytic subnetwork

Henceforth we will refer to the above partition of the species set for a MAS as the *food-waste-member-core* (FWMC) partition of the MAS. Moreover, we will refer to the submatrices generated by restricting to the food, waste, member, and core species, as \bar{S}_F , \bar{S}_W , \bar{S}_M , and \bar{S}_C , respectively. Finally, notice that MASs with identical FWMC partitions form an equivalence relation. Thus we will call MASs A and B to be equivalent if they have the

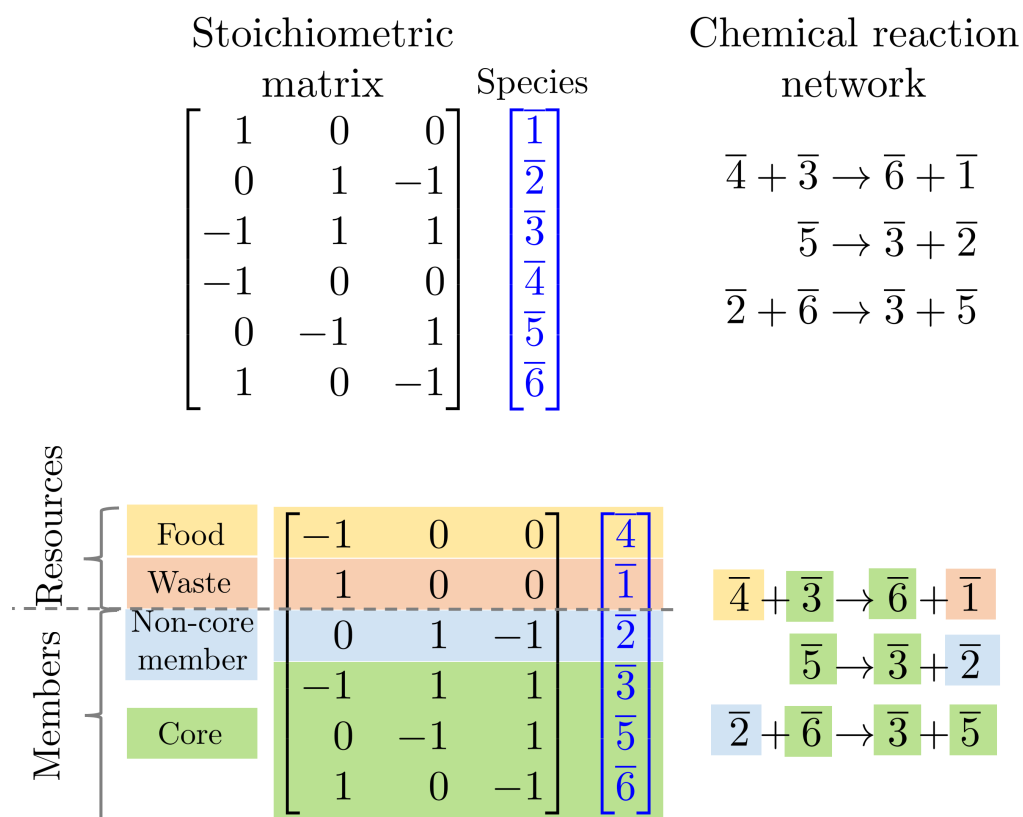


Figure 2.4: Food-waste-member-core and resource-member partition of a minimal autocatalytic subnetwork in the 1-constituent Cluster CRN (CCRN) with $L = 6$. Using the typology of Blokhuis et al. this subnetwork has a type two core.

same FWMC partition,

$$A \equiv B \iff \text{FWMC}(A) = \text{FWMC}(B).$$

Our *quadrartite* partition of the species set can be coarse-grained to a bipartite *resource-member* partition where the *resource* set is the union of

food and waste sets (and member set is the same as above) (see Fig. 2.4). Our partitioning of the species set by examining the coefficients of the stoichiometric matrix is cognate with the partitioning done by Avanzini et al. in (Avanzini et al., 2022), however, since their interest is in null cycles, they partition the set of species into a *resource-member* partition. It is also aligned with, and is in fact a refinement of, the *food-waste-member* partitioning by Peng et al. in (Peng et al., 2020), where it is shown that different compositions of distinct autocatalytic cycles can lead to various ecological dynamics (see Fig. 2.1).

Example 2.5. *In the network \mathcal{G}^* from Example 2.4, the MASs are:*

1. $A_1 = (\{A, B, C\}, \{A+B \rightarrow C, C \rightarrow 2A\})$ with FWMC partition $\{\{B\}, \{\}, \{\{A, C\}\}\}$.
2. $A_2 = (\{A, B, D\}, \{A+B \rightarrow D, D \rightarrow 2B\})$ with FWMC partition $\{\{A\}, \{\}, \{\{B, D\}\}\}$.

Notice that both these subnetworks have empty waste and non-core member species sets.

Flow, species and partition productive cones

Flow-productive cone: Recall from property 1 of an autocatalytic core that $\bar{\mathbb{S}}_C$ is invertible. The chemical interpretation of the inverse is that the k^{th} column is a flow vector that increases the k^{th} species by exactly one unit, making it an elementary mode of production. We define the *flow-productive cone*, denoted by \mathcal{F} , as the cone generated by the elementary

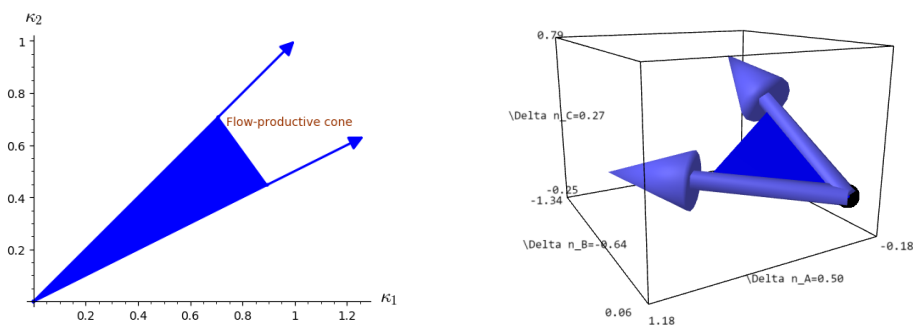


Figure 2.5: The flow and species productive cones for the Example 2.6 are shown in the left and right panel, respectively. The dimensions of the embedding space of the flow and species productive cones are the number of reactions and species in the minimal autocatalytic subnetwork, respectively. Notice that the food species B is strictly consumed (nonpositive) and the core species A, C are strictly produced (nonnegative) in the species-productive region. For this example, the partition-productive cone is identical to the species-productive cone.

modes of production of a minimal autocatalytic subnetwork (MAS),

$$\mathcal{F}(A) = \text{cone}(\text{cols}(\bar{\mathbb{S}}_C^{-1})). \quad (2.2)$$

By definition, an element in the interior of the flow-productive cone, the *flow productive region*, of A corresponds to a flow vector for which A is productive, i.e. the core member species of A are strictly produced. It should be noted that the inverse of $\bar{\mathbb{S}}_C$ might contain negative coefficients. Since we are only considering reversible networks, negative flows on a reaction simply correspond to the reverse reactions. In general, however,

for non reversible networks the productive cone may be defined as

$$\mathcal{F}(\mathcal{A}) = \text{cone}(\text{cols}(\bar{\mathbb{S}}_C^{-1})) \cap \mathbb{R}_{\geq 0}^{\mathcal{A}_C}.$$

Species-productive cone: We term the image of the flow-productive cone of a MAS under its stoichiometric matrix $\bar{\mathbb{S}}$ as the *species-productive cone* and denote it by \mathcal{P} ,

$$\begin{aligned} \mathcal{P}(\mathcal{A}) &= \bar{\mathbb{S}} \cdot \mathcal{F}(\mathcal{A}) \\ &= \text{cone}(\text{cols}(\bar{\mathbb{S}} \cdot \bar{\mathbb{S}}_C^{-1})). \end{aligned} \quad (2.3)$$

Therefore, each vector in the interior of the species-productive cone for a MAS, which we refer to as the *species-productive region*, describes a change in concentration of the participating species for which the subnetwork is productive in the core member species. We remark that the *species-consumptive cone* for the *drainable subnetwork* \mathcal{A}^{op} , which is obtained by reversing the edges of the autocatalytic subnetwork \mathcal{A} , is the opposite of the species-productive cone:

$$\mathcal{P}(\mathcal{A}^{\text{op}}) = -\mathcal{P}(\mathcal{A}).$$

Partition-productive cone: Consider a MAS with species set \mathcal{S} . Let us denote the FWMC partition of \mathcal{S} as $\mathbb{T} = \{\mathbb{T}_F, \mathbb{T}_W, \mathbb{T}_{M/C}, \mathbb{T}_C\}$. Let \mathbf{e}_i be a column vector of size \mathcal{S} with 1 at the i^{th} position and 0 elsewhere. We

define the set of basis vectors \mathcal{B}_T for the partition T to be

$$\mathcal{B}_T = \mathcal{B}_T^F \cup \mathcal{B}_T^W \cup \mathcal{B}_T^{M/C} \cup \mathcal{B}_T^C \quad (2.4)$$

where

$$\begin{aligned} \mathcal{B}_T^F &= \{-\mathbf{e}_f | f \in T_F\}, \\ \mathcal{B}_T^W &= \{\mathbf{e}_w | w \in T_W\}, \\ \mathcal{B}_T^{M/C} &= \{-\mathbf{e}_m \cup \mathbf{e}_m | m \in T_{M/C}\}, \\ \mathcal{B}_T^C &= \{\mathbf{e}_c | c \in T_C\}. \end{aligned}$$

Thus the basis \mathcal{B}_T for a partition T spans the negative half-space of each species in the food set, the positive half-space of each species in the waste and core set, and both negative and positive half-spaces of each species in the non-core member set. Let the reaction network have conservation laws $\{\mathbf{x}_i\}$, such that $\mathbf{x}_i \cdot \mathbb{S} = 0$. We define the *partition-productive cone* for the partition T , denoted by $\mathcal{Q}(T)$, to be the cone generated by the basis vectors restricted to the stoichiometric subspace

$$\mathcal{Q}(T) = \text{cone}(\mathcal{B}_T) \perp \{\mathbf{x}_i\}.$$

As noted in the previous subsection, belonging to an FWMC partition

is an equivalence relation. By definition, the species-productive cones of all equivalent autocatalytic subnetworks belonging to the partition T will lie in the partition-productive cone $\mathcal{Q}(T)$.

Example 2.6. *In this example, we will identify the flow-productive, species-productive and partition-productive cones for the MAS A_1 from Example 2.5 (shown in Fig. 2.5). For $A_1 = (\{A, B, C\}, \{A + B \rightarrow C, C \rightarrow 2A\})$,*

$$\begin{aligned}\bar{\mathcal{S}} &= \begin{bmatrix} -1 & 2 \\ -1 & 0 \\ 1 & -1 \end{bmatrix} \\ \bar{\mathcal{S}}_C &= \begin{bmatrix} -1 & 2 \\ 1 & -1 \end{bmatrix} \\ \bar{\mathcal{S}}_C^{-1} &= \begin{bmatrix} 1 & 2 \\ 1 & 1 \end{bmatrix}.\end{aligned}$$

The flow-productive cone is

$$\mathcal{F}(A_1) = \text{cone} \left(\begin{pmatrix} \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 2 \\ 1 \end{bmatrix} \end{pmatrix} \right),$$

and the species-productive cone is

$$\mathcal{P}(\mathcal{A}_1) = \text{cone} \left(\left(\begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ -2 \\ 1 \end{bmatrix} \right) \right).$$

Notice that the graph \mathcal{G}^* (from Example 2.4) has a positive conservation law $\mathbf{x} = [1, 1, 2]$. The partition-productive cone for the partition $\mathcal{T}_1 = \{\{\mathcal{B}\}, \{\}, \{\{\mathcal{A}, \mathcal{C}\}\}$ is

$$\begin{aligned} \mathcal{Q}(\mathcal{T}_1) &= \text{cone} \left(\left(\begin{bmatrix} 0 \\ -1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right) \perp \mathbf{x} \right) \\ &= \text{cone} \left(\left(\begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ -2 \\ 1 \end{bmatrix} \right) \right) \\ &= \mathcal{P}(\mathcal{A}_1). \end{aligned}$$

For this example, the species-productive and partition-productive cones are identical, which is an instance of a general result proved in Theorem 2.12.

2.3.2 Mathematical results

To understand the geometry of the different cones defined in the previous subsection, we prove some results about their behavior. In order to

organize the partitions, we give the conditions under which two minimal autocatalytic subnetworks (MASs) with different partitions will have non-intersecting species-productive regions in Proposition 2.7. Next, we explore conditions under which the species-productive cone is identical with the partition-productive cone for a MAS in Theorem 2.12. Under such conditions, the partition itself contains the information of the species-productive regions of all the MASs belonging to it. Finally, in Proposition 2.14, we understand the topological properties a CRN must possess in order for the species-productive cones of two MASs to intersect. In Sec. 2.3.3 we will use these results to construct a visualization scheme for the list of MASs in any CRN.

Proposition 2.7. *Two autocatalytic subnetworks with different food sets have disjoint species-productive regions if their non-core member species sets are empty.*

Proof. We will show that if two autocatalytic subnetworks A and B have distinct food sets and their non-core member species sets are empty, then the interiors of their species-productive cones (species-productive regions) do not intersect,

$$\text{int}(\mathcal{P}(A)) \cap \text{int}(\mathcal{P}(B)) = \emptyset.$$

Let $A_F/B_F \neq \emptyset$. Recall that the species-productive cone of the subnetworks lies within their partition-productive cone $\mathcal{Q}(\text{FWMC}(A))$ and

$\mathcal{Q}(\text{FWMC}(\text{B}))$. Let \mathbf{y} be the vector of length S such that

$$[\mathbf{y}]_i = \begin{cases} -1 & \text{if } i \in A_F/B_F \\ 0 & \text{otherwise.} \end{cases}$$

Then, by definition of the partition-productive cones, $\mathbf{y} \cdot \text{int}(\mathcal{Q}(\text{FWMC}(\text{A}))) > 0$ and $\mathbf{y} \cdot \text{int}(\mathcal{Q}(\text{FWMC}(\text{B}))) \leq 0$. Thus, \mathbf{y} defines a hyperplane separating the two species-productive regions, and the two regions do not intersect. \square

Remark 2.8. *If the non-core member species set is not empty, then the result is no longer true. For instance, see Example 2.13.*

Example 2.9. *Using Examples 2.5 and 2.6, the species-productive cones for A_1 and A_2 in the species set $\{A, B, C, D\}$ are*

$$\mathcal{P}(A_1) = \text{cone} \left(\begin{pmatrix} \begin{bmatrix} 1 \\ -1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ -2 \\ 1 \\ 0 \end{bmatrix} \right),$$

$$\mathcal{P}(A_2) = \text{cone} \left(\begin{pmatrix} \begin{bmatrix} -1 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} -2 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right).$$

If we let $\mathbf{y} = [0, -1, 0, 0]$, then $\mathbf{y} \cdot \text{int}(\mathcal{P}(A_1)) > 0$ and $\mathbf{y} \cdot \text{int}(\mathcal{P}(A_2)) < 0$. Thus \mathbf{y} defines a separating hyperplane, and the two species-productive regions are disjoint.

Lemma 2.10. *The restriction of the species-productive cones of all MASs to their core (member) species is the positive orthant in the core species.*

Proof. Recall that the species-productive cone of a MAS is defined as

$$\mathcal{P}(A) = \text{cone}(\text{cols}(\bar{\mathbb{S}} \cdot \bar{\mathbb{S}}_C^{-1})).$$

By definition of the matrix inverse, the restriction of the product to the core species is the identity matrix,

$$\bar{\mathbb{S}} \cdot \bar{\mathbb{S}}_C^{-1} \Big|_C = I_C.$$

Thus, the species-productive cone restricted to the core species is the positive orthant. \square

Lemma 2.11. *If a CRN has any positive conservation laws, the union of the food set and the non-core member set of any MAS must be non-empty.*

Proof. Let A be a MAS, with associated stoichiometric matrix $\bar{\mathbb{S}}$, of a CRN with a positive conservation law given by the vector \mathbf{x} . Suppose A does not have any element in the food set A_F or the non-core member set $A_{M/C}$. For any flow \mathbf{v} in the productive region of A , we know that $\bar{\mathbb{S}} \cdot \mathbf{v} \succ 0$. Also,

since \mathbf{x} is a positive conservation law, $\mathbf{x} \cdot \bar{\mathbf{S}} \cdot \mathbf{v} \stackrel{!}{=} 0$. Since the sum of positive values cannot add up to zero, we have a contradiction. Thus, there must be at least one element in the food set or the non-core member set of a MAS with positive conservation laws. \square

Theorem 2.12. *For a CRN with positive conservation laws, the species-productive cone of any MAS with exactly one more species than the core set is identical to its partition-productive cone.*

Proof. Let A be a MAS of a CRN with positive conservation laws $\{\mathbf{x}_i\}$ with exactly one more species than the core set, denoted by $C = \{C_1, \dots, C_C\}$. From Lemma 2.11, we know that the extra species must be either in the food set or non-core member set of A . In either case, the species must be net consumed by the subnetwork to respect the positive conservation laws, and we denote it by f .

Let us label the species set as $\{f, C_1, \dots, C_C\}$. Recall that the partition-productive cone is defined as $\mathcal{Q}(\text{FWMC}(A)) = \text{cone}(\mathcal{B}) \perp \{\mathbf{x}_i\}$, where

$$\mathcal{B} = \left\{ \begin{array}{c} \begin{bmatrix} -1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ \vdots \\ 0 \end{bmatrix}, \dots, \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix} \end{array} \right\}.$$

Note that the additional basis vector for the non-core member species with positive entries is omitted since we know that it must be consumed in

the species-productive region to respect the positive conservation laws. Moreover, using Lemma 2.10, we know that the restriction of the species-productive cone to the core species is the positive orthant in the core species. In particular, for the species-productive cone $\mathcal{P}(A) = \text{cone}(\text{cols}(\bar{\mathbb{S}} \cdot \bar{\mathbb{S}}_C^{-1}))$, we have

$$\bar{\mathbb{S}} \cdot \bar{\mathbb{S}}_C^{-1} = \begin{bmatrix} \mathbf{f} \\ I_C \end{bmatrix},$$

where \mathbf{f} is a row vector $[\mathbf{f}_1, \dots, \mathbf{f}_C]$ and I_C is the identity matrix of size C . Since the CRN also has conservation laws, these coefficients must satisfy

$$\begin{aligned} [\mathbf{x}_i]_f \cdot \mathbf{f}_1 + [\mathbf{x}_i]_{C_1} \cdot 1 &= 0 \\ &\vdots \\ [\mathbf{x}_i]_f \cdot \mathbf{f}_C + [\mathbf{x}_i]_{C_C} \cdot 1 &= 0, \end{aligned}$$

for every conservation law indexed by i . But these are simply the basis vectors of the partition-productive cone \mathcal{B} when made orthogonal to the conservation laws. Thus the partition-productive cone is identical to the species-productive cone. \square

Example 2.13. Consider two MASs (from Sec. 2.4.2 Fig. 2.3)

$$\begin{aligned}
 A_1 &= (\{\bar{1}, \bar{2}, \bar{3}, \bar{4}\}, \{\bar{1} + \bar{3} \rightarrow \bar{4}, \\
 &\quad \bar{1} + \bar{4} \rightarrow \bar{2} + \bar{3}, \bar{2} + \bar{2} \rightarrow \bar{3} + \bar{1}\}), \\
 A_2 &= (\{\bar{1}, \bar{2}, \bar{3}, \bar{4}\}, \{\bar{1} + \bar{3} \rightarrow \bar{4}, \\
 &\quad \bar{1} + \bar{4} \rightarrow \bar{2} + \bar{3}, \bar{2} + \bar{2} \rightarrow \bar{4}\}).
 \end{aligned}$$

The partitions of A_1 and A_2 are $\{\{\}, \{\}, \{\bar{1}, \bar{2}, \bar{3}, \bar{4}\}\}$ and $\{\{\bar{1}\}, \{\}, \{\bar{2}, \bar{3}, \bar{4}\}\}$, respectively. A simple calculation yields that their species-productive cones are in fact identical, and

$$\mathcal{P}(A_1) = \mathcal{P}(A_2) = \text{cone} \left(\begin{pmatrix} \begin{bmatrix} -2 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} -3 \\ 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -4 \\ 0 \\ 0 \\ 1 \end{bmatrix} \end{pmatrix} \right).$$

Proposition 2.14. If the species-productive cones of two MASs have an intersection, the reversible graph of their union has a null flow.

Proof. Let the productive cones of two MASs A and B , have a non-empty intersection, $\mathcal{P}(A) \cap \mathcal{P}(B) \neq \emptyset$. Let $A \cup B$ be the subnetwork obtained by taking the union of the two subnetworks, and let $\mathbb{S}_{A \cup B}$ be the associated stoichiometric matrix. Then, there must be non-identical flows \mathbf{v}_A and \mathbf{v}_B

with support in A and B , respectively, such that

$$\mathbb{S}_{A \cup B} \cdot \mathbf{v}_A = \mathbb{S}_{A \cup B} \cdot \mathbf{v}_B.$$

This implies that $\mathbb{S}_{A \cup B} \cdot (\mathbf{v}_A - \mathbf{v}_B) = 0$, and thus the kernel of the stoichiometric matrix is non-empty. In general, the difference of their flows $(\mathbf{v}_A - \mathbf{v}_B)$ can have negative coefficients, however, it yields the interpretation of a null flow on a network if the union of the two subnetworks $A \cup B$ is made reversible. \square

Remark 2.15. *The converse statement, if the reversible graph of the union of two MASs contains a null cycle then their species-productive cones must intersect, is not true. For example, consider the MASs (from Sec. 2.4.2 Fig. 2.3)*

$$\begin{aligned} A_1 &= (\{\bar{1}, \bar{2}, \bar{3}\}, \{\bar{1} + \bar{2} \rightarrow \bar{3}, \bar{1} + \bar{3} \rightarrow \bar{2} + \bar{2}\}), \\ A_2 &= (\{\bar{1}, \bar{2}, \bar{3}\}, \{\bar{2} \rightarrow \bar{1} + \bar{1}, \bar{3} + \bar{1} \rightarrow \bar{2} + \bar{2}\}). \end{aligned}$$

Notice that $\text{FWMC}(A_1) = \{\{\bar{1}\}, \{\}, \{\{\bar{2}, \bar{3}\}\}\}$ and $\text{FWMC}(A_2) = \{\{\bar{3}\}, \{\}, \{\{\bar{1}, \bar{2}\}\}\}$. Thus, using Proposition 2.7, the two species-productive cones do not intersect. However, the union of the two subnetworks yields

$$\begin{aligned} A_1 \cup A_2 &= (\{\bar{1}, \bar{2}, \bar{3}\}, \{\bar{1} + \bar{2} \rightarrow \bar{3}, \\ &\quad \bar{1} + \bar{3} \rightarrow \bar{2} + \bar{2}, \bar{2} \rightarrow \bar{1} + \bar{1}\}). \end{aligned}$$




Figure 2.6: The intersection data of partition-productive cones for all partitions of the complete $L = 5$ 1-constituent CCRN (described in Sec. 2.4.2).

which is a null cycle (with null flow $[1, 1, 1]^T$).

2.3.3 Geometry and visualization

In the last subsection, we proposed an algorithm that, given a stoichiometric matrix of a CRN, outputs a list of the MASs it contains. In this section, we will explain how to take the list of MASs and visualize their combina-

torics and geometry. Recall that for each MAS, we define a flow-productive cone in the flow space of the graph where each core member species is strictly produced, and a species-productive cone on the space of changes in concentration (population). Geometrically, the space of changes in chemical concentration (i.e., the velocity space of chemical concentrations) is the stoichiometric subspace of the CRN. Thus, the list of MASs can be seen as yielding a partial polyhedral decomposition of the stoichiometric space induced from the flows on the hypergraph (CRN). We remark that, for a more complete decomposition one must also consider the consumptive cones of the minimal drainable subnetworks. However, it is not necessary that the union of the autocatalytic and drainable subnetworks span the complete stoichiometric space (for example, see Sec. 2.4.2 Fig. 2.12).

In subsec. 2.3.1, we showed that each subnetwork can be assigned to its FWMC equivalence class. As explained in Proposition 2.7, while there are cases where the species-productive cones of different equivalence classes can be shown to not intersect, in general the productive regions of different equivalence classes can intersect. To visualize the list of equivalence classes to which the MASs belong, we run the algorithm for finding an intersection between each pair of partition-productive cones and obtain a two-dimensional square matrix, C_{pp} , of dimension equal to the number of equivalence classes. Let us denote the list of equivalence classes by \mathcal{L} . The

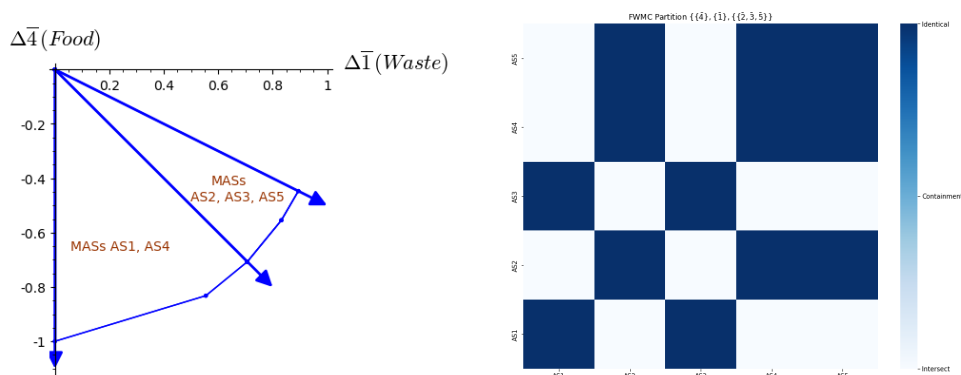


Figure 2.7: All MASs belonging to the $\{\{\bar{4}\}, \{\bar{1}\}, \{\{\bar{2}, \bar{3}, \bar{5}\}\}\}$ class in the complete $L = 5$ 1-constituent CCRN are considered. The projections of their species-productive cones on the food and waste species are shown in the left panel. The right panel displays the intersection information of MASs within a class analogous to Fig. 2.6. Note that there are no MASs in the class that are disjoint, and thus the ‘disjoint’ label is omitted.

entry of C_{pp} in the i^{th} row and j^{th} columns is given by,

$$[C_{pp}]_i^j = \begin{cases} 2 & \text{if } Q(\mathcal{L}_i) = Q(\mathcal{L}_j), \\ 1 & \text{if } Q(\mathcal{L}_i) \subsetneq Q(\mathcal{L}_j), \\ 0 & \text{if } \emptyset \neq Q(\mathcal{L}_i) \cap Q(\mathcal{L}_j) \subsetneq Q(\mathcal{L}_i), Q(\mathcal{L}_j), \\ -1 & \text{if } Q(\mathcal{L}_i) \cup Q(\mathcal{L}_j) = \emptyset. \end{cases}$$

Notice that any asymmetry in entries across the diagonal indicates that only one of the partition-productive cones completely contain the other. This matrix can then be visualized as a heat map, for example see Fig. 2.6.

Each equivalence class can contain several MASs. The species-productive cones of the MASs in a class are not always identical. While they will share

the same projection on the core species, the productive regions in the non-core species can be very different. For example, let us pick the equivalence class $\{\{\bar{4}\}, \{\bar{1}\}, \{\{\bar{2}, \bar{3}, \bar{5}\}\}\}$ from the list of classes shown in Fig. 2.6. From the figure, we know that it contains 5 MASs. We plot the projection of the species-productive region in the non-core species in the top panel of Fig. 2.7. In higher dimensions when more non-core members are involved, this representation can get rather cumbersome. Thus, we employ a similar visualization as C_{pp} to depict the intersection of MASs within an equivalence class. Whether or not there is an intersection between the species-productive cones can be ascertained using our algorithm outlined in the previous subsection. For an example of the resulting visualization for the same equivalence class considered above, see the bottom panel of Fig. 2.7. In the same manner, a visualization for the information of pairwise intersection of the productive cones of all MASs for a CRN can also be obtained, for example see Fig. 2.8.

Example 2.16. *Consider the complete $L = 5$ 1-constituent CCRN of order two. The list of all the FWMC classes, the number of MASs they contain, and their intersection information is summarized in Fig. 2.6. Consider the class*




Figure 2.8: The pairwise intersection information for all the species-productive cones of the $L = 5$ 1-constituent complete CCRN of order two. The reaction network and the list of MASs can be found in Ref. Gagrani et al. (2023a).

$\{\{\bar{4}\}, \{\bar{1}\}, \{\{\bar{2}, \bar{3}, \bar{5}\}\}\}$. It contains the MASs:

$$AS1 = \{\bar{2} + \bar{2} \rightarrow \bar{3} + \bar{1},$$

$$\bar{5} \rightarrow \bar{2} + \bar{3},$$

$$\bar{4} + \bar{3} \rightarrow \bar{2} + \bar{5}\},$$

$$AS2 = \{\bar{4} + \bar{2} \rightarrow \bar{1} + \bar{5},$$

$$\bar{5} \rightarrow \bar{2} + \bar{3},$$

$$\bar{4} + \bar{3} \rightarrow \bar{2} + \bar{5}\},$$

$$AS3 = \{\bar{3} + \bar{3} \rightarrow \bar{5} + \bar{1},$$

$$\bar{5} \rightarrow \bar{2} + \bar{3},$$

$$\bar{4} + \bar{2} \rightarrow \bar{3} + \bar{3}\},$$

$$AS4 = \{\bar{3} + \bar{3} \rightarrow \bar{5} + \bar{1},$$

$$\bar{4} + \bar{2} \rightarrow \bar{5} + \bar{1},$$

$$\bar{5} \rightarrow \bar{2} + \bar{3}\},$$

$$AS5 = \{\bar{3} \rightarrow \bar{2} + \bar{1},$$

$$\bar{4} + \bar{2} \rightarrow \bar{5} + \bar{1},$$

$$\bar{5} \rightarrow \bar{2} + \bar{3}\}.$$

The projection of the species-productive cones for the above MASs and their intersection information is summarized in Fig. 2.7.




Figure 2.9: Consider a CRN of polymers whose species set is arranged as a binary tree in two letters, shown to the left. Then the *cluster* resulting from each binary string is obtained by counting the number of occurrences of each alphabet. The clusters thus obtained form a lattice structure, shown to the right.

2.4 Cluster chemical reaction networks (CCRN) framework

All realistic chemical reactions, ignoring nuclear reactions, conserve the number of atoms of each type. Often, but not always, reactions involving polymers also conserve the number of each monomer. This motivates taking a *coarse-grained* perspective, where we forget the bond structure of each molecule or polymer and only count the number of atoms or monomers, respectively. Following terminology from Chapter 7 of (Kelly, 2011), we will refer to this coarse-grained description of each species as a *cluster* (see Fig. 2.9). This also allows us to coarse-grain a realistic CRN into a new CRN which has clusters as species and reactions between multisets of clusters induced from the reactions in the original CRN. We will call the resulting CRN, a *cluster chemical reaction network* (CCRN) obtained from

the original CRN.

There are three advantages of considering CCRN descriptions of CRNs. First, consider a CRN whose species set consists of all linear polymers formed of D monomers up to size N . Then, the CCRN induced by the CRN would reduce the number of species from exponential (D^N) to polynomial (N^D) in the size of the polymer, providing a significant computational advantage. Secondly, since each reaction in the original CRN induces a reaction in the CCRN with the same topology, the mapping would preserve the autocatalytic property. In other words, if a CRN is autocatalytic, its induced CCRN will also be autocatalytic. However, since it is a many-to-one map, multiple autocatalytic subnetworks of the CRN might map to the same autocatalytic subnetwork in its CCRN or new autocatalytic subnetworks in the CCRN might emerge (see Example 2.17). Thirdly, one can later re-introduce structure into the CCRN by adding more species with the same cluster counts (conserved quantities). In this way, the coarse-graining can be systematically refined to recover the original CRN, and allows a gradual complexification of the model.

Example 2.17. Consider the CRN

$$\mathcal{G} = (\mathcal{S}, \mathcal{R})$$

$$\mathcal{S} = \{A, B, BAA, ABB, ABBA, BAAB, ABBABAAB\}$$

$$\begin{aligned} \mathcal{R} = \{ & ABBA + BAA + B \rightarrow ABBABAAB, \\ & BAAB + ABB + A \rightarrow ABBABAAB, \\ & ABBABAAB \rightarrow ABBA + BAAB. \} \end{aligned}$$

The CCRN obtained by counting the number of As and Bs in each string and representing the resulting clusters as $\overline{n_A}, \overline{n_B}$ is

$$\text{CCRN}(\mathcal{G}) = (\mathcal{S}_C, \mathcal{R}_C)$$

$$\mathcal{S}_C = \{\overline{1,0}, \overline{0,1}, \overline{2,1}, \overline{1,2}, \overline{2,2}, \overline{4,4}\}$$

$$\mathcal{R}_C = \{\overline{2,2} + \overline{2,1} + \overline{0,1} \rightarrow \overline{4,4},$$

$$\overline{2,2} + \overline{1,2} + \overline{1,0} \rightarrow \overline{4,4},$$

$$\overline{4,4} \rightarrow \overline{2,2} + \overline{2,2}. \}$$

Notice that while the CRN has one minimal autocatalytic subnetwork (MAS), the CCRN has two MASs. This can be remedied by remembering that BAAB and ABBA are distinct by introducing a new species $\overline{2,2}^*$. A 'finer' CCRN would

then be

$$\begin{aligned} \text{CCRN}(\mathcal{G})^* &= (\mathcal{S}_C^*, \mathcal{R}_C^*) \\ \mathcal{S}_C^* &= \{\overline{1,0,0,1,2,1,1,2,2,2,2,2^*}, \overline{4,4}\} \\ \mathcal{R}_C^* &= \{\overline{2,2} + \overline{2,1} + \overline{0,1} \rightarrow \overline{4,4}, \\ &\quad \overline{2,2^*} + \overline{1,2} + \overline{1,0} \rightarrow \overline{4,4}, \\ &\quad \overline{4,4} \rightarrow \overline{2,2} + \overline{2,2^*}.\} \end{aligned}$$

Note that $\text{CCRN}(\mathcal{G})^$ has identical information to the original CRN and is the finest CCRN that can be obtained from the original CRN.*

We formally define a CCRN in Sec. 2.4.1, and systematically explore the properties of CCRNs with one conserved quantity (type of atom or monomer) in Sec. 2.4.2. We also provide the statistics of autocatalytic subnetworks found in such reaction networks and present a worked-out-example in Sec. 2.4.2. We briefly discuss the computational challenges in scaling the algorithm for fully connected models in Sec. 2.4.2 and introduce rule-generated CCRNs in Sec. 2.4.3.

2.4.1 Formalism

A cluster chemical reaction network (CCRN) is a CRN with the additional structure that, upon excluding the inflow and outflow reactions, the network has at least one nonnegative integer conservation law (for the defi-

inition of a conservation law, see Sec. 2.2.1). Each conservation law corresponds to a type of *constituent* that is conserved. We denote the number of distinct constituents by D and label them as A_1, \dots, A_D . The species of a CCRN, termed *clusters*, are multisets of constituents and denoted by an overline (notation chosen to be consistent with (Liu and Sumpter, 2018)) over a vector of nonnegative integers representing the number of each constituent in the cluster. For instance,

$$\bar{\mathbf{n}} := \overline{n_1, \dots, n_D}$$

denotes the cluster comprising n_1, \dots, n_D of constituents A_1, \dots, A_D , respectively. Notice that a unit of a constituent represents a unit of conserved quantity, thus a cluster is labelled by its set of conserved quantities. In case of multiple clusters with the same conserved quantities, we distinguish them with an asterisk $*$. For example, in Table 2.2, $\bar{1}$ and $\bar{1}^*$ refer to a monomer and an activated monomer, respectively. Finally, we define the *length* of a cluster to be the sum of its conserved quantities,

$$l(\mathbf{n}) := \sum_{i=1}^D n_i.$$

In a CCRN, *complexes* are multisets of clusters and denoted by

$$\mathbf{c}^\alpha := \sum_{\mathbf{n}} c_{\mathbf{n}}^\alpha \cdot \bar{\mathbf{n}}.$$

As an abuse of notation, we denote the stoichiometry of complex c^α also by c^α , where now it is the column vector with entries c_n^α . We define the *size* of a complex to be its vector sum of conserved quantities of each type, and denote it as

$$|c^\alpha| = \sum_{\mathbf{n}} c_n^\alpha \cdot \mathbf{n}.$$

We define the *width* of a complex to be total number of clusters in the multiset, and denote it by

$$w(c^\alpha) := \sum_{\mathbf{n}} c_n^\alpha.$$

In the CCRN, every reaction must be such that the size of the source and target complexes are identical. Thus, a reaction $c^\alpha \rightarrow c^\beta$ is allowed only if

$$|c^\alpha| = \sum_{\mathbf{n}} c_n^\alpha \cdot \mathbf{n} = \sum_{\mathbf{n}} c_n^\beta \cdot \mathbf{n} = |c^\beta|.$$

This restriction consistently ensures that the number of constituents across each reaction are conserved, as

$$\sum_{\mathbf{n}} \mathbf{n} \cdot (c^\beta - c^\alpha)_n = 0.$$

Let $\mathcal{H}_D = (\mathcal{S}, \mathcal{R})$ denote a CCRN with D constituents. Then for any

finite CCRN,

$$\begin{aligned}\mathcal{S} &\subset \{\bar{\mathbf{n}}|\mathbf{n} \in \mathbb{Z}_{\geq 0}^D\}, \\ \mathcal{R} &\subseteq \{c^\alpha \rightarrow c^\beta \mid |c^\alpha| = |c^\beta|\}.\end{aligned}$$

We define the length of a CCRN by the maximum length of its clusters, and denote it by

$$l(\mathcal{H}_D) = \max_{\mathcal{S}} |\mathbf{n}|.$$

We also define the *order* of a CCRN to be the maximum width of its complexes, denoted by

$$o(\mathcal{H}_D) = \max_{\mathcal{R}} w(c^\alpha).$$

We define a *complete CCRN* of length L and order w to consist of all species up to length L and all possible reactions that are allowed by the conservation laws with complexes of width up to w . Since any reaction consisting of more than two reactants or products can be written as chains of second-order reactions, for our analysis *we will restrict to CCRNs of order two*.

2.4.2 Complete CCRN

1 constituent CCRN

Let $\mathcal{H}_1^L = (\mathcal{S}, \mathcal{C}, \mathcal{R})$ denote a complete 1-constituent CCRN of length L and order two. 1-constituent CCRN of length L means that the species set




Figure 2.10: Number of minimal autocatalytic subnetworks (MASs) varying with L for the complete 1-constituent CCRN of order two. For each L , the distribution of MAS with the number of reactions is also shown. Notice that total MASs for each L , plotted as black dots, increases exponentially with L .

Max. size L	CCRN properties					# reactions in a MAS				
	$ \mathcal{C} $	$ \mathcal{R} $	ℓ	ι	δ	2	3	4	5	6
3	6	6	3	3	1	2	0	0	0	0
4	11	14	5	8	3	9	11	0	0	0
5	17	26	7	16	6	24	98	48	0	0
6	24	44	9	29	10	48	461	768	331	0
7	32	68	11	47	15	85	1549	6028	5673	1709

Table 2.1: The complete 1-constituent CCRN of size L and order two is considered for L from 3 – 7. The topological properties of the CCRN, such as number of linkage classes ℓ and deficiency δ , are shown in the left half (for calculations, see Appendix A.2). The number of MASs for a given number of reactions is shown in the right half (for a visual representation, see Fig. 2.10).




Figure 2.11: The first few reactions in the complete 1-constituent CCRN of order two. The reactions containing clusters only up to length four ($L = 4$) are shown in yellow boxes.

consists of clusters up to length L, denoted as $\mathcal{S} = \{\bar{1}, \bar{2}, \dots, \bar{L}\}$. By order two, we mean that all complexes are at most of width two, and the set of complexes is denoted as $\mathcal{C} = \{\bar{a}\} \cup \{\bar{a} + \bar{b}\}$ for $\bar{a}, \bar{b} \in \mathcal{S}$ and $b \leq a \leq L$. By a *complete CCRN*, we mean that the set of reactions contains all the possible

reactions that are allowed by the conservation law, i.e.

$$\mathcal{R} = \{\bar{a} + \bar{b} \rightleftharpoons \bar{c} + \bar{d} | a + b = c + d\} \cup \\ \{\bar{a} \rightleftharpoons \bar{c} + \bar{d} | a = c + d\}.$$

The first few linkage classes of the complete 1-constituent CCRN of order two are shown in Fig. 2.11 and the complete CCRN for $L = 4$ is highlighted in yellow. The information about the reaction networks and the MASs that they contain for the complete 1-constituent CCRNs of L ranging from 3 to 7 is collected in Table 2.1 and Fig. 2.10.

Consider the complete 1-constituent CCRN with $L = 4$ of order two, the reaction network for which is shown in Fig. 2.11. Notice that this subnetwork has 5 linkage classes, 11 complexes, and 14 reactions. An application of the algorithm for enumerating all the MASs of the CCRN discussed in Ref. Gagrani et al. (2023a) yields a list of 20 subnetworks, shown in Fig. 2.3. In the list, there are 9 MASs with two reactions and 11 MASs with three reactions. Since the CCRN has 4 species, and 1 conservation law, the stoichiometric subspace is of dimension 3. Moreover, due to the conservation law, from Lemma 2.11 each MAS must possess at least one non-core species. This means that there cannot be any 4 reaction MAS as it would require at least 5 distinct species, which would be a contradiction.

Upon obtaining a list of MASs for a given CRN, we would like to understand how their species-productive and partition-productive cones

intersect. For the $L = 4$ CCRN considered above, using Theorem 2.12, any MASs within a partition have identical species-productive cones and it is identical to the partition-productive cone. Thus, we do not show any plot for the species-productive cone intersection data. Moreover, any two partitions consisting of the same core of 3 species must have the same partition-productive cone, as the stoichiometric subspace is of dimension 3. For example, the partition-productive cones of $\{\{\bar{1}\}, \{\}, \{\{\bar{2}, \bar{3}, \bar{4}\}\}\}$ and $\{\{\}, \{\}, \{\bar{1}, \{\bar{2}, \bar{3}, \bar{4}\}\}\}$ are identical. The information of the intersection of partition-productive cones for different FWMC partitions are shown in the bottom panel of Fig. 2.3.

Unlike the $L = 4$ case, for the complete 1-constituent CCRN with $L = 5$ of order two, there can be MASs in the same equivalence class whose species-productive cones do not intersect. Moreover, it is not a priori clear whether or not two partition-productive cones of different partitions will intersect. Thus we employ our visualization scheme and display this information in Fig. 2.6. As described in Sec. 2.3.3, for a particular equivalence class we also show the intersection information of the species-productive cones in Fig. 2.7.


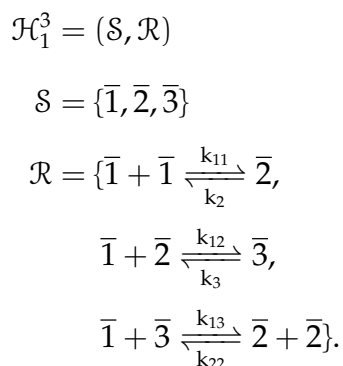


Figure 2.12: The 2-dimensional stoichiometric subspace along with the autocatalytic productive cones for the $L = 3$ complete 1-constituent CCRN are shown. The 6 rays emanating from the origin, labelled by their directions in the 3-dimensional space, form the edges of productive cones where one species is consumed and another is produced. For instance, the ray $[-3, 0, 1]$ corresponds to the direction in the stoichiometric space where $\bar{1}$ is consumed, $\bar{3}$ is produced, and $\bar{2}$ remains constant.

$L=3$ 1-constituent CCRN

Consider the complete 1-constituent CCRN with $L = 4$ of order two, given by



This network has two MASs, namely

$$\begin{aligned}
 A_1 = & \{ \bar{1} + \bar{2} \rightarrow \bar{3}, \\
 & \bar{1} + \bar{3} \rightarrow \bar{2} + \bar{2} \}, \\
 A_2 = & \{ \bar{2} \rightarrow \bar{1} + \bar{1}, \\
 & \bar{3} + \bar{1} \rightarrow \bar{2} + \bar{2} \}.
 \end{aligned}$$

Since the CCRN has one conservation law, the stoichiometric subspace is of dimension 2 and perpendicular to $[1, 2, 3]^T$. The two-dimensional stoichiometric subspace is shown in Fig. 2.12. The 6 rays emanating from the origin, labelled by their directions in the 3-dimensional space, form the edges of productive cones where one species is consumed and another is produced. The species-productive cones of the two autocatalytic cycles are also labelled. Note that the species-productive regions is the complete cone within the bounding rays, and the finite boundary is drawn simply to enhance visualization. For example, the species-productive cone of the MAS A_1 is bounded by the rays $[-2, 1, 0]$ and $[-3, 0, 1]$. Notice that the FWMC partition of A_1 is $\{\{\bar{1}\}, \{\}, \{\{\bar{2}, \bar{3}\}\}\}$, consistent with the species-productive region that strictly consumes species $\bar{1}$ and strictly produces species $\bar{2}$ and $\bar{3}$.

