

On a thermodynamic approach to biomolecular interaction networks

Ricardo Honorato-Zimmer

Doctor of Philosophy
Laboratory for Foundations of Computer Science
School of Informatics
University of Edinburgh
2017

Abstract

We explore the direct and inverse problem of thermodynamics in the context of rule-based modelling. The direct problem can be concisely stated as obtaining a set of rewriting rules and their rates from the description of the energy landscape such that their asymptotic behaviour when $t \rightarrow \infty$ coincide. To tackle this problem, we describe an energy function as a finite set of connected patterns \mathcal{P} and an energy cost function ϵ which associates real values to each of these energy patterns. We use a finite set of reversible graph rewriting rules \mathcal{G} to define the qualitative dynamics by showing which transformations are possible. Given \mathcal{G} and \mathcal{P} , we construct a finite set of rules $\mathcal{G}_{\mathcal{P}}$ which i) has the same qualitative transition system as \mathcal{G} and ii) when equipped with rates according to ϵ , defines a continuous-time Markov chain that has detailed balance with respect to the invariant probability distribution determined by the energy function. The construction relies on a technique for rule refinement described in earlier work and allows us to represent thermodynamically consistent models of biochemical interaction networks in a concise manner.

The inverse problem, on the other hand, is to i) check whether a rule-based model has an energy function that describes its asymptotic behaviour and if so ii) obtain the energy function from the graph rewriting rules and their rates. Although this problem is known to be undecidable in the general case, we find two suitable subsets of Kappa, our rule-based modelling framework of choice, where this question can be answered positively and the form of their energy functions described analytically.

Acknowledgements

I would like to thank everyone that has helped me along the way. It has been a marvelous journey with all its highs and lows and has involved much learning and growth in my scientific as well as personal life. This experience has transformed me in many ways.

I am eternally grateful to my supervisor Vincent Danos. He took me as a PhD student, forming me as a scientist and as a better person. I feel an immense sense of gratitude when I think of the many things he patiently taught me during all these years. It is not an exaggeration to say that he taught me how to think. I am lucky to have had a genius supervisor.

Many thanks go to my parents, Renate Zimmer and Ricardo Honorato, for having given me the life and energy, the teachings and nourishment, the good ideas, beliefs and love that prepared me to succeed in this adventure. Without them none of this would have been possible. Importantly, I would like to thank as well my sister Daniela and brother Francisco, who, together with my parents, shaped my thinking since an early age, encouraged me in my scientific career, and from whom I have learned to question my own beliefs and understanding. I have still much to learn from them.

My deepest feelings of gratitude go to my friends Philipp Thomas, Sandro Stucki, Sebastian Jaramillo, Ilias Garnier and Tobias Heindel for sharing their time and knowledge with me since I met them. They have also contributed enormously to develop the precision and clarity of mind that was necessary in a multitude of topics of relevance to this thesis.

I would like to thank my friends Aníbal Ravani, Gráinne Foster, Sophie Currie, Stuart Burrows, Aileen Dickson, Susan Robinson, Alejandro Granados, Sebastian, Benura Azeroglu, Philipp, Claudia Cianci, Sandro, Nicolas Behr, Laura Kreiling, Patric Fulop, Andreea Beica, Ilias, Huisong Li and Guillaume Terradot for the various forms of unconditional support, companionship and encouragement that kept me positive and made possible the realisation of this manuscript. Their reassuring empathy in the good and bad moments of this journey and kind hospitality have been the emotional fuel of the writing process.

I feel a great sense of gratitude towards Mark Miller for the tremendous gift that his teachings and loving presence have been. He has changed the direction of my life and the impact that this change has had on the making of the final version of this manuscript cannot be overstated. I am also grateful to Alfred Hofmann and Alexander Shulgin for their wonderful inventions that provided me with the beautiful insights, at times clearly perceived and other times just subconsciously at work, that started this transition in a subtle and sometimes not so subtle way.