We also want to make a few remarks about the details that are absent in Fig. 2.12. Firstly, for every MAS, there is a minimal drainable subnetwork obtained by reversing the reaction edges. The species-productive cone

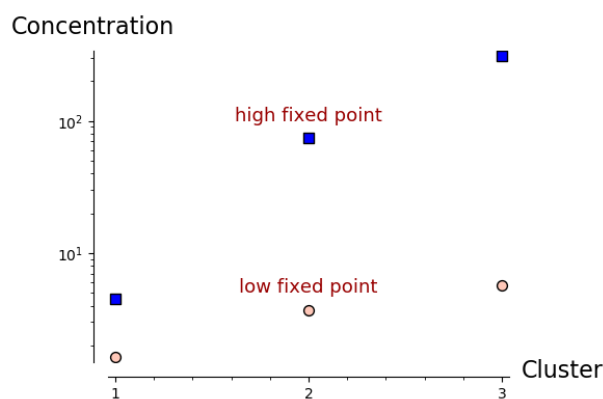


Figure 2.13: Multiple fixed points for the complete 1-constituent CCRN with $L = 3$ and order two obtained using Feinberg's deficiency one algorithm (see Appendix A.3). The parameters for the algorithm are chosen to be $\mu_{\bar{1}} = 1, \mu_{\bar{2}} = 3, \mu_{\bar{3}} = 4 = \eta$, and the rate constants for the model are $k_{11} = 1, k_2 = 1.001, k_{12} = 1.1, k_3 = 1.011, k_{13} = 0.146, k_{22} = 0.027$.

of these drainable networks will be the negative of their autocatalytic counterparts. Secondly, notice that there is no MAS with the partition $\{\{\bar{2}\}, \{\}, \{\{\bar{1}, \bar{3}\}\}\}$. In case there was one, as will be if we allow reactions of order 3, then its species-productive cone would be bound by the rays $[0, -3/2, 1]$ and $[1, -1/2, 0]$.

Notice from Table 2.1 that this CCRN has a deficiency of one. This means that when taken with mass-action kinetics, it has the potential to exhibit multistability. To investigate this, we used Feinberg's deficiency one algorithm (Feinberg, 1988) in Appendix A.3 and obtained rate constants for which the CCRN has multiple steady states. An example of multiple steady states and the rate constants using which they are obtained are reported in Fig. 2.13.

Number theoretic result for 1-constituent CCRN of order two

Since CCRNs are defined by number conservation laws, finding autocatalytic subnetworks in CCRN corresponds to number-valued solutions of systems of equations. In this subsection we will show that if there exists a two-reaction MAS in a partition of a 1-constituent CCRN, then it must be unique.

Lemma 2.18. *Let A be a 2-reactions MAS in a 1-constituent CCRN of order two. Then, there does not exist any other MASs in the CCRN with the same FWMC partition as A .*

Proof. Let $A_C = \{\bar{i}, \bar{j}\}$ be the core species in A (assume without loss of generality that $i > j$). The possible 2-reaction MASs with that set of core species are given by reactions of the form

$$\begin{cases} \bar{i} + \bar{l}_1 \rightarrow \bar{j} + \bar{l}_2, \\ \bar{j} + \bar{l}_3 \rightarrow 2\bar{i}. \end{cases}$$

$$\begin{cases} \bar{i} + \bar{m}_1 \rightarrow 2\bar{j}, \\ \bar{j} + \bar{m}_2 \rightarrow \bar{i} + \bar{m}_3. \end{cases}$$

Note that each non-core species can take the value 0, indicating the absence of a reactant. However, due to Lemma 2.11, the conservation law will require that at least one food species is non zero. Observe, also, that in both types of potential MASs there is a reaction which is uniquely

determined by \bar{i} and \bar{j} (being then $\bar{\ell}_3$ and \bar{m}_1 also unique for this choice of the core set). Concretely $\bar{\ell}_3 = \overline{2i - j}$ and $\bar{m}_1 = \overline{2j - i}$. Note that these reactions are different since it is assumed that $i \neq j$. Thus, the above reactions reduce to:

$$M_1 : \begin{cases} \bar{i} + \overline{j + \ell - i} \rightarrow \bar{j} + \bar{\ell}, \\ \bar{j} + \overline{2i - j} \rightarrow 2\bar{i}. \end{cases}$$

$$M_2 : \begin{cases} \bar{i} + \overline{2j - i} \rightarrow 2\bar{j}, \\ \bar{j} + \overline{i + m - j} \rightarrow \bar{i} + \bar{m}. \end{cases}$$

Now we will do a case-by-case analysis to show that each FWMC-partition will contain at most one solution of the systems of equations.

Note that the only possible partition of the species induced by the reactions in M_1 is: $\text{FWMC}(M_1) = \{\bar{j} + \bar{\ell} - \bar{i}, \overline{2i - j}\}, \{\bar{\ell}\}, \{\{\bar{i}, \bar{j}\}\}$ whereas the partition for M_2 is: $\text{FWMC}(M_2) = \{\overline{2j - i}, \bar{i} + \bar{m} - \bar{j}\}, \{\bar{m}\}, \{\{\bar{i}, \bar{j}\}\}$ From both types of partition, the only possibility that both coincide is that $\bar{\ell} = \bar{m}$ and $\bar{i} = \bar{j}$ contradicting the assumption that $i > j$.

□

Given two different species \bar{i}, \bar{j} with $i > j \geq 1$, the above result also provides a way to construct all the 2-reactions MASs with those species being the core species. The explicit reactions taking part of the MASs are in the form:

- In case $i \leq \frac{L}{2}$:

$$\text{MAS}_1 = \begin{cases} \text{R}_1: \bar{i} + \overline{j + \ell - i} \rightarrow \bar{j} + \bar{\ell}, \\ \text{R}_2: \bar{j} + \overline{2i - j} \rightarrow 2\bar{i}. \end{cases}$$

for all $\ell \in [i - j, L] \cap \mathbb{Z}$.

- In case $\frac{i}{2} \leq j \leq \frac{L}{2}$:

$$\text{MAS}_2 = \begin{cases} \text{R}_1: \bar{i} + \overline{2j - i} \rightarrow 2\bar{j}, \\ \text{R}_2: \bar{j} + \overline{i + m - j} \rightarrow \bar{i} + \bar{m}. \end{cases}$$

for all $m \in [0, L + j - i] \cap \mathbb{Z}$.

Computational challenges in scaling

Consider the complete 2-constituent CCRN of order two, defined as

$$\mathcal{H}_2 = (\mathcal{S}, \mathcal{R})$$

$$\mathcal{S} = \{\overline{\mathbf{a}}, \overline{\mathbf{b}} \mid \mathbf{a}, \mathbf{b} \in \mathbb{Z}_{\geq 0}\}$$

$$\mathcal{R} = \{\overline{\mathbf{a}}, \overline{\mathbf{b} + \mathbf{c}}, \overline{\mathbf{d}} \rightarrow \overline{\mathbf{e}}, \overline{\mathbf{f} + \mathbf{g}}, \overline{\mathbf{h}} \mid$$

$$\mathbf{a} + \mathbf{c} = \mathbf{e} + \mathbf{g} \text{ and } \mathbf{b} + \mathbf{d} = \mathbf{f} + \mathbf{h}\}.$$

We wish to remark that for \mathcal{H}_2 and other CRNs with a large number of reactions, generating the whole set of autocatalytic subnetworks is computationally challenging. As the number of reaction increases, the dimension

of the space where the MASs are to be found also increases, and even finding a single MAS in the network implies solving a difficult optimization problem that might be computationally costly. Even being each of the cycles polynomial-time solvable (which does not seem to be the case), the complexity of the enumeration algorithm may turn into exponential (see, e.g., (Garey and Johnson, 1979)) since the number of subnetworks increases considerable with the number of species and reactions. Thus for practical reasons, it maybe more feasible to consider sparser rule-generated networks, as explained in the next subsection and shown in Fig. 2.2.

2.4.3 Rule generated CCRNs

A naive computational modelling of artificial chemistry (Banzhaf and Yamamoto, 2015) often runs into issues of memory, storage and processing due to a combinatorial explosion of chemical species and reactions involved in the CRN. For instance, if one uses string chemistry (Moyer et al., 2020) (or polymer sequence chemistry) to model polymers of size up to N with D distinct monomers (D constituents), a straightforward calculation shows that one is required to track $\sum_{i=1}^N D^i > D^N$ species. Since for any physically relevant model $N \gg D$, so it is clear that the computation will soon become intractable as one increases the size of the string (polymer) N while also considering their relevant network of interactions. On the other hand, even a simple graph-grammar for generating artificial chemistry, with the application of (Andersen et al., 2016), can generate a

Monomer activation	$\bar{1} \rightleftharpoons \bar{1}^*$
Catalyzed activation	$\bar{C} + \bar{1} \rightleftharpoons \bar{1}^* + \bar{C}$
Polymerization	$\begin{aligned} \bar{1} + \bar{1}^* &\rightleftharpoons \bar{2} \\ \bar{2} + \bar{1}^* &\rightleftharpoons \bar{3} \\ &\vdots \end{aligned}$
1-constituent	$\bar{j} + \bar{1}^* \rightleftharpoons \overline{j+1}$
2-constituent	$\begin{cases} \overline{j, k} + \overline{1, 0}^* \rightleftharpoons \overline{j+1, k} \\ \overline{j, k} + \overline{0, 1}^* \rightleftharpoons \overline{j, k+1} \end{cases}$
\vdots	\vdots
Templated polymerization	$\begin{cases} \overline{j, k} + \overline{k-1, j+1, 0}^* \rightleftharpoons \overline{j, k+k, j} \\ \overline{j, k} + \overline{k, j-1+0, 1}^* \rightleftharpoons \overline{j, k+k, j} \end{cases}$
Fusion	$\bar{j} + \bar{k} \rightleftharpoons \overline{j+k}$
Fission	$\bar{j} \rightleftharpoons \overline{j-k+k}$


Table 2.2: Several examples of rules for generating a sparse CCRN relevant to biochemistry. Here, in accordance with the definition of a CCRN, the variables j, k must be chosen such that the resulting clusters have a non-negative integer quantity of each constituent.

reaction network that is computationally unfeasible. Moreover, even if one restricts the reaction network in the computationally feasible regime, an enumeration of autocatalytic subnetworks or motifs will generically be an extremely long list. In such a scenario, it might be unclear how to simplify the model without stripping away the essential details or adding artifacts.

We argue that the *Cluster framework*, or considering CCRNs induced by CRNs, can help alleviate some of the issues raised above. Any rule generated CRN, in string chemistry or built using graph-grammar methods, must induce rules on the CCRN. Since the induced CCRN can have an exponentially reduced species set, it offers a way to constrain the model to a more computationally feasible regime. Since autocatalytic networks are preserved under this coarsening (with some caveats mentioned in the introduction to this section), one can use the CCRN to obtain a smaller list of MASs. One can then gradually complexify the CCRN by adding more species with the same conserved quantities until the required behavior of the CRN is captured by the CCRN.

For example, consider polymer chemistry with two monomers A and B . For an arbitrary polymer ω , the addition of an A or B yields $\omega \cdot A$ or $\omega \cdot B$, respectively. In the cluster framework, if we map A , B , and ω to $\overline{1,0}$, $\overline{0,1}$, and $\overline{j,k}$, respectively, the polymerization reactions in the CCRN

become



In realistic chemistry, however, a monomer addition to a polymer is only done by an activated monomer and not an unactivated monomer. Thus, we may want to distinguish activated and unactivated monomers in our model, for which we can add extra species $\overline{1, 0}^*$ and $\overline{0, 1}^*$. The resulting polymerization reactions and other examples of rule generated CCRNs are shown in Table 2.2.

2.5 Discussion and future research

In this work, we began by reviewing the notions of autocatalysis in literature in Sec. 2.2.2. We term the autocatalysis that can be detected from the stoichiometric matrix as *stoichiometric autocatalysis*, and show that it is more constrained than *formal* or *exclusive* autocatalysis. Any system showing stoichiometric autocatalysis is formally and exclusively autocatalytic, which means that enumeration of stoichiometrically autocatalytic subnetworks can only underestimate the autocatalytic structure in a CRN.

In (Blokhuis et al., 2020), it was shown that there are exactly five types of minimal stoichiometrically autocatalytic cores, where a core is obtained

by only selecting a few species from the complete set of species that participate in the autocatalytic reactions. In Sec. 2.3, we extend their work by defining stoichiometric autocatalysis for autocatalytic subnetworks defined on the complete set of participating species in the reaction network. We show that for any such subnetwork, the species set can be uniquely assigned a *food-waste-member-core* partition, and define geometrically relevant quantities for them. We then derive mathematical results about these quantities, and provide an algorithm for exhaustively enumerating all minimal autocatalytic subnetworks (MASs). We also show how to take the list of MASs and obtain a visual summary of their combinatorics.

In Sec. 2.4, we define the *cluster chemical reaction network* (CCRN) framework. Since CCRNs are defined using conservation laws, we argue that this provides a natural framework for coarse-graining real CRNs. We then use the tools developed in Sec. 2.3 to organize a list of MASs for maximally connected 1-constituent CCRNs with L species and show that the number of MASs increases exponentially with L . While the fully connected CCRNs might not have physical relevance, in Sec. 2.4.3 we comment on how sparse CCRNs maybe obtained using rules induced by real polymer or biological chemistry (see Table 2.2).

Our work opens up a few lines of future research, namely:

1. For a simple network in Example 2.1, proven to be exclusively autocatalytic in Example 2.3, we proposed a modification such that the resulting graph was stoichiometrically autocatalytic in Example 2.4.

Is there a general algorithm to modify any exclusively autocatalytic motif to a stoichiometrically autocatalytic motif? Having such an algorithm would allow us to use the typology of (Blokhuys et al., 2020) and the results in this work to understand the autocatalytic properties of any CRN.

2. Within the CCRN framework, we restricted the clusters to have non-negative constituents. However, one can also consider the clusters defined on the complete integer lattice. The resulting reaction network could then be used to model real nuclear reactions. For example, if the positive conserved quantity corresponds to positive charge or subatomic particle type, the negative conserved quantity would refer to a negative charge or anti-particle, respectively.
3. The generation of MASs, for general CRNs, and especially for CCRNs, is computationally costly. Our optimization-based approach requires solving a series of binary mathematical programming problems with an increasing number of linear constraints. Further research includes the particular study of these optimization problems, in particular, its polyhedral properties that will allow to strengthen the formulation and solve them more efficiently. Exploiting the algebraic properties behind conservation laws of sequences of integer numbers would also lead us to understand the mathematical insights of these networks.

4. Realistic CRNs will not coarse-grain to a fully connected CCRN, thus appreciably reducing the number of reactions in the network, as well as the MASs that they contain. As a future investigation, similar to (Anderson and Nguyen, 2022), can we consider randomly generated CCRNs in an Erdos-Renyi framework and compute a bound on the number of autocatalytic cycles which persist as a function of the probability of a reaction being allowed in the network?
5. CCRNs are defined in terms of their conservation laws, and have more degrees of symmetry than generic CRNs. Having completely enumerated the stoichiometrically autocatalytic ecology of a complete CCRN, can we better understand and characterize their kinetics?

3 ACTION FUNCTIONAL GRADIENT DESCENT (AFGD)

ALGORITHM FOR ESTIMATING ESCAPE PATHS IN STOCHASTIC CHEMICAL REACTION NETWORKS

3.1 Introduction

Stochastic modelling has played an increasingly central role in science since the advent of high speed computing. There are, however, many categories of events that are vital to the organization of long-term dynamics for which the rate diminishes exponentially with system size ('rare' events) and computer simulations become unaffordable. A typical approach in these scenarios is to employ importance sampling approaches in the stochastic simulation to efficiently simulate the rare event of interest (see Hartmann and Schütte (2012); Cao and Liang (2013); Biondini (2015)). In this work, for the particular case of stochastic chemical reaction networks, we provide a deterministic alternative for estimating the likelihood of rare events.

Stochastic nonlinear dynamical processes exhibiting multistability, or multiple coexisting attractors, have found several scientific applications ranging from modelling climate change to population biology and the origin of life (Benzi et al., 1983; Dyson, 1982; Smith and Krishnamurthy, 2015; Smith and Morowitz, 2016), and are an active subject of interdisciplinary research. A question of critical importance for any practical application of such a system is, how often does it transition out of a stable attractor and

what are the least-improbable paths by which such a transition occurs? In literature, the dynamics that arise due to the system's transitioning from one stable attractor to another and the optimal path of transition is referred to as 'switching dynamics' and 'escape path', respectively.

Many phenomena in biology, chemistry, physics or engineering can be modelled as a chemical reaction network (CRN). It is also well known that CRNs can exhibit a host of dynamics, including limit cycles and multistability (Yu and Craciun, 2018). Perhaps less well-known is the role of Hamilton-Jacobi ray theory underlying stochastic CRN, and the interpretation of escape paths as particular characteristic curves that arise as a solution to the associated Hamilton-Jacobi equation. In our work, we review the necessary formalism and employ techniques from mathematical physics and numerical optimization to estimate the escape paths for a multistable CRN. More precisely, we first recognize that the escape paths are variational (locally least-action) solutions to the action functional, and subsequently use the functional to perform a gradient descent that converges on the desired path.

It must be noted that neither the recognition of the escape path as a variational solution nor using the action functional to perform gradient descent is novel to us (for instance, see Berne et al. (1998)), however, we differ from the earlier work in a few ways. In the past, the use of Hamilton-Jacobi theory to find the escape path for CRN has only been made using the 'shooting-method', which relies on integrating Hamilton's

equations of motion, rather than a functional gradient descent approach that we present here (see Dykman et al. (1994)). On the other hand, while the action functional has been used for a gradient descent, to our best knowledge, the form of the Hamiltonian assumed has always been separable in momentum and position (see Weinan et al. (2004); Bouchet et al. (2019)). This form of the Hamiltonian, with other constraints on the functional form, makes it amenable to finding an analytic form for the associated Lagrangian which drastically simplifies the gradient descent. While this assumption is typical for a Hamiltonian describing mechanical energy in physics, it is far too restrictive for a generic Hamiltonian in chemistry, as can be seen below:

$$\begin{aligned} H_{\text{ME}}(\mathbf{p}, \mathbf{q}) &= K(\mathbf{p}) + U(\mathbf{q}), \\ H_{\text{CRN}}(\mathbf{p}, \mathbf{q}) &= \sum_{\mathbf{y}_\alpha, \mathbf{y}_\beta} (e^{(\mathbf{y}_\beta - \mathbf{y}_\alpha) \cdot \mathbf{p}} - 1) k_{\mathbf{y}_\alpha \rightarrow \mathbf{y}_\beta} \mathbf{q}^{\mathbf{y}_\alpha}, \end{aligned} \quad (3.1)$$

where \mathbf{p} and \mathbf{q} are the momentum and position coordinates, $K(\mathbf{p})$ and $U(\mathbf{q})$ are the kinetic and potential energy, and H_{ME} and H_{CRN} (Eq. 3.16) are the Hamiltonians for mechanical energy and chemical reaction networks respectively. In our algorithm, however, we start from the Hamiltonian, numerically solve for the Lagrangian at each iteration and use it to calculate the descent direction. While our algorithm is designed for stochastic CRNs, it is amenable to generalization for other classes of stochastic Hamiltonian dynamical systems.

The layout of the chapter is as follows. In Section 3.2, we derive the Hamilton-Jacobi theory for stochastic processes starting from a master equation and apply it to stochastic CRN. We show that the Hamilton-Jacobi theory of the non-equilibrium potential (NEP) arises naturally as a result of a variational principle applied to the *action functional*. In Section 3.3, we propose an *action functional gradient descent* (AFGD) algorithm that finds the variational solution in the space of paths constrained at end points for a given Hamiltonian value constraint. We also explain how the algorithm can be used to find least-improbable escape paths out of a stable attractor and assign a value to the NEP along them. The details of the implementation are provided in Appendix B.4, and a MATLAB implementation is made available at (Gagrani, 2022). In Section 3.4, we demonstrate the applications of the AFGD algorithm on several CRNs. We first consider the Selkov model, and compare the result of the algorithm against the escape trajectory found by the ‘shooting-method’ (Figure 3.9). We then define a class of high dimensional birth-death models, namely ‘N-Schlögl model’, and compare the results of the algorithm against a stochastic simulation for the 2-Schlögl model (Figure 3.4). We also use the algorithm on the six dimensional 6-Schlögl model, and compare the result against the integration of Hamilton’s equations of motion (Figure B.7). Finally, in Section 3.5, we conclude with a discussion of our contribution and potential avenues of future research.

3.2 From master equation to Hamilton-Jacobi equation for CRN

The application of Hamilton-Jacobi theory to chemical reaction networks (CRNs) has a long history (Gang, 1987; Dykman et al., 1994). In this section, we review the necessary formulation needed to understand the switching dynamics of stochastic chemical reaction networks (Section 3.2.3) and to devise a variational algorithm for predicting the transitions (Section 3.3). In Section 3.2.1 we derive the non-equilibrium potential for continuous time Markov population processes and explain its role in estimating the probability of stochastic events. In Section 3.2.2, we derive the Hamiltonian for a CRN and prove that the NEP is a Lyapunov function along the deterministic trajectories under mass-action kinetics. In Appendix B.1, we investigate the relevance our work might have to a stochastic modelling practitioner by posing a general practical problem, giving an overall picture of the solution and explaining where our algorithm fits in. In Appendix B.2, we employ the Hamilton-Jacobi formalism to recover the ACK theorem and NEP for complex-balanced systems. In Appendix B.3, we define a ‘non-equilibrium action’, of which both the master equation and Schrödinger equation can be seen as a variational solution, derive the path integral formula and action functional for stochastic population processes, and calculate the first and second variational derivatives of the action functional. There are many resources that provide an

alternative treatment of the same subject, for instance see Snarski (2021); Smith (2020) and references therein.

3.2.1 Hamilton-Jacobi equation and Non-Equilibrium Potential (NEP) for stochastic dynamics

A continuous time Markovian population process is specified by its *master equation*,

$$\begin{aligned} \frac{\partial}{\partial t} \rho(\mathbf{n}, t) &= \sum_{\mathbf{n}'} \mathbb{T}_{\mathbf{n}\mathbf{n}'} \rho(\mathbf{n}', t) \\ \sum_{\mathbf{n}} \rho(\mathbf{n}, t) &= 1 \quad \text{for all } t \\ \sum_{\mathbf{n}} \mathbb{T}_{\mathbf{n}\mathbf{n}'} &= 0 \quad \text{for all } \mathbf{n}' \end{aligned} \tag{3.2}$$

where t is the time variable, \mathbf{n} is a D -dimensional discrete vector in the positive integer lattice $\mathbb{Z}_{\geq 0}^D$ denoting a position in the state space (to simplify notation we drop the arrow in $\vec{\mathbf{n}}$) and $\rho(\mathbf{n}, t)$ is a time-evolving probability distribution function (PDF) over the state space. \mathbb{T} is referred to as a transition operator, and in this chapter we only consider time independent transition operators.

For what follows, it is useful to recast Eq. 3.2 into the following form

and define a Hamiltonian operator \hat{H} that acts on the PDF ρ ,

$$\frac{\partial}{\partial t} \rho(n, t) = \hat{H} \left(-\frac{\partial}{\partial n}, n \right) \rho(n, t) \quad (3.3)$$

where $\partial/\partial n$ is the infinitesimal-shift operator¹.

The resemblance of Eq. 3.3 to the Schrödinger equation is not a mere coincidence, and we show in Section B.3.1 how they can both be derived from a common variational problem of extremizing, what is termed as, the ‘non-equilibrium action’ (NEA) in (Eyink, 1996). It should come as no surprise then that we can employ the same machinery developed in mathematical physics for quantum mechanics to derive results about stochastic dynamics. This line of reasoning has a long history and we refer the readers to (Doi, 1976; Baez and Biamonte, 2012; Smith and Krishnamurthy, 2015; Smith, 2019) for a ‘second-quantization’ treatment via abstract linear algebra.

We can integrate Eq. 3.3 starting from a PDF $\rho(n_0, 0)$ at time 0 to get a

¹A function in x can be shifted by y by the application of a *shift operator* given by $e^{y \frac{\partial}{\partial x}}$, as can be seen by

$$f(x + y) = \sum_{n=0}^{\infty} \frac{y^n}{n!} \frac{\partial^n f(x)}{\partial x^n} = e^{y \frac{\partial}{\partial x}} f(x).$$

distribution $\rho(\mathbf{n}_T, T)$ at time T , indexed by \mathbf{n}_0 and \mathbf{n}_T respectively.

$$\begin{aligned}\rho(\mathbf{n}_T, T) &= e^{\int_0^T \hat{\mathbf{H}} dt} \rho(\mathbf{n}_0, 0) \\ &= \int [\mathbf{d}\mathbf{n}] \int [\mathbf{d}\mathbf{p}] e^{-\mathcal{A}[\mathbf{n}(t), \mathbf{p}(t)]} \rho(\mathbf{n}_0, 0)\end{aligned}\quad (3.4)$$

where \mathcal{A} is the *action functional*

$$\mathcal{A}[\mathbf{n}(t), \mathbf{p}(t)] = \int_0^T \left[\mathbf{p} \cdot \frac{d\mathbf{n}}{dt} - \underline{\mathbf{H}}(\mathbf{p}, \mathbf{n}) \right] dt. \quad (3.5)$$

and \mathbf{p} is a momentum variable canonically conjugate to the position variable \mathbf{n} . The second line in Eq. 3.4 is the path-integral formula, a proof of which can be found in Appendix B.3.2, and $[\mathbf{d}\mathbf{n}] [\mathbf{d}\mathbf{p}]$ is the path-integral measure. Note that in Eq. 3.5, the Hamiltonian function $\underline{\mathbf{H}}(\mathbf{p}, \mathbf{n})$ has the same functional form as the Hamiltonian operator $\hat{\mathbf{H}}$ except the operator $-\partial/\partial\mathbf{n}$ is replaced by the variable \mathbf{p} .

In the first line of Eq. 3.4, the exponentiation of the Hamiltonian operator is to be understood as a time-ordered matrix product over small intervals dt . Its role is to time-evolve the initial distribution $\rho(0)$ and accumulate probability through all possible chains of states until time T so as to obtain a new distribution $\rho(T)$. In the second line, upon taking appropriate limits signifying a continuous time parameter, we obtain another description of the same process where we sum over all the paths starting at \mathbf{n}_0 at time 0 and ending at \mathbf{n}_T at time weighted with an appropriate

measure. The negative log of the measure of a particular path is given by the action functional (in other words, \mathcal{A} takes as input a path and returns its log-improbability), and integrating under the path-integral measure amounts to summing over all paths. For a didactic introduction to the topic and its application to stochastic dynamics, we refer the reader to (Baez and Biamonte, 2012).

It is useful to pass from a discrete to a continuous state space by descaling n with a scale factor V and considering the large V limit. In order to align our presentation here with the mathematical physics literature, we denote the continuous variable by $q = n/V$ and refer to the continuous state space as configuration space. There are several interpretations the variables can take; the one of interest to us is where n is a population vector (thus $n \geq 0$), q is the concentration vector and V is a scale of the the total population. Following Eq. 3.4, we can write the time evolution of a distribution in the transformed coordinate q as

$$\begin{aligned} \rho(q, T) &= \int [dq] \int [dp] e^{-V\mathcal{A}[q(t), p(t)]} \rho(q(0), 0) \\ \mathcal{A}[q(t), p(t)] &= \int_0^T \left[p \cdot \frac{dq}{dt} - H(p, q) \right] dt, \end{aligned} \quad (3.6)$$

where \mathcal{A} is the action functional in the new coordinates and has a different

Hamiltonian $H(p, q)$ which relates to $\underline{H}(p, n)$ by ²

$$VH(p, q) \asymp \underline{H}(p, qV) = \underline{H}(p, n). \quad (3.7)$$

It must be noted that Eq. 3.7 is an additional constraint on the form of the Hamiltonian, but is one that is satisfied by the CRN Hamiltonian which we are interested in for the scope of this chapter. Moreover, if the asymptotic equality was an equality then the condition is often referred to by saying ‘the Hamiltonian is extensive in the position coordinate’, or ‘the Hamiltonian is a homogenous function of degree one in the position coordinate’. We will return to this point in the next subsection, once we write the explicit Hamiltonian for a CRN.

Together the position and momentum coordinates (q, p) constitute a point in the *phase space* (Arnol’d, 2013). In order to evaluate the distribution at time T using Eq. 3.6, in principle we need to sum over all the phase space paths that end at q at time T . Evaluating the full sum, however, is not necessary to obtain the leading large- V asymptotics. Through Laplace’s or saddle-point method, for large V the path integral is dominated by the saddle point of the functional and the distribution at time T can be

²We say $f(V) \asymp g(V)$ (read as $f(V)$ is asymptotic to $g(V)$) if

$$\lim_{V \rightarrow \infty} \frac{f(V)}{g(V)} = 1.$$

approximated by

$$\rho(q, T) \simeq N(V) e^{-V\mathcal{A}[q^*(t), p^*(t)]} \rho(q^*(0), 0)$$

where $(q^*(t), p^*(t))$ is such that

$$\delta\mathcal{A}[q^*(t), p^*(t)] = 0 \text{ and } q^*(T) = q, \quad (3.8)$$

and $N(V)$ is a normalization factor expounded upon in the subsequent paragraph. The constraint in the last line specifies the stationarity or optimality condition on the saddle point path $(q^*(t), p^*(t))$, also referred to as the path of stationary action, and we henceforth refer to it simply as the *optimal path* (for a MinMax formulation of the saddle-point see Section 3.3.1). We require here that the Hamiltonian function is convex in the momentum variable p , which we show is the case for CRN Hamiltonians with mass-action kinetics in Section 3.2.2.

There are two caveats to Eq. 3.8 that we now point out. First, the normalization factor $N(V)$ in the first line depends not only on the scale factor V but also on the optimal path. There is an in-principle method, although costly in practice, to obtain $N(V)$ from the action functional itself, outlined in (Kirsten and McKane, 2003), Ch. 7 of (Coleman, 1988) and Ch. 4 of (Smith and Krishnamurthy, 2015) (for an application, see Manikandan and Krishnamurthy (2017)). Second, since the path integral is a sum over all paths, there are several stationary paths starting from different $q^*(0)$ that reach q at T , and one must perform a sum over all

of them. This becomes increasingly relevant in the limit $T \rightarrow \infty$, where the system can bounce back and forth multiple times between a $q^*(0)$ and q (see Smith and Krishnamurthy (2015); Coleman (1988)). There is, however, a two-fold resolution to this problem. First, the normalization factor $N(V)$ is sub-exponential in V , and can be well-approximated by unity large for V (also see the introduction to Section 3.3). Second, the initial distribution $\rho(q, 0)$ is typically taken to be peaked at a given value, thus yielding negligible contribution from every $q^*(0)$ that is not near the peak of the initial condition. Thus we only pick the least-improbable direct paths that start from near the peak of the initial distribution and reach q in Eq. 3.8.

In the first line of Eq. 3.8, the only free variables are configuration q and final time T , which together characterize a time evolving distribution. The optimality condition picks a path $(q^*(t), p^*(t))$, and it is only the probability mass at $q^*(0)$ in the initial distribution at time 0 i.e. $\rho(q^*(0), 0)$, weighted appropriately by the action functional, that contributes to the final distribution at configuration q and time T . The weight factor $e^{-V\mathcal{A}[q^*, p^]}$ accounts for the fact that not all the probability mass from $q(0)$ goes to q in time T , and in this sense takes the interpretation of a conditional probability.

Since the conditional probability term is only in the exponential, it is useful to write the probability distribution itself as an exponential function $\rho(q, t) = e^{-VS(q, t)}$. Here $S(q, t)$ is the *action function* (not be confused with

\mathcal{A} which is the action functional) and its time evolution is given as follows

$$\begin{aligned} S(q, T) &= \mathcal{A} [q^*(t), p^*(t)] + S(q^*(0), 0) \\ &= \int_0^T \left[p^* \cdot \frac{dq^*}{dt} - H(p^*, q^*) \right] dt + S(q^*(0), 0) \end{aligned} \quad (3.9)$$

where in the second line, we expand the action functional using Eq. 3.6. Thus the action function evolves by the optimal value of the action functional, and since it is a function we can also consider its total differential

$$\begin{aligned} S(q, T) &= \int_0^T dS + S(q^*(0), 0) \\ &= \int_0^T \left[\frac{\partial S(q^*, t)}{\partial q^*} \cdot \frac{dq^*}{dt} + \frac{\partial S(q^*, t)}{\partial t} \right] dt + S(q^*(0), 0). \end{aligned} \quad (3.10)$$

Comparing equations 3.9 and 3.10, we get

$$\begin{aligned} \frac{\partial S(q^*, t)}{\partial q^*} &= p^* \\ \frac{\partial S(q^*, t)}{\partial t} &= -H(p^*, q^*) \\ \frac{\partial S(q^*, t)}{\partial t} &= -H \left(\frac{\partial S}{\partial q^*}, q^* \right) \end{aligned} \quad (3.11)$$

where the first two lines are obtained by comparison and the last line is obtained by substituting the first line in the second line. The non-linear PDE in the last equation is called the **Hamilton-Jacobi equation**. For an introduction to the subject and its history, we refer the readers to (Goldstein et al., 2002; Arnol'd, 2013) (what we call as the action function is

also called Hamilton's principal function). We briefly remark that the definite integral of the action functional along the optimal path $\int_0^T dS$ also acts as a divergence function in information geometry (see Leok and Zhang (2017); Smith (2019)) and is denoted as $S(q^*(T)||q^*(0))$ yielding the relation $S(q_T^*) = S(q_T^*||q_0^*) + S(q_0^*)$, as mentioned in Eq. B.2.

The role of *momentum* for stochastic systems can be best understood by the Hamilton Jacobi equation. The first line in Eq. 3.11 shows that momentum is the gradient of the action function, which in turn is descaled negative log probability of the time evolving distribution given some initial conditions. Using this, we see that $p^*(0)$ must be the gradient of the action function S at $q^*(0)$, and that an optimal path in phase space corresponds to a contour in the probability distribution arising by scaling and exponentiating the action.

The third line in Eq. 3.2 ensures that the transition operator always has a cokernel, namely the row vector consisting of all ones $[1, \dots, 1]$. The Perron-Frobenius theorem for stochastic matrices proves that the operator also has a unique right kernel, referred to as the stationary distribution which we denote here by π . π thus satisfies $\sum_n \mathbb{T}_{n'n} \pi(n, t) = 0$ or equivalently

$$\hat{H}\pi(n, t) = 0.$$

From Eqs. 3.2, 3.3 one can see that the stationary distribution is indeed time independent as $\partial\pi(n, t)/\partial t = 0$, thus justifying its name.

The descaled log-improbability or action function of the stationary distribution is commonly referred to as the **non-equilibrium potential (NEP)** and denoted by \mathcal{V} (see Anderson et al. (2015)),

$$\pi(\mathbf{n}) \asymp e^{-\mathcal{V}(\mathbf{q})}, \quad (3.12)$$

where the normalization constant has been set to unity following the discussion below Eq. 3.8. As the action of a stationary distribution, it must satisfy

$$\begin{aligned} \frac{\partial \mathcal{V}}{\partial t} &= 0, \\ \text{H} \left(\frac{\partial \mathcal{V}}{\partial \mathbf{q}}, \mathbf{q} \right) &= 0 \end{aligned} \quad (3.13)$$

where the second line is the Hamilton-Jacobi equation from Eq. 3.11 that must be satisfied by the NEP.

In the remainder of this section, we will develop and demonstrate the applications of Hamilton-Jacobi theory for CRNs. Finding the NEP \mathcal{V} for a general CRN is of importance for getting any numerical estimates on its behavior, and it is the task that we will concern ourselves with in this chapter starting Section 3.3. As explained above, to calculate the difference of the action function or NEP between two points we first need to find the optimal path connecting them, which is precisely what the Action Functional Gradient Descent (AFGD) algorithm is designed to do.

3.2.2 Hamilton-Jacobi theory for stochastic Chemical Reaction Networks (CRN)

In this subsection, we start by introducing the mathematical formulation of CRNs and derive its Hamiltonian function in the concentration coordinate. Next, we show that the Hamiltonian function is convex in the momentum coordinate, a necessary condition for the existence of optimal paths, and write the equations that the optimal paths must satisfy. Then we classify the optimal paths constituting the stationary distribution into two categories, namely relaxation paths and escape paths, and show that the former yield the deterministic mass-action kinetics (MAK) and the latter yields the non-equilibrium potential (NEP). Finally we show that, for a general CRN with multiple fixed points, the NEP is always a Lyapunov function with respect to MAK.

A CRN is defined by the triple $\{\mathcal{S}, \mathcal{C}, \mathcal{R}\}$, where $\mathcal{S}, \mathcal{C}, \mathcal{R}$ are the set of species, complexes and reactions respectively.

$$\begin{aligned}\mathcal{S} &= \{S_1, \dots, S_i, \dots, S_{|\mathcal{S}|}\} \\ \mathcal{C} &= \{y_1, \dots, y_\alpha, \dots, y_{|\mathcal{C}|} : y_\alpha \in \mathbb{N}^{|\mathcal{S}|}\} \\ \mathcal{R} &= \{y_{ff} \xrightarrow{k_{y_{ff} \rightarrow y_{fi}}} y_{fi} : k_{y_\alpha \rightarrow y_\beta} \geq 0\},\end{aligned}$$

where Roman letters (i, j) and Greek letters (α, β) are used to denote

species and complex indices, respectively. A complex is a multi-set of species, and is denoted by the column vector \mathbf{y}_α representing the stoichiometry of the multi-set. The state of a CRN is characterized by a population vector $\mathbf{n} \in \mathbb{Z}_{\geq 0}^{|\mathcal{S}|}$, and the time evolution of a probability distribution over the state space is given by (for a rigorous microphysical derivation, see Gillespie (1992))

$$\begin{aligned} \frac{d\rho(\mathbf{n}, t)}{dt} &= \sum_{\mathbf{y}_\alpha, \mathbf{y}_\beta} \frac{k_{\mathbf{y}_\alpha \rightarrow \mathbf{y}_\beta}}{V^{y_\alpha - 1}} \cdot \left(\prod_{i=1}^{|\mathcal{S}|} \frac{(\mathbf{n}_i - \mathbf{y}_{\beta,i} + \mathbf{y}_{\alpha,i})!}{(\mathbf{n}_i - \mathbf{y}_{\beta,i})!} \rho(\mathbf{n} - \mathbf{y}_\beta + \mathbf{y}_\alpha, t) - \prod_{i=1}^{|\mathcal{S}|} \frac{\mathbf{n}_i!}{(\mathbf{n}_i - \mathbf{y}_{\alpha,i})!} \rho(\mathbf{n}, t) \right) \\ &= \sum_{\mathbf{y}_\alpha, \mathbf{y}_\beta} \frac{k_{\mathbf{y}_\alpha \rightarrow \mathbf{y}_\beta}}{V^{y_\alpha - 1}} \left(\frac{(\mathbf{n} - \mathbf{y}_\beta + \mathbf{y}_\alpha)!}{(\mathbf{n} - \mathbf{y}_\beta)!} \rho(\mathbf{n} - \mathbf{y}_\beta + \mathbf{y}_\alpha, t) - \frac{\mathbf{n}!}{(\mathbf{n} - \mathbf{y}_\alpha)!} \rho(\mathbf{n}, t) \right) \end{aligned} \quad (3.14)$$

where we get the last equation by making use of the multi-index notation.

Writing Eq. 3.14 in the form of equation Eq. 3.3 by making use of the shift operator (see footnote 1), we get

$$\begin{aligned} &\frac{d\rho(\mathbf{n}, t)}{dt} \\ &= \sum_{\mathbf{y}_\alpha, \mathbf{y}_\beta} \left(e^{-(\mathbf{y}_\beta - \mathbf{y}_\alpha) \cdot \frac{\partial}{\partial \mathbf{n}}} - 1 \right) \frac{k_{\mathbf{y}_\alpha \rightarrow \mathbf{y}_\beta}}{V^{y_\alpha - 1}} \frac{\mathbf{n}!}{(\mathbf{n} - \mathbf{y}_\alpha)!} \rho(\mathbf{n}, t) \\ &\equiv \hat{\mathbf{H}}_{\text{CRN}} \left(-\frac{\partial}{\partial \mathbf{n}}, \mathbf{n} \right) \rho(\mathbf{n}, t) \end{aligned}$$




Figure 3.1: Diagrammatic representation of 1-Schlögl, 2-Schlögl, 3-Schlögl and 4-Schlögl model CRNs from top left to bottom right. Species, complexes, stoichiometry and transitions are represented with solid circles, empty circles, solid lines and dashed arrows, respectively.

and obtain the following Hamiltonian operator for CRN

$$\hat{\underline{H}}_{\text{CRN}} = \sum_{y_\alpha, y_\beta} \left(e^{-(y_\beta - y_\alpha) \cdot \frac{\partial}{\partial \mathbf{n}}} - 1 \right) \frac{k_{y_\alpha \rightarrow y_\beta}}{V^{y_\alpha - 1}} \frac{\mathbf{n}!}{(\mathbf{n} - y_\alpha)!}. \quad (3.15)$$

Following the derivation in Section B.3.2, we can replace the differential operator $-\partial/\partial \mathbf{n}$ with momentum variable \mathbf{p} to obtain the Hamiltonian function

$$\underline{H}_{\text{CRN}}(\mathbf{p}, \mathbf{n}) = V \sum_{y_\alpha, y_\beta} \left(e^{(y_\beta - y_\alpha) \cdot \mathbf{p}} - 1 \right) \frac{k_{y_\alpha \rightarrow y_\beta}}{V^{y_\alpha - 1}} \frac{\mathbf{n}!}{(\mathbf{n} - y_\alpha)!}.$$

Finally, using the observation $\frac{(qV)!}{(qV - y)!} \asymp (qV)^y$, we can pass to the Hamil-

tonian function in the concentration variable $q = n/V$ using Eq. 3.7 to get

$$H_{\text{CRN}}(\mathbf{p}, \mathbf{q}) = \sum_{\mathbf{y}_\alpha, \mathbf{y}_\beta} (e^{(\mathbf{y}_\beta - \mathbf{y}_\alpha) \cdot \mathbf{p}} - 1) k_{\mathbf{y}_\alpha \rightarrow \mathbf{y}_\beta} \mathbf{q}^{\mathbf{y}_\alpha}. \quad (3.16)$$

For the Hessian of the Hamiltonian function in the momentum variable \mathbf{p} , we get

$$\frac{\partial H_{\text{CRN}}}{\partial \mathbf{p}^2} = \sum_{\mathbf{y}_\alpha, \mathbf{y}_\beta} (\mathbf{y}_\beta - \mathbf{y}_\alpha)(\mathbf{y}_\beta - \mathbf{y}_\alpha)^T e^{(\mathbf{y}_\beta - \mathbf{y}_\alpha) \cdot \mathbf{p}} k_{\mathbf{y}_\alpha \rightarrow \mathbf{y}_\beta} \mathbf{q}^{\mathbf{y}_\alpha}. \quad (3.17)$$

It can be seen that Hessian of the Hamiltonian in \mathbf{p} is positive definite as it is a sum of symmetric dyadics, each of which is positive definite. Thus, as commented in the discussion below Eq. 3.8, the Hamiltonian function for a CRN is indeed convex in the momentum variable.

As noted in the previous subsection, the convexity of the Hamiltonian in the momentum coordinates allows us to find stationary or optimal paths where the variation of the action functional vanishes. The optimal phase space paths $(\mathbf{q}(t), \mathbf{p}(t))$ satisfy the Hamilton's equations of motion (EoM)

$$\begin{aligned} \frac{d\mathbf{q}}{dt} &= \frac{\partial H}{\partial \mathbf{p}}, \\ \frac{d\mathbf{p}}{dt} &= -\frac{\partial H}{\partial \mathbf{q}} \end{aligned} \quad (3.18)$$

where we have dropped the * in denoting the optimal path for simplifying

notation. For a derivation, see Section B.3.3, where we also show that the Hamiltonian is constant along the optimal path and the equivalence of Hamilton's EoM with the Hamilton-Jacobi equations. Intuitively, the ODEs of Hamilton's EoM are the characteristic curves of the Hamilton-Jacobi PDE, and thus the union of all solutions to Hamilton's EoM form the complete solution to the Hamilton-Jacobi PDE.

Substituting the CRN Hamiltonian from Eq. 3.16 into Eq. 3.18 and omitting the species index for notational clarity, we get

$$\begin{aligned}\frac{dq}{dt} &= \sum_{y_\alpha, y_\beta} (y_\beta - y_\alpha) (e^{(y_\beta - y_\alpha) \cdot p}) k_{y_\alpha \rightarrow y_\beta} q^{y_\alpha} \\ \frac{dp}{dt} &= - \sum_{y_\alpha, y_\beta} (e^{(y_\beta - y_\alpha) \cdot p} - 1) k_{y_\alpha \rightarrow y_\beta} \frac{\partial q^{y_\alpha}}{\partial q},\end{aligned}\quad (3.19)$$

which are the set of ODEs satisfied by the optimal paths for a CRN.

Let us begin investigating the optimal paths specified by the above equations. For reasons mentioned in Eq. 3.13, we only consider the subclass of optimal paths in the $H_{\text{CRN}} = 0$ submanifold. For the remainder of this section, we drop the subscript CRN and refer to the CRN Hamiltonian simply by H .

Notice that for $p = 0$, the Hamiltonian is identically zero i.e. $H(p =$

$0, q) = 0$ for all q . Substituting this assignment in Eq. 3.19, we get

$$\begin{aligned} \left. \frac{dp}{dt} \right|_{p=0} &= 0, \\ \left. \frac{dq}{dt} \right|_{p=0} &= \sum_{y_\alpha, y_\beta} (y_\beta - y_\alpha) k_{y_\alpha \rightarrow y_\beta} q^{y_\alpha}. \end{aligned} \quad (3.20)$$

The first line ensures that, since the time derivative of p vanishes, $p = 0$ is a consistent assignment everywhere along the optimal path. The second line is nothing but the set of ODEs corresponding to the law of mass-action kinetics for a CRN. Thus the optimal paths with $p = 0$ correspond to the deterministic trajectories of a CRN, which we term **relaxation trajectories** and denote as $q_{\text{rel}}(t)$.

To understand why the $p = 0$ solution must correspond to the deterministic behaviour of CRN in the Hamilton-Jacobi formalism, let us calculate the change in action along the relaxation trajectory. From Eq. 3.9,

$$\begin{aligned} \int_0^T dS[q_{\text{rel}}] &= \int_0^T \left[0 \cdot \frac{dq_{\text{rel}}}{dt} - H(0, q_{\text{rel}}) \right] dt \\ &= 0. \end{aligned} \quad (3.21)$$

Recall from Eq. 3.8 that the exponential of the scaled change in action acted as a conditional probability up to a normalization factor. Since the change in action is identically zero, almost all of the probability measure

without any suppression evolves along the relaxation trajectory yielding

$$\rho(\mathbf{q}_{\text{rel}}(T), T) = N\rho(\mathbf{q}_{\text{rel}}(0), 0).$$

It is quite common for the deterministic dynamics of a CRN, Eq. 3.20, to exhibit multiple fixed points. Let us denote a fixed point by \underline{q} . Then each fixed point must satisfy

$$\begin{aligned} \frac{d\mathbf{q}}{dt} = 0 &= \sum_{\mathbf{y}_\alpha, \mathbf{y}_\beta} (\mathbf{y}_\beta - \mathbf{y}_\alpha) k_{\mathbf{y}_\alpha \rightarrow \mathbf{y}_\beta} \underline{q}^{\mathbf{y}_\alpha} \\ &= \sum_{\mathbf{y}_\alpha} \mathbf{y}_\alpha \sum_{\mathbf{y}_\beta} (k_{\mathbf{y}_\beta \rightarrow \mathbf{y}_\alpha} \underline{q}^{\mathbf{y}_\beta} - k_{\mathbf{y}_\alpha \rightarrow \mathbf{y}_\beta} \underline{q}^{\mathbf{y}_\alpha}). \end{aligned} \quad (3.22)$$

The existence of multiple steady states can be ruled out in many cases by certain topological considerations and is studied in deficiency theory. For an introduction and detailed derivation of some results, we refer the reader to (Feinberg and Horn, 1974; Feinberg, 2019; Smith and Krishnamurthy, 2017, 2021).

For our purposes in this chapter, we restrict our attention to generic CRNs with multiple fixed points. These points can be indexed by the number of repelling directions, which can easily be determined using the eigenvalues of the Jacobian of the flow field in Eq. 3.20. A fixed point with no repelling or all attracting directions is referred to as a *stable attractor*. In a purely deterministic setting, one would expect any evolving probability distribution to end up concentrated at the multiple stable attractors with

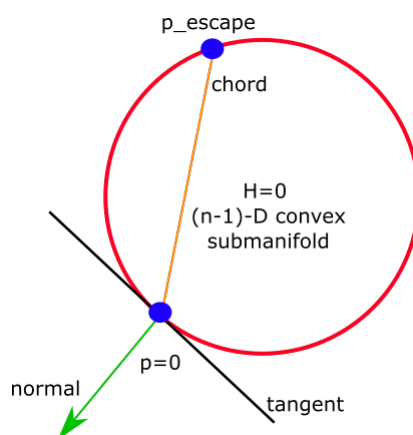


Figure 3.2: Diagrammatic proof of the Lyapunov property of the non-equilibrium potential with respect to mass action kinetics.

ratios determined by the initial conditions. In a stochastic setting however, there can be large fluctuations that take the system out from the vicinity of a stable attractor to a generic point $\underline{q} \neq \underline{q}$ in its basin of escape. The path along which the system arrives at \underline{q} with a leading exponential probability will be the optimal path that starts from \underline{q} and ends at \underline{q} , and the probability of this event is precisely given by the action function along that path using 3.8. We call the optimal paths out of a fixed point in the $H(p, q) = 0$ submanifold as **escape trajectories**, denoted by q_{esc} , and from eqs 3.12, 3.13 it is precisely the action function along paths that determine the value of the NEP,

$$\begin{aligned} \mathcal{V}(\underline{q}) - \mathcal{V}(\underline{q}) &= \int_{q_{\text{esc}}}^{\underline{q}} p_{\text{esc}}(q') \frac{dq'}{dt} dt \\ &= \int_{\underline{q}}^{\underline{q}} p_{\text{esc}}(q') dq' \end{aligned} \quad (3.23)$$

where $(q_{\text{esc}}, p_{\text{esc}})$ is the optimal path that connects \underline{q} and q . Note that the momentum assignment along the escape path must necessarily be non-zero, except at the fixed points, and the Hamiltonian function must evaluate to zero everywhere along the optimal path.

$$\begin{aligned} p_{\text{esc}}(q) &\equiv p(q_{\text{esc}}) \neq 0 \\ H(p_{\text{esc}}, q_{\text{esc}}) &= 0 \end{aligned} \tag{3.24}$$

To understand why the momentum assignment along escapes must be non-zero everywhere except the fixed points and to prove that the NEP is a Lyapunov function with respect to the relaxation trajectories (mass-action kinetics), let us investigate the geometry of the $H(p, q) = 0$ submanifold in the $2|\mathcal{S}|$ dimensional phase space. Since we are considering a CRN with fixed points, there must exist shift vectors $y_\beta - y_\alpha$ such that $H \rightarrow \infty$ when $|p| \rightarrow \infty$ in all directions. Also, recall from Eq. 3.17 that for a fixed concentration q , the Hamiltonian is a convex function in p . This means the $H(p, q) = 0$ submanifold at a given q must be a closed surface and bounded in the interior with $H < 0$. Next, let us consider the outward normal to the surface given by $\partial H / \partial p$. Recall from Hamilton's equations from Eq. 3.19 that the outward normal is equal to the rate of change of concentration. Since by the fixed point condition

$$\frac{\partial H(0, \underline{q})}{\partial p} = 0, \tag{3.25}$$




Figure 3.3: Escape paths and heteroclinic network for the 2-Schlögl model.

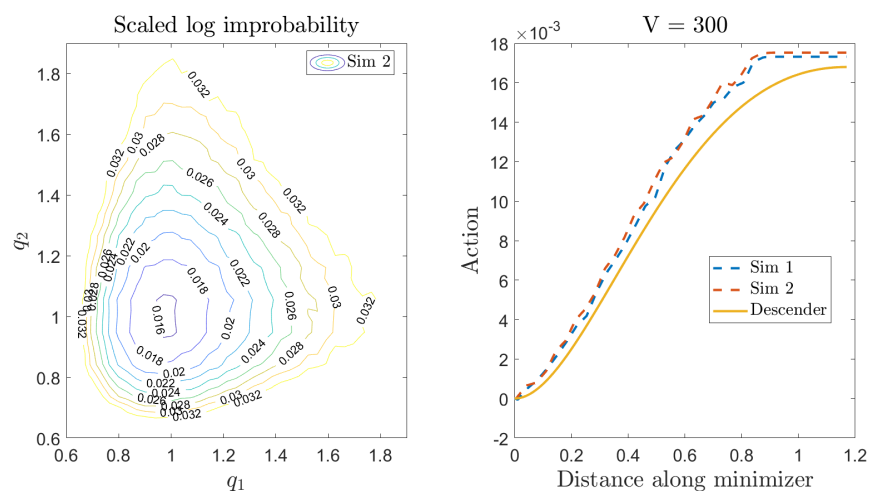


Figure 3.4: Comparing the algorithm against Gillespie stochastic simulation for the 2-Schlögl model (Section 3.4.2). Scaled log-improbability of the stationary distribution π is defined to be $(-1/V) \log(\pi)$.

this means that the outward normal to the surface vanishes and thus the $H(p, q) = 0$ submanifold must pinch off at the fixed points yielding $p = 0$ as the only solution in the submanifold.

Finally, to show that the NEP is a Lyapunov function along the deterministic MAK, let us consider its time derivative along a relaxation

trajectory.

$$\begin{aligned}
\frac{d\mathcal{V}(q_{\text{rel}}(t))}{dt} &= \frac{\partial\mathcal{V}(q)}{\partial q} \cdot \frac{dq}{dt} \Big|_{\text{rel}} \\
&= p_{\text{esc}}(q) \cdot \frac{dq}{dt} \Big|_{p=0} \\
&= (p_{\text{esc}} - 0) \cdot \frac{\partial H}{\partial p} \Big|_{p=0} \\
&\leq 0
\end{aligned} \tag{3.26}$$

where the first line follows from chain rule, the second line uses the Hamilton-Jacobi relation of momentum and NEP from Eq. 3.23, and the third line follows from Eq. 3.20. The last line follows from the fact that the gradient of a convex function along a chord in the convex subset (secant) is non-positive, or equivalently the outward normal always makes more than 90° with a secant, shown diagrammatically in Figure 3.2.

3.2.3 Switching dynamics for a multistable stochastic

CRN

Let us refer to the set of attractors (fixed points) of a CRN, when taken under MAK, by $\underline{q} \equiv \{q_1, q_2, \dots, q_N\}$. Now, label the attractors by the number of positive eigenvalues of the Jacobian of the MAK rates or the mixed-Hessian of the Hamiltonian $\partial^2 H(p, \underline{q}) / \partial p \partial q$ evaluated at $p = 0$. Recall that the positive eigenvalues give the number of repelling directions

around the attractor. Let us refer to the fixed points with zero and one repelling directions as *stable* and *saddle* attractors, respectively.

For large scale factor V , the switching dynamics of the stochastic CRN will be generally governed by the stable and saddle attractors. Recall from Section 3.2.1 that the most likely or optimal paths for a stochastic trajectory lie in the $H = 0$ submanifold. Moreover, in Section 3.2.2, we showed that there are both trajectories leading into a stable attractor (termed relaxation trajectories) and coming out of a stable attractor (escape paths). In particular then, corresponding to relaxation trajectories that emanate nearby a saddle attractor into a stable attractor, there will also be escape paths emanating from the stable attractor and terminating at the saddle attractor. The data obtained by collecting all the stable and saddle attractors and the relaxation and escape paths joining them is called as the *heteroclinic network* for the CRN. For example, the heteroclinic network obtained for the 2-Schlögl model from Figure 3.1 is shown in Figure 3.3 (also see Section 3.4.2).

Generically, one expects a stochastic CRN to spend a long time near a stable attractor before transitioning to a subsequent one by passing a saddle attractor in between. Recall that the NEP or action along an escape trajectory (see Section 3.2.1) quantifies the scaled log-improbability of the escape event. Moreover, as we have shown in Section 3.2.2, the NEP is a Lyapunov function along MAK. This means that the NEP must monotonically decrease along the relaxation trajectory, or correspondingly monotonically

increase along the escape trajectory. For example, the NEP for the 2-Schlögl model along all the escape heteroclinic orbits is shown in the right panel of Figure 3.3. Moreover, in Figure 3.4, the NEP is shown for a particular escape and compared against the occupation log-improbability probability obtained by a Gillespie simulation.

We briefly want to remark that since the NEP increases along an escape trajectory, the quantity must not be seen as an *entropy* for a stochastic system. This viewpoint will naturally lead to paradoxical violations of the second law of thermodynamics. However, in (Smith, 2020), it is shown that for an appropriate definition of entropy for stochastic CRN, its rate of change for any system can be asymptotically estimated by the rate of change of the NEP along a relaxation trajectory, thus yielding a rigorous *second law* for CRN.

3.3 Action Functional Gradient Descent (AFGD) algorithm

A central property of chemical reaction networks (CRNs) that makes them particularly useful for modeling complex systems is their ability to exhibit multiple attractors when taken under deterministic mass-action kinetics. Moreover, in a stochastic CRN, due to fluctuations in population sizes, it is possible for the system to transition out of a stable attractor into another with some probability. It is of practical value to quantify the

Object	Notation
Configuration space	$Q \subset \mathbb{R}^D$
Space of curves	$\mathcal{P}_Q = \{q : [0, 1] \rightarrow Q, q(0) = q_I, q(1) = q_F\}$
Curve	$q \in \mathcal{P}_Q$
Phase space	$T^*Q \subset \mathbb{R}^{2D}$
Space of trajectories	$\mathcal{P}_{T^*Q} = \{(q, p) : [0, T] \rightarrow T^*Q, q^*(0) = q_I, q^*(T) = q_F\}$
Trajectory	$(q, p) \equiv \gamma \in \mathcal{P}_{T^*Q}$

Table 3.1: Notation used for denoting paths in configuration and phase space.

transition events, termed as *escapes*, and can help in experiment design as well as network inference (Langary and Nikoloski, 2019). However, the escapes are rare, scaling exponentially in $-V$, compared to fluctuations in population that are polynomial in $1/V$, making stochastic simulation an inefficient approach to estimate the leading consequences of escapes, which dominate basin-switching.

In this section, we provide a deterministic alternative to estimating the escape paths between appropriate start and end points in chemical concentration space using a variational method. Specifically, we present a functional gradient descent algorithm that finds the optimal trajectory, that minimizes the action functional while satisfying the boundary conditions. A mathematical formulation of the optimization problem is posed in Section 3.3.1 and the algorithm is explained in Section 3.3.2, while relegating the technical details to Appendix B.4. A discussion of the algorithm’s performance costs and its relation to other methods is presented in Section 3.3.3.

Finally, we will briefly remark on obtaining the probability of system transition from the escape path. First, notice that if the boundary conditions are such that the trajectory emanates from a stable fixed point and terminates in a saddle fixed point, then the resulting escape trajectory does not lie in the solution set of the deterministic mass-action kinetics, and can only be interpreted in the stochastic framework (see Section 3.2.2). Let the phase space trajectory of such an escape be denoted by γ^* . For reasons explained below Eq. 3.8, the likelihood of a system of scale factor V starting from the stable fixed point and escaping to the saddle fixed point can be well-approximated simply by $e^{-V\mathcal{A}[\gamma^*]}$ (see Eq. 3.12). For example, we choose $V = 300$ in Figure 3.4, the authors consider values of V as low as 10 in Figure 4 of (Smith and Krishnamurthy, 2017), and in both cases the occupation probability from the stochastic simulation data is well in agreement with the exponential of the scaled action along the escape trajectory.

3.3.1 Formulation as a MinMax problem

We refer to the space of states in which we can find our chemical system as *configuration space*, denoted by Q . Suppose we have $|\mathcal{S}| = D$ species in our network, then the state of a system at a given time is a vector of concentration of each species $\in \mathbb{R}_{\geq 0}^D$. The space of paths that a system can take in configuration space is denoted by \mathcal{P}_Q , and we denote a path by $q(s)$. Notice that although the space of parametrized differentiable paths

lies in the tangent bundle of Q (see Milnor et al. (1963)), the parametrization is arbitrary and we do not yet have a notion of physical time and velocity. To introduce a notion of physical time, we have to introduce a Hamiltonian function (see Eqs. 3.6, 3.16) defined on the cotangent bundle of Q , commonly referred to as the *phase space*. The coordinates canonically conjugate to configuration coordinates q in the cotangent space are also called as momentum coordinates and denoted by p . Moreover, since the Hamiltonian is convex in momentum, every path in the configuration space has a unique lift in phase space parametrized by time variable t , which we denote by $\gamma \equiv (q(t), p(t))$. To avoid confusion, henceforth we refer to a path in configuration space in \mathcal{P}_Q as a *curve*, and phase space path in \mathcal{P}_{T^*Q} as a *trajectory*.

In the notation described in Table 3.1, we can succinctly formulate the problem that we wish to devise an algorithm to solve. We wish to find the optimal phase space trajectory, i.e. a trajectory that minimizes the action functional, constrained such that its projection to the configuration space curve begins at configuration q_I and ends at q_F , and the phase space trajectory is along the $H(p, q) = 0$ submanifold (reason for the $H(p, q) = 0$

constraint can be found in discussion around Eq. 3.13 in Section 3.2.1).

$$\begin{aligned}
 & \text{Find } \gamma^* \equiv (q^*, p^*) \in \mathcal{P}_{T^*Q} \\
 & \text{such that } \mathcal{A}[\gamma^*] = \min_{q(t)} \max_{p(t)} \int_0^T \left(p \cdot \frac{dq}{dt} - H(p, q) \right) dt, \\
 & H(p^*(t), q^*(t)) = 0, \\
 & q(0) = q_I, q(T) = q_F.
 \end{aligned} \tag{3.27}$$

Once the escape path from q_I to q_F is identified, the difference in the non-equilibrium potential between the two configurations is given by value of the action along the path, i.e.

$$\mathcal{V}(q_F) - \mathcal{V}(q_I) = \mathcal{A}[\gamma^*]. \tag{3.28}$$

For reasons discussed in Section 3.2.3, we are only interested in finding escape paths from a stable to a nearby saddle fixed point, thus we will choose the end points to be in the set of fixed points of the mass-action kinetics \underline{q} , as defined in Eq. 3.22. Moreover, from Eq. 3.25, we know that the $H(p, q) = 0$ phase space submanifold pinches off and the only solution to $H(p, \underline{q}) = 0$ is $p(\underline{q}) = 0$. Thus, for our purposes here, the momentum end point constraints are $p(q_I) = 0 = p(q_F)$.




Figure 3.5: Lifting a curve in configuration space to a trajectory in phase space by assigning a momentum and time difference to each segment.

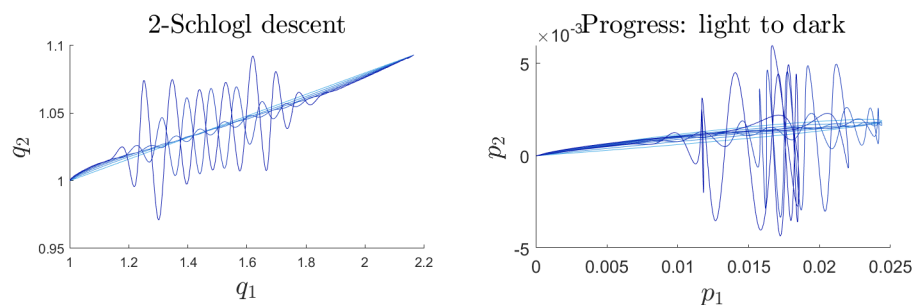


Figure 3.6: Functional gradient descent for 2-Schlogl model without filtering. Note that the low-frequency waves first amplify and then concentrate at higher frequencies near a central, badly-performing place on the trajectory.

3.3.2 Algorithm and pseudocode

Main idea

For a CRN exhibiting multiple fixed points, we start by picking a stable and a nearby saddle fixed point. The fixed points are determined by finding roots of the polynomial rate equations of mass-action kinetics, and

can be found for high dimensional systems using a numerical algebraic geometry software such as Bertini (Bates et al.). Once the end points have been determined, we obtain a configuration space *curve* by joining the end points with a straight line, which serves as an *initial condition* for the algorithm (see App. B.4.2). While in the actual algorithm the initial curve will be uniformly discretized at finitely many points, in this subsection, for the sake of explaining the main idea, consider that the curve is continuous.

Let the continuous configuration space curve at iteration I be denoted by q^I . Then, as explained in App. B.4.3 and shown in Figure 3.5, we *lift* the curve to a phase space *trajectory* by assigning a scalar dt and vector p^I at each point, such that the variation of the action functional in momentum is zero while the phase space trajectory is constrained to the $H(p, q) = 0$ submanifold,

$$\begin{aligned} \frac{dq^I}{dt} &= \left. \frac{\partial H}{\partial p} \right|_{(q^I, p^I)} \\ 0 &= H(p^I, q^I). \end{aligned}$$

Assigning momentum and time coordinates at each point along the discrete curve (see Figure 3.5) can be implemented using standard constraint optimization solvers (see B.4.3) and is the *costliest step* in the algorithm. Notice that since the Hamiltonian is convex in momentum, the solution exists. However, recall from Eq. 3.1 that the CRN Hamiltonian, as opposed to the generic Hamiltonian for mechanical energy, is not quadratic in mo-

mentum. Thus, the momentum coordinate generally cannot be solved for analytically and the difficulty of assigning it will depend on both the dimensionality of the system as well as the number of reactions in the network.

From Eq. B.25, the variation of the action functional along a phase space trajectory (q, p) is

$$\begin{aligned} \delta\mathcal{A}[q, p] &= \int dt \left(\frac{\delta\mathcal{A}}{\delta q} \cdot \delta q + \frac{\delta\mathcal{A}}{\delta p} \cdot \delta p \right) \text{ where} \\ \frac{\delta\mathcal{A}}{\delta q} &= - \left(\frac{dp}{dt} + \frac{\partial H}{\partial q} \right), \\ \frac{\delta\mathcal{A}}{\delta p} &= \left(\frac{dq}{dt} - \frac{\partial H}{\partial p} \right). \end{aligned} \quad (3.29)$$

Thus, as explained in App. B.4.4, we update the new configuration curve in the next iteration $q^{I+1} = q^I + \delta q^I$, where a variation δq^I is chosen to be in the **steepest descent direction** g^I (to first order), where

$$\begin{aligned} g^I &:= - \frac{\delta\mathcal{A}}{\delta q} = \left(\frac{dp}{dt} + \frac{\partial H}{\partial q} \right) \Big|_{(q^I, p^I)}, \\ \text{and } \delta q^I &= \epsilon g^I \end{aligned}$$

with a step-size $\epsilon > 0$ picked by employing a backtracking line search (see B.4.6, (Wright et al., 1999)). This assignment of δq^I ensures that, using

Eq. B.25, the variation of the action is negative semidefinite

$$\begin{aligned} \delta\mathcal{A} &= -\epsilon \frac{\delta\mathcal{A}}{\delta q} \cdot \frac{\delta\mathcal{A}}{\delta q} \\ &\leq 0, \end{aligned} \tag{3.30}$$

with equality only at the optimal curve.

The algorithm for finding the solution to Eq. 3.27 can thus be seen as performing a max by assigning optimal momentum values, followed by moving towards the min by taking a gradient in the descent direction in each iteration. The convergence of the algorithm relies on the convexity of the action functional around the optimal solution, on which we comment in App. B.3.4.

Noise and discretization

Since the optimal trajectory is a solution to a variational problem, it must be smooth. A naive implementation along the above account on a finitely sampled or discretized curve will fail on two ends. Firstly, there are two major sources of noise that self-amplify along the iterations, namely due to numerical solving of assigning momentum values (lifting the curve) and discretization. Secondly, any information in the optimal curve below the length scale of the discretization length will not be able to be captured. We elaborate on this in App. B.4.5, and for an illustration of the results with such a naive implementation, see Figure 3.6. **Our contribution** in this work

Algorithm 1 Pseudocode for Action Functional Gradient Descent (AFGD) algorithm (implementation (Gagrani, 2022))

```

I ← 1; ε ← εIC; fc ← fIC; ΔI ← ∞
qIC ← Initial_condition
qI ← Space_uniform_sampling[qIC]
(q, p, t)I ← Lift_curve[qI]           ▷ Lift curve to phase space trajectory
SI ← Action[(q, p)I]                 ▷ Calculate action
gI = Functional_gradient[(q, p, t)I]   ▷ Calculate functional gradient
gsI = Filter_gradient[gI, fc]       ▷ Low-pass filter gradient
ΔS ← 0
while I ≤ IMax or fc ≤ fMax or ΔI > Δthresh do
  I ← I + 1
  εI ← Step_Size(qI-1, gsI-1, εI-1)           ▷ Pick step size
  qI ← qI-1 + εI × gsI-1
  (q, p, t)I ← Lift_curve[qI]
  SI ← Action[(q, p)I]
  ΔS ← SI - SI-1
  if |ΔS| < ΔSthresh or ε < εthresh then
    fc ← fc + Δf; εI ← εIC           ▷ Loosen pass-band; refresh step-size
    qsI ← Filter_in_time_uniform[qI]       ▷ Optional: increase
sampling/anneal
    (q, p, t)I ← Lift_curve[qsI]
  end if
  ΔI ← Distance_Hamilton's_equations[(q, p)I] ▷ Verification: Least
distance from integrated EoM
  gI ← Functional_gradient[(q, p, t)I] ▷ Calculate functional gradient
  gsI ← Filter_in_space_uniform[gI, fc]   ▷ Low-pass filter gradient
end while

```

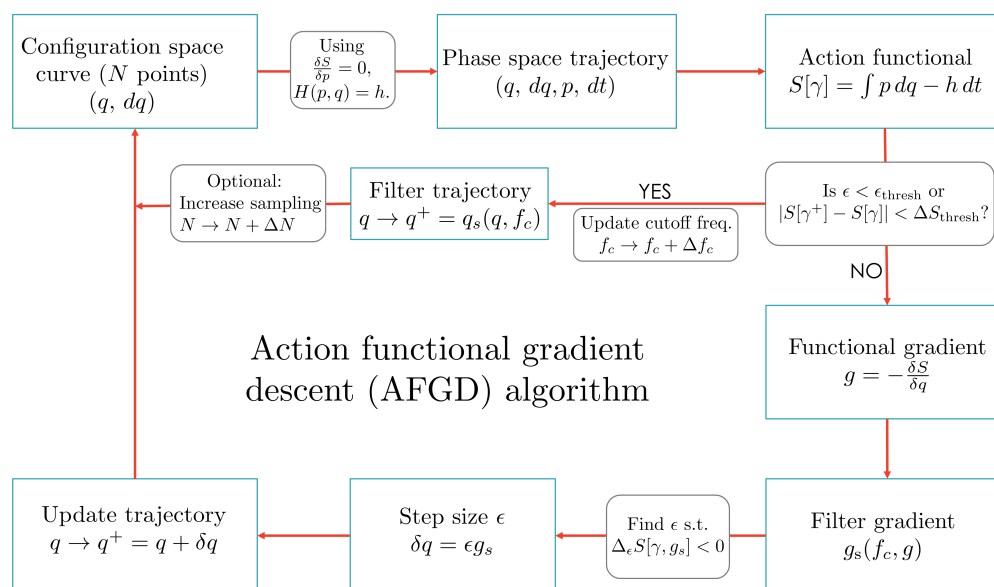


Figure 3.7: The AFGD algorithm in a nutshell. For the CRN Hamiltonian, when the end points are fixed points of the system then $h = 0$ must be chosen. For a detailed description, see Section 3.3.2.

then is to provide a principled algorithm for controlling these sources of noise and providing a recipe for how, in principle, one could approach the continuum limit in a controlled fashion where the algorithm is guaranteed to converge to the unique solution provided the action functional is convex.

Stuck in a local minima: The objective of the algorithm is to converge upon a phase space trajectory where the value of the action functional along it is minimized. If the difference in the action calculated in two subsequent iterations is below a threshold, or the step size is below a certain threshold, the descent is considered stuck.

Controlling noise: Firstly, in order to control the noise in the functional gradient g , we employ a low-pass filter and obtain a smoothed gradient g_s , as explained in App. B.4.5. To begin the descent, we choose a small low-pass frequency and allow the algorithm to proceed until it is stuck in a local minimum. The descent gets stuck away from the global minimum whenever the remaining frequencies to be relaxed in the trajectory fall above the pass-band shoulder. Thus we slightly increase the cut-off frequency of the low-pass filter, allowing more meaningful frequencies to go through and at the same time keeping noise from creeping in, and let the algorithm converge to a new curve (explained in App. B.4.7). Moreover, it must be noted that the algorithm uses a single step size for the complete curve, i.e. $q^{I+1} = q^I + \epsilon g_s$, where ϵ is a scalar. Thus we find it best to filter the gradient in a space-uniform parametrization of the curve (see Eq. B.51), which ensures that the magnitude of the gradient affects all the discrete segments in a uniform fashion.

Secondly, we also low-pass filter the configuration curve q^I in a time-uniform sampling (see Eq. B.52) each time we update the cutoff frequency. The reason for this is, in practice, although we smooth the gradient, some noise might still accumulate in the descended curve. This step ensures that the descended curve is smooth, and in practice, can often take the curve closer to the optimal curve even in the absence of a descent step (for instance, see Figure B.5 and the discussion in B.4.8).

Annealing/Reducing discretization interval: If the algorithm is stuck at a curve where no further descent is possible by increasing the low-pass frequency or filtering the trajectory, we increase the sampling of the curve by linearly interpolating on a smaller discretization interval. We also refer to this process as *annealing*, the details of which can be found in App. B.4.7. One might consider adding ≥ 1000 points in the trajectory, and restarting the algorithm from a new value of the low-pass cut-off frequency at this finely sampled trajectory. A typical reason why the algorithm does not descend further, despite not finding the optimal trajectory, is that the gradient might have meaningful information at length scale smaller than the discretization interval. Increasing the number of sample points then ensures that the meaningful information is not cancelled by the low-pass filter. Through examples considered in the next section, we show that this step can take the curve arbitrarily close to the optimal curve, as it theoretically should since we are approaching the continuum limit of the trajectory.

Verification protocol

At the optimal trajectory, we know that the functional gradient of the action in both position and momentum must be zero, in other words, the phase space trajectory must satisfy Hamilton's equations of motion. This gives a canonical verification method for the algorithm, namely numerical integration of Hamilton's equations forwards and backwards from each

point in the phase space trajectory and looking for how closely they approach the boundary points. The algorithm halts when we have an exact trajectory that starts and ends at the fixed points, and the deviation of the numerical integration from the end points gives us a way to quantify convergence. We explain the process by which one can obtain such a quantity in App. B.4.8. However, it must be noted that symplectic integration in high dimensions is often numerically unstable and we refer readers to (Hairer et al., 2006; McLachlan et al., 2009), for their careful and efficient implementation.

Algorithm in a nutshell

Starting with a uniformly discretized straight line between the end points, the Action Functional Gradient Descent (AFGD) algorithm consists of three nested subroutines:

1. Lift the curve to a trajectory and calculate the functional gradient. Filter the gradient in space-uniform parametrization and calculate a step size to descend. Descend to obtain a new curve.
2. If the descent is stuck, filter the curve in time-uniform parametrization and advance the space-uniform filter parameters. Go to routine 1.
3. If the descent is stuck, anneal or add more sample points and refresh the filter parameters. Go to routine 1.

Optionally, to quantify a distance from the global minimum, one can also repeatedly employ the verification protocol.

A pseudocode is presented in Algorithm 1 and the algorithm is diagrammed in Figure 3.7. The mathematical and implementation details of the subroutines can be found in App. B.4 and the code can be found in (Gagrani, 2022).

3.3.3 Performance costs and relation to other methods

From the last subsection, recall that the costliest step of our algorithm is the constrained optimization function needed to lift the curve. Let us refer to the cost of solving the optimization problem as $C_{c.opt.}$. Also, recall that the algorithm consists of three nested subroutines, namely obtaining the gradient and descending, advancing the filter-parameter when stuck, and advancing the number of sample points when stuck in the second routine. Let us refer to the number of points in the curve to be N_p , and the number of times the second and third routines are called as N_2 and N_3 , respectively. Moreover, suppose that each time the third routine is performed, ΔN_p points are added to the curve. The cost of running the algorithm then can simply be estimated as

$$\text{Cost of algorithm} \approx \sum_{i=0}^{N_3} (N_p + i\Delta N_p) C_{c.opt.} \times N_2.$$

As commented upon in Section B.4.3, a parallel implementation of calculating the lift of a curve can reduce the cost of computation by a factor proportional to the number of cores.

The proposed AFGD algorithm is a novel functional gradient descent method for solving two-point boundary value problems for a Hamiltonian system. Since it is known that the solutions satisfy the Hamilton equations of motion, another popular method for solving the same type of problem is the *shooting method* (Press et al. (2007)). In the shooting method, after selecting the two boundary points, one searches for the trajectory connecting them by integrating a brush of trajectories forwards and backwards from the starting and end points, respectively, in order to reach one or more fitting points in between. The technical difficulties associated with the shooting method method are two-fold. First, the search cost of the method increases exponentially with system dimension, as the (cross-sectional) brush of trajectories that must be sampled and refined has codimension 1 in the configuration space (starting position) and also codimension 1 in the tangent space (direction). Second, the method as a single algorithm can also fail for CRNs because the Hamiltonian diverges exponentially in the conjugate momentum variables (see Eq. 3.16). Stabilizing numerical integrators with diverging exponentials is challenging (Hairer et al., 2006; Duruisseaux et al., 2021), and if floating-point precision is magnified so that no trajectory can converge to a fixed point, the shooting method must be performed recursively along a path, introducing a further step of error

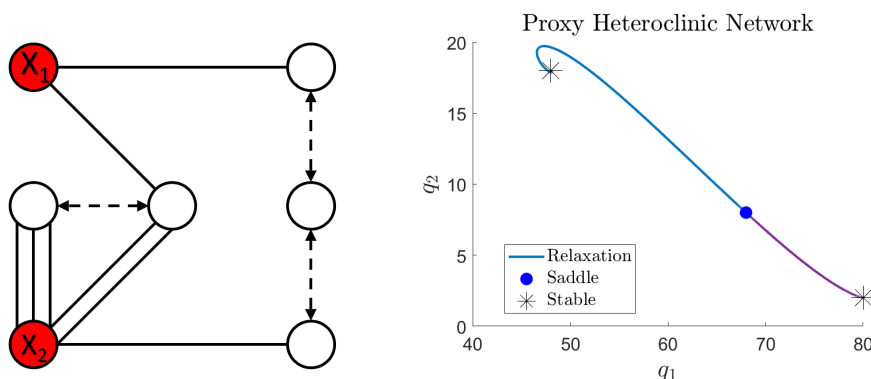


Figure 3.8: Diagrammatic representation of the reaction network (left) and a proxy-heteroclinic network using relaxation trajectories for the Selkov model.

bounding and path updating (for example, see Smith and Krishnamurthy (2021)).

The AFGD algorithm, on the other hand, does not suffer from the same issues. Firstly, it relies on computing derivatives, minimizing objective functions, and filtering, all of which are significantly cheaper and more robust than numerical integration. Secondly, due to the above mentioned reasons, it is also feasible in large dimensions. However, it must be noted that the computation time scales with proportionally to the number of points in the trajectory, and in high dimensions a rather fine sampling of the curve might be needed to get a meaningful escape. The demerit of the AFGD algorithm is that it does not find the exact escape path, as the action can reach its minimum value within tolerance without having completely descended upon the solution to the equations of motion. An example of this can be seen towards the ends of the trajectory near the

stable fixed points in Figure 3.9, where the trajectory converged upon by the AFGD does not agree completely with that from the shooting method. This suggests that, even in high dimensions, some hybrid of the shooting method and AFGD can be used if the exact escape trajectory is desired. However, if simply a reliable estimate of the action along the escape trajectory is required (to estimate the escape probability, for example), then the AFGD algorithm should be preferred.

3.4 Application and results

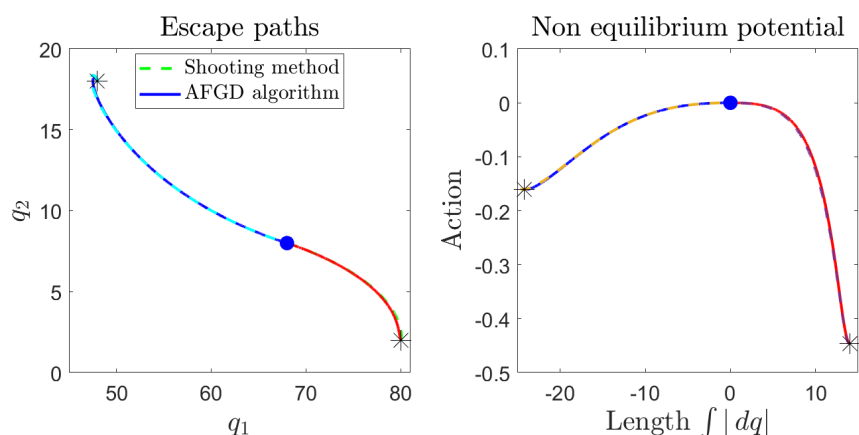


Figure 3.9: Comparing escape paths obtained from AFGD algorithm against the shooting method (left panel) and plotting the non equilibrium potential along the escape (right panel). The data is taken from (Smith and Krishnamurthy, 2021), and can be compared against Figures 9 and 15 of the same reference. (For details on the differences between the two methods, see the last paragraph in Section 3.3.3.)


In this section, we will demonstrate the applicability of the AFGD algorithm using three models with varying features. The first application

we consider will be to determine the escape paths and NEP of the Selkov model. The Selkov model, first introduced in 1968 by E. Sel'kov (Sel'Kov, 1968) to model self-oscillations in glycolysis, exhibit relaxation curves that spiral into the stable attractors. Since a spiralling curve, in principle, needs an exponential number of sample points near the fixed point, it poses a difficult challenge for the AFGD algorithm. The other two applications that we consider are higher dimensional analogs of the Schlögl model. The Schlögl model was introduced by F. Schlögl in 1972 to understand non-equilibrium phase transitions in chemical reaction systems (Schlögl, 1972). The Schlögl model is an example of a 1-dimensional birth-death process, for which the NEP can be analytically found using Hamilton-Jacobi theory, as we show in Section 3.4.2. We then define N dimensional analogs of the Schlögl model, that we term as N-Schlögl models, for which no analytic results are yet known. We then make use of the AFGD algorithm for the 2-Schlögl and 6-Schlögl model, and verify it against Gillespie simulation and Hamilton's equations of motion, respectively. It must be noted that throughout the examples considered in this section, we tune only the algorithm parameters while keeping the underlying algorithm identical to the pseudocode 1, thus demonstrating that the algorithm is agnostic to the dimensionality of the model. For a MATLAB implementation of the AFGD algorithm, using which we obtain the figures displayed in this chapter, see Gagrani (2022).

3.4.1 Selkov model

In this subsection we will consider the well-known Selkov model, which has been analyzed by eikonal methods in (Dykman et al., 1994; Smith and Krishnamurthy, 2021). A peculiarity of the model, that makes it both interesting and challenging for analysis, is that both relaxation and escape trajectories exhibit vorticity around the fixed points. In both (Dykman et al., 1994; Smith and Krishnamurthy, 2021), the authors identify the escape trajectories by using the *shooting method* (see Press et al. (2007)), which amounts to integrating a brush of optimal trajectories in phase space emanating from a point, selecting the one that approaches the desired end point most closely, and repeating. The AFGD algorithm, however, identifies the optimal trajectory using a functional gradient descent, and we summarize the main results of the implementation as well as compare it to the optimal trajectory obtained via the shooting method in (Smith and Krishnamurthy, 2021) here.

To explicitly define the model, our choice of rate constants is identical to (Smith and Krishnamurthy, 2021). The concrete model we consider is,



A diagrammatic representation of the reaction network can be seen in

the left panel of Figure 3.8. The resulting fixed points of the model are at $\underline{q} \in \{(80, 2), (68, 8), (48, 18)\}$, which we first characterize using the number of repelling directions and then find the *relaxation* trajectories that emanate close to a saddle and reach the stable point in the right panel of Figure 3.8, which we refer to as a ‘proxy-heteroclinic network’. Our goal is to find the *escape* trajectories that emanate from each of the two stable points and reach the saddle point, and thus obtain the ‘true heteroclinic network’. For the sake of brevity we will only demonstrate the workings of the algorithm on the top left escape path starting from $(48, 18)$ and ending at $(68, 8)$, however the same procedure can be used to obtain the other escape as well.

The first feature that one might notice in the top left relaxation path is that it spirals inwards towards the stable fixed point. This is not uncommon for a CRN and we discuss how to obtain initial conditions for such trajectories in App. B.4.2. We then sample the initial configuration space curve uniformly with 2500 points and run the algorithm, the results of which are summarized in Figure B.6. The top row of the figure displays how the descent progresses across the iterations, and the lower two rows summarize how the action, step size, cutoff frequency and minimum distance from the integrated EoM change along the descent.

Although the action seems to stabilize and the step size is zero roughly beyond iteration 150, the momentum assignment at the converged trajectory (top right panel) clearly looks unreliable. As a rule of thumb,

we expect the optimal trajectory to be smooth and containing only low frequency modes. Thus, one can see that despite the AFGD algorithm's effort to smooth the cusp in the momentum initial conditions, it has not yet converged even close to the true solution. To remedy this, we sample the converged curve in 4000 points, a process we call *annealing* (explained in Section 3.3 and App. B.4.7), run the algorithm again starting with this initial condition and plot the summary in Figure 3.10. This time we indeed find that the converged momentum assignment is without any cusp (right panel of first row), the action has converged to a yet lower value (second row) and the least distance from the integrated equations of motion Δ drops below 0.5 (third and fourth row), guaranteeing convergence near the true solution.

We now proceed by the same method to find the other escape trajectory and plot them against the optimal trajectory found by the shooting method in (Smith and Krishnamurthy, 2021) in the left panel of Figure 3.9. Using Eq. 3.23, we find the NEP along the escape trajectories and thus obtain the true heteroclinic network for the Selkov model, as displayed in the right panel of Figure 3.9.

3.4.2 N-D Schlögl model

1-D Schlögl model We begin our discussion by first considering 1-D birth-death processes, of which the 1-D Schlögl model is an example. A 1-species reaction network is called as a birth-death process if the difference

of the vectors denoting the target and source complex for each reaction is either positive or negative one, that is

$$|y_\beta - y_\alpha| = 1 \text{ for all reactions in } \mathcal{R},$$

$$\text{where } \mathcal{R} = \{y_{\text{ff}} \xrightarrow{k_{y_{\text{ff}} \rightarrow y_{\text{fi}}}} y_{\text{fi}} : k_{y_\alpha \rightarrow y_\beta} \geq 0\}.$$

The reactions where the stoichiometry of the target complex is one more or less than the source correspond to ‘birth’ or ‘death’ reactions, respectively (Anderson et al., 2015).

Using the form of the CRN Hamiltonian in Eq. 3.16, we can write the Hamiltonian function of a 1-D birth death process as

$$H_{1\text{-b.d.}}(p, q) = (e^p - 1)r_{+1}(q) + (e^{-p} - 1)r_{-1}(q), \quad (3.32)$$

where $r_{+/-1}$ are polynomials in q with coefficients as the corresponding rate constants appearing in the birth/death reactions. From the Hamiltonian, one can read that the deterministic rate of growth in the concentration of the species is $\dot{q} = r_{+1}(q) - r_{-1}(q)$ and the roots of \dot{q} correspond to fixed points of the system.

We can now proceed to find the NEP for such processes by solving for an escape momentum assignment in the $H_{1\text{-b.d.}} = 0$ submanifold. Since we have only one species and a two-dimensional phase space manifold, for a

birth-death process this constraint uniquely picks a $p_{\text{esc}} \neq 0$,

$$\begin{aligned} H_{1\text{-b.d.}}(p_{\text{esc}}, q) &= 0, \\ p_{\text{esc}}(q) &= \ln \left(\frac{r_{-1}(q)}{r_{+1}(q)} \right). \end{aligned} \quad (3.33)$$

Following the discussion on Hamilton-Jacobi theory in Section 3.2.2, integrating the escape momentum in Eq. 3.33 yields the NEP \mathcal{V} . Notice that at the fixed points, $p_{\text{esc}} = 0$, and correspondingly the NEP is at a local extrema. To ensure that the NEP is always greater than zero, we find a possibly non-unique fixed point \underline{q} , such that

$$\underline{q} = \arg \min_q \int_0^q \ln \left(\frac{r_{-1}(q)}{r_{+1}(q)} \right) dq,$$

using which the NEP is defined as

$$\mathcal{V}_{1\text{-b.d.}}(q) = \int_{\underline{q}}^q \ln \left(\frac{r_{-1}(q)}{r_{+1}(q)} \right) dq. \quad (3.34)$$

The ratio of the value of the stationary distribution at points n_2 and n_1 in a stochastic simulation of scaling volume V is then given by

$$\frac{\pi(n_2)}{\pi(n_1)} = \exp \left(-V \int_{q_1}^{q_2} \ln \left(\frac{r_{-1}(q)}{r_{+1}(q)} \right) dq \right). \quad (3.35)$$

The results concerning the NEP of 1-D birth death systems are well known, and an alternative derivation can be found in (Anderson et al.,

2015). A particular example of a birth-death process relevant to our purposes here is the **Schlögl model** (Schlögl, 1972). For a pedagogical exposition of an application of Hamilton-Jacobi theory to Schlögl model and its generalizations, see Lazarescu et al. (2019); Smith (2020). For this section, we use the concrete 1 species reaction network




where the rate constants are identical to example 10 of (Anderson et al., 2015). For a diagrammatic representation of the reaction network, see top left panel in Figure 3.1. For the choice of reaction network in Eq. 3.36, we have

$$\begin{aligned} r_{-1}(q) &= 11q + q^3 \\ r_{+1}(q) &= 6 + 6q^2 \\ \frac{dq_{\text{rel}}}{dt} &= \left. \frac{\partial H}{\partial p} \right|_{p=0} = r_{+1}(q) - r_{-1}(q) \\ &= -(q-1)(q-2)(q-3). \end{aligned} \quad (3.37)$$

Notice that the relaxation flow field has three roots, which means that the model exhibits three fixed points for the rate constants chosen $q \in \{(1), (2), (3)\}$. We display the velocity and momentum assignment of the variational solution along the $H(p, q) = 0$ submanifold for this model

in Figure 3.11.

2-D Schlögl model We define the 2-D Schlögl model to be a two species reaction network with 1-Schlögl reactions in each species and diffusion between the two species. For the results in this section and App. B.4, we use the concrete network



which we represent diagrammatically in the top right panel in Figure 3.1. For the choice of rate constants, the relaxation flow field has nine fixed points, out of which four are stable (with no repelling direction), four are saddle (with one repelling direction) and one is unstable (with all repelling directions).