This list of acknowledgements would not be complete without a special mention to all the people in the group, namely and in no particular order, Milana Filatenkova, Andrea Weise, William Waites, Guoli Yang, Matteo Cavaliere, Fredrik Dahlqvist, Matthias Sachs, Katharina Heil, Emilia Wysocka, David Sterratt, Oksana Sorokina, John Wilson-Kanamori, Hristiana Pashkuleva, Argyris Zardilis, Simon Lyons, Gaëlle Candel, Nick, Patric, Andreea, Guillaume, Sandro, Ilias, Tobias and Vincent, for the stimulating conversations, interesting ideas and sincere desire to understand that created the excellent environment for which our meetings are famous. I would also like to thank Russ Harmer, Jean Krivine and Jérôme Feret for the camaraderie, good feedback and helpful advice.

I am grateful to everyone that helped me get to the PhD. In particular, special thanks go to Tomas Perez-Acle and Juan Carlos Letelier. Last and perhaps most importantly to the present version of this manuscript, I would like to thank Paweł Sobociński and Ilias whose invaluable feedback and dedicated reading of the first version of this manuscript catalysed the improvements that gave life to this final version.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(Ricardo Honorato-Zimmer)

Contents

1	Introduction	1
1.1	Historical background	1
1.2	Kappa	8
1.3	Related work	14
2	The direct problem: From energy to rules	17
2.1	Minimal glueings	19
2.2	Refinements	25
2.3	Thermodynamic growth policy	29
2.4	Rates and detailed balance	35
2.5	Linear kinetic model	36
2.6	Example: Triangles all the way down	38
2.7	Example: Flagellum's engine	43
2.8	Non-linear energy functions	53
3	The inverse problem: From rules to energy	56
3.1	Cooperative assembly systems	57
3.2	Flipping and binding	60
	Bibliography	65

Chapter 1

Introduction

1.1 Historical background

In the history of natural sciences, there has been two main approaches to describe dynamical systems, which I call here *kinetics* and *thermodynamics*. The first approach goes all the way back to Newton's laws of motion (Newton, 1687). Loosely speaking, it describes a system by the positions and momenta of each particle. A description of this type is as explicit and detailed as it can get for the type of systems under consideration in classical mechanics, i.e. the movement of point particles. In the case of a classical mechanical system that has N particles, a state of the system is specified by a vector in $\mathbb{R}^{2 \cdot 3 \cdot N}$. In general, a kinetic description is the full description of the dynamics of the system in terms of the velocities of its processes.

The thermodynamic approach, on the other hand, gives a description based on an *energy function*. The energy function is defined on the states of the system and assigns a real value to each state, its energy. That is, a state is described by a single scalar regardless of how many particles it comprises. Naturally, this approach endowed the description of a dynamical system in classical mechanics with a remarkable conciseness, simplicity and elegance. It first appeared in the work of Lagrange (1811) and Hamilton (1834), and has been subsequently used as the basis for most of modern physics. Once in possession of the energy function, the kinetic description (i.e. the equations of motion) can be derived from it. However the converse is not true: in general a kinetic description might not have an energy function from which it can be derived (Santilli, 1978) because of non-conservative (dissipative) forces. Obtaining an

energy function from the equations of motion is referred to as the *inverse problem* and it was first attended to by Helmholtz (1887). Both the direct and the inverse problem are the interest of this thesis and we aim to answer these questions in the context of biomolecular interaction networks.

Half a century after Hamilton's work researchers like Maxwell, Boltzmann, and Gibbs applied the ideas of classical mechanics to *atoms* in order to describe physical properties of matter like pressure, the capacity to transfer heat, and others. This body of work came to be known as *statistical mechanics* and was used to explain Brownian motion by Einstein (1905), which after its experimental verification (Perrin, 1908) settled the debate about the existence of atoms. It tried to explain, however unsuccessfully, the second law of thermodynamics and thus how irreversible processes arise from reversible ones. Perhaps it failed because the second law does not hold in general, e.g. in small systems and short time scales (G. M. Wang et al., 2002). A theorem in dynamical systems generalises the second law and can explain these results (Evans and Searles, 2002). Here, however, we concern ourselves with equilibrium thermodynamics and all processes considered are reversible.