As explained in Section 3.2.1, in order to understand the switching dynamics from one stable fixed point to another, we need to only consider the least-improbable path of escape. We know that the least-improbable path between two nearby fixed points is through the saddle point between them, and thus we look for the equation of motion that emanates from a stable attractor and ends at any of the closest saddle points using the AFGD algorithm. Before looking for escape paths however, we find it useful to find the relaxation trajectories that emanate close to a saddle

point and reach a nearby stable point. This information yields a ‘proxy-heteroclinic network’ and tells us whether or not the system exhibits any vorticity in the areas of interest. For the position and classification of the fixed points and a ‘proxy-heteroclinic network’ for the 2-Schlögl model, see Figure 3.12.

We now proceed to identifying all the relevant escape paths to create a heteroclinic network for the 2-Schlögl model. As an illustrative example of an application of the AFGD algorithm, we first focus our attention to the bottom left escape path from the stable root $\underline{q}_I = (1, 1)$ to the saddle root $\underline{q}_F = (2.16, 1.09)$. From the criteria specified in App. B.4.2, we begin the descent from a straight line initial condition as shown in Figure B.2. We run our algorithm starting from a curve consisting of 500 points, and we summarize the progress of the descender in Figure B.5. The topmost row depicts how the the phase space trajectory changes in the configuration space (left) and momentum space (right) as the algorithm progresses. The last two rows summarize the progress of the action along the escape, step size, cutoff frequency and least distance from the end points of the integrated EoM during descent.

One can observe that the descent in the action has a step-like behavior which is mirrored in the step size as well as cutoff frequency. What is happening is that the algorithm descends at a given setting of a cutoff frequency as much as it can, by reducing the step size as long as it is above a certain threshold and the value of the action function is strictly

decreasing. Once no further step can be taken because the step size has decreased below the minimum threshold of ϵ_{thresh} , the cutoff frequency is slightly raised to allow the pass-band to inject more meaningful signal into the gradient. It can be seen in Figure 3.13 that the overall magnitude of the gradient decreases and its power spectrum becomes flatter as the descent proceeds through the iterations of the algorithm. For more details, see App. B.4.7.

The bottom right panel in Figure B.5 shows how the minimum distance Δ of the integrated Hamilton's EoM, as explained in App. B.4.8, changes along iterations of the algorithm. Curiously, the distance continues to decrease even when the step size is zero for $I > 55$. The reason for this is that our algorithm also smooths the configuration space curve in a time-uniform sampling, which for the 2-Schlögl demonstrably takes the trajectory towards the optimal solution. The distance, however, ceases to decrease after a point, and it is a signal that the number of points in the trajectory needs to be increased if further progress is to be made. We show the closest integrated trajectory overlaid on the descender trajectory for a few iterations in Figure 3.14. Since, in the previous example, we have already discussed the process of descending further towards the true solution by 'annealing' (increasing the number of sample points in the trajectory), we will not pursue it here.

Validation against Gillespie algorithm Now that we have verified the correctness of the descended trajectory against Hamilton’s EoM, we can ask how well does it perform against a stochastic modelling method such as the Gillespie algorithm (Gillespie et al., 2007)? To compare the two, we run two simulations with scaling volume $V = 300$ and find the scaled-log improbability of their stationary distribution π , i.e. we calculate $(-1/V) \log(\pi)$. Since the two simulations give similar results, we only display the contour plot resulting from the second simulation for clarity in the left panel of Figure 3.4. We then proceed to find the scaled log-improbability along only the escape trajectory output by the AFGD algorithm and compare it against the action or NEP found using Eq. 3.23. The result is displayed in the right panel of Figure 3.4. We can see that the scaled NEP obtained from the simulation is higher than the log-improbability obtained from the algorithm, as it theoretically should because the variational solution provides a lower bound, becoming exact in the $V \rightarrow \infty$ limit. For a similar plot for the 1-Schlögl model, see Figure 1 in (Anderson et al., 2015).

Having convinced ourselves of the correctness of the algorithm by two different means, i.e. via integrating Hamilton’s EoM and verifying against stochastic simulation, we proceed to find the true heteroclinic network for the 2-Schlögl model. The heteroclinic network consists of all the escape trajectories from every stable attractor, as well as the log improbability or NEP along each escape. We obtain the escape paths by running the

algorithm for every pair of nearby stable and saddle points, and display the resulting heteroclinic network in Figure 3.3.

N-D Schlögl model The AFGD algorithm is defined independently of the dimensions of the system, and thus works equally well for dimensions higher than 2. To demonstrate this point, we consider an N-dimensional generalization of the Schlögl model. Analogous to the 2-Schlögl model, we define a system with N species and include the Schlögl reactions for each species, as well as diffusion amongst the species.

For $N > 2$ we have to make a choice regarding the underlying diffusion network, and to simplify considerations we choose a fully connected diffusion network, i.e. each species is diffusing with all species. Other choices could also be made, such as the species could form a 1-D or a 2-D lattice with diffusion only between the nearest neighbors, but this choice is beside the point of our purposes here. A diagrammatic representation of the 3-D and 4-D Schlögl model can be seen in the lower line of Figure 3.1.

For an application of the AFGD algorithm to higher dimensional systems, we consider the 6-Schlögl model. To demonstrate the efficacy of the algorithm, the descended phase space trajectory is plotted against Hamilton's EoM in Figure B.7. The output trajectory is obtained using 2000 sample points and running for under 300 iterations at a relatively low cutoff frequency. Since other details are not particularly illuminating,

we omit the summary plots for brevity, but make it available in (Gagrani, 2022).

3.5 Discussion

Chemical reaction networks (CRNs) are essential for modeling a wide range of natural phenomena such as star formation, the origin of life, spatial or ecological patterns in living organisms, and climate (Smolin (1996); Smith and Morowitz (2016); Turing (1990); Benzi et al. (1983)). The widespread utility of CRNs stems from their ability to exhibit dynamic equilibria which, unlike the equilibrium at the top or bottom of a potential well where the velocities of objects are zero, are states of a system where the composition remains unchanging although the constituent species are being dynamically exchanged. The state of a star, organism, ecosystem or climate, when modeled as a stochastic CRN, can undergo transitions from one dynamic equilibrium to another, and it is the probability of such a transition occurring that we give an algorithm to numerically estimate in this work. It must be noted that since the probability of transitions is exponentially suppressed in the number of simulated species, finding cheaper ways of estimating them is of practical importance.

Our main contribution has been to employ the Hamilton-Jacobi formalism to rigorously formulate the problem of finding transitions between steady states (or fixed points) of a CRN as a MinMax problem, and to con-

struct a principled algorithm to solve it. The functional whose value along its optimal points are the desired transitions is called the *action* functional, due to which we name our algorithm as the *Action Functional Gradient Descent* (AFGD) algorithm. Our algorithm only requires computation of derivatives, solving function optimization problems, and basic tools from signal processing. Moreover, while the algorithm itself does not rely on numerical integration, its validity can be readily verified by integrating the equations of motion starting anywhere along the converged output (for more detailed, see Section 3.3). Finally, in Section 3.4, we explore applications of the algorithm on several high dimensional problems and validate them against other methods of obtaining transition paths and probabilities.

While in this work, we only use the algorithm to calculate transitions between two fixed points, in principle it can be used to find a transition between any two points which are guaranteed to have a direct optimal phase space trajectory connecting them. In particular, rather than escaping from a stable fixed point to a nearby saddle point, one might be interested in finding the probability of an escape to any point within the basin of escape of the stable point in the long time limit. As explained in Section 3.2, the optimal trajectories that connect the stable fixed point to the desired point will also be in the $H(p, q) = 0$ submanifold, and can readily be found (by leaving the momentum at the end point unspecified) by the AFGD algorithm. In this way, by connecting arbitrary points to their nearest stable

fixed points, in principle, one can assign a transition probability to each point on the state space and recover the occupation probability distribution that one would otherwise obtain by running a stochastic simulation of a CRN for a very long time. We will then leave it as future work to take as input the time-series data or occupation probability distribution obtained from a stochastic simulation and to learn the CRN from which it was generated. Due to the widespread utility of CRN in modeling real-world phenomena, a machine learning algorithm to infer the CRN from its simulated data would have several significant scientific applications.

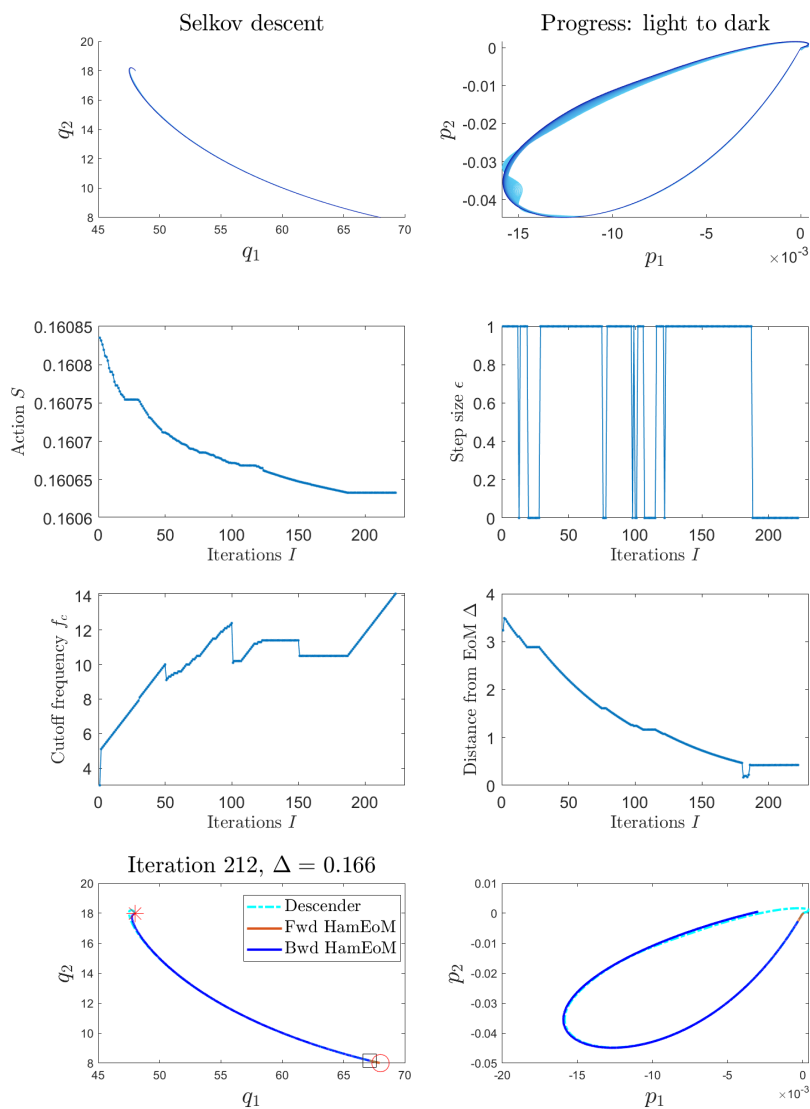


Figure 3.10: Descent progress, summary and proof of convergence for the Selkov model after annealing to a trajectory with 4000 points. For initial descent with 2500 points, see Figure B.6.




Figure 3.11: $H(p, q) = 0$ submanifold in the tangent and cotangent space for the 1-D Schlögl model.

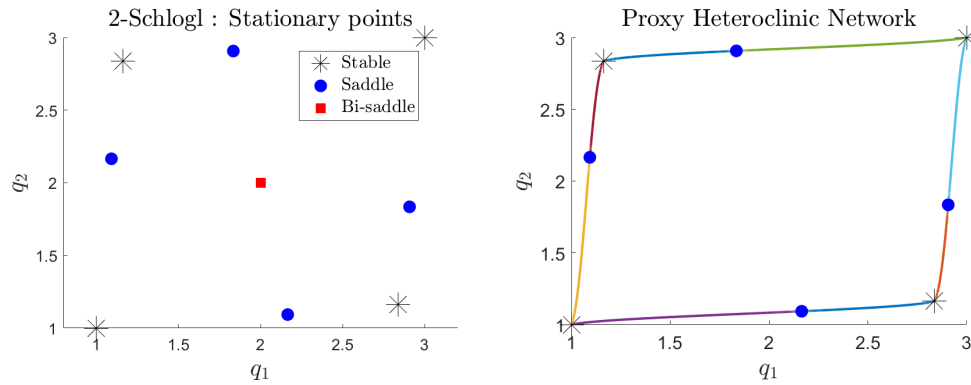


Figure 3.12: Fixed points labelled by repelling directions (left) and proxy-heteroclinic network using relaxation trajectories (right) for the 2-Schlögl model.

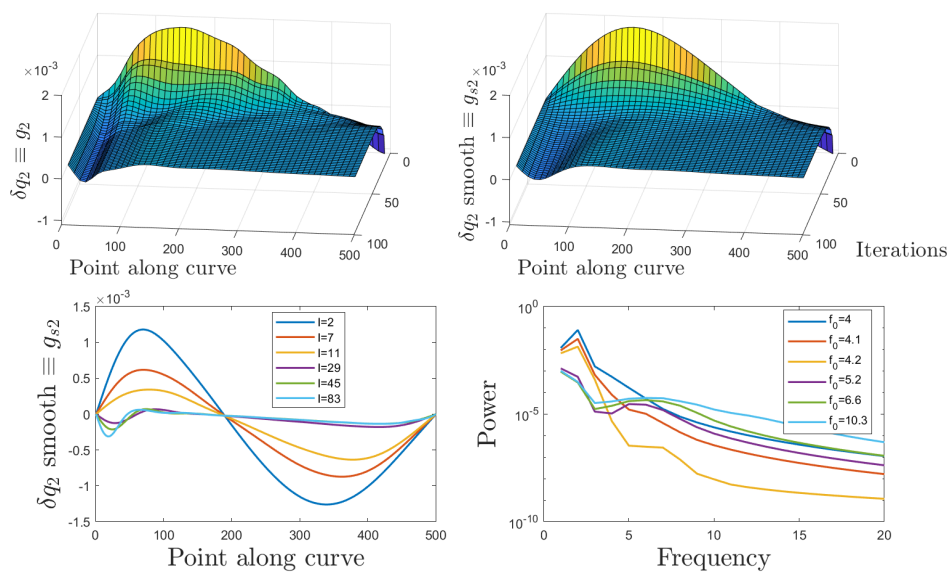


Figure 3.13: Extracting functional gradient from phase space trajectory and smoothing it (top panel). Smoothed gradient and power spectrum at selected iterations (bottom panel) for the 2-Schlögl model.

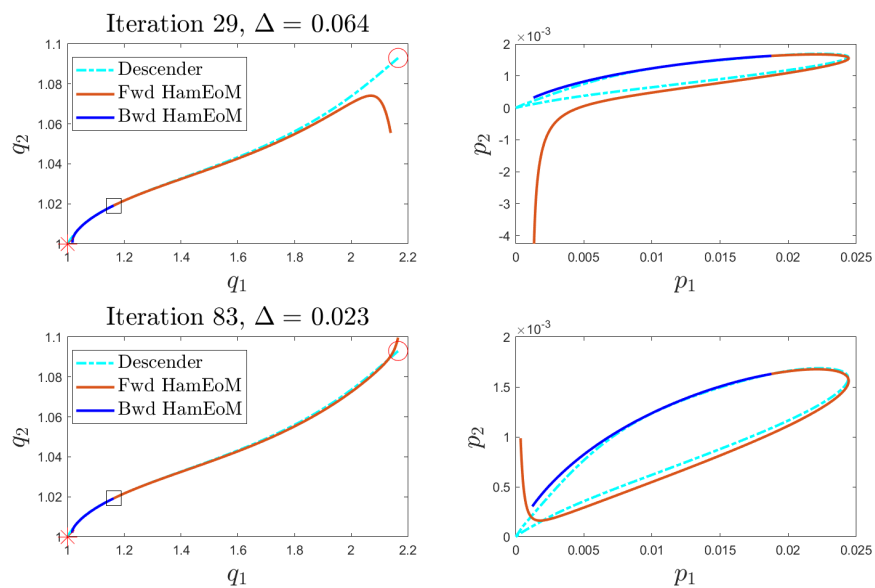


Figure 3.14: Progress of integrated Hamilton's equations of motion against the AFGD algorithm at selected iterations for the 2-Schlögl model.

4 MATHEMATICS OF EVOLVING ECOLOGIES IN STOCHASTIC CHEMICAL REACTION NETWORKS AND THE ORIGINS OF BIOCHEMICAL LIFE

4.1 A mathematical formalism for modeling evolving ecologies in chemical reaction networks

In this section, we propose a mathematical formalism for a bottom-up approach to modeling an evolving ecology using CRNs. We articulate the role of mathematics in the formalism in subsec. 4.1.1, and explain the physical assumptions leading to a CRN description of the phenomenon under investigation in subsec. 4.1.1. In subsec. 4.1.2, by purely topological considerations, we define our notion of an *ecology*. Then, in subsec. 4.1.3, we give a mathematical characterization of various types of *evolution* for a CRN with chemical kinetics.

4.1.1 The role of mathematics in our formalism and the physical assumptions leading to our basic set-up.

A primary utility of mathematics in the sciences is its ability to unambiguously discern what is impossible, thus imposing limits on what is

possible. For example, elementary geometric considerations yield that it is physically impossible to inscribe a square of side length d in a circle of diameter d . Similarly in the physical sciences concerning the motion of particles, variational analysis is used to derive the optimal paths between two points (Goldstein et al., 2002), imposing a limit on some measurable quantity (time for light, or ‘action’ for massive particles) in any experiment satisfying the relevant constraints.

In our formalism, we take a bottom-up approach to modelling a system under investigation. We define a *system* by defining the types of all pragmatically irreducible entities, and their network of interactions with each other and the environment. The role of mathematics in our formalism, then, is to systematically investigate the variabilities that the system can exhibit, and derive limits on these variabilities. In particular, we give a precise mathematical criterion for what constitutes an ecology, and whether or not it is capable of evolution. Whether or not the modeled systems are physically feasible, and their variabilities admitted in nature, are practical questions left unresolved by mathematical considerations alone.

For a bottom-up approach to modeling a phenomenon under investigation, our formalism requires that:

- The types of all, possibly infinite, pragmatically irreducible entities in the system is known.¹ We refer to the types as *species*, labelled by

¹The notion of *irreducibility* is not to be taken in the physical sense, like a description in terms of subatomic particles, rather in a pragmatic sense, like a molecule or a cell

s , and the list of species as the *species set*, denoted by $\mathcal{S} = \{s_1, \dots, s_S\}$.

- The list of all stoichiometric interactions, termed *reactions*, by which the count of each entity changes are known. Each reaction changes the composition of the species population by whole numbers, and their list, denoted by \mathcal{R} , includes the interactions of all the species with one another and with their environments.

When taken together, the information $(\mathcal{S}, \mathcal{R})$ is said to define a CRN. In the remaining subsection, we will describe the assumptions that need to be made to obtain a CRN description.

The first assumption that we make to obtain a CRN description is that each particle of a species is indistinguishable, and we can only access the count or population of each species, termed *species composition*, of the system. Second, we assume that the mechanism of interactions is memoryless, or satisfies the *Markovian* property. By this we mean that the state of a system in the future only depends on the current state, and is not affected by its past. In Appendix C.1 and Ch. 3, we show that these processes can be formalized by way of a *stochastic Hamiltonian* and derive the Hamilton-Jacobi theory for stochastic population processes.

The third assumption that we make to obtain a CRN description is on the stochastic Hamiltonian. To obtain a CRN description, we assume that the stochastic dynamics can be specified such that each stochastic entity may be considered to be an irreducible entity for a chemical or biological experiment, respectively.

transition can be expressed as a list of rules that transmute a collection of species into another. In other words, we require that the stochastic Hamiltonian be written as a sum over reaction channels in the set \mathcal{R} , as described in Appendix C.2 and Ch. 2. We remark that CRNs can also be used to describe cosmological and astrophysical processes by considering all the subatomic particles as the species set and their interactions afforded by the Standard Model (Group, 1998) as the reaction set, and collect some examples in Table 4.1 and Fig. 1.1.


Species	Reactions	Description
electron e^- , neutron n ,	$n \rightleftharpoons p^+ + \bar{\nu}_e + e^-$	Beta minus (β^-) decay
electron e^- , hydrogen H ,	$e^- + p^+ \rightleftharpoons H + \gamma$	Recombination
positron, neutron, photon, neutrino, proton	See Fig. 1.1	Fusion in the Sun

Table 4.1: Examples of chemical reaction networks in cosmology, astrophysics, and particle physics.


Finally, we make a fourth assumption to obtain a kinetic description from CRNs. Stochastic chemical kinetics assumes that the particles come in discrete amounts and are in a well-mixed reactor. Moreover, it assumes that each reaction proceeds stochastically proportional to the ways of sampling the input reactants from the system. For a detailed explanation see subsec. C.3 or Ref. Gillespie et al. (2007). While these kinetic assumptions are not always justified (see Michaelis–Menten kinetics (Qian and Bishop, 2010)), they do have a microphysical basis. For instance, the stochastic chemical kinetics assumed here have been shown to be exact for any gas-phase chemical system that is kept well-stirred and thermally equilibrated in Ref. Gillespie (1992). We remark that the assumption of well-stirred does not mean that CRNs cannot account for spatial inhomogeneity. To model diffusion systems that differentiate between locations, we simply need to add new species at each site and explicitly track their location index (for example, see Ref. Muñuzuri and Pérez-Mercader (2022); Plum and Baum (2022)).

4.1.2 A topological definition of an ecology

A stoichiometric interaction, or a reaction, in the set \mathcal{R} is of the form



where the coefficients y_i and y'_i are nonnegative integers ($\in \mathbb{Z}_{\geq 0}$) for all species s_i in the set \mathcal{S} . Any collection $w_1 s_1 + \dots + w_S s_S$ with nonnegative coefficients w_i is referred to as a *complex*, and the vector w is called as the stoichiometry of the complex. In what follows, we will simply denote a complex by its stoichiometric vector w . Thus, notationally, the reaction above simplifies to



and y and y' are called the input and output complex, respectively.

Let us label the reactions in the set \mathcal{R} as

$$\mathcal{R} = \{y_1 \rightarrow y'_1, \dots, y_{|\mathcal{R}|} \rightarrow y'_{|\mathcal{R}|}\},$$

where all complexes y_i and y'_i for $i \in [1, |\mathcal{R}|]$ are not necessarily distinct.


We denote *species composition* at a given instant by

$$n := (n_1, \dots, n_S),$$

where each n_i is a non-negative integer. The effect of a reaction $y \rightarrow y'$ on a species composition n , then, is to remove the complex y and replace it with the complex y' . Equivalently, a reaction $y \rightarrow y'$ transforms the composition from n to $n - y + y'$, denoted as $n \rightarrow n - y + y'$. For a bottom-up approach to modeling, we require that every such elementary stoichiometric interaction by which the species composition can be changed must be

in the set \mathcal{R} . Conversely, we require that \mathcal{R} only consist of such elementary interactions.

In CRNs, the firing of a collection of elementary reactions in the set \mathcal{R} can also be expressed as a reaction, however this reaction does not necessarily belong to the set \mathcal{R} and is termed a *composite reaction*. For example, in a network with two reactions $\{y_1 \rightarrow y'_1, y_2 \rightarrow y'_2\}$, firing the first and second reaction f_1 and f_2 times, respectively, we obtain the composite reaction




The firings are formalized as **flows** on a CRN in Ch. 2, and are explained further in Appendix C.2.

A topological property of a CRN of particular interest for our formalism is *autocatalysis*. Autocatalytic cycles have been used for modelling the metabolic cycles in an organism, the life cycle of an organism (Smith, 2021), and in general can be seen as a minimal model for a living organism (Hordijk and Steel, 2018). The concept of autocatalysis, dating back to the late 1800s (Ostwald, 1890; Peng et al., 2022b), has several distinct formulations in CRN theory (Andersen et al.; Deshpande and Gopalkrishnan, 2014; Blokhuis et al., 2020). Here, by autocatalysis we mean what is referred to as *exclusive autocatalysis* in Ref. Andersen et al.. A discussion of autocatalysis and related concepts can be found in Appendix C.2 and Ch. 2.

A subset of reactions $\mathcal{R}' \subset \mathcal{R}$ will be set to be an *autocatalytic subnetwork* (AS) if it satisfies the following properties:

1. There exists a flow on the subnetwork given by \mathcal{R}' such that the resulting composite reaction is of the form:



where $\mathbf{0} \prec \mathbf{m} \prec \mathbf{n}$.² Here \mathbf{m} and \mathbf{n} are the stoichiometries of the set \mathcal{C} in the input and output of the composite reaction, respectively, and $\mathbf{o}\mathcal{C} = \sum_i \mathbf{o}_i \mathcal{C}_i$.

2. For every reaction $y \rightarrow y' \in \mathcal{R}'$, the input complex y and the output complex y' contain at least one species from the set \mathcal{C} . This ensures that there is no flow if the population of every species in the set \mathcal{C} is zero.

For an AS that satisfies the above properties, \mathcal{C} will be said to be its *autocatalytic set*.

Furthermore, generalizing the formulation in Ch. 2, for each AS, we can define:

- A food set F defined as the union of all the species that occur only as reactants in the reaction set \mathcal{R}' .

² $A \prec B$, read as *A precedes B*, if all entries of $A - B$ are strictly negative. Analogously, $A \succ B$, read as *A succeeds B*, if all entries of $A - B$ are strictly positive.

- A waste set W defined as the union of all the species that occur only as products in the reaction set \mathcal{R}' .
- A member set M defined as the species which occur as both reactants and products in the reaction set \mathcal{R}' .

We refer to the assignment of food, waste, and member sets for an AS as its *food-waste-member* or FWM partition. We briefly remark that, by construction, each AS can be uniquely assigned a FWM partition, and their mathematical properties are developed further in Ch. 2.

Ecology: We define an **ecology** to be a collection of ASs in a CRN.³ Our use of the terminology is justified by the observation that, as explained above, an organism can be modelled using ASs in a CRN. The ecology as a whole, then, can be seen as a dynamic network of *positive* and *negative interactions* between different organisms. Moreover, since each AS has its own set of food and waste species, an ecology can be assigned a *food set* as the union of the food sets and elements in the member set that only appear as reactants in the composite reactions of all the ASs while excluding the union of all their waste sets, and elements in the member set that only appear as products in the composite reactions. The food set thus defined imposes a limit on the least number of species required by the ASs to exhibit their autocatalytic character. Notice that certain CRNs can exhibit

³Given a CRN, a linear-programming algorithm to enumerate all minimal ASs which have *distinct* species in the input and output complexes is given in Ch. 2.

ecologies that do not have a food set, and we term them *closed ecologies*. Arguably, after the dominance of photosynthetic recycling, barring rare astronomical events, there is no appreciable exchange of atoms from outer space. Thus the CRN of Earth with atoms and molecules as species can be said to exhibit a closed ecology.

4.1.3 A kinetic definition of evolution

The effect of the reactions on the species composition dynamics is obtained by selecting a model for how the interactions proceed and choosing values for kinetic parameters. Our model of the effect of reactions on the system is that, each reaction in the set \mathcal{R} occurs with a probability proportional to the number of ways of selecting the collection \mathbf{y} from the current composition. For a reaction $\mathbf{y} \rightarrow \mathbf{y}'$, this choice of dynamics models the scenario that a collection of entities specified by the reactant complex \mathbf{y} come together and transmute into entities specified by the product complex \mathbf{y}' . This assumptions constitutes *stochastic chemical kinetics* in a continuous-flow stirred tank reactor. For a detailed derivation, see Appendix C.3 and Ch. 3.

A CRN \mathcal{G} , when taken with stochastic chemical kinetics, is given by the triple $(\mathcal{S}, \mathcal{R}, \mathcal{K})$, where

$$\mathcal{K} = \{k_{\mathbf{y} \rightarrow \mathbf{y}'} \in \mathbb{R}_{>0} \mid \mathbf{y} \rightarrow \mathbf{y}' \in \mathcal{R}\}.$$

In stochastic chemical kinetics, the unnormalized probability of each reaction $y \rightarrow y'$ being selected randomly is $k_{y \rightarrow y'} n^y$, where n is the discrete species composition in the integer lattice $\mathbb{Z}_{\geq 0}^{|S|}$ and

$$n^y = \prod_{i=1}^{|S|} \frac{n_i!}{(n_i - y_i)!}.$$

Instead of tracking the discrete species composition n , it is useful to track the continuous species composition $q \in \mathbb{R}_{\geq 0}^{|S|}$ obtained by dividing the discrete composition with a scale factor V ,

$$q = \frac{n}{V}.$$

For example, in molecular dynamics, we use the Avogadro's number $V = 6.022 \times 10^{23}$ to obtain a molar concentration q from molecular count n .

It can be shown that, for large scale factors V , the effect of stochasticity on the (continuous) species composition dynamics are exponentially suppressed. In fact, in the limit $V \rightarrow \infty$ the system is governed by deterministic equations, which are commonly referred to as *Mass-Action Kinetics* (MAK). A central feature of MAK, that makes it especially useful for describing non-linear phenomena, is its ability to exhibit multiple types of *attractors* in the composition space, namely: discrete attractors, limit cycles, and chaotic attractors (Yu and Craciun, 2018; Willamowski and Rössler,

1980). A CRN exhibiting multiple attractors is said to be *multistable*. Each type of attractor poses its own set of theoretical and computational challenges, and for the purpose of this work, we restrict ourselves only to *discrete attractors*.

Formally, the (discrete) attractors are defined as the stable fixed points of the MAK such that any perturbation around the fixed point results in the system relaxing towards it. There are, however, other fixed points where, if the initial conditions are along a particular line called their repelling dimension, the system relaxes away from the point. We term the fixed points with exactly one repelling dimension as *transient points*, and generically refer to any fixed point with strictly more than zero repelling dimensions as *unstable points*. A detailed mathematical description of this taxonomy can be found in Appendix C.3. We state without proof that, every two attractors must have at least one transient point that separates them.

In systems with a small scale factor V , the underlying stochasticity of the model leads to an appreciable deviation of the dynamics from MAK. An effect of particular interest for our formalism is that of *basin-switching*. In a multistable system, while a system stays at an attractor composition indefinitely under MAK, stochastic dynamics allows for transition from one attractor to another, via a transient point. The paths that the system takes to escape from an attractor into nearby transient points are called *escape paths*. The least-improbable (or most probable) escape paths are

essential to the calculation of the *non-equilibrium potential* (NEP) function (for an exposition, see Appendix C.1). The NEP, denoted by \mathcal{V} , yields a direct measure of the negative descaled log-probability of the system of being in a composition, \mathbf{q} , up to a global additive constant⁴ (see Eq. C.10)

$$\mathcal{V}(\mathbf{q}) + c \asymp -\frac{1}{V} \log(\text{Probability}(\mathbf{q})).$$

In Ch. 3, we provide a functional gradient descent algorithm to compute the escape paths and the associated NEP for any CRN.

Cost: Consider the set of species composition attractors, and denoted it by $\underline{\mathbf{q}}$, where

$$\underline{\mathbf{q}} := \{\underline{\mathbf{q}}_1, \underline{\mathbf{q}}_2, \dots, \underline{\mathbf{q}}_N\}.$$

To simplify terminology, and in accordance with connections between Hamilton-Jacobi theory, optimal control theory, and chemical reaction networks (see Ref. Snarski (2021)), we refer to the value of the NEP (or the negative descaled log probability) at an attractor as its **cost**.

$$\text{Cost}(\underline{\mathbf{q}}) = \mathcal{V}(\underline{\mathbf{q}})$$

⁴We say $f(V) \asymp g(V)$ (read as $f(V)$ is asymptotic to $g(V)$) if

$$\lim_{V \rightarrow \infty} \frac{f(V)}{g(V)} = 1.$$

Note that, alike any other potential, the value of the cost has only a relative significance. In particular, the difference in the cost between two attractors is a measure of the ratio of their log probability, where a less costly attractor is exponentially more probably occupied given a long enough time. Conversely, it is less probable for the system to be at a more costly attractor composition, than another one of lesser cost.

Evolvedness: The notion of **evolvedness** is one that cannot be given a purely mathematical definition since it depends on what the species are, and how unlikely they are to found in equilibrium. Thus, to define evolvedness, we first require that each species be assigned a positive *weight vector*. The weight of a species can quantify, for instance, how difficult it is to assemble the species (Marshall et al., 2021) from the building blocks available in the environment, or the difficulty of discovering that species in equilibrium. The different coefficients of the weight vector for each species measure evolvedness along different dimensions. Let the number of dimensions along which we measure evolvedness be denoted by e . Then, all the weight vectors can be collected in a weight matrix of dimensions $e \times |S|$, which we denote by \mathbf{w} , where

$$\text{weight}(s_i) = \mathbf{w}^i.$$

Once the weights are assigned, then we define the evolvedness of an attractor composition as the weighted sum of the composition,

$$\text{Evolvedness}(\underline{q}) = \sum_{i=1}^S \mathbf{w}^i q_i.$$

For example, let the species be polymers in one monomer. Define the species set as $\{\bar{1}, \bar{2}, \dots, \bar{M}\}$, where \bar{j} denotes a polymer of length j . Suppose the system is coupled to an environment modeled as a chemical reservoir of monomers. Then, we may assign the weight matrix $\mathbf{w} = [1, 2, \dots, M]$, where the weight of each species is a scalar quantity. The evolvedness of an attractor composition \underline{q} then evaluates to

$$\text{Evolvedness}(\underline{q}) = \sum_{i=1}^M i q_i.$$

Defining evolution: Using the terminology developed above, we will define the terms *directional evolution*, *open-ended evolution*, and *historical contingency* for a stochastic CRN.

1. Given a stable (composition) attractor of a stochastic CRN, if there exists a more evolved attractor that is less costly than the current one, then the attractor will be said to be capable of **directional evolution**.
2. If a CRN has infinitely many attractors, and at each attractor, there is at least one more attractor that is less costly and more evolved

than the current one, then such a stochastic CRN will be said to be capable of **open-ended evolution**.

3. If there are at least two more evolved and less costly attractors between which the system can only (predominantly) transition through less evolved and costlier attractors, then the stochastic CRN will be said to exhibit **historical contingency**. This follows because, starting from a less evolved more costly attractor, the trapping of the system in any of the more evolved less costly attractor is a contingent choice which affects its future evolution.

4.2 CRN models of evolving ecologies

In this section, we provide examples of different CRN ecologies that exhibit evolution. We briefly review the formalism for ascertaining evolution in Sec. 4.2.1. In Sec. 4.2.2, we explore a CRN in one species and demonstrate that, with an appropriate choice of rate constants and reaction set, it can be made to exhibit N attractors for any N . Considering the distance of an attractor from the origin as a measure of its evolvedness, we prove in Ref. Gagrani et al. (2023b) that there exist rate constant assignments such that for any attractor, the cost of each more evolved attractor is strictly lesser than its cost. By our characterization of the different types of evolution in Sec. 4.1.3, this system is capable of directional evolution from all but the most evolved attractor. Moreover, we conjecture that a rate constant

assignment can also be found in the limit $N \rightarrow \infty$, in which case this system will be said to be capable of open-ended evolution.

In Sec. 4.2.3, we consider a polymer model in one monomer. Against a background of polymerization and degradation reactions, we add additional reactions that select a polymer to play the functional role of a catalyst for monomer activation, and refer to the polymer as a *kinase*. In the presence of such a kinase, we show that the polymer model forms an ecology, and under appropriate rate constants, the system can exhibit multistability. Furthermore, we show that the analysis of multistability is similar to the previously considered one-dimensional CRN. Using the algorithm from Ch. 3, we show that this model can exhibit directional evolution from any but the most evolved attractor. Finally, in Sec. 4.4, we consider a polymer model in two monomers. To each polymer, we assign a two-dimensional evolvedness using the abundance of each monomer in it. Then, by selecting certain polymers as kinases and appropriate rate constants for the system, we show that the model has the potential to exhibit two evolved attractors, albeit in different dimensions. We show that it is highly unlikely that the system transitions between the two evolved attractors, except by going through a more costly and less evolved attractor. Thus, we conclude that our polymer model in two monomers can exhibit historical contingency.

4.2.1 Procedure for ascertaining evolution

Consider a CRN with kinetics $\mathcal{G}_K = (\mathcal{S}, \mathcal{R}, K)$. Since all CRNs in this section will be taken with kinetics, to avoid notational clutter we will simply denote the CRN as \mathcal{G} . The Hamiltonian for the CRN is written as (see Eq. C.19 in Appendix C.3)

$$H_{\text{CRN}}(\mathbf{p}, \mathbf{q}) = \sum_{\mathcal{G}} (e^{(y'-y)\mathbf{p}} - 1) k_{y \rightarrow y'} \mathbf{q}^y. \quad (4.1)$$

The equations of mass-action kinetics for this CRN are given by (see Eq. C.20)

$$\frac{d\mathbf{q}}{dt} = \sum_{\mathcal{G}} (y' - y) k_{y \rightarrow y'} \mathbf{q}^y. \quad (4.2)$$

Let us denote the fixed-point or attractors of the CRN by \underline{q} , where

$$\left. \frac{d\mathbf{q}}{dt} \right|_{\underline{q}} = 0.$$

Each attractor can be assigned stability by the criterion explained in Appendix C.3.2.

The non-equilibrium potential (NEP) for the system, $\mathcal{V}(\mathbf{q})$, is found by solving (see Eq. C.11)

$$H \left(\frac{\partial \mathcal{V}}{\partial \mathbf{q}} \neq 0, \mathbf{q} \right) = 0, \quad (4.3)$$

and using

$$\mathcal{V}(\mathbf{q}) = \int \frac{\partial \mathcal{V}}{\partial \mathbf{q}} d\mathbf{q}.$$

This step is computationally difficult, and analytic solutions are infeasible in all but the simplest cases. We provide a functional gradient descent algorithm to obtain the NEP for a general CRN in Ch. 3.

For notational clarity, let us denote each stable root by the row vector \mathbf{t}_i and refer to the set of stable roots with \mathbf{t} , where

$$\mathbf{t} := \{\mathbf{t}_1, \dots, \mathbf{t}_{|\mathbf{t}|}\}.$$

Define a weight matrix \mathbf{w} quantifying the evolvedness of each species. The evolvedness of each stable root is then given by

$$\text{Evolvedness}(\mathbf{t}_i) = \mathbf{w}\mathbf{t}_i^T.$$

The cost of the attractors is given by $\mathcal{V}(\mathbf{t}_i)$

$$\text{Cost}(\mathbf{t}_i) = \mathcal{V}(\mathbf{t}_i).$$

Finally, using the criteria specified at the end of Sec. 4.1.3, we ascertain which type of evolution, if any, is the CRN capable of.




Figure 4.1: The CRN \mathcal{G}_2 is shown with two rate constants that admit directed evolution. The rate of change given by mass-action kinetics (MAK) and an arbitrarily scaled non-equilibrium potential (NEP) are overlaid in the display. The rate constants $\{k_0, k_1, k_2, k_3\}$ for the left and right panel are $\{10, 17, 8, 1\}$ and $\{12, 20, 9, 1\}$, respectively.

4.2.2 Open ended evolution in a one-dimensional CRN

Let us consider a stochastic chemical kinetic system with 1 species, denoted by X .

One reversible reaction: Suppose that there is an external reservoir from which X can flow directly into and out of our system with rate constants k_0 and k_1 , respectively. This yields our first CRN, which we refer to as \mathcal{G}_1 , where

$$\mathcal{G}_1 = \left(\{X\}, \left\{ \emptyset \xrightleftharpoons[k_1]{k_0} X \right\} \right).$$

Let us denote the concentration of X simply by q , and its conjugate momentum by p . Using Eq. 4.1, the Hamiltonian for the CRN \mathcal{G}_1 is given

by

$$H_1(p, q) = (e^p - 1)k_0 + (e^{-p} - 1)k_1q.$$

Using Eq. 4.2, the relaxation trajectory for the system is given by

$$\frac{dq}{dt} = k_0 - k_1q.$$

As can be seen from the above equation, it has a unique complex-balanced attractor at

$$\underline{q} = \frac{k_0}{k_1}.$$

Since there is a unique complex-balanced attractor for the CRN, as proved in Appendix B of Ch. 3, the

$$H\left(\frac{\partial \mathcal{V}}{\partial q}, q\right) = 0, \quad \frac{\partial \mathcal{V}}{\partial q} \neq 0$$

solution for the system is given by the momentum assignment


$$\frac{\partial \mathcal{V}_1}{\partial q} = \ln\left(\frac{q}{\underline{q}}\right) = \ln\left(\frac{k_1q}{k_0}\right).$$

This momentum assignment yields the NEP \mathcal{V}_1 as

$$\begin{aligned} \mathcal{V}_1 &= \int_{\underline{q}}^q dq' \ln\left(\frac{q'}{\underline{q}}\right) \\ &= q \ln\left(\frac{q}{\underline{q}}\right) - (q - \underline{q}), \end{aligned} \tag{4.4}$$

where $\underline{q} = k_0/k_1$. Since this CRN has only one type of attractor, no evolution is possible.

Two reversible reactions: Now, let us add another reversible reaction channel



to our CRN. Let us denote the resultant reaction network by \mathcal{G}_2 , where


$$\mathcal{G}_2 = \left(\{X\}, \left\{ \emptyset \xrightleftharpoons[k_1]{k_0} X, 2X \xrightleftharpoons[k_3]{k_2} 3X \right\} \right).$$

This is the well-studied *Schlögl-model* (Wilhelm, 2009), and, as we will see, it has the potential to exhibit multistability.

Following terminology defined in Appendix C.2, notice that \mathcal{G}_2 has 4 complexes, 2 linkage classes, and a stoichiometric subspace of dimension 1. Using Eq. C.18, it has a deficiency δ of one,

$$\delta = 4 - 2 - 1 = 1.$$

It can be seen that it has one independent *minimal delta null subnetwork* (MDNS)



a null-flow for which is given by $[1, 1]^T$. Moreover, using the definitions detailed in Sec. 4.1.2, \mathcal{G}_2 has one *minimal autocatalytic subnetwork* (MAS),

given by

$$\text{MAS} = \{2X \rightarrow 3X\},$$

with the autocatalytic set $\{X\}$, and one *minimal drainable subnetwork* (MDS), given by

$$\text{MDS} = \{3X \rightarrow 2X\},$$

with the drainable set $\{X\}$. *Since this CRN has an autocatalytic subnetwork, it forms an ecology.*

The Hamiltonian for \mathcal{G}_2 is

$$H_2(p, q) = (e^p - 1)(k_0 + k_2 q^2) + (e^{-p} - 1)(k_1 q + k_3 q^3).$$

The relaxation trajectories are given by

$$\frac{dq}{dt} = (k_0 + k_2 q^2) - (k_1 q + k_3 q^3).$$

Notice that the above equation can have up to three roots in the positive orthant, two of which may be stable. The relationship of the roots to the rate constants is simply given by the symmetric polynomials in three variables using Viète's formula (Korniłowicz and Pał, 2017).

As can be solved from the $H\left(\frac{\partial \mathcal{V}}{\partial q}, q\right) = 0$ condition, the non-zero momentum for the escape trajectories is given by

$$\frac{\partial \mathcal{V}}{\partial q} = \ln\left(\frac{k_1 q + k_3 q^3}{k_0 + k_2 q^2}\right),$$

yielding the NEP \mathcal{V} as

$$\mathcal{V}_2(q) = \int_{q^*}^q dq \ln \left(\frac{k_1 q + k_3 q^3}{k_0 + k_2 q^2} \right),$$

where

$$q^* = \arg \min_q \int_0^q dq \ln \left(\frac{k_1 q + k_3 q^3}{k_0 + k_2 q^2} \right).$$


Since we have only one species in the system, we simply assign the weight vector [1] as a measure of its evolvedness. For the two attractors $\{t_1, t_2\}$, their degree of evolvedness is thus equal to their difference from the origin, where the attractor further from the origin is more evolved. Notice that the rate constants can be picked to move the attractors around such that, for

$$0 < t_1 < t_2,$$

$$\mathcal{V}_2(t_1) > \mathcal{V}_2(t_2).$$

Some examples of such assignments can be found in Fig. 4.1. *Thus, the ecology given by this CRN is capable of directed evolution from the least evolved attractor.*

N reversible reactions: In general, we can add more reversible reaction channels to our system of the form



to our CRN. Since each reaction channel only adds or removes a single X species, the resultant network is called a *birth-death* process. Denoting the resultant reaction network by \mathcal{G}_N , where

$$\mathcal{G}_N = \left(\{X\}, \left\{ \emptyset \xrightleftharpoons[k_1]{k_0} X, 2X \xrightleftharpoons[k_3]{k_2} 3X, \dots, (2N-2)X \xrightleftharpoons[k_{2N-1}]{k_{2N-2}} (2N-1)X \right\} \right).$$

The Hamiltonian for \mathcal{G}_N is

$$H_N(p, q) = (e^p - 1)r_+(q) + (e^{-p} - 1)r_-(q),$$

where

$$\begin{aligned} r_+ &= k_0 + k_2 q^2 + \dots + k_{2N-2} q^{2N-2}, \\ r_- &= k_1 q + k_3 q^3 + \dots + k_{2N-1} q^{2N-1}. \end{aligned}$$

Notice that $dq/dt = r_+ - r_-$ can have $2N - 1$ roots in the positive orthant with appropriate rate constants.

In general, for this model, considering more *birth-death* reaction channels (Lazarescu et al., 2019; Anderson et al., 2015), simply adds terms to the r_+ and r_- polynomials. Using the same reasoning outlined above, the NEP \mathcal{V}_N becomes

$$\mathcal{V}_N(q) = \int_{q^*}^q dq \ln \left(\frac{r_-(q)}{r_+(q)} \right).$$

Labelling the attractors by t , where

$$0 < t_1 < t_2 < \dots < t_N,$$

we **prove** in Ref. Gagrani et al. (2023b) that a rate constant assignment can be found such that

$$\mathcal{V}_N(t_1) > \mathcal{V}_N(t_2) > \dots > \mathcal{V}_N(t_N).$$

Thus, the ecology given by this CRN containing N reactions is capable of directional evolution from all but the most evolved attractor.