Statistical mechanics had to be extended in order to explain the chemical interactions and reactions that molecules undergo. It was greatly helped by the axiomatisation of probability theory by Kolmogorov (1933) and the further developments by Doob (1937) and Feller (1940), who, among others, established the theoretical framework for continuous-time Markov chains (CTMCs). Below you can find the definition of (time-homogeneous) CTMCs and q -matrices that will be used here.

Definition 1. A (stable and conservative) q -matrix Q on an at most countable set of states \mathcal{S} is an $\mathcal{S} \times \mathcal{S}$ matrix with elements $q_{ij} \in \mathbb{R}$, $i, j \in \mathcal{S}$ such that $0 \leq q_{ij} < \infty$ when $i \neq j$ and $q_{ii} = -\sum_{j \neq i} q_{ij} > -\infty$.¹

The q -matrix plays the role of the time derivative of the transition probabilities at time 0 and induces the evolution of a probabilistic state according to the Kolmogorov backward equation,

$$\frac{d}{dt}P(t) = QP(t), \quad P(0) = I \quad (1.1)$$

where $P(t)$ is the $\mathcal{S} \times \mathcal{S}$ matrix with elements $p_{ij}(t) \in \mathbb{R}$ the probability that we were in state i at time 0 and are in state j at time t . When the q -matrix is

¹ If unstable q_{ii} can be $-\infty$ and if non-conservative $q_{ii} \leq -\sum_{j \neq i} q_{ij}$.

stable and conservative there exists a unique minimal² solution $P(t)$ of Eq. 1.1 (Anderson, 1991). We shall work with this type of q -matrices and assume there is a transition function $P(t)$ whenever we have a q -matrix Q and vice versa.

Given a probability distribution $\xi(0)$ on \mathcal{S} (seen as a row vector) the state at $t = 0$, the probability distribution $\xi(t)$ after time t is given by $\xi(t) = \xi(0)P(t)$. We say the q -matrix is *irreducible* if every state is reachable regardless of the initial state, i.e. $p_{ij}(t) > 0$ for all $i, j \in \mathcal{S}$ and some $t \geq 0$.

Definition 2 (CTMC). A continuous-time Markov chain is a tuple $(\mathcal{S}, \xi(0), Q)$ with \mathcal{S} an at most countable set of states, $\xi(0)$ a probability distribution on \mathcal{S} representing the initial state and Q the q -matrix of the Markov chain.

A few important properties of CTMCs for the present work are given below.

Definition 3 (detailed balance). A q -matrix Q on \mathcal{S} is said to be time reversible iff there is a probability distribution π on \mathcal{S} such that

$$\pi_i q_{ij} = \pi_j q_{ji} \quad (1.2)$$

for all $i, j \in \mathcal{S}$. Then Q is said to have detailed balance with respect to π .

Detailed balance was first proposed, in a slightly stronger form that requires every path going from i to j to have a reverse path with which it is in equilibrium, by Wegscheider (1901) in the context of chemical kinetics. Its validity for other physical systems was argued by Lewis (1925) and Tolman (1925). Tolman called the generalised principle *microscopic reversibility*.

Definition 4. A probability distribution π on \mathcal{S} is invariant for a q -matrix Q iff $\pi Q = 0$, i.e.

$$-\pi_i q_{ii} = \pi_i \sum_{j \neq i} q_{ij} = \sum_{j \neq i} \pi_j q_{ji}$$

That is to say, π is invariant whenever the action of Q on π does not change π or equivalently when π is a fixpoint of Q .

Lemma 1. Suppose the q -matrix Q has detailed balance with respect to π . Then π is invariant for Q .

² If $P'(t)$ is any non-negative solution of Eq. 1.1, then $p_{ij}(t) \leq p'_{ij}(t)$ for all $i, j \in \mathcal{S}$ and $t \geq 0$.