Infinite reversible reactions: It is clear, from the arguments above, that N can be made arbitrarily large. By adding infinitely many such reactions, labelling the attractors by t , where

$$0 < t_1 < t_2 < \dots < t_c < \dots,$$

we conjecture in Ref. Gagrani et al. (2023b) that, under some constraints, a rate constant assignment can be found such that

$$\mathcal{V}_\infty(t_1) > \mathcal{V}_\infty(t_2) > \dots > \mathcal{V}_\infty(t_c) > \dots$$

If the above is true, then the ecology given by this CRN containing infinite reactions will be said to be capable of open-ended evolution.




Figure 4.2: Some polymerization reactions of polymers in a single monomer. The monomer is labelled as $\bar{1}$, the polymers are labelled with the number of monomers they contain, and different geometries of polymers are distinguished by primes (').

4.2.3 Directional evolution in polymer models in one monomer

Consider a continuously-stirred tank reactor that is connected to an external reservoir of monomers, from which the monomers, labelled as $\bar{1}$, can flow into and out of the reactor with rates k_i and k_o , respectively,

$$\mathcal{G}_{\text{Env.}} = \{\{\bar{1}\}, \left\{ \emptyset \xrightleftharpoons[k_o]{k_i} \bar{1} \right\}\}.$$

Suppose the monomers can form bonds between themselves to form longer polymers. As shown in Fig. 4.2, the structure of the polymers can in general be open (see $\bar{3}$), closed (see $\bar{3}'$), or can have multiple intersections (not shown). For demonstrating a simple polymer ecology that can exhibit

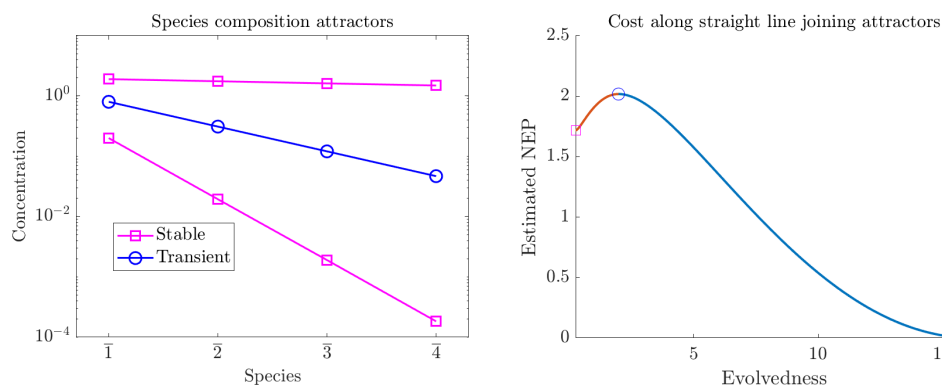


Figure 4.3: An example of rate constant assignments for the CRN \mathcal{G}_c that exhibits directional evolution. The rate constants $\{k_i, k_o, l_c, m_c\}$ are chosen to be $\{0.1476, 1, 2.9, 1\}$. All spontaneous addition and degradation rate constants are chosen to be identical, with $s_i = s = 1$ and $d_i = d = 2.06$. The species composition attractors along with their stability are displayed in the left panel. An estimate of the NEP using the AFGD algorithm by finding the Action along straight lines joining the attractors are shown in the right panel. The measure of evolvedness along the joining curve is shown along the x -axis.

directional evolution, we only consider open polymers that do not have any self intersections.


Let us label the (open) polymer with n monomers as \bar{n} . Consider the space of all reactions that the polymers can participate that conserve the total number of monomers in the system. Let us denote the reactant and product complex for such a reaction by c^α , where

$$c^\alpha := \sum_n c_n^\alpha \cdot \bar{n}.$$

Then, a reaction $c^\alpha \rightarrow c^\beta$ is allowed only if

$$\sum_n c_n^\alpha \cdot n = \sum_n c_n^\beta \cdot n.$$

The reaction network for such a system includes reactions of the type



Such reaction networks, that are defined through conservation laws, are termed cluster chemical reaction networks (CCRN) and developed further in Ch. 2. Observe that the CRN of polymers, when not constrained to a subset of reactions, contains many autocatalytic subnetworks and thus forms an ecology. An enumeration of the autocatalytic subnetworks in the CCRN with only up to second-order reactions can be found in Ch. 2, and it is also shown that this count grows exponentially in the size of the longest polymer in the system.

Null-model: To simplify matters, let us restrict our CRN to only reactions where the polymers elongate via *end additions* of a monomer.

We denote the CRN, with up to polymers of size n , thus obtained as

$\mathcal{G}_{\text{End}} = (\mathcal{S}, \mathcal{R}, \mathcal{K})$, where

$$\begin{aligned} \mathcal{S} &= \{\bar{1}, \bar{2}, \dots, \bar{n}\}, \\ (\mathcal{R}, \mathcal{K}) &= \left\{ \begin{aligned} &\bar{1} + \bar{1} \xrightleftharpoons[d_1]{s_1} \bar{2}, \\ &\bar{1} + \bar{2} \xrightleftharpoons[d_2]{s_2} \bar{3}, \\ &\bar{1} + \bar{3} \xrightleftharpoons[d_3]{s_3} \bar{4}, \\ &\dots \\ &\overline{n-1} + \bar{1} \xrightleftharpoons[d_{n-1}]{s_{n-1}} \bar{n} \end{aligned} \right\}. \end{aligned}$$

Here s_i and d_i are the rate constants of spontaneous addition and degradation of a monomer to and from a polymer of size i , respectively.

Our model of background processes, or *null-model*, is the CRN $\mathcal{G}_{\text{Null}}$, where

$$\mathcal{G}_{\text{Null}} = \mathcal{G}_{\text{Env.}} \cup \mathcal{G}_{\text{End.}}$$

While we can write a Hamiltonian for this model, it is not particularly illuminating. Instead, we directly investigate the attractors for this model. Denoting the concentration of species \bar{i} by q_i , and its rate of change by \dot{q}_i ,

the MAK for the CRN is

$$\begin{aligned}
 \dot{q}_n &= -d_{n-1}q_n + s_{n-1}q_1q_{n-1}, \\
 \dot{q}_{n-1} &= -(d_{n-2} + s_{n-1}q_1)q_{n-1} + (s_{n-2}q_1q_{n-2} + d_{n-1}q_n), \\
 &\dots \\
 \dot{q}_2 &= -(d_1 + s_2q_1)q_2 + (s_1q_1^2 + d_2q_3), \\
 \dot{q}_1 &= (-k_oq_1 + k_i) + \sum_{i=1}^{n-1} (-s_iq_iq_1 + d_iq_{i+1}) \\
 &\quad + (-s_1q_1^2 + d_1q_2).
 \end{aligned}$$


Elementary considerations yield that there is a single species composition attractor for the CRN, and it is given by

$$\begin{aligned}
 \underline{q}_1 &= \frac{k_i}{k_o}, \\
 \underline{q}_i &= \prod_{j=1}^{i-1} \left[\frac{s_j}{d_j} \right] (\underline{q}_1)^i \text{ for } i \in [2, n].
 \end{aligned}$$

Since there is only a single attractor, the null-model is not an evolving CRN.


Catalysts for monomer formation It is common for polymers to play a functional role in the system from which they emerge. Suppose that the polymer of size c , \bar{c} , acts as a catalyst for monomer formation. In accordance with literature, We refer to such a polymer as a *kinase* (see Ref. Kannaiyan and Mahadevan (2018)). Introducing a kinase then amounts

to adding the reactions of the form



to our CRN.

Let us denote the CRN resulting from taking the union of the null-model and reaction R_c as \mathcal{G}_c . Observe that the reactions



form an autocatalytic subnetwork with food set $\{\overline{c-1}\}$ and member set $\{\bar{1}, \bar{c}\}$. Thus, the resulting CRN \mathcal{G}_c is an ecology.

It is also easy to see that \mathcal{G}_c has non-zero deficiency (see Appendix C.2.1) and can exhibit multistability for an appropriate choice of rate constants. The rates of polymers of size greater than 1 is left unchanged by the addition of R_c , and thus they retain their functional form

$$\underline{q}_i = \prod_{j=1}^{i-1} \left[\frac{s_j}{d_j} \right] (q_1)^i \text{ for } i \in [2, n].$$

However, the attractor concentration of $\bar{1}$, denoted by \underline{q}_1 , is given by roots

of the equation

$$0 = (-k_o \underline{q}_1 + k_i) + \prod_{j=1}^{c-1} \left[\frac{s_j}{d_j} \right] (l_c - m_c \underline{q}_1) (\underline{q}_1)^c.$$

Notice that the above equation is of the same form as that considered in the previous subsection, Sec. 4.2.2. In particular then, we can add any number of such catalysts to obtain any number of roots, where the position of the roots are determined by the rate constants.

Consider the case with only one kinase \bar{c} . The rate constants can then be chosen such that the above equation has up to three roots. Analogous to the setup in Ref. Gagrani et al. (2023b), let us pick the three roots such that the root closer to the origin has multiplicity two. It is easy to see that the root furthest from the origin is an attractor, and the root closer to the origin is an inflection point. Notice that, starting from any point to the right of the inflection point, under MAK, the system will relax to the attractor (or the fixed point furthest from the origin). As shown in Ch. 3, since the NEP is a Lyapunov function in the MAK, the value of the NEP is monotonically decreasing from the inflection point to the attractor. Now, if we slightly displace one of the lower roots to a value lower than the inflection point, then the root closest to the origin is an attractor and the inflection point becomes a transient point. Due to the same reason as above, the value of the NEP is maximum at the transient point and monotonically decreases to its left until the attractor. However, since the value of the NEP changes

continuously with the position of the roots, the lower root can be chosen to be close enough to the middle root such that the NEP at the root closest to the origin is higher than the NEP at the further root.

To each polymer, let us assign an evolvedness which is equal to the number of monomers it is consisted of. Thus, the evolvedness of an attractor composition t is given by

$$\text{Evolvedness}(t) = \sum_{j=1}^n j t_j.$$

In order to ascertain whether such a model is capable of directional evolution, we need to solve the Hamilton-Jacobi equation and find the NEP \mathcal{V}_c . For the particular case shown in Fig. 4.3, we employ the AFGD algorithm from Ch. 3 and calculate \mathcal{V}_c . While these are not exact solutions to the Hamilton-Jacobi equations, the difference in their magnitudes, and a few iterations of the algorithm suggest that the actual NEP in the more evolved attractor is lesser than the less evolved state for this choice of rate constants. As is clear from the figure and the argument in the preceding paragraph, there exists some choice of rate constants such that, if the attractors are such that t_1 is less evolved than t_2 , then

$$\mathcal{V}_c(t_1) > \mathcal{V}_c(t_2).$$

Thus, this CRN is capable of directional evolution from its least evolved attractor.

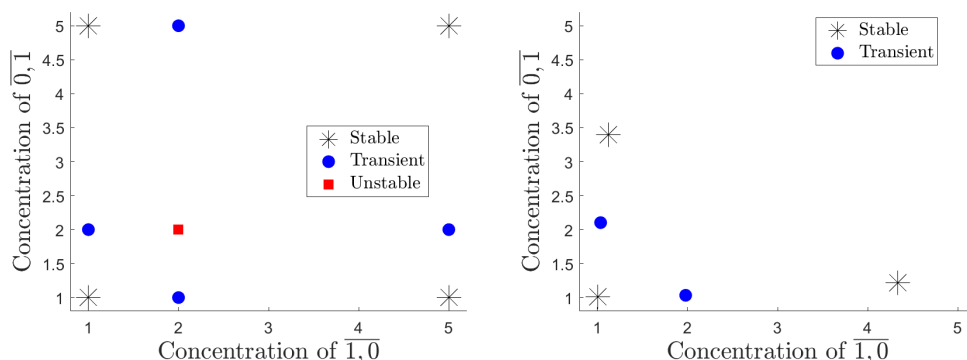


Figure 4.4: Attractors for CRNs \mathcal{G}_L and \mathcal{G}_R from Sec. 4.2.4 are shown on the left and right panel, respectively. In the left panel, there is a unique most evolved attractor and the other evolved attractors are directly connected to it through a transient point. Thus, this system is capable of directional evolution. In the right panel, there is no unique most evolved attractor, and a transition between the most evolved attractors can only occur through a less-evolved attractor. Thus, this system is capable of exhibiting historical contingent evolution.

Reasoning by analogy from the 1-dimensional model, we argue that this CRN can be made to be directionally evolving from all but its most evolved attractor for any number of attractors. Moreover, we conjecture that by adding infinite catalysts, the rate constants can be chosen to make the CRN exhibit open-ended evolution.



4.2.4 Historically contingent evolution in polymer models in two monomers

Let us consider polymers composed of two different types of monomers, labelled A and B. Since we are restricting to only open polymers, or polymers that avoid self-intersection, we can identify each polymer by its sequence of monomers. For instance, the complete list of polymers in two

monomers of size three is:

$$\{AAA, AAB, ABA, BAA, ABB, BAB, BBA, BBB\}.$$

Let ω_2^n denote an arbitrary polymer of size n in two monomer. Then the end addition reactions are of the type:



In all of the different cases, the reaction is of the form $\omega^n \rightarrow \omega^{n+1}$.

Notice that there are 2^n polymers of size n in two monomers,

$$|\omega_2^n| = 2^n.$$

While we can enumerate the complete reaction network for our model, it is not feasible to perform analysis with an exponentially growing list of species. To overcome this challenge, we introduce a coarse-graining CCRN with polynomially growing species in Ch. 2. While the formalism is developed there, the basic construction is this. For each polymer, instead of tracking the exact sequence of monomers, we only count the number of monomers of each type. These letter-count sets are instances of clusters (as letters are conserved), and we refer to the species and CRN thus

obtained as *clusters* and *two-constituent CCRN*, respectively. For example, the polymer AABA maps to the cluster (3A, 1B). In what follows, we use the basis (A, B) and denote the clusters as $\overline{i, j}$, where

$$(iA, jB) \rightarrow \overline{i, j}.$$

Polymerization CRN: Let our simplified CRN of the polymer in two monomers, or the two-constituent CCRN, be denoted as $\mathcal{G}_{\text{Poly.}} = (\mathcal{S}, \mathcal{R}_{\text{Poly.}})$. Then the species set is given by

$$\mathcal{S} = \{\overline{1, 0}, \overline{0, 1}, \overline{1, 1}, \dots, \overline{i, j}, \dots\},$$

and the reaction set is given by

$$\mathcal{R}_{\text{Poly.}} = \left\{ \left\{ \overline{i, j} + \overline{1, 0} \xrightleftharpoons[\text{dA}_{ij}]{\text{sA}_{ij}} \overline{i+1, j}, \right. \\ \left. \overline{i, j} + \overline{0, 1} \xrightleftharpoons[\text{dB}_{ij}]{\text{sB}_{ij}} \overline{i, j+1} \right\} \left| \forall i, j \in \mathbb{Z}_{\geq 0}, i+j \geq 1 \right\}.$$

In general, the attractors of this CRN cannot be analytically found. To simplify analysis, consider the case that all spontaneous addition and

degradation rate constants are identical to one another,

$$sA_{ij} = sB_{ij} = s \quad \forall i, j,$$

$$dA_{ij} = dB_{ij} = d \quad \forall i, j.$$

Then the attractor concentration for any cluster can be found in terms of the attractor concentration of the monomers, and is of the form


$$\underline{q}_{i,j} = \left(\frac{s}{d}\right)^{i+j-1} \left(\underline{q}_{1,0}\right)^i \left(\underline{q}_{0,1}\right)^j. \quad (4.5)$$

Monomer formation and catalysts: Suppose the system was in contact with a chemical reservoir from which monomers $\overline{1,0}$ and $\overline{0,1}$ could flow in and out of the system. Let us denote the CRN containing reactions with the environment as $\mathcal{G}_{\text{Env.}} = (\mathcal{S}_{\text{Env.}}, \mathcal{R}_{\text{Env.}})$, where


$$\begin{aligned} \mathcal{S}_{\text{Env.}} &= \{\overline{1,0}, \overline{0,1}\} \\ \mathcal{R}_{\text{Env.}} &= \left\{ \emptyset \xrightleftharpoons[k_{A_o}]{k_{A_i}} \overline{1,0}, \right. \\ &\quad \left. \emptyset \xrightleftharpoons[k_{B_o}]{k_{B_i}} \overline{0,1} \right\}. \end{aligned}$$

Furthermore, suppose some polymers acted as catalysts for monomer formation. As argued in the last subsection, the addition of each catalyst yields a new autocatalytic subnetwork, making such a CRN an *ecology*.

Let us denote the set of catalysts and reversible reactions that catalyze monomer formation as \mathcal{C} and $\mathcal{R}_{\text{Cat.}}$, respectively, where $\mathcal{R}_{\text{Cat.}}$ contains reactions of the form



or



for each $\overline{e, f} \in \mathcal{C}$. Notice that not all rate constants have to be positive, thus the catalysts can selectively catalyze one but not the other monomer formation. Let us denote the CRN of catalyzation reactions as $\mathcal{G}_{\text{Cat.}}$, where $\mathcal{G}_{\text{Cat.}} = (\mathcal{C}, \mathcal{R}_{\text{c}})$.

Consider the CRN \mathcal{G} that contains in-flow and out-flow reactions $\mathcal{G}_{\text{Env.}}$, polymerization reactions $\mathcal{G}_{\text{Poly.}}$, and reactions that catalyze monomer concentration $\mathcal{G}_{\text{Cat.}}$,

$$\mathcal{G} = \mathcal{G}_{\text{Env.}} \cup \mathcal{G}_{\text{Poly.}} \cup \mathcal{G}_{\text{Cat.}}$$

Analogous to analysis in the last subsection, the attractors of the CRN can

be found by using the roots of the system:

$$\begin{aligned}
 0 &= (kA_i - kA_o) \underline{q}_{1,0} + \\
 &\quad \sum_{e,f \in \mathcal{C}} \left(\frac{s}{d}\right)^{e+f-1} (lA_{e,f} - mA_{e,f} \underline{q}_{1,0}) (\underline{q}_{1,0})^d (\underline{q}_{0,1})^e, \\
 0 &= (kB_i - kB_o) \underline{q}_{0,1} + \\
 &\quad \sum_{e,f \in \mathcal{C}} \left(\frac{s}{d}\right)^{e+f-1} (lB_{e,f} - mB_{e,f} \underline{q}_{0,1}) (\underline{q}_{1,0})^d (\underline{q}_{0,1})^e,
 \end{aligned}$$

along with Eq. 4.5. Observe that the roots of the above system of polynomials yield attractors in the monomer concentrations, and the polymer concentration are fixed by the $\mathcal{G}_{\text{Poly}}$.

Having obtained attractor compositions, we need to define a notion of their evolvedness. Since the species of this system are composed of two monomers, both of which are of equal importance, we assign a vector valued weight to each species. In particular, to the cluster species $\overline{i,j}$, let us assign the weight vector $[i,j]^T$,

$$\text{Weight}(\overline{i,j}) = [i,j]^T.$$

Thus, the evolvedness of an attractor composition \underline{q} will be given by

$$\text{Evolvedness}(\underline{q}) = \sum_{i,j} q_{i,j} [i,j]^T.$$

Notice that the evolvedness is a vector quantity, each dimension quanti-

fying the evolvedness in a given monomer. For the case considered in this work, where the attractor concentration of the large clusters takes a product form of the monomer concentrations, the evolvedness are monotonically increasing in the concentrations of the monomers. Thus, we will simply use the concentration of the monomers $\overline{1,0}$ and $\overline{0,1}$ at the attractor as a proxy for the vector-valued evolvedness of an attractor.

Now we will consider two concrete examples of CRNs of the above type, denoted as $\mathcal{G}_L = (\mathcal{S}, \mathcal{R}_L)$ and $\mathcal{G}_R = (\mathcal{S}, \mathcal{R}_R)$. In both cases \mathcal{S} is the space of all two constituent clusters, which can be identified with the positive orthant of the integer lattice in two dimensions.

Let the first reaction set be given by :

$$\begin{aligned} \mathcal{R}_L = & \mathcal{R}_{\text{Poly.}|s=1,d=5.01} \cup \\ & \left\{ \begin{aligned} & \emptyset \xrightleftharpoons[17]{10} \overline{1,0}, \\ & \overline{2,0} \xrightleftharpoons[5.01]{40.08} \overline{2,0} + \overline{1,0}, \\ & \emptyset \xrightleftharpoons[17]{10} \overline{0,1}, \\ & \overline{0,2} \xrightleftharpoons[5.01]{40.08} \overline{0,2} + \overline{0,1} \end{aligned} \right\}. \end{aligned}$$

The attractors in the monomer concentrations are shown in the left panel of Fig. 4.4. It can be seen that the CRN has four attractors $\{(1, 1), (1, 5), (5, 1), (5, 5)\}$. As explained in the last paragraph, we can use the concentration of the monomers as a proxy for the vector-valued evolvedness of the attractors.

Thus, clearly, $(1, 1)$ and $(5, 5)$ are the least and most evolved attractors, respectively. Due to symmetry of the model and similar considerations made earlier in this subsection, the ordering of the cost of the attractors is such that:

$$\text{Cost}(1, 1) < \text{Cost}(1, 5) = \text{Cost}(5, 1) < \text{Cost}(5, 5).$$

Hence we conclude that \mathcal{G}_L is capable of directional evolution.

Let the second reaction set be given by :

$$\begin{aligned} \mathcal{R}_L = & \mathcal{R}_{\text{Poly.}}|_{s=1, d=5.01} \cup \\ & \left\{ \emptyset \xrightleftharpoons[17]{10} \overline{1, 0}, \\ & \overline{2, 0} \xrightleftharpoons[5.01]{40.08} \overline{2, 0} + \overline{1, 0}, \\ & \overline{3, 2} \xrightleftharpoons[13.54]{31.5} \overline{3, 2} + \overline{1, 0}, \\ & \emptyset \xrightleftharpoons[17]{10} \overline{0, 1}, \\ & \overline{0, 2} \xrightleftharpoons[5.01]{40.08} \overline{0, 2} + \overline{0, 1}, \\ & \overline{2, 3} \xrightleftharpoons[36.82]{56.7} \overline{2, 3} + \overline{0, 1} \right\}. \end{aligned}$$

The attractors in the monomer concentrations are shown in the right panel of Fig. 4.4. It can be seen that the CRN has three attractors, roughly around $\{(1, 1), (4, 1), (1, 3)\}$. Although, $(1, 1)$ is the least evolved attractor, the other attractors $(4, 1)$ and $(1, 3)$ are both more evolved albeit in different dimen-

sions. Also, observe that a transient point separates each of the more evolved attractor from the lesser evolved one. Thus, from considerations explained in Ch. 3, suppose the system were to reach an evolved attractor, the least improbable path out of the attractor would be through the transient point to a lesser evolved attractor. In other words, for large enough population sizes, the system will almost never transition from one of the more evolved attractors directly into the other evolved attractor, and will always go through the lesser evolved attractor. *Thus, we conclude that the ecology given by \mathcal{G}_R is capable of exhibiting historically contingent evolution.*

4.3 Discussion

Stochastic processes are central to modeling complex systems with very many degrees of freedom. Stochastic CRNs, a particular type of stochastic process, models the scenario where the species are well-mixed in the system and reactions proceed randomly by bringing different species into contact. In this work, we provide a mathematically rigorous formalism for ascertaining when a CRN forms an ecology, and when a stochastic CRN is capable of directional, open-ended, or historically contingent evolution. We then employ the formalism on polymer models in one and two monomers and show that there are kinetic parameters such that the stochastic system is guaranteed to evolve towards attractors with greater abundance of longer polymers.

The polymer models, worked in Sec. 4.2, are intended to provide answers to the questions raised in the introduction (Sec. 1.1). They show that, in an environment which only provides small molecules (monomers) which have an ability to create bigger molecules (polymers), there can exist conditions under which an ecology reliably evolves to a composition with a higher mean molecule size and lower cost. Moreover, evolutionary transitions can also be potentially open-ended because a system can repeatedly evolve to other compositions with even higher mean lengths and lower costs. The general formalism and the precise definitions of *composition*, *ecology*, *evolvedness*, and *cost* were explained in Sec. 4.1.

Limitations and future work: Although we have defined a notion for evolution in stochastic CRNs, the mathematical considerations here are not sufficiently developed to explain the origins of biochemical life in its modern form.

First, the the stochastic CRNs considered here only model a homogeneous species distribution, and makes the *continuously-stirred tank reactor* assumption. On the other hand, Darwinian selection and evolution, characterized by inheritance, differential success, and mutation, relies on a notion of compartmentalization and spatial heterogeneity. As remarked in Sec. 4.1.1, while spatial heterogeneity can be modelled by introducing new sets of species labelled by a spatial index (see Refs. (Muñuzuri and Pérez-Mercader, 2022; Plum and Baum, 2022)), this approach very quickly

becomes infeasible for modelling even modestly sized reaction networks at each spatial site. Moreover, work remains to clearly define a notion of individuality, inheritance, and differential success for a spatially heterogeneous evolving ecology of CRNs.

Second, a key feature of biochemical life is, what is known as the *central dogma* (Li and Xie, 2011), which posits that genetic *information* only flows from nucleic acids to protein. Equivalently, the space of possibilities or the *entropy* of the bonds in a protein is controlled by the RNA from which it was translated in its distant *hereditary* history (Smith, 2008). It is widely understood that different topologies of CRNs can achieve various computations (Ma, 2009; Varghese et al., 2015; Dittrich, 2005). In fact, they are shown to be Turing complete (Fages et al., 2017) and can implement neural networks (Anderson et al., 2021). Thus, we posit that CRNs can model and implement the types of control systems needed to explain the emergence of the central dogma. In our formalism, we recognize the role of autocatalysis in organizing CRN dynamics and forming an ecology (Peng et al., 2020). However, further work is needed to formalize how these ecologies themselves could evolve into more complex ecologies that could eventually yield a control system resembling biochemical life as we know it.

A APPENDIX TO CHAPTER 2

A.1 Deficiency theory and multistability in CRNs

In (Feinberg, 2019), a CRN is defined by the triple $\{\mathcal{S}, \mathcal{C}, \mathcal{R}\}$, where \mathcal{S} , \mathcal{C} , and \mathcal{R} are the set of species, complexes, and reactions respectively. Recall from Sec. 2.2.1 that a CRN is a hypergraph with hypervertices in the set \mathcal{C} and hyperedges in \mathcal{R} . The stoichiometric matrix \mathbb{S} is defined such that (see Eq. 2.1)

$$\text{cols}(\mathbb{S}) = \{\mathbf{y}' - \mathbf{y} \mid \mathbf{y} \rightarrow \mathbf{y}' \in \mathcal{R}\}.$$

Alternatively, \mathbb{S} can be written as the product $\mathbf{Y} \cdot \mathbb{M}$, where the columns of \mathbf{Y} are the stoichiometries of the hypervertices

$$\text{cols}(\mathbf{Y}) = \{\mathbf{y} \mid \mathbf{y} \in \mathcal{C}\}$$

and \mathbb{M} is the oriented vertex-edge incidence matrix of the hypergraph (see Balandin (1940); Feinberg (2019))

$$\text{cols}(\mathbb{M}) = \{\mathbf{e}_{\mathbf{y}'} - \mathbf{e}_{\mathbf{y}} \mid \mathbf{y} \rightarrow \mathbf{y}' \in \mathcal{R}\},$$

where $\mathbf{e}_{\mathbf{y}}$ is the vector with 1 at hypervertex \mathbf{y} and 0 otherwise. Note that \mathbf{Y} and \mathbb{M} are of sizes $|\mathcal{S}| \times |\mathcal{C}|$ and $|\mathcal{C}| \times |\mathcal{R}|$, respectively, yielding \mathbb{S} of size $|\mathcal{S}| \times |\mathcal{R}|$.

Since \mathbb{S} is a linear operator, using the rank-nullity theorem we get

$$\dim(\text{domain}(\mathbb{S})) = \dim(\text{im}(\mathbb{S})) + \dim(\text{ker}(\mathbb{S})). \quad (\text{A.1})$$

Let E, V, s denote the number of reactions (hyperedges), complexes (hypervertices) and the dimension of the stoichiometric subspace S , respectively. Then, by definition,

$$\begin{aligned} \dim(\text{domain}(\mathbb{S})) &= E, \\ \dim(\text{im}(\mathbb{S})) &= s. \end{aligned}$$

Using $\mathbb{S} = Y \cdot \mathbb{M}$, we have

$$\dim(\ker(\mathbb{S})) = \dim(\ker(\mathbb{M})) + \dim(\ker(Y) \cap \text{im}(\mathbb{M})). \quad (\text{A.2})$$

Notice that the complexes and reactions, without referring to the underlying species, form a graph, the vertex-edge incidence matrix of which is given by \mathbb{M} . Using properties of the vertex-edge incidence matrix, if the number of independent loops in the graph are ι , then

$$\dim(\ker(\mathbb{M})) = \iota.$$

Denoting $\dim(\ker(Y) \cap \text{im}(\mathbb{M}))$ by δ , also referred to as the *deficiency* of the CRN, substituting the above and Eq. A.2 in Eq. A.1, we get

$$\delta = E - \iota - s. \quad (\text{A.3})$$

Furthermore, let us denote the number of connected components or *linkage classes* of the complex-reaction graph by ℓ . Notice that ℓ and ι are the zeroth and first Betti numbers of a graph, respectively (Hatcher, 2005). Using the formula for the Euler characteristic of a graph, we have the relation

$$V - E = \ell - \iota, \quad (\text{A.4})$$

where V is the number of vertices in the graph. Substituting the above in Eq. A.3, we obtain the more widely used expression for the deficiency

$$\delta = V - \ell - s. \quad (\text{A.5})$$

For direct derivations of Eq. A.5 without resorting to a counting of independent loops, see Smith and Krishnamurthy (2017); Feinberg (2019).

Intuitively, deficiency is a topological property of the reaction network that informs about the steady state dynamics. In particular, a non-zero deficiency is a necessary, but not sufficient, condition for the dynamics to exhibit *multistability* or multiple steady states. In certain cases, such as for deficiency one networks (Feinberg, 1988), there exist algorithms that can help classify whether or not multiple steady states can exist for a given CRN and identify the rate constants such that under mass-action kinetics it exhibits multistability. Once a model exhibiting multistability




Figure A.1: Representation of the complete 1-constituent CCRN with $L = 6$. Complex label (a, b) on a lattice point corresponds to the complex $\bar{a} + \bar{b}$. All the complexes connected by a red line correspond to a fully connected graph in the reaction network. For example, the red line connecting $\{(5, 0), (4, 1), (3, 2)\}$, labelled $a + b = 5$, corresponds to the reactions $\{\bar{5} \rightleftharpoons \bar{4} + \bar{1} \rightleftharpoons \bar{3} + \bar{2} \rightleftharpoons \bar{5}\}$ in the CCRN.

is finalized, one can also use algorithms from (Gagrani and Smith, 2022; Dykman et al., 1994) to estimate the most-likely path the system will take, along with its probability, to escape from one steady state to another.

A.2 Kernel of the stoichiometric matrix for the complete 1-constituent CCRN

Let $\mathcal{H}_1^L = (\mathcal{S}, \mathcal{C}, \mathcal{R})$ denote a complete 1-constituent CCRN of length L and order two. 1-constituent CCRN of length L means that the species set consists of clusters up to length L , denoted as $\mathcal{S} = \{\bar{1}, \bar{2}, \dots, \bar{L}\}$ (for a detailed definition, see Sec. 2.4.2.) For a diagrammatic representation of the reaction network, see Fig. A.1. Here, following Appendix A.1, we will perform the necessary calculations to calculate the kernel of the stoichiometric matrix for a general length L .

First, we will count the number of complexes in a complete 1-constituent CCRN of length L . Let us denote the set of complexes of size N as \mathcal{C}_L^N . Re-

call that size of a complex is defined as the sum of the conserved quantities in each cluster. For $2 \leq N \leq L$,

$$\mathcal{C}_L^N = \{\overline{N}, \overline{N-1} + \overline{1}, \dots, \overline{N/2} + \overline{N/2}\}.$$

For $L < N < 2L - 2$,

$$\mathcal{C}_L^N = \{\overline{N-L} + \overline{L}, \dots, \overline{N/2} + \overline{N/2}\}.$$

Thus, the number of complexes as a function of N and L are

$$|\mathcal{C}_L^N| = \begin{cases} \lfloor N/2 \rfloor + 1 & \text{for } 2 \leq N \leq L, \\ L + 1 - \lceil N/2 \rceil & \text{for } L < N \leq 2L - 2. \end{cases}$$

The total number of complexes for a CCRN of length L is thus

$$|\mathcal{C}_L| = \sum_N |\mathcal{C}_L^N| = \frac{(L+1)(L+2)}{2} - 4.$$

This can be found by either summing the individual contributions or by observing from Fig. A.1 that the complexes occupy the lower triangle including the diagonal for a lattice with $L + 1$ rows and columns excluding 4 points.

Let us denote the set of reactions of complexes of size N by \mathcal{R}_L^N . Since the reaction network of the complete CCRN is obtained by connecting all allowed reactions, for complexes of size N , the number of reactions are

$$|\mathcal{R}_L^N| = \begin{cases} 2^{\binom{\lfloor N/2 \rfloor + 1}{2}} & \text{for } 2 \leq N \leq L, \\ 2^{\binom{L+1 - \lceil N/2 \rceil}{2}} & \text{for } L < N \leq 2L - 2, \end{cases}$$

where the factor of 2 is multiplied since all reactions are reversible. Since, each N contributes exactly one linkage class (or connected component), the total number of linkage classes ℓ for the complete CCRN is $2L - 3$.


Recall from Appendix A.1, the kernel of the stoichiometric matrix is the sum of the independent loops ι and the deficiency δ , where using

Equations A.4 and A.5 we have

$$\begin{aligned} \iota &= |\mathcal{R}_L| - |\mathcal{C}_L| + \ell & (\text{A.6}) \\ &= \sum_{N=2}^{2L-2} |\mathcal{R}_L^N| - \frac{L^2 - 5L - 2}{4} \end{aligned}$$

$$\begin{aligned} \delta &= |\mathcal{C}_L| - \ell - s & (\text{A.7}) \\ &= \frac{(L+1)(L+2)}{2} - 4 - (2L-3) - (L-1) \\ &= \frac{(L-1)(L-2)}{2}. \end{aligned}$$

The intersection $\text{Im}(\mathbb{M}) \cap \text{Ker}(Y)$, the dimension of which is the deficiency δ , counts the null flows of the CCRN which are not simple loops in the network. While an explicit basis can be obtained using subnetworks of the form




we will not be pursuing it since we do not need it for any result in this work.

A.3 Deficiency-one algorithm for $L=3$ 1-constituent CCRN

The complete $L = 3$ 1-constituent CCRN is the simplest complete CCRN, and as we will see, it possesses non-trivial dynamical properties. In particular, from Table 2.1, it can be seen that the deficiency δ for the CCRN is one. This hints at the possibility that the CRN, when taken with mass-action kinetics, exhibits multistability. In this section, using Feinberg's deficiency-one algorithm from (Feinberg, 1988, 2019), we will confirm that the network indeed has the capability of exhibiting multiple steady states, and find a relation between the rate constants and concentrations where that is the case. The algorithm works by converting the CRN into a

system of linear inequalities, and if they have a solution, they can be used to obtain a relation between the rate constants and multiple steady states.

The $L = 3$ 1-constituent CCRN, denoted by $\mathcal{G} = (\mathcal{S}, \mathcal{C}, \mathcal{R})$, is given by the species set $\mathcal{S} = \{\bar{1}, \bar{2}, \bar{3}\}$, complex set $\mathcal{C} = \{\bar{1} + \bar{1}, \bar{2}, \bar{1} + \bar{2}, \bar{3}, \bar{1} + \bar{3}, \bar{2} + \bar{2}\}$ and reaction set



Notice that the above reaction network is a regular (Feinberg, 1988) deficiency-one reaction network, and is thus a valid candidate for the deficiency-one algorithm. To proceed with our calculation, we will follow closely the example in Chapter 17 of (Feinberg, 2019) and Sec. 5 in (Feinberg, 1988).

The confluence vector $g \in \mathbb{R}^c$ for our network can be calculated to be (up to sign)

$$\begin{aligned}g_{\bar{2}} &= -1, & g_{\bar{1}+\bar{1}} &= 1, \\ g_{\bar{1}+\bar{2}} &= -1, & g_{\bar{3}} &= 1, \\ g_{\bar{1}+\bar{3}} &= -1, & g_{\bar{2}+\bar{2}} &= 1.\end{aligned}$$

For our construction, we will use the following upper-middle-lower partition

$$\begin{aligned}\mathbf{U} &:= \{\bar{1} + \bar{3}, \bar{2} + \bar{2}\}, \\ \mathbf{M} &:= \{\bar{1} + \bar{2}, \bar{3}\}, \\ \mathbf{L} &:= \{\bar{1} + \bar{1}, \bar{2}\}.\end{aligned}$$

The partition induces the following system of inequalities on the chemical potential,

$$2\mu_{\bar{2}} > \mu_{\bar{1}} + \mu_{\bar{3}} > \mu_{\bar{3}} = \mu_{\bar{1}} + \mu_{\bar{2}} > \mu_{\bar{2}} > 2\mu_{\bar{1}}.$$

The inequalities are satisfied by

$$\begin{aligned}\mu_{\bar{1}} &> 0, \\ \mu_{\bar{2}} &> 2\mu_{\bar{1}}, \\ \mu_{\bar{3}} &= \mu_{\bar{1}} + \mu_{\bar{2}},\end{aligned}$$

which clearly have a solution, for e.g. $\mu_{\bar{1}} = 1, \mu_{\bar{2}} = 3, \mu_{\bar{3}} = 4$. Thus, the CCRN can exhibit multiple steady states.

To find the rate constants at which the system exhibits multistability, we use subsection 5.3 in (Feinberg, 1988). Following the reference, let us denote the mass-action kinetics rate vector by $\kappa_{y \rightarrow y'} := k_{y \rightarrow y'} x^y$, where x is the concentration vector and $x^y := \prod_{i \in S} x_i^{y_i}$. Choosing the monotonically increasing function to be the exponential $\phi(x) = e^x$, and choosing η so that $\mu_{\bar{1}} + \mu_{\bar{3}} < \eta < \mu_{\bar{2}}$ we get

$$\begin{aligned}\kappa_{\bar{1}+\bar{3} \rightarrow \bar{2}+\bar{2}} &= 1 \cdot \frac{e^{2\mu_{\bar{2}}} - e^{\eta}}{e^{2\mu_{\bar{2}}} - e^{\mu_{\bar{1}}}e^{\mu_{\bar{3}}}} = k_{13}x_{\bar{1}}x_{\bar{3}}, \\ \kappa_{\bar{2}+\bar{2} \rightarrow \bar{1}+\bar{3}} &= 1 \cdot \frac{e^{\mu_{\bar{1}+\bar{3}}} - e^{\eta}}{e^{2\mu_{\bar{2}}} - e^{\mu_{\bar{1}}}e^{\mu_{\bar{3}}}} = k_{22}x_{\bar{2}}^2, \\ \kappa_{\bar{1}+\bar{2} \rightarrow \bar{3}} &= k_{12}x_{\bar{1}}x_{\bar{2}}, \\ \kappa_{\bar{3} \rightarrow \bar{1}+\bar{2}} &= k_{12}x_{\bar{1}}x_{\bar{2}} - 1 = k_3x_{\bar{3}}, \\ \kappa_{\bar{2} \rightarrow \bar{1}+\bar{1}} &= 1 \cdot \frac{e^{2\mu_{\bar{1}}} - e^{\eta}}{e^{2\mu_{\bar{1}}} - e^{\mu_{\bar{2}}}} = k_2x_{\bar{2}}, \\ \kappa_{\bar{1}+\bar{1} \rightarrow \bar{2}} &= 1 \cdot \frac{e^{\mu_{\bar{2}}} - e^{\eta}}{e^{2\mu_{\bar{1}}} - e^{\mu_{\bar{2}}}} = k_{11}x_{\bar{1}}^2.\end{aligned}$$

From (Feinberg, 1988), it is also known that if $x := (x_{\bar{1}}, x_{\bar{2}}, x_{\bar{3}})$, then so is $x^* := (x_{\bar{1}}e^{\mu_{\bar{1}}}, x_{\bar{2}}e^{\mu_{\bar{2}}}, x_{\bar{3}}e^{\mu_{\bar{3}}})$, thus yielding multiple steady states for our system.

B APPENDIX TO CHAPTER 3

B.1 Experimental applications of AFGD

Consider a dynamical experiment where the individual populations of a list of different types or ‘species’ of objects is tracked through time. Supposing all objects are discrete and objects of the same species are indistinguishable, we can label the different species as $\{S_1, S_2, \dots\} \equiv \mathcal{S}$ and obtain a population vector $(n_1, n_2, \dots) \equiv \vec{n}$ which represents the state of the experiment. Furthermore, we will denote the entry in the time series data at time t as (\vec{n}_t, t) or simply \vec{n}_t .

In a stochastic experiment, starting from the same initial condition does not guarantee the same observation after some time t has elapsed. In this case, if we are to completely understand the dynamics of the system under investigation, we have to consider a collection or ‘ensemble’ of experiments starting from some initial condition $(\vec{n}_0, 0)$ and see how the relative frequencies of the different observations in the experiment ensemble changes in time. More precisely, from the data of the experiment ensemble, we get the probability that the experiment is found in \vec{n}_t at time t given that it was in \vec{n}_0 at time 0 , denoted as $\mathbb{P}[(\vec{n}_t, t) \cap (\vec{n}_0, 0)]$.

Recall from the definition of conditional probability that

$$\mathbb{P}[\vec{n}_t \cap \vec{n}_0] = \mathbb{P}[\vec{n}_t | \vec{n}_0] \mathbb{P}[\vec{n}_0], \quad (\text{B.1})$$

where $\mathbb{P}[\vec{n}_t | \vec{n}_0]$ denotes the probability of observing \vec{n}_t conditioned on having observed \vec{n}_0 . This formula can be useful to us in two ways. First, if the conditional probability $\mathbb{P}[\vec{n}_t | \vec{n}_0]$ is known, given an observation at the time of initialization (detection), we can make a probabilistic prediction (retrodiction) of what will be (was) the state of the system at a future (past) time t . Second, if we have access to both $\mathbb{P}[\vec{n}_t \cap \vec{n}_0]$ and $\mathbb{P}[\vec{n}_0]$, then we can improve our model of the system. In the remainder of the paper we focus on the first of these and comment on the second use in Section 3.5.

In Section 3.2.1 (Eq. 3.8), we give a detailed exposition of how the Hamilton-Jacobi formalism yields asymptotic estimates to the conditional probability term $\mathbb{P}[\vec{n}_t | \vec{n}_0]$. The basic idea can be summarized as follows. Whenever the dynamics of the ensemble of experiments can be modelled

as a Hamiltonian dynamical system satisfying certain conditions, in the limit of a large number of objects in each experiment (making sample fluctuations negligible) the conditional probability $\mathbb{P}[\vec{n}_t|\vec{n}_0]$ can be estimated by finding the optimal path that takes the system from \vec{n}_0 to \vec{n}_t . The optimality condition can take various interpretations such as most probable, least improbable, least costly, etc. (Snarski, 2021), but in all these interpretations there is a variational principle at play which derives from the Hamiltonian structure of the dynamics. Instead of yielding the conditional probability, Hamilton-Jacobi theory gives a descaled conditional log improbability that is referred to as the action function and denoted by

$$S(\vec{q}_t|\vec{q}_0) \equiv -\frac{1}{V} \log(\mathbb{P}[V\vec{q}_t|V\vec{q}_0]), \quad (\text{B.2})$$

where $\vec{q} = \vec{n}/V$. In technical terms, the action function is a solution to the Hamilton-Jacobi PDE, and the optimal paths given by ODEs of Hamilton's equations of motion are its characteristic curves. While historically this formalism has been used to describe motion of celestial or terrestrial objects and is the workhorse of classical mechanics, it can indeed also be used to quantify population space dynamics such as in chemical reactions, population genetics or economics (Baez and Biamonte, 2012; Smith and Morowitz, 2016; Smith, 2021; Shubik and Smith, 2016).


A particular type of measurement procedure that is of practical relevance is letting the stochastic experiment evolve under its own rules for an extremely long time i.e. several orders of magnitude more time than any characteristic time scale in the system. Due to a theorem by Perron and Frobenius (Pillai et al., 2005), the probability distribution of the ensemble of experiments will always approach a stationary distribution independent of the initial conditions, which we denote by $\pi(\vec{n})$. By definition then, $\mathbb{P}[(\vec{n}', \infty) \cap (\vec{n}, 0)] = \pi(\vec{n}')$, and the ratio of the stationary distribution at two population states \vec{n} and \vec{n}' , will be given by

$$\frac{\pi(\vec{n}')}{\pi(\vec{n})} = \mathbb{P}[(\vec{n}', \infty)|(\vec{n}, 0)] \equiv \pi(\vec{n}'|\vec{n}). \quad (\text{B.3})$$

Thus, for the stationary distribution, the conditional probability between two events $\pi(\vec{n}'|\vec{n})$ is simply equal to the ratio of their stationary probabilities $\pi(\vec{n}')/\pi(\vec{n})$. Using Hamilton-Jacobi theory, we will show in Section 3.2.1 that $\pi(\vec{n}'|\vec{n})$ is asymptotically estimated by finding the optimal path

connecting \vec{n} and \vec{n}' in the $H(p, q) = 0$ submanifold of the phase space. Thus, the complete solution to the Hamilton-Jacobi PDE in the $H(p, q) = 0$ manifold, also referred to in literature as Non-Equilibrium Potential (NEP), plays a central role for a certain set of questions. We will return to this point and explain its relevance to our algorithm after we introduce the basic idea behind stochastic chemical reaction network theory.