Proof. From Eq. 1.2 we obtain

$$\sum_{i \in \mathcal{S}} \pi_i q_{ij} = \sum_{i \in \mathcal{S}} \pi_j q_{ji} = -\pi_j q_{jj},$$

as $\sum_{i \in \mathcal{S}} q_{ji} = -q_{jj}$ for any fixed state j . ■

Once an invariant probability distribution is reached by the Markov chain it stays there forever. We would therefore like to know when an invariant probability distribution is realised by the Markov chain.

Definition 5 (ergodicity). *A q -matrix Q is ergodic when there is a probability distribution π on \mathcal{S} such that for all $i, j \in \mathcal{S}$,*

$$\lim_{t \rightarrow \infty} P_{ij}(t) = \pi_j$$

This is equivalent to say that the Markov chain will converge to the probability distribution π regardless of the initial state $\zeta(0)$.

Lemma 2. *Suppose the q -matrix Q is irreducible and has an invariant probability distribution π . Then Q is ergodic and converges to π .*

The proof for this lemma can be found in part 2 of theorem 1.6 in chapter 5 of Anderson's book (1991, pages 160–161).

CTMCs have a strong kinetic flavour as they describe stochastic processes in terms of probability flows happening at a certain rate. It is natural to wonder then how the thermodynamic approach looks like in the stochastic world. It turns out the energy function has a very clear interpretation in this setting, namely, that of defining the probability π_i that the system finds itself in state $i \in \mathcal{S}$ as follows (McQuarrie, 2000, page 40).³

$$\pi_i = \frac{e^{-E(i)}}{\sum_{j \in \mathcal{S}} e^{-E(j)}} \quad (1.3)$$

This is known as the *Boltzmann distribution*. Note that (i) when given the probability distribution π the energy function is defined uniquely only up to an additive constant;⁴ (ii) by convention the sign of the energy is inverted so lower energies represent more favourable states; and (iii) in the case of detailed balance, we obtain $e^{E(j)-E(i)} = q_{ji}/q_{ij}$ by combining Eq. 1.3 and Eq. 1.2.

³ We express the energy in units of $1/k_B T$ to avoid writing this term explicitly.

⁴ In other words, if we change the energy of each state by adding a fixed constant we obtain the same probability distribution π .

How do we construct a CTMC from an energy function? The first formulation to shed light on this problem was proposed by Metropolis et al. (1953). The algorithm asks for an energy function and an *a priori* one-step transition probability matrix A where each element a_{ij} ($i, j \in \mathcal{S}$) denotes the probability that we choose to jump to state j when we are at state i . Hence $\sum_{j \in \mathcal{S}} a_{ij} = 1$ for any fixed i and we write a_{i-} for this probability distribution. The A matrix is assumed to be symmetric, i.e. $a_{ij} = a_{ji}$ for all $i, j \in \mathcal{S}$, although this is not strictly necessary and the algorithm has been later generalised to work under a weaker assumption ($a_{ij} = 0$ iff $a_{ji} = 0$) by Hastings (1970).

Note that when addressing the direct problem for CTMCs by using the Metropolis algorithm we require an extra ingredient — the A matrix — which was not needed in classical mechanics. This is because in classical mechanics there are implicit assumptions of continuity on \mathcal{S} that supply this information. The state space is $\mathbb{R}^{2 \cdot 3}$ and, intuitively, an allowed transition in this continuous space is a differential change in any direction, i.e. dx, dy, dz . On the other hand, the method that will be presented in §2 does not ask for *a priori* transition probabilities but only which reversible transitions are possible at all.

The construction gives a *discrete-time* Markov chain that converges to the probability distribution π in Eq. 1.3. The algorithm works as follows. Given any state $i \in \mathcal{S}$ we pick a neighbour state j at random according to the probability distribution a_{i-} . We evaluate the energy function at i and j to compute $\Delta E = E(j) - E(i)$ and proceed with the transition with probability 1 if $\Delta E < 0$ and probability $e^{-\Delta E}$ if $\Delta E > 0$. Otherwise we stay at state i . In both cases time (a natural number) is increased by 1. We repeat for state j if the transition was successful and i otherwise.

To see that π , as defined in Eq. 1.3, is the invariant probability distribution of the discrete-time Markov chain we show that it has (the discrete-time version of) detailed balance with respect to π . The probability p_{ij} of jumping from i to j is a combination of the *a priori* probability a_{ij} and the probability of accepting that transition, which depends on ΔE .

$$p_{ij} = a_{ij} \min(1, e^{-\Delta E})$$

By taking the ratio of p_{ij} and p_{ji} we obtain

$$\frac{p_{ij}}{p_{ji}} = \frac{a_{ij} \min(1, e^{E(i)-E(j)})}{a_{ji} \min(1, e^{E(j)-E(i)})} = \frac{\min(1, e^{E(i)-E(j)})}{\min(1, e^{E(j)-E(i)})}$$

since $a_{ij} = a_{ji}$ by symmetry of A . Suppose $E(i) - E(j) > 0 > E(j) - E(i)$, then

$$\frac{p_{ij}}{p_{ji}} = e^{E(i)-E(j)} = \frac{e^{-E(j)}}{e^{-E(i)}} = \frac{\pi_j}{\pi_i}$$

It is easy to see that when $E(j) - E(i) > 0 > E(i) - E(j)$ we obtain the same equation. Hence the discrete-time Markov chain has detailed balance with respect to π . Provided the a priori transition probability matrix A makes it possible to reach any state from any other state, the Markov chain will converge to π as $t \rightarrow \infty$.

The Metropolis-Hastings algorithm can be generalised to the continuous-time case (Diaconis and Miclo, 2009). However, the algorithm require us to either (i) compute the energy of all states to obtain the probabilities p_{ij} (or transition rates q_{ij} in the continuous-time case), or (ii) do rejection sampling, as outlined above. Option (i) can be very time-consuming when \mathcal{S} is large or it's costly to evaluate the energy function. Option (ii) can be inefficient when the rejection rate is high. For these reasons we explore an alternative method in this thesis. We partition the state space in regions of equal energy and group transitions according to these regions. This is made possible by assuming extra structure on \mathcal{S} (to be introduced in §1.2).

Let us go back to the stochastic modelling of chemical interactions mentioned above. The theory of CTMCs allows one to frame the dynamics of chemical reaction systems. A stochastic approach to such systems was pioneered by Delbrück (1940) and has been common practice for decades (McQuarrie, 1967). The physical conditions under which this approach is plausibly valid has been argued by Gillespie (1976).

Since the number of molecules of a species is a priori unbounded and thus \mathcal{S} might be infinite, one would like to have a way to express these systems in a finite and simple form. A language that could do this was designed by Petri (1962). This language, later called *Petri nets*, sees reactions as transformations of multisets of chemical species.

Definition 6. A multiset M over a set X is a map from X to the naturals assigning to each element $x \in X$ the number of copies $M(x) \in \mathbb{N}$ of that element in the multiset.

There is a natural partial order \leq on multisets over X . We say $M \leq N$ if for each element $x \in X$, $M(x) \leq N(x)$. We write $\mathbf{M}(X)$ for the set of all multisets over X .

Definition 7. Given a set of species Σ , a reaction r is a pair (L, R) with L and R multisets over Σ . We refer to L and R as the left- and right-hand side of r .

Definition 8. A Petri net is a pair (Σ, \mathfrak{R}) of sets of species Σ and reactions \mathfrak{R} .

A state of a Petri net is a multiset over Σ , usually called a *marking*. A reaction can occur in a given state M only if its left-hand side $L \leq M$.

Definition 9. A match of the left-hand side L of a reaction on a state M is an injective function from L to M that identifies each copy of species s in L with a copy of s in M .

We write $[L; M]$ for the set of matches from L to M . From this definition we have that the number of matches $|[L; M]|$ from L to M is

$$|[L; M]| = \prod_{s \in \Sigma} \binom{L(s)}{M(s)}$$

A reaction is said to be elementary iff its rate is proportional to the number of matches of its left-hand side. This is known as the *law of mass action* in chemistry. Here we consider only elementary reactions.