To explain the type of systems that can be modelled using stochastic chemical reaction networks, we first need an understanding of a ‘reaction’. In a system described by the population of its individual species, a reaction is a rule for a transition that replaces a collection or multiset of objects with a different one. For example, suppose in a system consisting of only two species, a reaction



will transition the population vector (n_1, n_2) to $(n_1 - 1, n_2 + 2)$. A ‘reaction network’ is a set of such reactions, each of which are assigned a rate constant determining a proportionality factor in the propensity or how often these reactions are to occur at random. In the particular case of a stochastic ‘chemical’ reaction network (CRN), the reaction rates are also proportional to the concentrations of the species and model proportional sampling without replacement. Although the dynamics of chemical reaction networks can be seen rigorously as arising from an underlying physical process (Gillespie, 1992), we will concern ourselves only with the modelling aspects of a stochastic CRN.

Chemical reaction network theory has a long history, and for an excellent review we point the readers to (Yu and Craciun, 2018). The most notable feature of chemical reaction systems is that they can exhibit a wide array of dynamics, with single or multiple attractors, limit cycles, etc. Although there are a host of applications that CRN have, to continue along our practical problem let us pick a particular one. Supposing that our system under investigation can be modelled as a multi-attractor CRN, we want to design an experiment to physically observe the different attractors. This amounts to finding how many experiments our ensemble should consist of so as to detect a significant amount in each attractor, which, as we will explain, is a type of problem that our algorithm can help get numerical estimates to.

Let us refer to the set of stable attractors in the population space by $\underline{\vec{n}} \equiv$

$\{\underline{n}_1, \underline{n}_2, \dots\}$. To estimate the number of experiments our ensemble must consist of, it is sufficient to estimate the ratio of the stationary distribution $\pi(\underline{n}_k | \underline{n}_j)$ at any two stable attractors \underline{n}_j and \underline{n}_k . From CRN theory and other considerations mentioned in the main text, any optimal path in the $H(p, q) = 0$ submanifold that emanates from the stable attractor \underline{n}_j must go through an adjoining saddle point, before going on to the next stable attractor, and so on in the process of reaching some \underline{n}_k . The optimal paths that take the system out of a stable attractor to an adjoining saddle attractor are called ‘escape paths’, and their collection consists of what we call a ‘heteroclinic network’. The heteroclinic network then determines the asymptotic stationary distribution of the stochastic system at the stable attractors, and the escape paths that constitute the heteroclinic network is precisely what our algorithm is designed to find (for e.g. see Figure 3.3).

We conclude this section by explaining how one can also use the algorithm to estimate the ratio $\pi(\vec{n} | \underline{n}_1)$ for a general multi-attractor CRN at any point \vec{n} . First, consider the case where the system exhibits a unique stable attractor \underline{n}_1 . In that case, the ratio $\pi(\vec{n} | \underline{n}_1)$ is found simply by determining the optimal path that connects the stable attractor \underline{n}_1 to the point \vec{n} , which is how we recover the Horn-Jackson potential using Hamilton-Jacobi theory in Section B.2. Next, in the case where the system has multiple attractors, we can use the algorithm to determine the stable attractor \underline{n}_k from which there is a direct escape path (that does not go through any other stable attractor) to \vec{n} . Once the optimal path joining \underline{n}_k and \vec{n} , or equivalently $\pi(\vec{n} | \underline{n}_k)$ is found, we can use the heteroclinic network (as discussed in the preceding paragraph) and

$$\pi(\vec{n} | \underline{n}_1) = \pi(\vec{n} | \underline{n}_k) \pi(\underline{n}_k | \underline{n}_1)$$

to determine the desired ratio. In technical terms, since the stationary distribution is asymptotically estimated by the solution of the Hamilton-Jacobi PDE in the $H(p, q) = 0$ submanifold, the conditional probability between any two points \vec{n}_1 and \vec{n}_2 is obtained by finding the difference in the values of the solution while traversing only along any of the optimal paths or characteristic curves. The geometry of the optimal paths in the $H(p, q) = 0$ submanifold is such that each stable attractor has a region within which each point is connected to the attractor by an escape path (escape basin (Smith and Krishnamurthy, 2021)), and neighboring regions

are connected by a saddle point which lies on the boundaries of two adjacent basins (for a detailed exposition, see Smith (2020)). Thus one can use our algorithm to find any number of desired characteristic curves of the Hamilton-Jacobi PDE in the $H(p, q) = 0$ submanifold and construct the complete stationary distribution for a chemical reaction network with multiple attractors.

B.2 Rederiving previous results in the Hamilton-Jacobi formalism

In this subsection we illustrate the usage of the formalism developed up to this point by rederiving some well known results in CRN theory. We start by giving a diagrammatic representation of CRNs, through the help of which we define the complex-balanced condition. Then we find the stationary distribution for such CRN by finding zero-eigenvectors of the corresponding Hamiltonian operators, which recovers the Anderson-Craciu-Kurtz (ACK) theorem (Anderson et al., 2010). Next, we find the NEP for such a distribution using Hamilton-Jacobi theory which recovers the Horn-Jackson potential and finally show that the negative descaled logarithm of the stationary distribution is indeed the NEP by comparison. We return to the techniques in this section by defining and finding the analytic NEP of another class of models, namely one dimensional (1-D) birth-death models in Section 3.4.2. There we also consider its generalization to N-dimensional (N-D) birth-death models, and use our algorithm to numerically estimate the corresponding NEPs.

Recall from Section 3.2.2 that a CRN is defined by the triple consisting the set of species \mathcal{S} , set of complexes \mathcal{C} and set of reactions \mathcal{R} . For a diagrammatic representation, we represent each species with a solid circle and each complex with an empty circle. Since a complex is a multi-set of species, we connect each complex to its constituting species with solid lines, where the number of lines denote the stoichiometry of the complex (denoted by column vector y). Finally, a reaction is a directed edge with a pair of complexes as source (y_α) and target (y_β), which we denote through a dashed line with an arrow indicating direction of reaction. For examples of such a representation, see Figures 3.1 and 3.8 or diagrams in (Krishnamurthy and Smith, 2017; Smith and Krishnamurthy, 2021).

In this diagrammatic representation, the deterministic flow of mass-action-kinetics can be visualized by assigning weights to the directed edges (dashed lines) equal to $k_{y_\alpha \rightarrow y_\beta} q^{y_\alpha}$. An equilibrium is said to be *complex balanced* (Horn and Jackson, 1972) at \underline{q} if the total flow directed out of each complex equals the total flow inwards or the net flow out of each complex is zero, i.e.

$$\sum_{y_\beta} \left(k_{y_\beta \rightarrow y_\alpha}(\underline{q})^{y_\beta} - k_{y_\alpha \rightarrow y_\beta}(\underline{q})^{y_\alpha} \right) = 0 \text{ for all } \alpha. \quad (\text{B.4})$$

Given a CRN that exhibits a complex-balanced steady state, we will now find the stationary distribution $\pi(n)$, such that $\hat{H}(-\partial/\partial n, n)\pi(n) = 0$. To identify the stationary distribution, we first need to make the same change of coordinates as we did for deriving the path integral formula in Section B.3.2 Eq. B.20, in which the Hamiltonian operator takes the form

$$\begin{aligned} & \hat{H}_{\text{CRN}} \left(z, \frac{\partial}{\partial z} \right) \\ &= \sum_{y_\alpha, y_\beta} (z^{y_\beta} - z^{y_\alpha}) \frac{k_{y_\alpha \rightarrow y_\beta}}{V^{y_\alpha - 1}} \left(\frac{\partial}{\partial z} \right)^{y_\alpha} \\ &= \sum_{y_\alpha} z^{y_\alpha} \sum_{y_\beta} \left[\frac{k_{y_\beta \rightarrow y_\alpha}}{V^{y_\beta - 1}} \left(\frac{\partial}{\partial z} \right)^{y_\beta} - \frac{k_{y_\alpha \rightarrow y_\beta}}{V^{y_\alpha - 1}} \left(\frac{\partial}{\partial z} \right)^{y_\alpha} \right] \end{aligned} \quad (\text{B.5})$$

where in the last line we have switched indices to rewrite the Hamiltonian in a form similar to Eq. B.4, also termed the *complex representation* in (Smith and Krishnamurthy, 2017). As mentioned in Section B.3.2, the change of coordinates corresponds to evolving the \mathcal{Z} -transform or the moment-generating function (MGF) rather than the distribution itself. Observe that the coordinate change preserves the commutation relations before and after the transformation,¹

¹This transformation is analogous to the Dirac transformation for simple harmonic oscillator in quantum mechanics, where we change from position and momentum to raising and lowering operators. z and $\partial/\partial z$ precisely play the role of the raising and lowering operators respectively, however, since stochastic dynamics preserves the ℓ_1 norm, the normalization is different. For the relevance of these operators in statistics, see Baez and Biamonte (2012).

$$\left[n, -\frac{\partial}{\partial n} \right] = \mathbb{I} = \left[\frac{\partial}{\partial z}, z \right]. \quad (\text{B.6})$$

We will refer to z as the raising operator and $\partial/\partial z$ as the lowering operator. For a representation of the above in abstract linear algebra making use of Fock space operators a, a^\dagger with $[a, a^\dagger] = 1$, see Ch. 4 of (Smith and Krishnamurthy, 2015).

Next, we find the eigenvectors of the lowering operator $\partial/\partial z$. Let us denote the eigenvector with eigenvalue Vc by $C(z)$. Then we have,

$$\begin{aligned} \frac{\partial C(z)}{\partial z} &= Vc C(z) \\ C(z) &= Ne^{cz}, \end{aligned}$$

where N is a normalization factor (determined in the next line). Recall that the \mathcal{Z} -transform or MGF of a probability distribution must be such that $C(z=1) = 1$, which means that $N = e^{-c}$ yielding

$$C(z) = e^{-Vc+Vcz} = e^{Vc(z-1)}. \quad (\text{B.7})$$

Let us denote the eigenvector of the lowering operator with eigenvalue Vq by $\underline{Q}(z)$. Then using Eq. B.5,

$$\begin{aligned} &\hat{\mathbb{H}}_{\text{CRN}} \underline{Q}(z) \\ &= \sum_{y_\alpha} z^{y_\alpha} \sum_{y_\beta} \left[\frac{k_{y_\beta \rightarrow y_\alpha}}{V^{y_\beta-1}} \left(\frac{\partial}{\partial z} \right)^{y_\beta} - \frac{k_{y_\alpha \rightarrow y_\beta}}{V^{y_\alpha-1}} \left(\frac{\partial}{\partial z} \right)^{y_\alpha} \right] \underline{Q}(z) \\ &= V \sum_{y_\alpha} z^{y_\alpha} \underline{Q}(z) \sum_{y_\beta} \left(k_{y_\beta \rightarrow y_\alpha}(\underline{q})^{y_\beta} - k_{y_\alpha \rightarrow y_\beta}(\underline{q})^{y_\alpha} \right) \\ &= 0, \end{aligned} \quad (\text{B.8})$$

where in second line we make use of the fact that \underline{Q} is an eigenvector of $\partial/\partial z$ and in the last line we use the complex-balanced condition from Eq. B.4.

Finally, to obtain the distribution $\pi(n)$ from $\underline{Q}(z)$, we can either take

the inverse- \mathcal{Z} transform or simply expand the exponential as a summation,

$$\begin{aligned}\underline{Q}(z) &= \sum_{n=0}^{\infty} z^n e^{-V\underline{q}} \frac{(V\underline{q})^n}{n!} \\ &= \sum_{n=0}^{\infty} z^n \pi(n).\end{aligned}$$

By reading the coefficients of the series, we obtain for the stationary distribution

$$\pi(n) = e^{-V\underline{q}} \frac{(V\underline{q})^n}{n!} \quad (\text{B.9})$$

which is the ACK or multi-Poisson distribution as derived in (Anderson et al., 2010; Smith and Krishnamurthy, 2017).

Having determined the stationary distribution using the Hamiltonian operator, we proceed to making use of Hamilton-Jacobi theory to find the NEP for complex-balanced system. Recall from 3.2.2 that to find the NEP, we need to find a momentum assignment $\mathbf{p}_{\text{esc}}(\mathbf{q}) \neq 0$ at every configuration \mathbf{q} , such that $H(\mathbf{p}_{\text{esc}}(\mathbf{q}), \mathbf{q}) = 0$ and the Hamilton's equations are satisfied for any initial condition. For finding the momentum assignment, namely escape momentum, along the $H(\mathbf{p}, \mathbf{q}) = 0$ submanifold it is easiest to recast the CRN Hamiltonian from Eq. 3.16 in the following form

$$\begin{aligned}H(\mathbf{p}, \mathbf{q}) &= \sum_{y_\alpha} (e^{\mathbf{p}})^{y_\alpha} \left[\sum_{y_\beta} \left(k_{y_\beta \rightarrow y_\alpha} (e^{-\mathbf{p}} \mathbf{q})^{y_\beta} - k_{y_\alpha \rightarrow y_\beta} (e^{-\mathbf{p}} \mathbf{q})^{y_\alpha} \right) \right].\end{aligned}$$

We will now show that $\mathbf{p}_{\text{esc}} = \ln(\mathbf{q}/\underline{\mathbf{q}})$ is such an assignment. Observe that,

$$\begin{aligned}H(\mathbf{q}, \mathbf{p}_{\text{esc}}) &= \sum_{y_\alpha} \left(\frac{\mathbf{q}}{\underline{\mathbf{q}}} \right)^{y_\alpha} \left[\sum_{y_\beta} \left(k_{y_\beta \rightarrow y_\alpha}(\underline{\mathbf{q}})^{y_\beta} - k_{y_\alpha \rightarrow y_\beta}(\underline{\mathbf{q}})^{y_\alpha} \right) \right] \\ &= 0,\end{aligned}$$

where the last line follows from the complex-balanced condition. Next, we calculate the total time derivative of p_{esc} and confirm the consistency of such a momentum assignment against the equations of motion.

$$\frac{dp_{\text{esc}}}{dt} = \frac{1}{q} \frac{dq}{dt} = \frac{1}{q} \frac{\partial H}{\partial p} = -\frac{\partial H}{\partial q},$$

where the last equality can be verified by a simple calculation.

Thus the NEP \mathcal{V} for CRN that exhibit a complex-balanced steady state is given by

$$\mathcal{V}(q) = \int_{\underline{q}}^q p_{\text{esc}} dq \quad (\text{B.10})$$

$$= q \ln \left(\frac{q}{\underline{q}} \right) - (q - \underline{q}), \quad (\text{B.11})$$

where an integration constant has been chosen such that $\mathcal{V}(\underline{q}) = 0$, and we recover the Horn-Jackson potential that appears in (Horn and Jackson, 1972).

Finally, we show the equivalence of the stationary distribution in Eq. B.9 and NEP in Eq. B.11. Using Eq. B.9 and substituting $n = Vq$, we have

$$\begin{aligned} -\frac{1}{V} \log(\pi(n)) &= -\frac{1}{V} (-V\underline{q} + Vq \log(V\underline{q}) - \log n!) \\ &\asymp q \log \left(\frac{q}{\underline{q}} \right) - (q - \underline{q}) \end{aligned}$$

where the last line is obtained by using Stirling's approximation. We recognize the last equation to be the same as the NEP in Eq. B.11, which completes our expository example of the equivalence of the Hamiltonian operator techniques and Hamilton-Jacobi theory for complex-balanced systems.

	Stochastic dynamics	Quantum dynamics
Space of states of system $q \in Q$	Discrete $n \in \mathbb{Z}_{\geq 0}^D$	Continuous $x \in \mathbb{R}^D$
Space of states of ensemble $\mathcal{F} \in L^p$	Probability distribution $\rho^R(n) \in \ell^1$	Wave-function $\psi(x) \in L^2$
Dual state space $\mathcal{F} \in L^{(1-\frac{1}{p})^{-1}}$	Sampling protocol $\rho^L(n) \in \ell^\infty$	State of detector $\psi(x) \in L^2$
Dual pairing or Inner product	$\langle \rho_1^L, \rho_2^R \rangle$ $\sum_n \rho_1^L(n) \rho_2^R(n)$	$\langle \psi_1, \psi_2 \rangle$ $\int \psi_1^*(x) \psi_2(x) dx$
Hamiltonian operator	ℓ^1 preserving dynamics Infinitesimal stochastic $\sum_n \hat{H}_{nn'} = 0$ for all n'	L^2 preserving dynamics Infinitesimal unitary $\hat{H} = -\hat{H}^\dagger$ $\hat{H} = -\frac{i}{\hbar} \hat{M}, \hat{M}^\dagger = \hat{M}$ (\hat{M} is Hermitian)
Forward equation	Master equation $\frac{\partial \rho^R(n)}{\partial t} = \hat{H} \rho^R(n)$ $\frac{\partial \rho^R(q)}{\partial t} = V \hat{H} \rho^R(q)$	Schrödinger equation $\frac{\partial \psi}{\partial t} = -\frac{i}{\hbar} \hat{M} \psi$
Adjoint representation $\langle p \Psi \rangle = \tilde{\Psi}(p)$	\mathcal{Z} -transform $\tilde{\rho}(z) = \sum_n z^n \rho(n)$ Laplace-transform $\tilde{\rho}_V(p) = \int_0^\infty e^{Vpq} \rho(q) dq$	Fourier transform $\tilde{\psi}(p) = \int dx e^{-\frac{i}{\hbar} px} \psi(x)$
Momentum operator $\langle p \hat{P} \Psi \rangle = p \langle p \Psi \rangle$	$\hat{P} = -\frac{\partial}{\partial n}$ $\hat{P} = -\frac{1}{V} \frac{\partial}{\partial q}$	$\hat{P} = \frac{\hbar}{i} \frac{\partial}{\partial x}$
Scaling limit	$V \rightarrow \infty$	$\hbar \rightarrow 0$
Large-deviation theory	$\rho^R(q, t) \asymp e^{-VS(q,t)}$	$\psi(x, t) \asymp e^{\frac{i}{\hbar} S(x,t)}$

Table B.1: Correspondence between stochastic and quantum dynamics

B.3 Non-equilibrium action, action functional and its first and second variational derivatives

B.3.1 Deriving stochastic and quantum dynamics from the non-equilibrium action (NEA)

Hamiltonian dynamical systems is a mathematical framework for formulating and analyzing dynamics of physical systems defined through a variational principle. The space of states (state-space) of a physical system is the space of all possible values one can observe upon measurements at the finest resolution. In physics, the measurable can be position, spin, energy of a particle or a group of particles, while in chemistry or biology, the measurable can be the count, concentration, etc. of a given species of molecule, organelle or organism. An experiment consists of an ensemble of systems, each obeying the same set of rules. These systems can either evolve simultaneously (like particles in a fluid, molecules in a solution, organisms in an ecosystem) or can be spatio-temporally separated (like a series of independent quantum or biological experiments) or both. In all cases, one can assign a distribution over the state-space and refer to it as the state of the experiment, which we denote by $|\Psi^R(t)\rangle$, as well as define a sampling protocol that takes a distribution as input and give a scalar quantity, which we denote by $\langle\Psi^L(t)|$.²

Following Eyink's construction in (Eyink, 1996), given a Hamiltonian operator \hat{H} we define the non-equilibrium action (NEA) functional Γ to be

$$\Gamma[\Psi^L(t), \Psi^R(t)] = \int \langle\Psi^L(t), (\partial_t - \hat{H})\Psi^R(t)\rangle dt$$

that takes a time-evolving state $|\Psi^R\rangle$ and sampling protocol $\langle\Psi^L|$ as input and yields a scalar quantity. If we expand the NEA around Hilbert-space

²Probability distribution and quantum amplitudes are both distributions, formally defined as vectors in L_1 and L_2 normed Hilbert space, respectively (Baez and Biamonte, 2012).

vectors ϕ^R and ϕ^L ,

$$\begin{aligned}\Psi^R &= \phi^R + \delta\Psi^R \\ \Psi^L &= \phi^L + \delta\Psi^L\end{aligned}$$

we get the variation in the NEA to be

$$\begin{aligned}\Gamma[\Psi^L(t), \Psi^R(t)] &= \Gamma[\phi^L(t), \phi^R(t)] \\ &+ \int \langle \delta\Psi^L(t), (\partial_t - \hat{H})\phi^R(t) \rangle dt \\ &+ \int \langle (-\partial_t - \hat{H}^\dagger)\phi^L(t), \delta\Psi^R(t) \rangle dt + \mathcal{O}(\delta^2).\end{aligned}$$

The stationary condition then yields the following Hamilton's equations of motion on the Hilbert space vectors

$$\begin{aligned}\partial_t \phi^R(t) &= \hat{H}\phi^R(t), \\ \partial_t \phi^L(t) &= -\hat{H}^\dagger \phi^L(t).\end{aligned}\tag{B.12}$$

The norm of the state and the sampling protocol must be preserved during the dynamics, which yields extra conditions on the Hamiltonian operator \hat{H} . In particular, when \hat{H} preserves the L_1 norm, then we get the Master equation from Eq. 3.2.

It is also straightforward to verify that the Hamiltonian operator that preserves the L_2 norm must correspond to an anti-Hermitian operator, implying $\hat{H} = -\hat{H}^\dagger$. Since the dual of an L_2 function is an L_2 function, there is no difference between the evolution of the state or the sampling protocol, thus reducing the two equations in Eq. B.12 to one, yielding the Schrödinger equation. We summarize these observations in table B.1.

The two equations B.12 can be collected in one by defining a density operator $\hat{\phi} = |\phi^R\rangle\langle\phi^L|$ and finding its total time derivative

$$\begin{aligned}\frac{d\hat{\phi}}{dt} &\equiv \frac{d|\phi^R\rangle\langle\phi^L|}{dt} = \frac{d|\phi^R\rangle}{dt}\langle\phi^L| + |\phi^R\rangle\frac{d\langle\phi^L|}{dt} \\ &= \hat{H}|\phi^R\rangle\langle\phi^L| - |\phi^R\rangle\langle\phi^L|\hat{H}^\dagger \\ &= [\hat{H}, \hat{\phi}].\end{aligned}\tag{B.13}$$

This is also known as the Von-Neumann equation and the operator $[\hat{H}, \cdot]$ is also referred to as the Liouville operator (Manzano, 2020). The statistical observation of the state of an experiment under a sampling protocol is given by $\langle O \rangle_\phi \equiv \text{Tr} \{ \hat{O} \hat{\phi} \}$. The time evolution of the statistical observable O along the variational solutions of the non-equilibrium action is then given by

$$\frac{d\langle O \rangle_\phi}{dt} = \text{Tr} \{ [\hat{O}, \hat{H}] \hat{\phi} \} + \text{Tr} \left\{ \frac{\partial \hat{O}}{\partial t} \hat{\phi} \right\}. \quad (\text{B.14})$$

For an excellent introduction to operator techniques in quantum statistical mechanics and quantum computation, see Kadanoff (2000); Nielsen and Chuang (2002).

B.3.2 Derivation of the path integral formula

In this section we will answer the question, given the state of an experiment (PDF) at time $t = 0$ and a stochastic Hamiltonian with which the system evolves following Eqs. B.12 and 3.2, what is its state at an arbitrary time $t = T$?

Let us revisit the time-evolution, from Eq. 3.3, for a PDF $\rho(\mathbf{n}, t)$, where $\mathbf{n} \in \mathbb{Z}_{\geq 0}^D$ denotes the position in state space and t is the time,

$$\frac{\partial}{\partial t} \rho(\mathbf{n}, t) = \hat{\underline{H}} \left(-\frac{\partial}{\partial \mathbf{n}}, \mathbf{n} \right) \rho(\mathbf{n}, t). \quad (\text{B.15})$$

Since \mathbf{n} takes only nonnegative integer values, we can consider the moment-generating function (MGF) or \mathcal{Z} -transform of the PDF $\rho(\mathbf{n}, t)$, given by $\tilde{\rho}(z, t)$, and its inverse

$$\tilde{\rho}(z, t) = \sum_{\mathbf{n}} z^{\mathbf{n}} \rho(\mathbf{n}, t) \quad (\text{B.16})$$

$$\rho(\mathbf{n}, t) = \frac{1}{2\pi i} \oint dz \frac{1}{z^{\mathbf{n}+1}} \tilde{\rho}(z, t). \quad (\text{B.17})$$

We refer to the first line as \mathcal{Z} -transform and the second as inverse \mathcal{Z} -transform (for an introduction, see Oppenheim et al. (1997)). Here $z \in \mathbb{C}^D$ and the contour integral is done over any contour that encloses $z = 0$.

Since Eq. B.15 in general defines the time evolution of any distribution, in particular we can consider the time evolution of the MGF or \mathcal{Z} -transform of a distribution

$$\begin{aligned}\frac{\partial \tilde{\rho}(z, t)}{\partial t} &= \sum_{\mathbf{n}} z^{\mathbf{n}} \hat{\mathbf{H}} \left(-\frac{\partial}{\partial \mathbf{n}}, \mathbf{n} \right) \rho(\mathbf{n}, t) \\ &= \hat{\mathbf{H}} \left(z, \frac{\partial}{\partial z} \right) \tilde{\rho}(z, t),\end{aligned}\quad (\text{B.18})$$

where $\hat{\mathbf{H}}$ is the time evolution operator for the \mathcal{Z} -transformed distribution, which we will henceforth refer to as the \mathcal{Z} -Hamiltonian operator.

Using the observations

$$\begin{aligned}\sum_{\mathbf{n}} z^{\mathbf{n}} e^{-y \frac{\partial}{\partial \mathbf{n}}} \rho(\mathbf{n}, t) &= z^y \sum_{\mathbf{n}} z^{\mathbf{n}} \rho(\mathbf{n}, t) \\ \sum_{\mathbf{n}} z^{\mathbf{n}} \frac{\mathbf{n}!}{(\mathbf{n} - y)!} \rho(\mathbf{n}, t) &= z^y \left(\frac{\partial}{\partial z} \right)^y \sum_{\mathbf{n}} z^{\mathbf{n}} \rho(\mathbf{n}, t),\end{aligned}$$

it is easy to see that the \mathcal{Z} -Hamiltonian operator for CRN from Eq. 3.15

$$\begin{aligned}\hat{\mathbf{H}}_{\text{CRN}} \left(-\frac{\partial}{\partial \mathbf{n}}, \mathbf{n} \right) \\ = \sum_{y_{\alpha}, y_{\beta}} \left(e^{-(y_{\beta} - y_{\alpha}) \cdot \frac{\partial}{\partial \mathbf{n}}} - 1 \right) \frac{k_{y_{\alpha} \rightarrow y_{\beta}}}{V^{y_{\alpha} - 1}} \frac{\mathbf{n}!}{(\mathbf{n} - y_{\alpha})!}.\end{aligned}\quad (\text{B.19})$$

takes the following form (for a more insightful derivation, see Baez and Biamonte (2012))

$$\hat{\mathbf{H}}_{\text{CRN}} \left(z, \frac{\partial}{\partial z} \right) = \sum_{y_{\alpha}, y_{\beta}} (z^{y_{\beta}} - z^{y_{\alpha}}) \frac{k_{y_{\alpha} \rightarrow y_{\beta}}}{V^{y_{\alpha} - 1}} \left(\frac{\partial}{\partial z} \right)^{y_{\alpha}}. \quad (\text{B.20})$$

Returning to the question posed in the beginning of this subsection, given a Hamiltonian operator $\hat{\mathbf{H}}$ and $\rho(\mathbf{n}, 0)$, we wish to find $\rho(\mathbf{n}, T)$. The approach that we will take is the following (for a pedagogical introduction to these methods, see Altland and Simons (2010); Peskin (2018)). First, we discretize time from 0 to T into N intervals of length Δt , such that $N\Delta t = T$. Next, we label the random variable denoting the count of the

system at time $i\Delta t$ by n_i . Since, we know that the system is evolving with the given Hamiltonian, there is a relation between $\rho(n_{i+1}, (i+1)\Delta t)$ and $\rho(n_i, i\Delta t)$. To find $\rho(n_{i+1}, (i+1)\Delta t)$, we first compute the \mathcal{Z} -transform of $\rho(n_i, i\Delta t)$, evolve it for Δt by B.18 and compute the inverse \mathcal{Z} -transform at n_{i+1} .




Figure B.1: Evolving distribution by Δt

Using the above procedure, which we represent in diagrammatic form in Figure B.1, we can read the distribution at time $(i+1)\Delta t$ to be

$$\begin{aligned}
 \rho(n_{i+1}, (i+1)\Delta t) & \quad (B.21) \\
 &= \frac{1}{2\pi i} \oint dz_i \frac{1}{z_i^{n_{i+1}+1}} e^{\Delta t \hat{\mathcal{H}}(z, \frac{\partial}{\partial z})} \sum_{n_i} z_i^{n_i} \rho(n_i, i\Delta t) \\
 &= \frac{1}{2\pi i} \oint \frac{dz_i}{z_i} \sum_{n_i} z_i^{-n_{i+1}} e^{\Delta t \hat{\mathcal{H}}(z, \frac{\partial}{\partial z})} z_i^{n_i} \rho(n_i, i\Delta t) \\
 &= \frac{1}{2\pi i} \oint \frac{dz_i}{z_i} \sum_{n_i} z_i^{-n_{i+1}+n_i} e^{\Delta t \mathcal{H}(z_i, n_i)} \rho(n_i, i\Delta t)
 \end{aligned}$$

where in the last line, the function $\mathcal{H}(z_i, n_i)$ is obtained by an application of the operator $\hat{\mathcal{H}}(z, \frac{\partial}{\partial z})$ on $z_i^{n_i}$. In particular for CRN, using B.20, we obtain

$$\mathcal{H}_{\text{CRN}}(z, n) = \sum_{y_\alpha, y_\beta} (z^{y_\beta - y_\alpha} - 1) \frac{k_{y_\alpha \rightarrow y_\beta}}{V^{y_\alpha - 1}} \frac{n!}{(n - y_\alpha)!}. \quad (B.22)$$

Finally, by repeated application of the above procedure at the initial PDF N times and taking the limit $N \rightarrow \infty$ and $\Delta t \rightarrow 0$, while their product

is held fixed $N\Delta t = T$, we obtain

$$\begin{aligned}
& \rho(\mathbf{n}_N, N\Delta t) \\
&= \prod_{i=0}^{N-1} \left(\frac{1}{2\pi i} \oint \frac{dz_i}{z_i} \sum_{\mathbf{n}_i} z_i^{-n_{i+1}+n_i} e^{\Delta t \mathcal{H}(z_i, \mathbf{n}_i)} \right) \rho(\mathbf{n}_0, 0) \\
&\text{using } z_i = e^{p_i} \\
&= \prod_{i=0}^{N-1} \left(\int \frac{dp_i}{2\pi i} \sum_{\mathbf{n}_i} e^{-p_i(n_{i+1}-n_i)} e^{\Delta t \underline{\mathbf{H}}(p_i, \mathbf{n}_i)} \right) \cdot \rho(\mathbf{n}_0, 0) \\
&\xrightarrow[\Delta t \rightarrow 0]{N \rightarrow \infty} \int \left[\frac{dp}{2\pi i} \right] \sum [dn] e^{-\int_0^T dt (p\dot{n} - \underline{\mathbf{H}}(p, \mathbf{n}) dt)} \rho(\mathbf{n}_0, 0),
\end{aligned}$$

where $[dp]$ and $[dn]$ are the path-integral measures and $\underline{\mathbf{H}}(p, \mathbf{n})$. The last line yields us the path-integral formula and we define the **action functional** as

$$\mathcal{A}[\mathbf{n}(t), \mathbf{p}(t)] = \int_0^T dt (p\dot{n} - \underline{\mathbf{H}}(p, \mathbf{n}) dt). \quad (\text{B.23})$$

For an alternative derivation of the functional integral in the Doi-Peliti formalism using coherent states, see Smith (2011); Smith and Krishnamurthy (2015); Smith (2019).

Note that For CRN, using B.22, the Hamiltonian function takes the particular form

$$\underline{\mathbf{H}}_{\text{CRN}}(\mathbf{p}, \mathbf{n}) = \sum_{\mathbf{y}_\alpha, \mathbf{y}_\beta} (e^{p \cdot (\mathbf{y}_\beta - \mathbf{y}_\alpha)} - 1) \frac{k_{\mathbf{y}_\alpha \rightarrow \mathbf{y}_\beta}}{V^{\mathbf{y}_\alpha - 1}} \frac{\mathbf{n}!}{(\mathbf{n} - \mathbf{y}_\alpha)!}. \quad (\text{B.24})$$

which is the same as the functional form of the Hamiltonian operator in Eq. B.19, with the operator $-\frac{\partial}{\partial \mathbf{n}}$ replaced by the momentum variable p , thus justifying our notation.

B.3.3 Optimality condition: Hamilton's equations of motion and their relation to the Hamilton-Jacobi equation

Analogous to Eq. B.23, in Eq. 3.6 we defined an action functional in the concentration and momentum coordinates (q, p) to be

$$\mathcal{A}[q(t), p(t)] = \int_0^T \left[p \cdot \frac{dq}{dt} - H(p, q) \right] dt,$$

where the relation of the Hamiltonian function $H(p, q)$ to $\underline{H}(p, n)$ was discussed in Eq. 3.17. As discussed around Eq. 3.8, the path integral is dominated by the value of the action functional around the optimal or stationary path, i.e. the path for which the first variation of the action is zero.

To calculate the first variation of the action, we consider the difference of the action functional between a path $(q(t), p(t))$ and $(q(t) + \delta q(t), p(t) + \delta p(t))$. In the following equation, we will suppress the time dependence to simplify notation and keep only the first order terms.

$$\begin{aligned} & \delta \mathcal{A}[q, p] \\ &= \mathcal{A}[q + \delta q, p + \delta p] - \mathcal{A}[q, p] \\ &= \int dt \left\{ \delta p \left(\frac{dq}{dt} - \frac{\partial H}{\partial p} \right) + \left(p \frac{d\delta q}{dt} - \frac{\partial H}{\partial q} \delta q \right) \right\} \\ &= \int dt \left\{ \delta p \left(\frac{dq}{dt} - \frac{\partial H}{\partial p} \right) - \delta q \left(\frac{dp}{dt} - \frac{\partial H}{\partial q} \right) + p \delta q \Big|_0^T \right\} \quad (\text{B.25}) \end{aligned}$$

where the last line is obtained by integrating by-parts. If we fix the end-points of the path in configuration space, then we can set $\delta q(0) = \delta q(T) = 0$, getting rid of the last term in the above equation.

Using Eq. B.25 and

$$\delta \mathcal{A}[q(t), p(t)] = \frac{\delta \mathcal{A}}{\delta q(t)} \delta q(t) + \frac{\delta \mathcal{A}}{\delta p(t)} \delta p(t)$$

we can read off the variation of the action functional in configuration and

momentum around $(q(t), p(t))$ to be

$$\begin{aligned}\frac{\delta \mathcal{A}}{\delta q(t)} &= - \left(\frac{dp}{dt} + \frac{\partial H}{\partial q} \right), \\ \frac{\delta \mathcal{A}}{\delta p(t)} &= \left(\frac{dq}{dt} - \frac{\partial H}{\partial p} \right).\end{aligned}\tag{B.26}$$

The optimality condition $\mathcal{A}[q^*, p^*] = 0$ for a trajectory $(q^*(t), p^*(t))$ then yields

$$\begin{aligned}\frac{dp^*}{dt} &= - \frac{\partial H}{\partial q^*}, \\ \frac{dq^*}{dt} &= \frac{\partial H}{\partial p^*}.\end{aligned}\tag{B.27}$$

The last two lines are also referred to as Hamilton's equations of motion (EoM) and are the equations that any optimal trajectory must satisfy.

There is a rich and deep mathematical structure underlying these equations, which is the subject of symplectic geometry. For instance, the change in value of an observable function $f(p, q)$ along an optimal trajectory (q^*, p^*, t) can be given in terms of its Poisson-bracket commutator with the Hamiltonian,

$$\begin{aligned}\frac{df}{dt} &= \frac{\partial f}{\partial t} + \frac{\partial f}{\partial q} \frac{dq}{dt} + \frac{\partial f}{\partial p} \frac{dp}{dt} \\ &= \frac{\partial f}{\partial t} + \frac{\partial f}{\partial q} \frac{\partial H}{\partial p} - \frac{\partial f}{\partial p} \frac{\partial H}{\partial q} \\ &= \frac{\partial f}{\partial t} + \{f, H\}.\end{aligned}\tag{B.28}$$

This immediately shows that the value of a time-independent Hamiltonian as well as any time independent operator that commutes with the Hamiltonian stays constant along equations of motion (notice the resemblance of the above equation with Eq. B.14). For a classical introduction to the subject, see Goldstein et al. (2002); Arnol'd (2013).

To end this subsection, following (Courant and Hilbert, 2008), we provide a proof of the equivalence between Hamilton's equations of motion and Hamilton-Jacobi equation on paths where the velocity and the

momentum are related by a Legendre transform.

$$\begin{aligned}
 p &= \frac{\partial S}{\partial q} \\
 \dot{p} &= \frac{d}{dt} \left(\frac{\partial S}{\partial q} \right) \\
 &= \frac{\partial^2 S}{\partial q \partial q} \dot{q} + \frac{\partial^2 S}{\partial q \partial t} \\
 &= \frac{\partial^2 S}{\partial q \partial q} \frac{\partial H}{\partial p} - \frac{\partial H}{\partial q} - \frac{\partial H}{\partial p} \frac{\partial p}{\partial q} \\
 &= \frac{\partial p}{\partial q} \frac{\partial H}{\partial p} - \frac{\partial H}{\partial p} \frac{\partial p}{\partial q} - \frac{\partial H}{\partial q} \\
 &= -\frac{\partial H}{\partial q}
 \end{aligned}$$

where the first and fourth line use the Hamilton-Jacobi equation. The fourth line also makes use of the relation between optimal momentum and velocity, or the Legendre-transform condition.

B.3.4 Second variational derivative of action functional: Onsager-Machlup action and convexity of instanton

In this subsection, we will first calculate the second variational derivative of the action functional. Next, we provide a derivation of the Onsager-Machlup action, followed by comments on the convexity of the action functional around any optimal trajectory.

To simplify notation we use the following convention. For vectors v, w , we represent their dot product $v^T w$ as vw . For a quadratic form M , we represent $M(v, v) \equiv v^T M v$ simply as Mv^2 . Finally, the time derivative and perturbation of a function x is given by

$$\begin{aligned}
 \dot{x} &= \frac{dx}{dt}, \\
 x^+ &= x + \delta x,
 \end{aligned}$$

respectively.

Recall, in Eq. 3.6, the action functional was defined to be

$$\mathcal{A}[q, p] = \int_0^T dt \{p\dot{q} - H\}.$$

Thus, the variation in the action functional up to terms second order in the variation is

$$\begin{aligned} \delta\mathcal{A} &= \mathcal{A}[q^+, p^+] - \mathcal{A}[q, p] \\ &= \int dt \left\{ \delta p \left(\dot{q} - \frac{\partial H}{\partial p} \right) - \delta q \left(\dot{p} + \frac{\partial H}{\partial q} \right) \right\} + \delta p \delta \dot{q} - \frac{1}{2} [\delta p \quad \delta q] \begin{bmatrix} \frac{\partial^2 H}{\partial p \partial p} & \frac{\partial^2 H}{\partial p \partial q} \\ \frac{\partial^2 H}{\partial q \partial p} & \frac{\partial^2 H}{\partial q \partial q} \end{bmatrix} \begin{bmatrix} \delta p \\ \delta q \end{bmatrix} + p \delta q \Big|_0^T \\ &= \int dt \left\{ \delta p \left(\dot{q} - \frac{\partial H}{\partial p} \right) - \delta q \left(\dot{p} + \frac{\partial H}{\partial q} \right) - \frac{1}{2} [\delta p \quad \delta q] \begin{bmatrix} \frac{\partial^2 H}{\partial p \partial p} & \left(-d_t + \frac{\partial^2 H}{\partial p \partial q} \right) \\ \left(d_t + \frac{\partial^2 H}{\partial q \partial p} \right) & \frac{\partial^2 H}{\partial q \partial q} \end{bmatrix} \begin{bmatrix} \delta p \\ \delta q \end{bmatrix} \right\} \\ &= \int dt \left\{ \delta p \left(\dot{q} - \frac{\partial H}{\partial p} \right) - \delta q \left(\dot{p} + \frac{\partial H}{\partial q} \right) - \frac{1}{2} \frac{\partial^2 H}{\partial p \partial p} \left[\delta p - \left(\frac{\partial^2 H}{\partial p \partial p} \right)^{-1} \left(\delta \dot{q} - \frac{\partial^2 H}{\partial p \partial q} \delta q \right) \right]^2 \right. \\ &\quad \left. + \frac{1}{2} \left[\left(\frac{\partial^2 H}{\partial p \partial p} \right)^{-1} \left[\left(d_t - \frac{\partial^2 H}{\partial p \partial q} \right) \delta q \right]^2 - \frac{\partial^2 H}{\partial q \partial q} \delta q^2 \right] \right\} + \left(p \delta q + \frac{1}{2} \delta p \delta q \right) \Big|_0^T, \end{aligned} \tag{B.29}$$

where we have implicitly assumed that the original and perturbed phase space trajectories lie in the same $H = h$ submanifold.

Now, if we assume that the two trajectories are constrained to the same end points in configuration space, i.e. $\delta q(0) = \delta q(T) = 0$, then the boundary terms vanish. Moreover, if we assume that the momentum assignment before and after the perturbation is optimal, i.e. p and p^+ are assigned to q and q^+ such that the phase space point is at its Legendre

transform before and after the perturbation, then we have

$$\begin{aligned} & \text{For } (q, p) \\ 0 &= \left(\dot{q} - \frac{\partial H}{\partial p} \right) \end{aligned} \quad (\text{B.30})$$

$$\begin{aligned} & \text{For } (q^+, p^+) \\ \dot{q}^+ &= \frac{\partial H(p^+, q^+)}{\partial p} \\ \delta \dot{q} &= \frac{\partial^2 H}{\partial p \partial q} \delta q + \frac{\partial^2 H}{\partial p \partial p} \delta p \\ 0 &= \delta p - \left(\frac{\partial^2 H}{\partial p \partial p} \right)^{-1} \left(\delta \dot{q} - \frac{\partial^2 H}{\partial p \partial q} \delta q \right) \end{aligned} \quad (\text{B.31})$$

Substituting eqs B.30 and B.31 in Eq. B.29 and making use of the assumptions, we get the variation in the action functional to be

$$\begin{aligned} \delta \mathcal{A} &= \int dt \left\{ -\delta q \left(\dot{p} + \frac{\partial H}{\partial q} \right) + \frac{1}{2} \left[\left(\frac{\partial^2 H}{\partial p \partial p} \right)^{-1} \left[\left(d_t - \frac{\partial^2 H}{\partial p \partial q} \right) \delta q \right]^2 - \frac{\partial^2 H}{\partial q \partial q} \delta q^2 \right] \right\} \\ &= \int dt \left\{ -\delta q \left(\dot{p} + \frac{\partial H}{\partial q} \right) + \frac{1}{2} \left[\left(\frac{\partial^2 H}{\partial p \partial p} \right)^{-1} \left[d_t \delta q + \frac{\partial H}{\partial p} - \frac{\partial H}{\partial p} - \frac{\partial^2 H}{\partial p \partial q} \delta q \right]^2 - \frac{\partial^2 H}{\partial q \partial q} \delta q^2 \right] \right\} \\ &= \int dt \left\{ -\delta q \left(\dot{p} + \frac{\partial H}{\partial q} \right) + \frac{1}{2} \left[\left(\frac{\partial^2 H}{\partial p \partial p} \right)^{-1} \left[\dot{q}^+ - \dot{q} - \frac{\partial^2 H}{\partial p \partial q} \delta q \right]^2 - \frac{\partial^2 H}{\partial q \partial q} \delta q^2 \right] \right\}, \end{aligned} \quad (\text{B.32})$$

where we add and subtract the same quantity in the second line and make use of the Legendre transform condition in the third line.

Onsager-Machlup Action $\delta \mathcal{A}$ measures the log-conditional improbability of trajectory q^+ given base trajectory q . If q is Hamilton's equation

of motion along $p = 0$, then we have

$$\begin{aligned} \left(\dot{p} + \frac{\partial H}{\partial q} \right) &= 0, \\ \frac{\partial^2 H}{\partial q \partial q} \Big|_{p=0} &= 0, \\ \frac{\partial^2 H}{\partial q \partial p} \Big|_{p=0} &= 0, \end{aligned} \tag{B.33}$$

and we thus recover the Onsager-Machlup action

$$\begin{aligned} \mathcal{A}_{\text{OM}}[q^+] &\equiv \mathcal{A}[(q^+, 0)] - \mathcal{A}[(q, 0)] \\ &= \int dt \left\{ \frac{1}{2} \left[\left(\frac{\partial^2 H}{\partial p \partial p} \right)^{-1} [\dot{q}^+ - \dot{q}]^2 \right] \right\}. \end{aligned} \tag{B.34}$$

Recall from the discussion around Eq. 3.20 that the equations of motion along which $p = 0$ are the deterministic trajectories of the system. Thus the Onsager-Machlup action measures the log improbability of a trajectory q^+ conditioned on the deterministic or relaxation trajectory q for the stochastic system. For a similar derivation and applications to population biology, see Smith (2021).

Convexity of action functional around an optimal path If q is an escape curve for a CRN Hamiltonian, only the first equation of Eqs. B.33 holds since $p(q_{\text{esc}}) \neq 0$ (see Eqs. 3.24). The variation in action around the optimal path then becomes

$$\delta \mathcal{A} = \int dt \left\{ \frac{1}{2} \left[\left(\frac{\partial^2 H}{\partial p \partial p} \right)^{-1} \left[\dot{q}^+ - \dot{q} - \frac{\partial^2 H}{\partial p \partial q} \delta q \right]^2 - \frac{\partial^2 H}{\partial q \partial q} \delta q^2 \right] \right\}. \tag{B.35}$$

Notice that due to convexity of the Hamiltonian in p , the first term is positive definite. However, nothing can be said a priori about the second term since the Hessian of the Hamiltonian in q for multi-stationary networks is in general non-convex. For the applications demonstrated in the paper, the variation is indeed positive definite and a global minimum exists in all

cases. Generically however, as can be seen from the equation above, not all optimal curves are global minimizers of the action. For an example of an infinite dimensional stochastic system exhibiting a saddle optimal solution, we refer the readers to Chapter 8, (Kamenev, 2011). The authors are not aware of a similar example for a finite dimensional chemical reaction network at this point, and leave it as a question for future investigation.

B.4 Details of AFGD

B.4.1 Notation

As noted in Table 3.1, starting Section 3.3.1, we use the following notation.

Configuration space : $Q \subset \mathbb{R}^D$

$\mathcal{P}_Q = \{q : [0, 1] \rightarrow Q : q(0) = q_I, q(1) = q_F\}$

Configuration space curve : $q \in \mathcal{P}_Q$

Phase space : $T^*Q \subset \mathbb{R}^{2D}$

$\mathcal{P}_{T^*Q} = \{(q, p) : [0, T] \rightarrow T^*Q : q^*(0) = q_I, q^*(T) = q_F\}$

Phase space trajectory : $\gamma \in \mathcal{P}_{T^*Q}$

We refer to paths in configuration space as ‘curves’, and paths in phase space as ‘trajectories’. This is an unusual choice intended to avoid confusion and simplify terminology, and will be used consistently through this section. Notice that curves are parametrized on the unit interval $[0, 1]$, which is the normalized arc length. On the other hand, trajectories are parametrized with time, on the interval $[0, T]$ (for more details, see Section 3.3.1).

Since computationally we have to work with discrete paths, we now introduce notation for discretization. Say we sample a continuous path $x \in X$ on N points, we denote its discrete counterpart with $\tilde{x} \in \tilde{X}$ where

$$\tilde{x}(n) = x\left(\frac{n}{N}\right) \quad \text{where } n \in [1, N] \subset \mathbb{Z}.$$

We denote the I^{th} iteration of the algorithm with a superscript I . For instance, the configuration space curve in the first iteration is denoted by

\tilde{q}^1 , where

$$\tilde{q}^1(\mathbf{n}) = q^{\text{IC}} \left(\frac{\mathbf{n}}{N} \right).$$




Figure B.2: Initial conditions for the AFGD algorithm.

B.4.2 Initial condition

To begin the descent towards an optimal trajectory, in principle, one can start with an arbitrary curve connecting the end-points q_I and q_F that does not pass through or go around another fixed point and remains in q_I 's basin of escape. In practice however, the descent might require a large number of sample points and high numerical accuracy to obtain the optimal trajectory. In this subsection we provide a method to classify the optimal trajectory and pick an appropriate initial condition with the desired number of sample points effectively.

We classify the optimal trajectory near a fixed point by linearizing Hamilton's equations of motion (EoM) and analyzing the eigenvalues of the Hessian. If the eigenvalues have a non-zero imaginary part then the optimal trajectory spirals into or out of the fixed point. An exposition of this technique can be found in Chapter 4 of (Kamenev, 2011) (their 'activation trajectory' is our 'escape trajectory'), but we repeat the relevant construction here.

Let the fixed point be denoted by q . Recall that the momentum value at the fixed point in the $H(p, q) = 0$ submanifold is identically zero i.e.

$p(\underline{q}) = 0$. Linearizing Hamilton's equations of motion around the fixed point, we get

$$\begin{aligned} \begin{bmatrix} \frac{dq}{dt} \\ \frac{dp}{dt} \end{bmatrix} &= \begin{bmatrix} \frac{\partial H}{\partial p} \\ -\frac{\partial H}{\partial q} \end{bmatrix}_{(\underline{q},0)} + \begin{bmatrix} \frac{\partial H}{\partial p \partial q} & \frac{\partial H}{\partial p^2} \\ -\frac{\partial H}{\partial q^2} & -\frac{\partial H}{\partial q \partial p} \end{bmatrix}_{(\underline{q},0)} \begin{bmatrix} (q - \underline{q}) \\ (p - 0) \end{bmatrix} \\ &= \begin{bmatrix} \frac{\partial H}{\partial p \partial q} & \frac{\partial H}{\partial p^2} \\ 0 & -\frac{\partial H}{\partial q \partial p} \end{bmatrix}_{(\underline{q},0)} \begin{bmatrix} (q - \underline{q}) \\ (p - 0) \end{bmatrix} \end{aligned} \quad (\text{B.36})$$

where the second line is obtained by using the fact that \underline{q} is a fixed point and $p(\underline{q}) = 0$. To simplify notation, we denote the column vector $(q - \underline{q}, p - 0)^T$ by $\Delta\gamma$. Then we can rewrite the above equation as

$$\begin{aligned} \frac{d\Delta\gamma}{dt} &= M\Delta\gamma \\ M &= \sum_{i=1}^{2n} \lambda_i v_i v_i^T \end{aligned} \quad (\text{B.37})$$

where M is the Hessian with eigenvectors v_i and their corresponding eigenvalues λ_i .

Let the initial condition be denoted by q^{IC} .

Systems without vorticity If all λ_i s are real then we know that the optimal trajectory will not spiral into the saddle or out of the stable fixed point. In this case, we pick the initial condition to be a straight line starting from the stable fixed point towards the saddle fixed point.

$$q^{\text{IC}}(s) = q_F + (1 - s)(q_I - q_F) \quad (\text{B.38})$$

The discrete initial curve for the algorithm will then be given by,

$$\tilde{q}^1(n) = q^{\text{IC}}\left(\frac{n}{N}\right). \quad (\text{B.39})$$

Systems with vorticity If some λ_i s are complex, then the trajectory spirals outwards or inwards around the stable or saddle fixed point respectively. For the purpose of this subsection, let us assume that the eigenvalues are complex around the stable fixed point q_I and real around the saddle fixed point q_F .

All the trajectories that depart from the fixed point, including the escape trajectory in the $H(p, q) = 0$ submanifold, lie in the D -dimensional submanifold³ (in the 2D-dimensional phase space) given by the superposition of the eigenvectors with non-zero eigenvalues. Thus, we can write the escape trajectory near the fixed point as

$$\begin{aligned}\Delta\gamma(0) &= \sum \mu_i v_i \\ \Delta\gamma(t) &= \sum_i \mu_i e^{\lambda_i t} v_i\end{aligned}\tag{B.40}$$

where we must pick μ_i such that

$$\begin{aligned}\mu_i &\in [-1, 1] \subset \mathbb{R} \quad \text{for all } i, \\ \mu_i &= 0 \quad \text{if } \text{Re}(\lambda_i) < 0, \\ \text{and } |\mu_i| &= |\mu_j| \quad \text{if } \bar{\lambda}_i = \lambda_j\end{aligned}$$

which ensures that we only consider the repelling subspace and the trajectory has purely real coordinates. There is still some arbitrariness in the selection of the magnitude and sign of the coefficients which must be resolved through experimentation on a case-by-case basis, the goal being to select coefficients such that the curve spirals outwards ($\text{Re}(\lambda_i) > 0$) from q_I towards q_F .

We know that the trajectory obtained by integrating the linearization is only valid in some small distance around the fixed point. To end the curve at the saddle fixed point, we find time t such that the first derivative of the spiral matches the slope of the joining line, i.e. solve for $t = t^* > 0$

³Also called the Lagrangian submanifold (Kamenev, 2011)

such that

$$\frac{\frac{dq_i}{dt}}{\frac{dq_1}{dt}} = \frac{(q(t) - q_F)_i}{(q(t) - q_F)_1} \quad \text{for all } i \in [2, n]. \quad (\text{B.41})$$

We obtain the complete initial condition by taking the union of the spiral part and straight part,

$$\begin{aligned} q_{\text{spiral}}(\mathbf{u}) &= q_I + \Delta\gamma(\mathbf{u}) \quad \text{for } \mathbf{u} \in [-T, t^*] \\ q_{\text{straight}}(\mathbf{u}) &= q_F + (1 - \mathbf{u})(q_I + \Delta\gamma(t^*) - q_F) \quad \text{for } \mathbf{u} \in [0, 1] \\ q^{\text{IC}} &= q_{\text{spiral}} \cup q_{\text{straight}}. \end{aligned} \quad (\text{B.42})$$

We then use the length along the curve to parametrize q^{IC} , and obtain the initial condition

$$\tilde{q}^1(n) = q^{\text{IC}}\left(\frac{n}{N}\right). \quad (\text{B.43})$$

B.4.3 Lift curve to trajectory

In every iteration, we begin with a configuration space curve q^I and lift it to a phase space trajectory $(q, p, t)^I$, which we then use to obtain a variation curve δq^I , which gives us a curve for the next iteration $q^{I+1} = q^I + \delta q^I$. In this subsection, we define the *lift* map and comment on methods for its implementation.

For a discretized curve \tilde{q} , for each segment $\tilde{q}(i)$ to $\tilde{q}(i+1)$ we assign an optimal momentum value $p(i)$ at the center of the segment and a time interval $\Delta t(i)$ denoting the time taken for transition from $q(i)$ to $q(i+1)$ (see Figure 3.5), using the following equations. The process of assigning an optimal momentum value using a convex Hamiltonian is an instance of a Legendre-transformation, for a detailed exposition see Touchette (2005).

$$\begin{aligned} 0 &= H\left(p(n), \tilde{q}(n) + \frac{\Delta\tilde{q}(n)}{2}\right) \\ \frac{\Delta\tilde{q}(n)}{\Delta t} &= \left. \frac{\partial H}{\partial p} \right|_{(p(n), \tilde{q}(n) + \frac{\Delta\tilde{q}(n)}{2})} \end{aligned} \quad (\text{B.44})$$

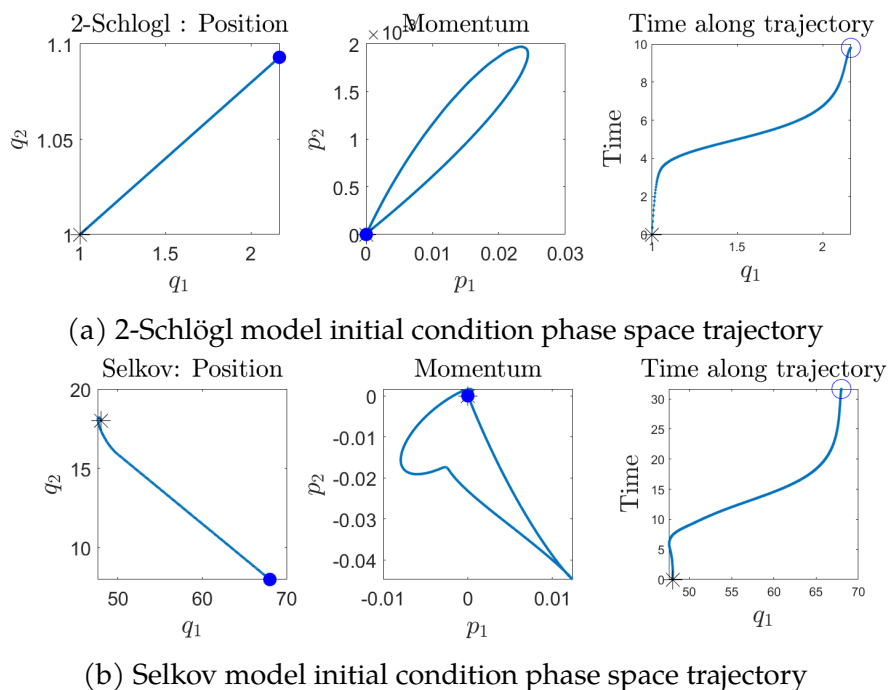


Figure B.3: Phase space trajectory obtained by lifting the initial conditions shown in Figure B.2.

or equivalently

$$\begin{aligned}
 \mathbf{p}(n), \Delta t(n) &= \arg \max_{\mathbf{p}} \min_{\Delta t} \\
 &\left(\mathbf{p} \cdot \Delta \tilde{\mathbf{q}}(n) - \left(H \left(\mathbf{p}, \tilde{\mathbf{q}}(n) + \frac{\Delta \tilde{\mathbf{q}}(n)}{2} \right) - 0 \right) \Delta t \right) \quad (\text{B.45})
 \end{aligned}$$

The $D+1$ values are assigned so as to simultaneously satisfy the $H(\mathbf{p}, \mathbf{q}) = 0$ constraint (1 equation) and the Legendre transform condition (D equations), as shown in B.44. In fact, the two problems can equivalently be assimilated in a single optimization problem, as shown in Eq. B.45, where Δt plays the role of a Lagrange-multiplier. Since the Hamiltonian is convex in momentum, the objective function is concave in \mathbf{p} , so a unique max exists. At the optimal momentum assignment, the gradient of H in \mathbf{p} and $d\mathbf{q}$ are in the same direction, and the magnitude of the two vectors is made equal via Δt .

Finally, we appropriately assign the momentum and time assignment along the configuration curve as shown below and obtain a discrete phase space trajectory.

$$\begin{aligned}
 t(1) &= 0 \\
 t(n) &= \sum_{m=1}^{n-1} \Delta t(m) \\
 \tilde{p}(1) &= 0 = \tilde{p}(N) \\
 \tilde{p}(n) &= \frac{p(n-1) + p(n)}{2} \\
 \tilde{\gamma}(n) &= (\tilde{q}, \tilde{p}, t)(n)
 \end{aligned} \tag{B.46}$$

We call this procedure as the *lift* of a configuration curve to phase space trajectory. Denoting the lift map by Λ , we define it as

$$\begin{aligned}
 \Lambda : \tilde{\mathcal{P}}_Q &\rightarrow \tilde{\mathcal{P}}_{T^*Q} \\
 \tilde{q} &\mapsto \tilde{\gamma} \equiv (\tilde{q}, \tilde{p}, t)
 \end{aligned} \tag{B.47}$$

We denote the inverse or the projection map from phase space trajectory to configuration space curve by Π , and define it to be

$$\begin{aligned}
 \Pi : \tilde{\mathcal{P}}_{T^*Q} &\rightarrow \tilde{\mathcal{P}}_Q \\
 \tilde{\gamma} &\mapsto \tilde{q}.
 \end{aligned} \tag{B.48}$$

Comments on implementation Eqs. B.44 can be solved numerically using a nonlinear least-squares problem solver, such as MATLAB's *lsqnonlin*. Eq. B.45 can be solved using a constrained optimization solver, such as MATLAB's *fmincon*.

Eq. B.44 can also be solved analytically by first doing a change of coordinates, like the one in Eq. B.22, in which the Hamiltonian function and its

derivatives are polynomials rather than consisting of exponential terms.

$$\begin{aligned} p &\rightarrow z = e^p \\ \mathcal{H}(q, z) &= 0 \\ \frac{\Delta q_i}{\Delta q_1} &= \frac{\frac{\partial \mathcal{H}}{\partial z_i} z_1}{\frac{\partial \mathcal{H}}{\partial z_1} z_i} \quad i \in [2, n] \subset \mathbb{Z} \end{aligned}$$

This gives D equations for D components of z . We can then use this to solve for Δt

$$\Delta t = \Delta q_1 \left(z_1 \frac{\partial \mathcal{H}}{\partial z_1} \right)^{-1}$$

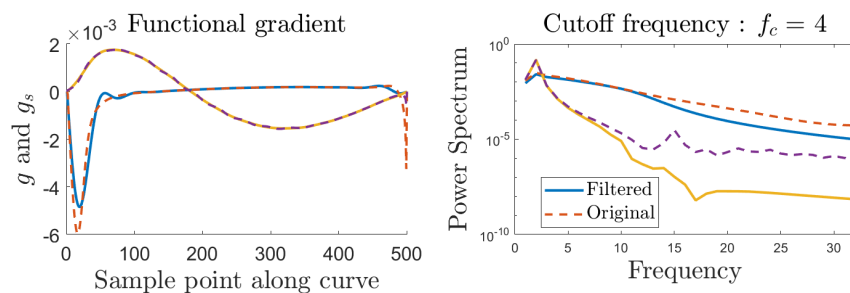
It must be noted that numerically solving for roots of polynomials can get costly in high dimensions. Also, the numerical precision decreases as we go from z to p coordinates, and it is the latter that we need for integrating Hamilton's equations of motion.

Irrespective of the implementation, for calculating momentum and time assignments along the complete trajectory, the problem B.45 or B.44 must be solved individually for each segment. Thus, a parallel implementation over multiple cores is natural and can save a lot of computational time. Also, we recommend smoothing the momentum values along the trajectory using a moving average at this step in order to partially reduce noise introduced by discretization and even out the effect of numerical errors. For a concrete MATLAB implementation, see Gagrani (2022).


B.4.4 Functional gradient - obtaining and filtering

In the action functional gradient descent (AFGD) algorithm, we perform a functional gradient descent using negative of the functional gradient of the action functional, derived in Eq. B.26. Denoting the functional gradient by g , we get

$$\begin{aligned} g &= -\frac{\delta \mathcal{A}}{\delta q} \\ &= -\left(\frac{dp}{dt} + \frac{\partial H}{\partial q} \right) \end{aligned}$$



(a) 2-Schlögl model original and filtered functional gradient at the first iteration.



(b) Selkov model original and filtered functional gradient at the first iteration.

Figure B.4: Extracting functional gradient from phase space trajectory and smoothing it. In all panels, the unfiltered and filtered curves are represented by dashed and solid lines, respectively.

Thus we define the discrete functional gradient \tilde{g} as

$$\tilde{g}(\mathbf{m}) = \left(2 \frac{\mathbf{p}(\mathbf{m}) - \mathbf{p}(\mathbf{m} - 1)}{\Delta t(\mathbf{m}) + \Delta t(\mathbf{m} - 1)} + \frac{\partial H}{\partial \mathbf{q}} \Big|_{(\tilde{\mathbf{q}}(\mathbf{m}), \tilde{\mathbf{p}}(\mathbf{m}))} \right) \quad (\text{B.49})$$

The result of naively performing a functional gradient descent using the above gradient is displayed in Figure 3.6. It can be seen that any numerical inaccuracies in solving for the momentum will be amplified by taking the time derivative, resulting in self-amplifying noise and instability. In order to smooth the noisy signal thus produced, we employ a filtering routine, that we explain in the next subsection, and obtain a smooth discrete function \tilde{g}_s that we use for updating the algorithm.

B.4.5 Filtering and resampling routines

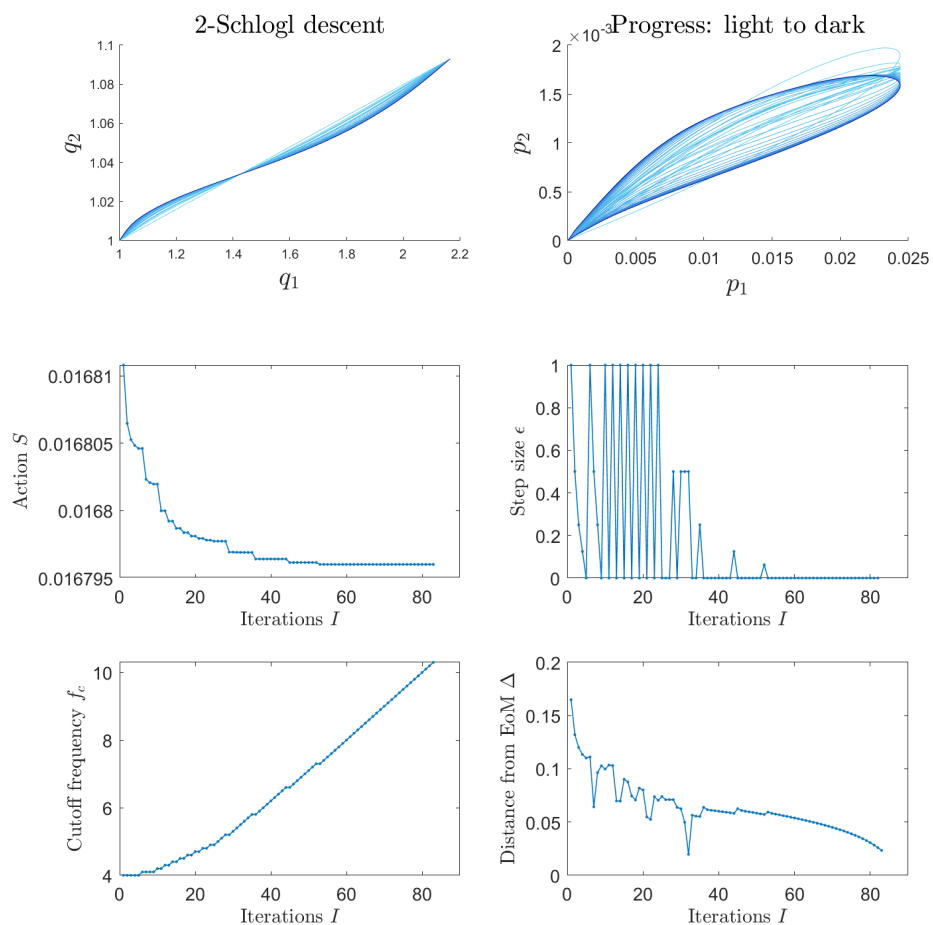


Figure B.5: Descent progress and summary for the 2-Schlogl model for a trajectory with 500 points.

Filtering

In order to filter a discrete function \tilde{y} which takes value at N points, we employ the following routine.

First, we subtract a straight line joining the end points from the function,

to get a new function $\delta\tilde{y}$ which is identically zero at the end point.

$$\begin{aligned}\tilde{y}_{\text{st}}(n) &= \tilde{y}(N) + \left(1 - \frac{n-1}{N-1}\right) (\tilde{y}(1) - \tilde{y}(N)) \\ \delta\tilde{y} &= \tilde{y} - \tilde{y}_{\text{st}}\end{aligned}$$

Next, we define a concatenated function $\delta\tilde{y}_c$ of size $2N - 1$ which is obtained by juxtaposing a flipped copy with a negative sign next to the original signal and by removing the duplicate point at N .

$$\delta\tilde{y}_c(u) = \begin{cases} -\delta\tilde{y}(N - u + 1) & \text{if } u \in [1, N - 1] \\ \delta\tilde{y}(u - N + 1) & \text{if } u \in [N, 2N - 1] \end{cases}$$

Next, we apply a Butterworth lowpass filter $\mathcal{B}(\cdot, f_c)^4$, with some cutoff frequency f_c , to the concatenated function and obtain a filtered function $\delta\tilde{y}_{\text{cs}}$. Notice $\delta\tilde{y}_c$ is an odd function across the mid-point N , and thus $\delta\tilde{y}_{\text{cs}}$ will also be of the same form. Concatenation before filtering is a common technique employed in signal processing, without which the end-points are not guaranteed to remain at zero after applying the lowpass filter.

$$\begin{aligned}f_c &:= \text{Cutoff frequency} \\ \delta\tilde{y}_{\text{cs}} &= \mathcal{B}(\delta\tilde{y}_c, f_c)\end{aligned}$$

Finally, we obtain the desired signal \tilde{y}_s by picking out only the second half of the smoothed concatenated function $\delta\tilde{y}_{\text{cs}}$ and adding that back to the straight line from the first step

$$\begin{aligned}\delta\tilde{y}_s(n) &= \delta\tilde{y}_{\text{cs}}(n + N - 1) \quad \text{for } n \in [1, N] \\ \tilde{y}_s &= \tilde{y}_{\text{st}} + \delta\tilde{y}_s.\end{aligned}\tag{B.50}$$

Resampling

To resample the discrete curve \tilde{q} , we first need a parametrization. Consider the continuous curve q obtained by interpolation. Then there are two choices of parametrization canonically available to us, namely

⁴For our implementation we choose Butterworth filter of order 4, but this choice is arbitrary and can be experimented with.

1. Arc-length parametrization

$$s(x) = \frac{\int_{q_I}^x |dq|}{\int_{q_I}^{q_F} |dq|} \quad (\text{B.51})$$

2. Time parametrization

$$s(x) = \frac{t(x)}{t(q_F)} \quad (\text{B.52})$$

Now we obtain a new discrete function \tilde{q}_u uniform in a chosen parametrization $s(x)$ by finding a point x where $s(x) = m/N$, i.e.

$$\tilde{q}_u(m) = \bar{x} \quad \text{such that } s(\bar{x}) = \frac{m}{N} \text{ for } m \in [1, N]. \quad (\text{B.53})$$

We will call a discrete curve \tilde{q}_u space-uniform or time-uniform sampled if we use the arc-length parametrization in Eq. B.51 or time parametrization in Eq. B.52 respectively.

B.4.6 Pick step size

Once we have a descent direction \tilde{g}_s obtained by filtering the functional gradient, we need to pick a step size that ensures that the value of the action functional is strictly decreasing. More precisely, we need to pick $\epsilon > 0$ such that

$$\begin{aligned} \mathcal{A}[\tilde{\gamma}^\epsilon] &< \mathcal{A}[\tilde{\gamma}], & (\text{B.54}) \\ \text{where } \tilde{\gamma} &= \Lambda(\tilde{q}) \\ \text{and } \tilde{\gamma}^\epsilon &= \Lambda(\tilde{q} + \epsilon\tilde{g}_s). \end{aligned}$$

Ideally, in order to maximize descent, we want to pick the largest $\epsilon > 0$ such that the above conditions are satisfied. In practice however, this will require us to solve another optimization problem which can be rather time consuming. Thus, for ease of implementation we employ the *backtracking line search* method in which one starts from a large value for ϵ and keeps making it smaller until the conditions are satisfied. For an exposition of

the method and more sophisticated ‘line search’ algorithms for picking a step size, see Wright et al. (1999).

If the step size is below a threshold $\epsilon < \epsilon_{\text{thresh}}$ or change in the value of the action is too small $|\mathcal{A}[\tilde{\gamma}^\epsilon] - \mathcal{A}[\tilde{\gamma}]| < \Delta S_{\text{thresh}}$, then we end the search and assign a step size $\epsilon = 0$. This indicates that the algorithm can not descend further with the given conditions and takes it to the next phase of either updating cutoff frequency or increasing the number of sample points.

As an illustration of convergence, see the fourth panel in Figures B.5, B.6 and 3.10. Notice that in all figures, the step size is identically zero for a few iterations, thus indicating that the algorithm must either be terminated or taken to the next phase.

B.4.7 Updating cutoff frequency and increasing sample points during descent

Updating cutoff frequency

The low-pass filter on the functional gradient serves the purpose of controlling noise due to discretization and numerical solving. However, it also cuts off meaningful signal in the gradient, especially in the beginning of the descent when the cutoff frequency f_c takes a very small value. To remedy this, at a given cutoff frequency, we let the algorithm converge till it cannot take a further step in the descent direction, and then increase the cutoff frequency to $f_c + \Delta f$. The value of Δf can be chosen by experimentation and we choose it to be 0.1 for our implementation.

We also define a maximum cutoff frequency f_{Max} as a considerable fraction of the Nyquist frequency. However, it must be noted that, in practice, the algorithm will stop descending at a much lower cutoff frequency than f_{Max} . In other words, for no step size will the filtered gradient yield a smaller value of the action functional than its current value. A reason for this might be that the gradient is dominated by noise which is being allowed to go through the pass-band. This is when, for descending further, we employ the *annealing* subroutine.

For an illustration of how the cutoff frequency updates with with iterations, see the fifth panel in Figures B.5, B.6 and 3.10. Notice that in all of these, there are a few iterations where the cutoff frequency remains the same before increasing to a slightly higher value. As explained earlier,

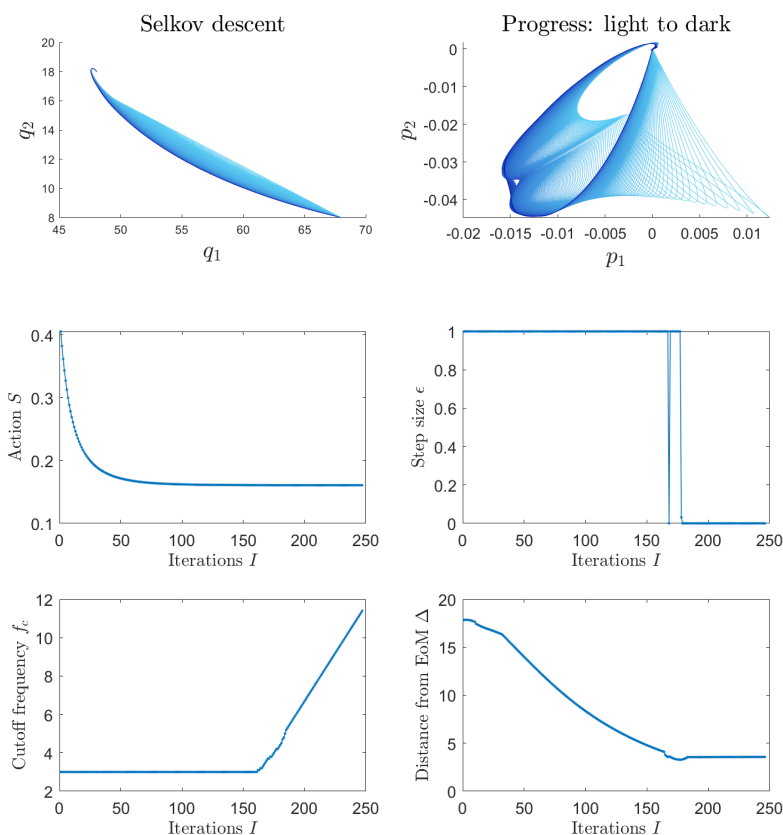


Figure B.6: Descent progress and summary for the Selkov model for a trajectory with 2500 points.

during these iterations the algorithm takes a non-zero step size and the value of the action functional steadily decreases.

Annealing or increasing sample points

At a given iteration the curve only has a finite number of points, say N . However, in principle, the optimal trajectory that minimizes the action functional is a continuous function i.e. we need the limit $N \rightarrow \infty$ to accurately represent it. We get around this problem by first descending with a small value for N until the algorithm converges, and then updating the number of points in the trajectory to $N + \Delta N$.

To illustrate the importance of this subroutine, we consider an appli-

cation of the AFGD algorithm to the Selkov model. In Figure B.6 it can be seen that the algorithm converges to some trajectory that exhibits non-differentiability in the momentum coordinates. We know that the true optimal path must be a smooth function, and thus the algorithm has not converged to the true solution. To remedy this, we take the converged trajectory and resample in the time-uniform sampling, defined in Section B.4.5, with 4000 points. As is evident from Figure 3.10 and the discussion in Section 3.4, this indeed takes the algorithm towards the optimal trajectory.

In principle, one must employ the annealing subroutine infinitely many times, since the true optimal trajectory is a continuous function. In practice however, since we can monitor the progress of the algorithm by integrating Hamilton's EoM, as we explain below, after a desired accuracy is reached one can terminate the AFGD algorithm and, if needed, use the shooting-method (Press et al., 2007; Dykman et al., 1994).

B.4.8 Integrating Hamilton's equations of motion and convergence criteria

The objective of the AFGD algorithm is to converge at the optimal curve constrained at the end points, such that the lifted trajectory is a solution to the Hamilton-Jacobi PDE in Eq. 3.13. In App. B.3.3, we prove that any solution of Hamilton-Jacobi must also satisfy Hamilton's equations of motion (EoM), thus the optimal lifted trajectory must also satisfy them. Since Hamilton's EoM is a system of coupled ODEs, it is easier to find their solution starting from an initial condition as opposed to solving HJ equations. We will now use this property of the optimal trajectory to define a 'distance' from the true solution at a given iteration and a convergence criterion, as explained below.

At a given iteration I , let us denote the discrete phase space trajectory as $\tilde{\gamma}$ and parametrize it with discrete index $m \in [1, N]$. Now, integrate Hamilton's EoM forwards and backwards starting from each point m until the configuration space distance from the saddle and stable fixed points, respectively starts diverging, having passed through its point of closest approach. We take the minimum Euclidean distance in configuration space near the stable and saddle points, and add them to obtain a $\Delta(m)$ for each point m . Finally, we find the minimum over all m to assign a

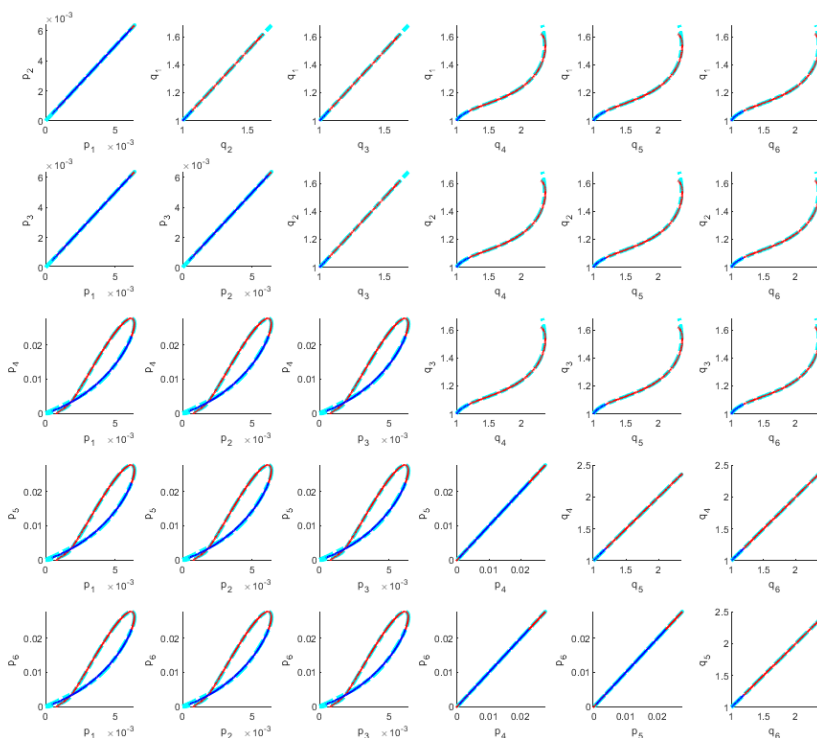


Figure B.7: Integrating Hamilton's equations of motion against AFGD algorithm output for the 6-D Schlögl model after 300 iterations. The convention is the same as used in Figure 3.14, and is omitted for neatness.

distance of the phase space trajectory, denoted by $\Delta^I = \min_m \Delta(m)$, and use it as a 'measure' for distance from optimality.

Note that since the optimal curve passes through both the saddle and stable fixed point, the distance Δ of the optimal trajectory must be zero by definition. In practice, however, one can define a Δ_{thresh} such that when $\Delta \leq \Delta_{\text{thresh}}$, we will declare the algorithm to have converged.

For an illustration of how this measure changes with iterations for the Selkov model, see the bottom right panel in Figures B.6, B.5 or last three panels in Figure 3.10. To see how Δ changes during the descent of the 2-Schlögl model, see Figure 3.14. In some of these figures, it can be seen that Δ continues to decrease even though the step size is 0. The reason

for this is the additional time-uniform filtering on the curve each time the cutoff frequency is updated. Since distance from Hamilton's EoM is a theoretically rigorous measure of convergence, we can use this to convince ourselves of convergence for higher dimensional models, where alternative verification methods such as the shooting method or Gillespie simulations can be rather expensive. For example of the application of this measure to a 6-D model, see Figure B.7.

C APPENDIX TO CHAPTER 4

C.1 Hamilton-Jacobi theory for stochastic population processes

Description	Denoted by	Belongs to
State	\mathbf{n}	$\mathbb{R}^D, \mathbb{Z}^D, \mathbb{Z}_{/p}^D$
Event	(\mathbf{n}, t)	$t \in \mathbb{R}$
Probability	$\rho(\mathbf{n}, t)$	$\rho \geq 0, \sum_{\mathbf{n}} \rho = 1$
Transition rate matrix	$W_{\mathbf{n} \rightarrow \mathbf{n}'}$	$\sum_{\mathbf{n}'} W_{\mathbf{n} \rightarrow \mathbf{n}'} = 0$

In this subsection, we briefly outline the arguments leading from a stochastic population process to its Hamilton-Jacobi theory for obtaining conditional probabilities of events.

C.1.1 States, events, and probability

Consider a phenomenon under scientific investigation, where the types of all (possibly infinite) pragmatically irreducible entities in an experiment are known. The notion of irreducibility is not to be taken in the physical sense, like a description in terms of subatomic particles, rather in a pragmatic sense, like a molecule or a cell may be considered to be an irreducible entity for a chemical or biological experiment, respectively. Let us refer to each irreducible entity as a *particle*, each type of an irreducible entity as a *species*, and denote it by s , the list of types as the *species set*, and denote it by $\mathcal{S} = \{s_1, \dots, s_S\}$, and the number of particles of a given species s_i as the *count* of that species, and denote it by n_i .

States: Each particle can possess internal degrees of freedom, such as chirality, charge, spin, or mass, depending on the species of which it is a member, and external degrees of freedom, such as relative position with respect to other particles. Typically, a state of the experiment is used to refer to the complete information of all the degrees of freedom of each particle in a system (Pathria, 2016). We are primarily interested in the *Population Processes* (PP) of species, for chemical and biological applications, and each particle of a type will be considered indistinguishable. Hence, by a

state, we simply mean the information of the counts of each species in the experiment and denote it by

$$\mathbf{n} := (n_1, \dots, n_S) \in \mathbb{Z}_{\geq 0}^S.$$

We refer to the population state as *species composition* in the main text.

Events: Consider a time parameter t external to the experiment, where $t \in \mathbb{R}$. The state of an experiment at some time t constitutes an *event* E , which we define as

$$E := (n_E, t_E).$$

Event E_1 is said to precede event E_2 if $t_{E_1} < t_{E_2}$.

Probability: Let us denote the probability of event E by $\rho(E)$, where

$$\rho(E) := \rho(n_E, t_E).$$

By the definition of probability, and our restriction to experiments where the species set is completely known,

$$\sum_{n_E} \rho(n_E, t_E) = 1 \quad \forall t_E.$$

Let us denote the joint probability of events E_1 and E_2 as $\rho(E_2 \cup E_1)$. Then, by the definition of conditional probability, we have

$$\rho(E_2 \cap E_1) = \rho(E_2|E_1)\rho(E_1), \tag{C.1}$$

where $\rho(E_2|E_1)$ is the probability that event E_2 occurs conditioned on event E_1 .

C.1.2 Event transition rates and probabilities

Consider events E_1 , E_2 , and E_3 , such that $t_{E_1} < t_{E_2} < t_{E_3}$. Then, by application of Eq. C.1, we have

$$\rho(E_3 \cap E_2 \cap E_1) = \rho(E_3|E_2 \cap E_1)\rho(E_2 \cap E_1).$$

If the phenomenon under investigation is such that the probability of a future event only depends on the current state of the experiment as opposed to prior events, i.e. the process satisfies the Markov property (Kelly, 2011), then the above relation simplifies to

$$\rho(E_3 \cap E_2 \cap E_1) = \rho(E_3|E_2)\rho(E_2 \cap E_1),$$

which, when summed over n_{E_2} , yields the Chapman-Kolmogorov equation

$$\rho(E_3 \cap E_1) = \sum_{n_{E_2}} \rho(E_3|E_2)\rho(E_2 \cap E_1). \quad (\text{C.2})$$

Event transition rates: Let us choose the events E_1, E_2 and E_3 such that $|t_{E_3} - t_{E_2}| \ll |t_{E_2} - t_{E_1}|$, and denote them as

$$E_1 = (n_i, t_i), E_2 = (n', t - \delta t), E_3 = (n, t),$$

where $t_i \ll t$ and $\delta t > 0$. Let us denote the probability of transition from state $n' \xrightarrow{\delta t} n$ in some small time δt to be $W_{n' \rightarrow n} \delta t$,

$$\rho(n, t \cap n', t - \delta t) = W_{n' \rightarrow n} \delta t,$$

and refer to $W_{n' \rightarrow n}$ as the *transition rate matrix*. Substituting the relation in Eq. C.2 by using

$$\rho(n, t | n', t - \delta t) = \left(1 - \delta t \sum_{n''} W_{n' \rightarrow n''} \right) \delta(n - n') + \delta t W_{n' \rightarrow n}$$

and taking the limit $\delta t \rightarrow 0$, we obtain the so-called continuous-time *master equation* (see Ref. Altland and Simons (2010))

$$\frac{\partial \rho(n, t \cap n_i, t_i)}{\partial t} = \sum_{n'} [W_{n' \rightarrow n} \rho(n', t) - W_{n \rightarrow n'} \rho(n, t)]. \quad (\text{C.3})$$

Notice that $W_{n' \rightarrow n}$ is the *rate of transition* from state n' to n and each coefficient has the dimension of inverse time. The master-equation is a partial differential equation (PDE) for the time-evolution of the joint probability $\rho(n, t \cap n_i, t_i)$ under a continuous-time Markov process. In what

follows, for notational simplicity, we will often omit the initial conditions (n_i, t_i) from the arguments of the probability distribution, and implicitly assume that any solution of the PDE found by providing initial condition represents the joint distribution.

Hamiltonian function and the Hamilton-Jacobi formalism: As has been shown in Ref. Kamenev (2023) and Ref. Gagrani and Smith (2023), it is useful to define Hamiltonian and momentum operators $\underline{\hat{H}}$ and \hat{P} , respectively, and recast the master-equation, Eq. C.3, as

$$\begin{aligned} \frac{\partial \rho(\mathbf{n}, t)}{\partial t} &= \sum_r (e^{-r \frac{\partial}{\partial \mathbf{n}}} - 1) W_{\mathbf{n} \rightarrow \mathbf{n}+r} \rho(\mathbf{n}, t) \\ &= \underline{\hat{H}} \left(\hat{P}, \mathbf{n} \right) \rho(\mathbf{n}, t), \end{aligned} \quad (\text{C.4})$$

where,¹

$$\hat{P} = -\frac{\partial}{\partial \mathbf{n}}.$$

Furthermore, if one passes from a discrete count variable n to a continuous *concentration* variable $q = n/V$, where V is a scale of the total population, then an asymptotic solution (in V) to the above equation can be obtained by employing large-deviations theory and the Hamilton-Jacobi formalism. While we summarize the main results below, for a complete derivation we refer the readers to Sec. II A of Ref. Gagrani and Smith (2023) and Chap. 4 of Ref. Smith and Krishnamurthy (2015).

From $\underline{\hat{H}} \left(\hat{P}, \mathbf{n} \right)$, we first obtain the Hamiltonian function $\underline{H}(p, \mathbf{n})$,

$$\underline{H}(p, \mathbf{n}) = \sum_r (e^{r p} - 1) W_{\mathbf{n} \rightarrow \mathbf{n}+r},$$

¹Formally extending n from an integer to the real domain, a function in n can be shifted to $n + y$ by the application of a *shift operator* given by $e^{y \frac{\partial}{\partial x}}$, as can be seen by

$$f(x + y) = \sum_{n=0}^{\infty} \frac{y^n}{n!} \frac{\partial^n f(x)}{\partial x^n} = e^{y \frac{\partial}{\partial x}} f(x).$$

by replacing the momentum operator \hat{P} with the momentum function p ,

$$\hat{P} \rightarrow p, \quad \underline{H}(\hat{P}, n) \rightarrow \underline{H}(p, n).$$

(Notice that the Hamiltonian function is always convex in the momentum coordinates p .) Next, we pass from a count variable to a concentration variable

$$n \rightarrow qV, \quad \underline{H}(p, n) \rightarrow \underline{H}(p, qV),$$

and define a new Hamiltonian function in the concentration coordinates $H(p, q)$

$$H(p, q) = \sum_r (e^{rp} - 1) w_{q \rightarrow q+r}, \quad (\text{C.5})$$

where $w_{q \rightarrow q'}$ is chosen such that²

$$\underline{H}(p, Vq) \asymp VH(p, q).$$

Under these transformations, either using path-integral methods explained in Ref. Gagrani and Smith (2023), or large deviation theory (see Fig. C.1), the joint probability of an event (q, t) succeeding (q_i, t_i) asymptotically equals

$$\rho(q, t \cap q_i, t_i) \asymp e^{-V \int_{q_i, t_i}^{q, t} dS(q', t')} \rho(q_i, t_i), \quad (\text{C.6})$$

where $S(q', t')$ satisfies the Hamilton-Jacobi equations (compare with Eq. C.4)

$$\begin{aligned} p^* &= \frac{\partial S(q, t)}{\partial q}, \\ \frac{\partial S(q, t)}{\partial t} &= H(p^*, q), \end{aligned} \quad (\text{C.7})$$

²We say $f(V) \asymp g(V)$ (read as $f(V)$ is asymptotic to $g(V)$) if

$$\lim_{V \rightarrow \infty} \frac{f(V)}{g(V)} = 1.$$

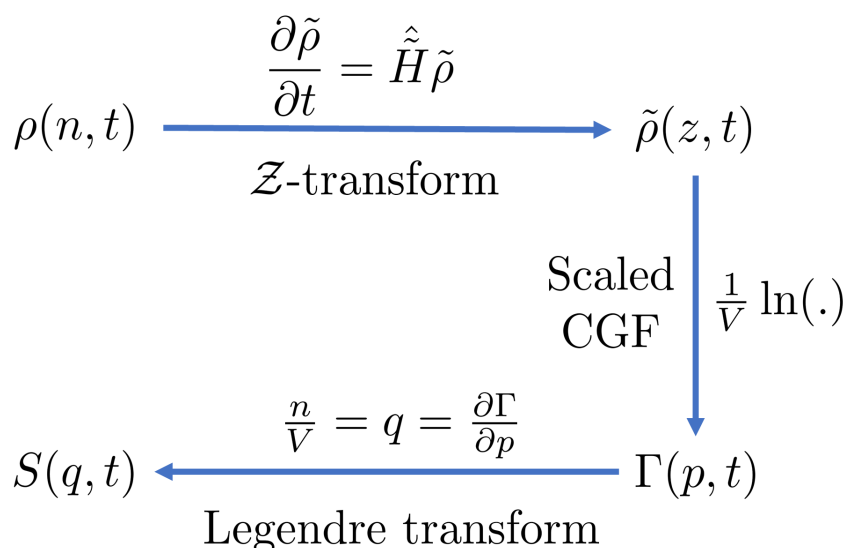


Figure C.1: Extracting a large-deviations function from a probability distribution for a stochastic population process (see Refs. (Smith and Krishnamurthy, 2015; Rassoul-Agha and Seppäläinen, 2015; Touchette, 2009; Ellis, 2006)). As an application of Cramér’s theorem, we get that $\rho(qV, t) \asymp e^{-VS(q,t)}$.

subject to the boundary condition

$$S(q_i, t_i) = -\frac{1}{V} \ln(\rho(q_i, t_i)).$$

$S(q, t)$ is called the *action* function and can be equivalently formulated as (see Sec. III A of Ref. Gagrani and Smith (2023))

$$\int_{q_i, t_i}^{q, t} dS(q', t') = \inf_{q'(t')} \sup_{p(t')} \int_{t_i}^t dt' (p \dot{q}' - H(p, q')), \quad (\text{C.8})$$

constrained to the conditions

$$q'(t_i) = q_i, \quad q'(t) = q.$$

Finally, as explained in Refs. (Arnol’d, 2013; Goldstein et al., 2002; Courant and Hilbert, 2008) and derived in Appendix C 3 of Ref. Gagrani

and Smith (2023), the phase space path $(q(t), p(t))$ where the ‘inf-sup’ (or ‘min-max’) condition is satisfied, henceforth referred to as the *optimal path*, is a solution to the Hamilton’s equations of motion

$$\begin{aligned}\frac{dq}{dt} &= \frac{\partial H(p, q)}{\partial p}, \\ \frac{dp}{dt} &= -\frac{\partial H(p, q)}{\partial q},\end{aligned}\tag{C.9}$$

constrained to the conditions

$$q'(t_i) = q_i, \quad q'(t) = q.$$

Event transition probability: It can be read from Eqs. C.1 and C.6 that the probability of event (q, t) conditioned on the event (q_i, t_i) is given by

$$\rho(q, t|q_i, t_i) \asymp e^{-\mathcal{V} \int_{q_i, t_i}^{q, t} dS(q', t') + c}.\tag{C.10}$$

Thus, intuitively, the Action function or a solution to the Hamilton-Jacobi equation can be seen as the time-evolving descaled log-improbability of observing the concentration q at time t conditioned on having observed concentration q_i at time t_i .