Petri nets can be given a stochastic interpretation in terms of a CTMC. Given a Petri net (Σ, \mathfrak{R}) , an initial marking M_0 and a rate map $k : \mathfrak{R} \rightarrow \mathbb{R}_{\geq 0}$ that assigns rates to reactions, we construct a CTMC $(\mathcal{S}, \zeta(0), Q)$ as follows.

$$\begin{aligned} \mathcal{S} &= \mathbf{M}(\Sigma) \\ \zeta(0)(x) &= \begin{cases} 1 & \text{if } x = M_0 \\ 0 & \text{if } x \neq M_0 \end{cases} \\ q_{MN} &= \sum_{\substack{r \in \mathfrak{R} \\ r=(L,R)}} k(r) \gamma_{MN}(L, R) \end{aligned}$$

with

$$\gamma_{MN}(L, R) = \begin{cases} |[L; M]| & \text{if } M - L + R = N \\ 0 & \text{otherwise} \end{cases}$$

Danos and Oury (2013) have solved the *direct* and *inverse* problem for Petri nets, that is, they have shown the conditions the set of reactions and rate map have to fulfil for a Petri net to have an energy function and what is the structure of said energy function.

Petri nets have limitations when we take into consideration what happens inside molecules in a chemical reaction. The chemical transformation taking

place amounts to a change in the way electrons are shared by atoms resulting in a relocation of chemical bonds. In other words, (non-radioactive) reactions are all about the binding and unbinding of atoms, how they establish connections and break them. This is poorly captured by a conversion of species as it is modelled by Petri nets. A consequence of this lack of a formal representation for molecular bonds is that certain systems of chemical reactions cannot be described in a finite way using Petri nets, e.g. unbounded polymerisation (think of a molecular chain that can always attach new links).

Recently, a formal language to describe biochemical interactions using rewriting rules, where molecules not just react but can also bind other molecules has been proposed by Danos and Laneve (2004). In the next section we introduce this language, called Kappa. The language shall give us a formal foundation from which we can address the *direct* and *inverse* problems mentioned above, namely, the problem of generating a set of rewriting rules from an energy function and vice versa.

1.2 Kappa

Kappa represents interactions among proteins, nucleic acids and other biomolecules as connections in a biomolecular network. In these networks, nodes stand for the biomolecules while connections represent transient molecular bonds (e.g. non-covalent interactions like hydrogen bonds). This network is constantly changing as molecules travel and interact with other molecules in a cell, which is viewed as the constant destruction and creation of the connections that make up the network.

Due to spatial constraints, molecules can physically interact with just so many other molecules at once. Exactly how many will depend on multiple factors like the size of the two interacting molecules and the region where they come in contact. These regions, known in molecular biology by the names of domains, motifs or binding sites, are simply called *sites* in Kappa. Any such site can bind at most one other site at a time. These sites belong to the nodes of the graph, which Kappa calls *agents*. In the same way a molecule is of a certain species, agents can be of different types. These types also live in a network, a static «network of possibilities» which informs us of the set of sites an agent can have and the possible connections that sites can form.

To make these ideas formal we use a simpler version of the category-theoretical approach introduced by Danos, Harmer and Winskel (2013). We first define the networks where types live and then use them as a basis to construct the actual biomolecular networks. Graphs are used as a mathematical model of networks and site graphs in particular for biomolecular interaction networks.

Definition 10. A site graph G consists of a finite set of agents \mathcal{A}_G , a finite set of sites \mathcal{S}_G , a map $\sigma_G : \mathcal{S}_G \rightarrow \mathcal{A}_G$ that assigns sites to agents and a symmetric edge relation \mathcal{E}_G on \mathcal{S}_G .

The pair $\mathcal{S}_G, \mathcal{E}_G$ form an undirected graph. Clearly, the definition of site graphs does not impose a bound on the number of connections a site can have. Indeed there is no restriction at all so far. This is the network where types live.

Sites not in the domain of \mathcal{E}_G are said to be *free*. One says G is *realisable* iff (i) no site has an edge to itself and (ii) sites have at most one incident edge. Each realisable site graph represents a (possibly partially specified⁵) state in which our biomolecular network can be. However it contains no typing information. We give a type to each agent and site in the graph by assigning to it an agent and site in the type graph. More precisely, we use a map from a realisable site graph to a site graph. Below we introduce such maps.