For very long time scales, in the limit $|t - t_i| \rightarrow \infty$, the evolving probability distribution given by the solution to the master-equation (Eq. C.3) converges upon the stationary distribution $\pi(q)$ (for reasons discussed towards the end of Sec. II A in Ref. Gagrani and Smith (2023)). The Action or scaled log-improbability of the stationary distribution $\pi(q)$ is given by the *Non-Equilibrium Potential* (NEP), denoted by $\mathcal{V}(q)$. By definition of the stationary distribution,

$$\frac{\partial \mathcal{V}}{\partial t} = 0.$$

Substituting it in Eq. C.7, we get that the NEP must satisfy

$$\mathcal{H}\left(p^* = \frac{\partial \mathcal{V}}{\partial q}, q\right) = 0.\tag{C.11}$$

Thus, from Eq. C.10, the probability of observing the event (q, ∞) condi-

tioned on the event $(q_i, 0)$ is given by

$$\rho(q, \infty | q_i, 0) \asymp e^{-\mathcal{V} \int_{q_i}^q d\mathcal{V}(q')}. \quad (\text{C.12})$$

From Eq. C.5, notice that for $p = 0$,

$$H(0, q) = 0 \quad \forall q.$$

Also, this solution satisfies the Hamilton's equation of motion (Eq. C.9) for momentum evaluation,

$$\left. \frac{dp}{dt} \right|_{p=0} = 0,$$

and yields

$$\left. \frac{dq}{dt} \right|_{p=0} = \sum_r r w_{q \rightarrow q+r} \quad (\text{C.13})$$

for the rate of change in concentration. Since the NEP \mathcal{V} along optimal paths with $p = 0$ is identically zero,

$$\mathcal{V}(q)|_{p=0} = 0,$$

the conditional probability in Eq. C.10 has the most contribution from trajectories that satisfy Eq. C.13. For the above mentioned reasons, and those explained in Sec. II C of Ref. Gagrani and Smith (2023), the solutions to Eq. C.13 are termed as the *deterministic* or *relaxation* dynamics of the stochastic system under consideration.

Finally, we remark that there are also optimal paths with momentum $p^* \neq 0$ on the $H(p^*, q) = 0$ submanifold. We will call these assignments of momenta as the *escape momenta*. These optimal paths correspond to the *least-improbable* paths constrained to the end-points, and are termed the *escape trajectories* when the initial fixed point is a stable fixed point of the relaxation dynamics given by Eq. C.13. Notice that these trajectories are due to the inherent stochasticity of the system, and completely missed by the deterministic description of the system given by Eq. C.13. For more details, see Sec. II C in Ref. Gagrani and Smith (2023).

C.2 Species interaction networks and their properties

From the master equation, Eq. C.3, we notice that, at the smallest time scale, the information about state transitions is in the transition matrix $W_{\mathbf{n}' \rightarrow \mathbf{n}}$. Recall that each entry in the counts vector \mathbf{n} denotes the population of the corresponding type of species. Since we are interested in the population dynamics of the particles of a species, rather than their geometric or spatial description, the interaction between the species cannot be a priori specified, and is governed by the phenomenon under investigation. In this subsection, we will consider the case where the transition matrix has a special structure that it inherits from the network of species interactions.

This interaction network is the purely topological data on the dynamics, one that can be said prior to the particular dynamics. When we take a reaction network, we can decompose it into subnetworks, by criteria defined in this subsection. It is to be noted that while the classification is done using linear algebra and polyhedral geometry, the decomposition into subnetworks is a topological operation on the interaction network and does not directly determine anything about the dynamics of the system. These considerations do, however, tell us where a system cannot be because each system must be in the same stoichiometric subspace as its initial boundary constraint. They can also help specify the polyhedral decomposition admitted by the space of flows on the hypergraph into null-spaces and positively or negatively flow-productive and species-productive cones (defined in Ref. Gagrani et al. (2023a)).

C.2.1 Chemical reaction network

Recall that the set of all species is denoted by \mathcal{S} . Suppose that each interaction is of the form $\mathbf{y} \rightarrow \mathbf{y}'$, where it takes a multiset (collection admitted with a non-negative multiplicity) of species $\mathbf{y} = \{\mathbf{y}_i \in \mathbb{Z}_{\geq 0} | s_i \in \mathcal{S}\}$ as a *reactant*, and produces a different multiset $\mathbf{y}' = \{\mathbf{y}'_i \in \mathbb{Z}_{\geq 0} | s_i \in \mathcal{S}\}$ as a *product*. The reactant and product multisets are termed *complexes*, and the set of all complexes is denoted by \mathcal{C} . A network of such interactions is formalized as a chemical reaction network (CRN) (Feinberg, 2019; Yu and

Craciun, 2018), and specified as a hypergraph

$$\mathcal{G} = (\mathcal{S}, \mathcal{R})$$

with each species $s \in \mathcal{S}$ and reaction $y \rightarrow y' \in \mathcal{R}$ forming a vertex and a hyper-edge, respectively. While it is traditional (Feinberg, 2019; Anderson et al., 2015) to define a CRN as the triple $(\mathcal{S}, \mathcal{C}, \mathcal{R})$, but as explained in Ref. Gagrani et al. (2023a) and Ref. Deshpande and Gopalkrishnan (2014), the two are equivalent, and our use of either will be dictated by convenience.

The transition rate matrix element $w_{q \rightarrow q+r}$ ($W_{n \rightarrow n+r}$) induced by a CRN \mathcal{G} is obtained by summing over all *flows* $\mathcal{K}_{y \rightarrow y'}$ ($\underline{\mathcal{K}}_{y \rightarrow y'}$) on reaction channels $y \rightarrow y'$ whose net contribution is to transform the concentration (count) from q to $q+r$ (n to $n+r$). A flow on CRN \mathcal{G} yields the transition rate matrix

$$\begin{aligned} W_{n \rightarrow n+r} &= \sum_{y \rightarrow y' \in \mathcal{G}} \delta(r - (y' - y)) \underline{\mathcal{K}}_{y \rightarrow y'}(n), \\ w_{q \rightarrow q+r} &= \sum_{y \rightarrow y' \in \mathcal{G}} \delta(r - (y' - y)) \mathcal{K}_{y \rightarrow y'}(q). \end{aligned} \quad (\text{C.14})$$

Using Eq. C.5, the Hamiltonian functions $\underline{H}_{\text{CRN}}(\mathbf{p}, \mathbf{n})$ and $H_{\text{CRN}}(\mathbf{p}, \mathbf{q})$ in the count and concentration for this transition rate matrix is

$$\begin{aligned} \underline{H}_{\text{CRN}}(\mathbf{p}, \mathbf{n}) &= \sum_{\mathcal{G}} (e^{(y'-y)\mathbf{p}} - 1) \underline{\mathcal{K}}_{y \rightarrow y'}(\mathbf{n}), \\ H_{\text{CRN}}(\mathbf{p}, \mathbf{q}) &= \sum_{\mathcal{G}} (e^{(y'-y)\mathbf{p}} - 1) \mathcal{K}_{y \rightarrow y'}(\mathbf{q}). \end{aligned} \quad (\text{C.15})$$

Using Eq. C.13, the deterministic dynamics under a CRN \mathcal{G} is given by

$$\frac{d\mathbf{q}}{dt} = \sum_{\mathcal{G}} (y' - y) \mathcal{K}_{y \rightarrow y'}(\mathbf{q}). \quad (\text{C.16})$$

As explained in Sec. II A of Ref. Gagrani et al. (2023a), it is useful to recast the above equation in a matrix form

$$\frac{d\mathbf{q}}{dt} = \mathbb{S} \mathcal{K}, \quad (\text{C.17})$$

where \mathbb{S} is the *stoichiometric matrix* with dimensions $|\mathcal{S}| \times |\mathcal{R}|$ and \mathcal{K} is the flow vector of size $|\mathcal{R}|$.

Stoichiometric subspace: We remark that Eq. C.17 ensures that the rates of change in concentration (or changes in count) for a CRN can only lie within the range of the stoichiometric matrix \mathbb{S} . This subspace is termed as the *stoichiometric subspace*, the dimension of which we denote by \mathfrak{s} . It can be seen that the stoichiometric subspace has dimensions

$$\begin{aligned}\mathfrak{s} &= \text{Rank}(\mathbb{S}) \\ &= |\mathcal{S}| - |\text{Kernel}(\mathbb{S}^T)|.\end{aligned}$$

We also want to briefly mention that elements in the left null space of the stoichiometric matrix, or $\text{Kernel}(\mathbb{S}^T)$, correspond to conserved quantities under the stoichiometry of the CRN (see Refs. (Feinberg, 2019)).

Subnetwork: We define a *subnetwork* \mathcal{G}' of a CRN \mathcal{G} to be the subgraph $\mathcal{G}' = (\mathcal{S}_{|\mathcal{R}'}, \mathcal{R}')$ where $\mathcal{R}' \subset \mathcal{R}$ and $\mathcal{S}_{|\mathcal{R}'}$ indicates that there is no omission of species in the subnetwork that participates in the set of reactions \mathcal{R}' of the original graph \mathcal{G} . We will denote the restriction of \mathbb{S} to \mathcal{G}' by

$$\bar{\mathbb{S}} := \mathbb{S}_{|\mathcal{G}'}$$

C.2.2 Minimal null subnetwork

We will define a subnetwork $\mathcal{G}' \subset \mathcal{G}$ to be a *null subnetwork* (NS) if there exists a flow $\mathcal{K} \in \text{Kernel}(\bar{\mathbb{S}})$ such that

$$\bar{\mathbb{S}} \mathcal{K} = 0.$$

Such a flow, one that is in the kernel of the stoichiometric matrix of a CRN, will be termed as a *null flow* for the CRN. Furthermore, we define a *minimal null subnetwork* (MNS) to be a NS which does not contain a smaller NS in it.

Minimal delta null subnetworks: A hypergraph \mathcal{G} can also be seen as a *directed ordinary graph* with vertices and edges that are the complexes and reactions from the sets \mathcal{C} and \mathcal{R} , respectively. Let us denote the number

of connected components in the ordinary graph of the CRN, also termed *linkage classes*, by ℓ . Furthermore, as explained in the appendix of Ref. Gagrani et al. (2023a), the stoichiometric matrix can be written as the product

$$\mathbb{S} = Y \mathbb{M},$$

where the columns of Y are the stoichiometries of the complexes

$$\text{cols}(Y) = \{\mathbf{y} \mid \mathbf{y} \in \mathcal{C}\}$$

and \mathbb{M} is the oriented vertex-edge incidence matrix of the ordinary graph (see Refs. (Balandin, 1940; Feinberg, 2019))

$$\text{cols}(\mathbb{M}) = \{\mathbf{e}_{\mathbf{y}'} - \mathbf{e}_{\mathbf{y}} \mid \mathbf{y} \rightarrow \mathbf{y}' \in \mathcal{R}\},$$

where $\mathbf{e}_{\mathbf{y}}$ is the vector with 1 at complex \mathbf{y} and 0 otherwise. From this representation of the stoichiometric matrix, it is clear that there may exist null flows \mathcal{K} which belong to the intersection of kernel of $Y\mathbb{M}$ with the orthogonal complement to kernel of \mathbb{M} in the domain of \mathbb{M} ,

$$\mathcal{K} \in \text{Kernel}(Y\mathbb{M}) \cap \text{Kernel}(\mathbb{M})^\perp.$$

Following the terminology from Ref. (Smith and Krishnamurthy, 2017), we will refer to these null flows as δ -flows. We define a minimal delta null subnetwork (MDNS) as a minimal null subnetwork that admits a δ -flow.

As proved in Ref. Feinberg (2019) or Appendix A of Ref. Gagrani et al. (2023a), the existence of a δ -flow is a necessary condition for a CRN to show multistability. The number of independent δ -flows is given by

$$\delta = |\mathcal{C}| - \ell - \mathfrak{s}, \tag{C.18}$$

where δ characterizes a topological property, termed as *deficiency*, of the network.

C.2.3 Siphon subnetwork

Extending the definitions of Deshpande et al. in Ref. Deshpande and Gopalkrishnan (2014) to subnetworks, a *siphon subnetwork* is defined as a subnetwork $\mathcal{G}' \subset \mathcal{G}$ if:

1. there exists a set $\mathcal{M} \in \mathcal{S}$, such that, there exists at least one species from \mathcal{M} in each reactant and product complex of the reaction set \mathcal{R}' .

Minimal autocatalytic subnetwork: We term a siphon subnetwork $\mathcal{G}' \subset \mathcal{G}$ as an *autocatalytic subnetwork* (AS), if there exists a species set $\mathcal{M} \in \mathcal{S}$ and a flow \mathcal{K} such that

1. all species in the set \mathcal{M} are strictly produced

$$\left. \frac{dq}{dt} \right|_{\mathcal{M}} = \bar{\mathbb{S}} \mathcal{K} \Big|_{\mathcal{M}} > \mathbf{0}.$$

We term an AS that does not contain any strictly smaller autocatalytic subnetwork as a minimal autocatalytic subnetwork (MAS).

Our definition of autocatalytic is identical to, what Andersen et al. in Ref. Andersen et al. define as, *exclusively autocatalytic*. For a detailed comparison of other definitions of autocatalysis, see Ref. Andersen et al. or Sec II B of Ref. Gagrani et al. (2023a). While we do not use their result in this work, we want to mention that Blokhuis et al. in Ref. (Blokhuis et al., 2020), have proved that for an even more restrictive notion of autocatalysis, what we term as *stoichiometrically autocatalysis*, all MASs can be categorized into five types depending on their autocatalytic core (for more details, see Sec. III A of Ref. Gagrani et al. (2023a)).

Finally, we remark that, often one is concerned with whether an introduction of an AS can potentially introduce new steady states in the system. To that end, Eq. C.18 provides a straightforward formula using which it can be checked whether addition of a new AS (or for that matter any subnetwork) increases the deficiency of the network.

Minimal drainable subnetwork: We term a siphon subnetwork $\mathcal{G}' \subset \mathcal{G}$ as a *drainable subnetwork* (DS), if there exists a species set $\mathcal{M} \in \mathcal{S}$ and a flow \mathcal{K} such that

1. all species in the set \mathcal{M} are strictly consumed (or drained)

$$\left. \frac{dq}{dt} \right|_{\mathcal{M}} = \bar{\mathbb{S}} \mathcal{K} \Big|_{\mathcal{M}} < \mathbf{0}.$$

We term an DS that does not contain any strictly smaller autocatalytic subnetwork as a minimal drainable subnetwork (MDS). In Ref. Deshpande and Gopalkrishnan (2014), Deshpande et al. explain the relevance of understanding both the autocatalytic and the drainable siphons in a comprehensive characterization of the dynamics of a CRN.

C.3 Modeling a mechanism and obtaining quantitative predictions

To obtain any quantitative predictions about the collective behavior of particles in an experiment, we need to propose a *mechanism* by which the interactions (modelled in the interaction network) are undertaken and provide the relevant *kinetic parameters*. For our purposes in this work, let us assume that the mechanism of an interaction is the following. We suppose that the particles under consideration are in a well-mixed mixture with a homogeneous particle density. Furthermore, a reaction $y \rightarrow y'$ proceeds with a rate proportional to the chance that the particles constituting the multi-set y collide with each other (mechanism), and a *rate constant* (kinetic parameter). We remark that, while these assumptions are characteristic of chemical kinetics in nature (see Refs. (Gillespie et al., 2007; Feinberg, 2019)), they have to be generalized further for biochemical processes (Johnson and Goody, 2011; Qian and Bishop, 2010).

C.3.1 Stochastic chemical kinetics:

A CRN $\mathcal{G} = (\mathcal{S}, \mathcal{R})$ is said to be taken with kinetics

$$K = \{k_{y \rightarrow y'} \in \mathbb{R}_{>0} \mid y \rightarrow y' \in \mathcal{R}\}$$

if for each reaction $y \rightarrow y' \in \mathcal{R}$, there is an associated positive *rate constant* $k_{y \rightarrow y'} \in \mathbb{R}_{>0}$ in the set K . Together, we represent the CRN with kinetics as the triple

$$\mathcal{G}_K = (\mathcal{S}, \mathcal{R}, K).$$

Since, as explained in the beginning of this subsection, we have assumed chemical kinetics, the flows on the reactions of the CRN are determined by the concentration (count) and the rate constant. Explicitly,

assuming a scale factor V , the flows in the count variable in Eq. C.14 can be written as (Anderson et al., 2015; Gagrani and Smith, 2023; Smith and Krishnamurthy, 2015)

$$\underline{\mathcal{K}}_{y \rightarrow y'} = \frac{k_{y \rightarrow y'}}{V^{y-1}} \prod_{i=1}^s \frac{n_i!}{(n_i - y_i)!} 1_{\{n_i \geq y_i\}}.$$

Making use of the scaling limit explained below Eq. C.5 and the multi-index notation, the flows in the concentration variable then becomes

$$\mathcal{K}_{y \rightarrow y'} = k_{y \rightarrow y'} q^y.$$

Substituting the above flow in Eq. C.15, we get that the Hamiltonian for a CRN with chemical kinetics is

$$H_{\text{CRN}}(p, q) = \sum_{\mathfrak{g}} (e^{(y'-y)p} - 1) k_{y \rightarrow y'} q^y. \quad (\text{C.19})$$

Since we will be assuming chemical kinetics in the rest of our work, henceforth we drop the subscript 'CRN', and refer to the Hamiltonian $H_{\text{CRN}}(p, q)$ as simply $H(p, q)$.

Deterministic mass-action kinetics: Using Eqs. C.13 and C.16, the relaxation trajectories of chemical kinetics, termed the *deterministic mass-action kinetics* (MAK), are given by

$$\begin{aligned} \frac{dq}{dt} &= \sum_{\mathfrak{g}} (y' - y) k_{y \rightarrow y'} q^y \\ &= \sum_{y' \in \mathfrak{g}} y' \sum_{y \in \mathfrak{g}} (k_{y \rightarrow y'} q^y - k_{y' \rightarrow y} q^{y'}). \end{aligned} \quad (\text{C.20})$$

Notice that the right hand-sides of the above set of ordinary differential equations (ODEs) are polynomials in the concentration variables, due to which they are also said to define a *polynomial dynamical system* (Craciun, 2019).

C.3.2 Stability, transitions, and cost of discrete attractors

Stability: In our work, we restrict to dynamical systems with a set of discrete fixed-points. Suppose there are N discrete fixed points, denoted by $\underline{q} := \{\underline{q}_1, \dots, \underline{q}_N\}$, then they satisfy

$$\left. \frac{dq}{dt} \right|_{\underline{q}} = \sum_{\mathfrak{g}} (y' - y) k_{y \rightarrow y'} \underline{q}^y = 0.$$

The stability of a fixed-point can be determined by solving for the eigenvalues of the Jacobian matrix resulting from the right-hand sides of the above set of equations. In particular, for a system of size $|\mathcal{S}|$, we will term a fixed-point with $|\mathcal{S}|$ strictly negative eigenvalues as a *attractor*, and fixed-point with $|\mathcal{S}| - 1$ strictly negative but one strictly positive eigenvalue as a *transient point*. In general, all fixed points with at least one strictly positive eigenvalue will be termed *unstable points*. For more details, see Sec. II C in Ref. Gagrani and Smith (2023).

We briefly remark that a root \underline{q} , such that

$$\sum_{y \in \mathcal{S}} \left(k_{y \rightarrow y'} \underline{q}^y - k_{y' \rightarrow y} \underline{q}^{y'} \right) = 0$$

is said to be *complex-balanced* (Baez and Biamonte, 2012; Anderson et al., 2010). A network with non-zero deficiency, under appropriate kinetics, has the potential to exhibit multiple attractors, neither of which are complex-balanced.

Transition between attractors: While determining whether or not a system can have multiple attractors is a challenging question (Feinberg, 1988, 1989), cases where there are multiple attractors are primarily of interest to us. In particular, we would like to know, given that a system is at one attractor at time $t = 0$, what is its probability of being in a different attractor after a very long time $t \rightarrow \infty$? As explained earlier in Sec. C.1.2, the descaled log-improbability of such an event is the difference of the NEP at the two steady states, where the value of the NEP is determined by finding a partial solution to the Hamilton-Jacobi equation. In Ref. Gagrani and Smith (2023), we develop a functional gradient descent algorithm that is designed to estimate such solutions for stochastic CRNs.

To conclude this subsection, we briefly explain the significance of the NEP and the role it plays in related fields. Colloquially, the NEP or the Action function along a trajectory (explained in Sec. C.1.2) can be said to quantify the *cost* of the trajectory (Snarski, 2021). This terminology is intuitive as, in the long-time limit, the log-probability of occupation of a state is negative of the cost of the state, thus exponentially favoring a lower cost than a higher cost. Moreover, any transition from one state to another is most likely to occur through the least costly trajectory connecting them, and these are precisely the solutions to Hamilton's equations of motion.

REFERENCES

-
- Aad, Georges, Tatevik Abajyan, B Abbott, J Abdallah, S Abdel Khalek, Ahmed Ali Abdelalim, R Aben, B Abi, M Abolins, OS AbouZeid, et al. 2012. Observation of a new particle in the search for the standard model higgs boson with the atlas detector at the lhc. *Physics Letters B* 716(1): 1–29.
- Altland, Alexander, and Ben D Simons. 2010. *Condensed matter field theory*. Cambridge university press.
- Andersen, Jakob L, Christoph Flamm, Daniel Merkle, and Peter F Stadler. 2016. A software package for chemically inspired graph transformation. In *Graph transformation: 9th international conference, icgt 2016, in memory of hartmut ehrig, held as part of staf 2016, vienna, austria, july 5-6, 2016, proceedings 9*, 73–88. Springer.
- Andersen, Jakob L, Christoph Flamm, Daniel Merkle, and Peter F Stadler. Defining autocatalysis in chemical reaction networks.
- Anderson, David F, Gheorghe Craciun, Manoj Gopalkrishnan, and Carsten Wiuf. 2015. Lyapunov functions, stationary distributions, and non-equilibrium potential for reaction networks. *Bulletin of mathematical biology* 77(9):1744–1767.
- Anderson, David F, Gheorghe Craciun, and Thomas G Kurtz. 2010. Product-form stationary distributions for deficiency zero chemical reaction networks. *Bulletin of mathematical biology* 72(8):1947–1970.
- Anderson, David F, Badal Joshi, and Abhishek Deshpande. 2021. On reaction network implementations of neural networks. *Journal of the Royal Society Interface* 18(177):20210031.
- Anderson, David F, and Tung D Nguyen. 2022. Prevalence of deficiency-zero reaction networks in an erdős–rényi framework. *Journal of Applied Probability* 59(2):384–398.
- Arnol'd, Vladimir Igorevich. 2013. *Mathematical methods of classical mechanics*, vol. 60. Springer Science & Business Media.

- Avanzini, Francesco, Nahuel Freitas, and Massimiliano Esposito. 2022. Circuit theory for chemical reaction networks. *arXiv preprint arXiv:2210.08035*.
- Baez, John C, and Jacob Biamonte. 2012. Quantum techniques for stochastic mechanics. *arXiv preprint arXiv: 1209.3632*.
- Balandin, AA. 1940. Structural algebra in chemistry. *Acta Physicochim*.
- Banzhaf, Wolfgang, and Lidia Yamamoto. 2015. *Artificial chemistries*. MIT Press.
- Barenholz, Uri, Dan Davidi, Ed Reznik, Yinon Bar-On, Niv Antonovsky, Elad Noor, and Ron Milo. 2017. Design principles of autocatalytic cycles constrain enzyme kinetics and force low substrate saturation at flux branch points. *Elife* 6:e20667.
- Bates, Daniel J., Jonathan D. Hauenstein, Andrew J. Sommese, and Charles W. Wampler. Bertini: Software for numerical algebraic geometry. Available at bertini.nd.edu with permanent doi: dx.doi.org/10.7274/R0H41PB5.
- Benzi, Roberto, Giorgio Parisi, Alfonso Suter, and Angelo Vulpiani. 1983. A theory of stochastic resonance in climatic change. *SIAM Journal on applied mathematics* 43(3):565–578.
- Berne, Bruce J, Giovanni Ciccotti, and David F Coker. 1998. *Classical and quantum dynamics in condensed phase simulations: Proceedings of the international school of physics*. World Scientific.
- Biondini, Gino. 2015. An introduction to rare event simulation and importance sampling. In *Handbook of statistics*, vol. 33, 29–68. Elsevier.
- Blokhuis, Alex, David Lacoste, and Philippe Nghe. 2020. Universal motifs and the diversity of autocatalytic systems. *Proceedings of the National Academy of Sciences* 117(41):25230–25236.
- Bouchet, Freddy, Joran Rolland, and Eric Simonnet. 2019. Rare event algorithm links transitions in turbulent flows with activated nucleations. *Physical review letters* 122(7):074502.

- Cao, Youfang, and Jie Liang. 2013. Adaptively biased sequential importance sampling for rare events in reaction networks with comparison to exact solutions from finite buffer dcme method. *The Journal of chemical physics* 139(2):07B605_1.
- Cardelli, Luca. 2011. Strand algebras for dna computing. *Natural Computing* 10:407–428.
- Coleman, Sidney. 1988. *Aspects of symmetry: selected erice lectures*. Cambridge University Press.
- Commons, Wikimedia. 2015. Fusion in the sun. File: Sunfusion2.jpg.
- Courant, Richard, and David Hilbert. 2008. *Methods of mathematical physics: partial differential equations*. John Wiley & Sons.
- Craciun, Gheorghe. 2019. Polynomial dynamical systems, reaction networks, and toric differential inclusions. *SIAM Journal on Applied Algebra and Geometry* 3(1):87–106.
- Deshpande, Abhishek, and Manoj Gopalkrishnan. 2014. Autocatalysis in reaction networks. *Bulletin of mathematical biology* 76(10):2570–2595.
- Dittrich, Peter. 2005. Chemical computing. *Unconventional programming paradigms* 97–97.
- Doi, Masao. 1976. Second quantization representation for classical many-particle system. *Journal of Physics A: Mathematical and General* 9(9):1465.
- Dolson, Emily, and Charles Ofria. 2021. Digital evolution for ecology research: a review. *Frontiers in Ecology and Evolution* 9:750779.
- Duruiseaux, Valentin, Jeremy Schmitt, and Melvin Leok. 2021. Adaptive hamiltonian variational integrators and applications to symplectic accelerated optimization. *SIAM Journal on Scientific Computing* 43(4): A2949–A2980.
- Dykman, Mark I, Eugenia Mori, John Ross, and PM Hunt. 1994. Large fluctuations and optimal paths in chemical kinetics. *The Journal of chemical physics* 100(8):5735–5750.

- Dyson, Freeman J. 1982. A model for the origin of life. *Journal of Molecular Evolution* 18(5):344–350.
- Eigen, Manfred, and Peter Schuster. 1977. A principle of natural self-organization: Part a: Emergence of the hypercycle. *Naturwissenschaften* 64:541–565.
- . 1978a. The hypercycle: A principle of natural self-organization part b: The abstract hypercycle. *Naturwissenschaften* 65:7–41.
- . 1978b. The hypercycle: a principle of natural self-organization part c: The realistic hypercycle. *The Science of Nature* 65(7):341–369.
- Eisenstein, Michael, et al. 2021. Artificial intelligence powers protein-folding predictions. *Nature* 599(7886):706–708.
- Ellis, Richard S. 2006. *Entropy, large deviations, and statistical mechanics*, vol. 1431. Taylor & Francis.
- Epstein, Irving R, and John A Pojman. 1998. *An introduction to nonlinear chemical dynamics: oscillations, waves, patterns, and chaos*. Oxford university press.
- Eyink, Gregory L. 1996. Action principle in nonequilibrium statistical dynamics. *Physical Review E* 54(4):3419.
- Fages, François, Guillaume Le Guludec, Olivier Bournez, and Amaury Pouly. 2017. Strong turing completeness of continuous chemical reaction networks and compilation of mixed analog-digital programs. In *Computational methods in systems biology: 15th international conference, cmsb 2017, darmstadt, germany, september 27–29, 2017, proceedings 15*, 108–127. Springer.
- Feinberg, Martin. 1988. Chemical reaction network structure and the stability of complex isothermal reactors—ii. multiple steady states for networks of deficiency one. *Chemical Engineering Science* 43(1):1–25.
- . 1989. Necessary and sufficient conditions for detailed balancing in mass action systems of arbitrary complexity. *Chemical Engineering Science* 44(9):1819–1827.
- . 2019. Foundations of chemical reaction network theory.

Feinberg, Martin, and Friedrich JM Horn. 1974. Dynamics of open chemical systems and the algebraic structure of the underlying reaction network. *Chemical Engineering Science* 29(3):775–787.

Gagrani, Praful. 2022. AFGD-for-CRN-escapes. URL <https://github.com/praful12/Descender-for-CRN-escapes>.

Gagrani, Praful, Victor Blanco, Eric Smith, and David Baum. 2023a. The geometry and combinatorics of an autocatalytic ecology in chemical and cluster chemical reaction networks. 2303.14238.

Gagrani, Praful, and Eric Smith. 2022. Action functional gradient descent algorithm for estimating escape paths in stochastic chemical reaction networks. *arXiv preprint arXiv:2210.15419*.

———. 2023. Action functional gradient descent algorithm for estimating escape paths in stochastic chemical reaction networks. *Physical Review E* 107(3):034305.

Gagrani, Praful, Vladimir Sotirov, David Baum, and Eric Smith. 2023b. Mathematics of evolving ecologies in chemical reaction networks and the origins of biochemical life (under preparation).

Gang, Hu. 1987. Stationary solution of master equations in the large-system-size limit. *Physical Review A* 36(12):5782.

Gardner, Jonathan P, John C Mather, Mark Clampin, Rene Doyon, Matthew A Greenhouse, Heidi B Hammel, John B Hutchings, Peter Jakobsen, Simon J Lilly, Knox S Long, et al. 2006. The james webb space telescope. *Space Science Reviews* 123:485–606.

Garey, Michael R, and David S Johnson. 1979. *Computers and intractability*, vol. 174. Freeman San Francisco.

Gillespie, Daniel T. 1992. A rigorous derivation of the chemical master equation. *Physica A: Statistical Mechanics and its Applications* 188(1-3):404–425.

Gillespie, Daniel T, et al. 2007. Stochastic simulation of chemical kinetics. *Annual review of physical chemistry* 58(1):35–55.

- Goldstein, Herbert, Charles Poole, and John Safko. 2002. Classical mechanics.
- Gopalkrishnan, Manoj. 2011. Catalysis in reaction networks. *Bulletin of mathematical biology* 73(12):2962–2982.
- Group, Particle Data. 1998. Review of particle physics. *The European Physical Journal C-Particles and Fields* 3(1-4):1–783.
- Hairer, Ernst, Marlis Hochbruck, Arieh Iserles, and Christian Lubich. 2006. Geometric numerical integration. *Oberwolfach Reports* 3(1):805–882.
- Hartmann, Carsten, and Christof Schütte. 2012. Efficient rare event simulation by optimal nonequilibrium forcing. *Journal of Statistical Mechanics: Theory and Experiment* 2012(11):P11004.
- Hatcher, Allen. 2005. *Algebraic topology*.
- Higgs, Paul G. 2021. When is a reaction network a metabolism? criteria for simple metabolisms that support growth and division of protocells. *Life* 11(9):966.
- Hordijk, Wim, and Mike Steel. 2004. Detecting autocatalytic, self-sustaining sets in chemical reaction systems. *Journal of theoretical biology* 227(4):451–461.
- . 2018. Autocatalytic networks at the basis of life's origin and organization. *Life* 8(4):62.
- Horn, Fritz, and Roy Jackson. 1972. General mass action kinetics. *Archive for rational mechanics and analysis* 47(2):81–116.
- Jaynes, Edwin T. 2003. *Probability theory: The logic of science*. Cambridge university press.
- Johnson, Kenneth A, and Roger S Goody. 2011. The original michaelis constant: translation of the 1913 michaelis–menten paper. *Biochemistry* 50(39):8264–8269.
- Kadanoff, Leo P. 2000. *Statistical physics: statics, dynamics and renormalization*. World Scientific.

Kamenev, Alex. 2011. *Field theory of non-equilibrium systems*. Cambridge University Press.

———. 2023. *Field theory of non-equilibrium systems*. Cambridge University Press.

Kannaiyan, Radhamani, and Daruka Mahadevan. 2018. A comprehensive review of protein kinase inhibitors for cancer therapy. *Expert review of anticancer therapy* 18(12):1249–1270.

Kauffman, Stuart, and Simon Levin. 1987. Towards a general theory of adaptive walks on rugged landscapes. *Journal of theoretical Biology* 128(1): 11–45.

Kelly, Frank P. 2011. *Reversibility and stochastic networks*. Cambridge University Press.

King, GAM. 1978. Autocatalysis. *Chemical Society Reviews* 7(2):297–316.

Kirsten, Klaus, and Alan J McKane. 2003. Functional determinants by contour integration methods. *Annals of Physics* 308(2):502–527.

Korniłowicz, Artur, and Karol Pał. 2017. Vieta's formula about the sum of roots of polynomials. *Formalized Mathematics* 25(2):87–92.

Krakauer, David, Nils Bertschinger, Eckehard Olbrich, Jessica C Flack, and Nihat Ay. 2020. The information theory of individuality. *Theory in Biosciences* 139:209–223.

Krishnamurthy, Supriya, and Eric Smith. 2017. Solving moment hierarchies for chemical reaction networks. *Journal of Physics A: Mathematical and Theoretical* 50(42):425002.

Langary, Damoun, and Zoran Nikoloski. 2019. Inference of chemical reaction networks based on concentration profiles using an optimization framework. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 29(11): 113121.

Lazarescu, Alexandre, Tommaso Cossetto, Gianmaria Falasco, and Massimiliano Esposito. 2019. Large deviations and dynamical phase transitions in stochastic chemical networks. *The Journal of Chemical Physics* 151(6): 064117.

- Leok, Melvin, and Jun Zhang. 2017. Connecting information geometry and geometric mechanics. *Entropy* 19(10):518.
- Li, Gene-Wei, and X Sunney Xie. 2011. Central dogma at the single-molecule level in living cells. *Nature* 475(7356):308–315.
- Liu, Yu, and David JT Sumpter. 2018. Mathematical modeling reveals spontaneous emergence of self-replication in chemical reaction systems. *Journal of Biological Chemistry* 293(49):18854–18863.
- Ma, W. 2009. Trusina a. el-samad h. lim wa, and tang c.:'. *Defining network topologies that can achieve biochemical adaptation*, *Cell* 138(4).
- Manikandan, Sreekanth K, and Supriya Krishnamurthy. 2017. Asymptotics of work distributions in a stochastically driven system. *The European Physical Journal B* 90(12):1–19.
- Manzano, Daniel. 2020. A short introduction to the lindblad master equation. *Aip Advances* 10(2):025106.
- Marshall, Stuart M, Cole Mathis, Emma Carrick, Graham Keenan, Geoffrey JT Cooper, Heather Graham, Matthew Craven, Piotr S Gromski, Douglas G Moore, Sara I Walker, et al. 2021. Identifying molecules as biosignatures with assembly theory and mass spectrometry. *Nature communications* 12(1):3033.
- McLachlan, RI, GRW Quispel, and PSP Tse. 2009. Linearization-preserving self-adjoint and symplectic integrators. *BIT Numerical Mathematics* 49(1):177–197.
- Milnor, John Willard, Michael Spivak, Robert Wells, and Robert Wells. 1963. *Morse theory*. 51, Princeton university press.
- Moyer, Devlin, Alan R Pacheco, David B Bernstein, and Daniel Segrè. 2020. Stoichiometric modeling of artificial string chemistries. *bioRxiv*.
- Muñuzuri, Alberto P, and Juan Pérez-Mercader. 2022. Unified representation of life's basic properties by a 3-species stochastic cubic autocatalytic reaction-diffusion system of equations. *Physics of Life Reviews*.
- Nielsen, Michael A, and Isaac Chuang. 2002. Quantum computation and quantum information.

Nowak, Martin A, and Hisashi Ohtsuki. 2008. Prevolutionary dynamics and the origin of evolution. *Proceedings of the National Academy of Sciences* 105(39):14924–14927.

Oono, Yoshitsugu. 2012. *The nonlinear world: Conceptual analysis and phenomenology*. Springer Science & Business Media.

Oppenheim, Alan V, Alan S Willsky, Syed Hamid Nawab, Gloria Mata Hernández, et al. 1997. *Signals & systems*. Pearson Educación.

Ostwald, Wilhelm. 1890. Über autokatalyse. *Berichte über die Verhandlungen der Königlich-Sächsischen Gesellschaft der Wissenschaften zu Leipzig, Mathematisch-Physische Klasse* 42:189–191.

Pathria, Raj Kumar. 2016. *Statistical mechanics*. Elsevier.

Peng, Zhen, Jeff Linderoth, and David A Baum. 2022a. The hierarchical organization of autocatalytic reaction networks and its relevance to the origin of life. *PLOS Computational Biology* 18(9):e1010498.

Peng, Zhen, Klaus Paschek, and Joana C Xavier. 2022b. What wilhelm ostwald meant by “autokatalyse” and its significance to origins-of-life research: Facilitating the search for chemical pathways underlying abiogenesis by reviving ostwald’s thought that reactants may also be autocatalysts. *BioEssays* 44(9):2200098.

Peng, Zhen, Alex M Plum, Praful Gagrani, and David A Baum. 2020. An ecological framework for the analysis of prebiotic chemical reaction networks. *Journal of theoretical biology* 507:110451.

Personal conversations with Philippe Nghe.

Peskin, Michael E. 2018. *An introduction to quantum field theory*. CRC press.

Pillai, S Unnikrishna, Torsten Suel, and Seunghun Cha. 2005. The perron-frobenius theorem: some of its applications. *IEEE Signal Processing Magazine* 22(2):62–75.

Plum, Alex M, and David A Baum. 2022. Aces in spaces: Autocatalytic chemical ecosystems in spatial settings. *arXiv preprint arXiv:2212.14445*.

Press, William H, Saul A Teukolsky, William T Vetterling, and Brian P Flannery. 2007. *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge university press.

Qian, Hong, and Lisa M Bishop. 2010. The chemical master equation approach to nonequilibrium steady-state of open biochemical systems: Linear single-molecule enzyme kinetics and nonlinear biochemical reaction networks. *International journal of molecular sciences* 11(9):3472–3500.

Rassoul-Agha, Firas, and Timo Seppäläinen. 2015. *A course on large deviations with an introduction to gibbs measures*, vol. 162. American Mathematical Soc.

Ruiz-Mirazo, Kepa, Carlos Briones, and Andrés de la Escosura. 2017. Chemical roots of biological evolution: the origins of life as a process of development of autonomous functional systems. *Open biology* 7(4): 170050.

Schlögl, Friedrich. 1972. Chemical reaction models for non-equilibrium phase transitions. *Zeitschrift für physik* 253(2):147–161.

Schneider, Thomas D. 2010. A brief review of molecular information theory. *Nano communication networks* 1(3):173–180.

Schuster, Peter. 2019. What is special about autocatalysis? *Monatshefte für Chemie-Chemical Monthly* 150:763–775.

Segré, Daniel, Doron Lancet*, Ora Kedem, and Yitzhak Pilpel. 1998. Graded autocatalysis replication domain (gard): kinetic analysis of self-replication in mutually catalytic sets. *Origins of Life and Evolution of the Biosphere* 28:501–514.

Sel’Kov, EE. 1968. Self-oscillations in glycolysis 1. a simple kinetic model. *European Journal of Biochemistry* 4(1):79–86.

Shubik, Martin, and Eric Smith. 2016. *The guidance of an enterprise economy*. Mit Press.

Smith, Eric. 2008. Thermodynamics of natural selection i: Energy flow and the limits on organization. *Journal of theoretical biology* 252(2):185–197.

———. 2011. Large-deviation principles, stochastic effective actions, path entropies, and the structure and meaning of thermodynamic descriptions. *Reports on Progress in Physics* 74(4):046601.

———. 2019. The information geometry of 2-field functional integrals. *arXiv preprint arXiv:1906.09312*.

———. 2020. Intrinsic and extrinsic thermodynamics for stochastic population processes with multi-level large-deviation structure. *Entropy* 22(10):1137.

———. 2021. Beyond fitness: selection and information flow through the constructive steps in lifecycles. *bioRxiv*.

Smith, Eric, and Supriya Krishnamurthy. 2015. *Symmetry and collective fluctuations in evolutionary games*. IOP Publishing.

———. 2017. Flows, scaling, and the control of moment hierarchies for stochastic chemical reaction networks. *Physical Review E* 96(6):062102.

———. 2021. Eikonal solutions for moment hierarchies of chemical reaction networks in the limits of large particle number. *Journal of Physics A: Mathematical and Theoretical* 54(18):185002.

Smith, Eric, and Harold J Morowitz. 2016. *The origin and nature of life on earth: the emergence of the fourth geosphere*. Cambridge University Press.

Smolin, Lee. 1996. Galactic disks as reaction-diffusion systems. *arXiv preprint astro-ph/9612033*.

Snarski, Michael. 2021. Hamilton-jacobi-bellman equations for chemical reaction networks. *arXiv preprint arXiv:2108.00514*.

Touchette, Hugo. 2005. Legendre-fenchel transforms in a nutshell. URL <http://www.maths.qmul.ac.uk/~ht/archive/lfth2.pdf>.

———. 2009. The large deviation approach to statistical mechanics. *Physics Reports* 478(1-3):1–69.

Turing, Alan Mathison. 1990. The chemical basis of morphogenesis. *Bulletin of mathematical biology* 52(1-2):153–197.

Varghese, Shaji, Johannes AAW Elemans, Alan E Rowan, and Roeland JM Nolte. 2015. Molecular computing: paths to chemical turing machines. *Chemical science* 6(11):6050–6058.

Veloz, Tomas, Pablo Razeto-Barry, Peter Dittrich, and Alejandro Fajardo. 2014. Reaction networks and evolutionary game theory. *Journal of mathematical biology* 68:181–206.

Watson, James D, and Francis HC Crick. 1953. The structure of dna. In *Cold spring harbor symposia on quantitative biology*, vol. 18, 123–131. Cold Spring Harbor Laboratory Press.

Weinan, E, Weiqing Ren, and Eric Vanden-Eijnden. 2004. Minimum action method for the study of rare events. *Communications on pure and applied mathematics* 57(5):637–656.

Wilhelm, Thomas. 2009. The smallest chemical reaction system with bistability. *BMC systems biology* 3:1–9.

Willamowski, K-D, and OE Rössler. 1980. Irregular oscillations in a realistic abstract quadratic mass action system. *Zeitschrift für Naturforschung A* 35(3):317–318.

Wright, Stephen, Jorge Nocedal, et al. 1999. Numerical optimization. *Springer Science* 35(67-68):7.

Yu, Polly Y, and Gheorghe Craciun. 2018. Mathematical analysis of chemical reaction systems. *Israel Journal of Chemistry* 58(6-7):733–741.