Definition 11. A homomorphism $h : G \rightarrow G'$ of site graphs is a pair of functions, $h_S : \mathcal{S}_G \rightarrow \mathcal{S}_{G'}$ and $h_A : \mathcal{A}_G \rightarrow \mathcal{A}_{G'}$, such that for all $s, s' \in \mathcal{S}_G$ we have (i) $h_A(\sigma_G(s)) = \sigma_{G'}(h_S(s))$ and (ii) if $s \mathcal{E}_G s'$ then $h_S(s) \mathcal{E}_{G'} h_S(s')$.

$$\begin{array}{ccc} \mathcal{S}_G & \xrightarrow{h_S} & \mathcal{S}_{G'} \\ \sigma_G \downarrow & & \downarrow \sigma_{G'} \\ \mathcal{A}_G & \xrightarrow{h_A} & \mathcal{A}_{G'} \end{array}$$


Put simply, homomorphisms preserve site ownership and connections. The diagram to the right is the corresponding commutative diagram in the category **Set** of sets and total functions to condition (i) in the definition. We say the homomorphism $g : G \rightarrow C$ is a *contact map* over C iff (i) G is realisable and (ii) whenever $g_S(s_1) = g_S(s_2)$ and $\sigma_G(s_1) = \sigma_G(s_2)$, then $s_1 = s_2$. Condition (ii) means that every agent in G has at most one copy of each site of its corresponding agent in C . We refer to C as the contact graph. Contact maps act as the typing map mentioned above. In particular, C specifies the types of

⁵ Below you can find the definition of a fully specified state, which we call a mixture.

agents that can exist in G , the sites that they may possess, and which of the $|\mathcal{S}_C|(|\mathcal{S}_C| + 1)/2$ possible edge types are actually valid.

Site graphs and homomorphisms form a category \mathbf{SG} . The composition of two homomorphisms $h_1 : G_1 \rightarrow G_2$ and $h_2 : G_2 \rightarrow G_3$ is a homomorphism $h : G_1 \rightarrow G_3$ with $h_S = h_{2,S} \circ h_{1,S}$ and $h_A = h_{2,A} \circ h_{1,A}$. It is easy to see that composition defined in this way is associative. The identity arrow $1_G : G \rightarrow G$ in \mathbf{SG} is defined using the identity functions of the corresponding sets.


A homomorphism $\psi : G \rightarrow G'$ is an *embedding* iff (i) ψ_A and ψ_S are injective and (ii) if s is free in G , so is $\psi_S(s)$ in G' . Injectivity of ψ_A and ψ_S implies that whenever $\psi : G \rightarrow G'$ is an embedding and G' is realisable then G is also realisable. An embedding $\psi : G \rightarrow G'$ between realisable site graphs can be lifted to a morphism between contact maps $g : G \rightarrow C$ and $g' : G' \rightarrow C$ iff the diagram on the right commutes in \mathbf{SG} .




Contact maps over C and embeddings form a category \mathbf{rSGe}_C . Composition and identity are defined in a similar manner to \mathbf{SG} . We write $[g; g']$ for the set of embeddings between g and g' in \mathbf{rSGe}_C and refer to g as a *pattern* to be matched in g' . We have a functor $|\cdot|$ from \mathbf{rSGe}_C to \mathbf{SG} which forgets types. In particular, if $g : G \rightarrow C$ is a contact map, we write $|g|$ for its domain G .

As an example, consider the site graph T for a triangle.

$$\begin{aligned}\mathcal{A}_T &= \{1, 2, 3\}, \quad \mathcal{S}_T = \{l_1, r_1, l_2, r_2, l_3, r_3\}, \\ \sigma_T &= \{s_a \mapsto a \mid s \in \{l, r\}, a \in \mathcal{A}_T\}, \\ \mathcal{E}_T &= \{(r_1, l_2), (l_2, r_1), (r_2, l_3), (l_3, r_2), (r_3, l_1), (l_1, r_3)\}\end{aligned}$$



Let us use T as the contact graph for a contact map $g : G \rightarrow T$ where




$$\begin{aligned}\mathcal{A}_G &= \{x, y, z\} & \sigma_G &= \{s_a \mapsto a \mid s \in \{l, r\}, a \in \mathcal{A}_G\} \\ \mathcal{S}_G &= \{l_x, r_x, l_y, r_y, l_z, r_z\} & \mathcal{E}_G &= \{(r_x, l_y), (l_y, r_x), (r_y, l_z), (l_z, r_y)\}\end{aligned}$$

and

$$\begin{aligned}g_A &= \{x \mapsto 1, y \mapsto 2, z \mapsto 3\} \\ g_S &= \{s_a \mapsto s_{a'} \mid s \in \{l, r\}, a \in \mathcal{A}_G, a' = g_A(a)\}\end{aligned}$$

Sites l_x and r_z in G are free, which we denote graphically by a stub coming out of the site. T and G are realisable since no site is bound to itself or bound to more than one other site. Note however that the codomain of a contact map (T in this case) does not have to be realisable in general.

To ease the definition of concrete contact maps, we colour agents according to their type and annotate sites by their name in C . The contact map $g : G \rightarrow T$ above can then be defined more succinctly by the following drawing.




where we have assigned colours orange, blue and green to agent types 1, 2, 3 in C . We have written l and r for sites l_1, l_2, l_3 and r_1, r_2, r_3 as the subscript can be deduced from the colour of the agent as well.

Whenever a contact map $g : G \rightarrow C$ specifies all sites that its type C permits for all its agents, that is, if for all $a \in \mathcal{A}_G$, $g_S(\sigma_G^{-1}(a)) = \sigma_C^{-1}(g_A(a))$, then we say g is a *mixture*. We write $\mathbf{M}(C)$ for the set of all mixtures in \mathbf{rSGe}_C . In the above example, g is a mixture. What other mixtures are there that have T as contact graph? We can have chains of any length and closed cycles of length some multiple of three like triangles, hexagons, etc. We can have any disjoint sum of them as well.

Mixtures, being fully specified biomolecular networks with respect to the type C , are a natural choice for the states of our dynamical system. We jump from state to state by the applications of *rules*.

Definition 12. A rule r is a pair of contact maps $r_L : L \rightarrow C$, $r_R : R \rightarrow C$ which differ only in their edge structures, i.e. $\mathcal{A}_L = \mathcal{A}_R$, $\mathcal{S}_L = \mathcal{S}_R$, $\sigma_L = \sigma_R$, $r_{L,A} = r_{R,A}$ and $r_{L,S} = r_{R,S}$.

In the context of the contact graph T , we obtain a rule that binds agents of type 1 with agents of type 2 as follows.



$$\begin{aligned}
 \mathcal{A}_L &= \mathcal{A}_R = \{u, v\} \\
 \mathcal{S}_L &= \mathcal{S}_R = \{r_u, l_v\} \\
 \sigma_L &= \sigma_R = \{r_u \mapsto u, l_v \mapsto v\} \\
 r_{L,A} &= r_{R,A} = \{u \mapsto 1, v \mapsto 2\} \\
 r_{L,S} &= r_{R,S} = \{r_u \mapsto r_1, l_v \mapsto l_2\} \\
 \mathcal{E}_L &= \emptyset \\
 \mathcal{E}_R &= \{(r_u, l_v), (l_v, r_u)\}
 \end{aligned}$$

Note that there is no site l in u and no site r in v . Hence r_L and r_R are not mixtures as they are only partially specified. Intuitively, this means that the rule can be applied regardless of whether those sites are bound or free.

When a rule r is applied to an embedding $\psi : r_L \rightarrow m$ it induces a *rewrite* of the mixture m by modifying the edge structure of the image of ψ from that of r_L to that of r_R . The result of rewriting is a new mixture $m^{(r,\psi)}$ (or simply

$$\begin{array}{ccc} r_L & \xrightarrow{\quad} & r_R \\ \psi \downarrow & & \downarrow \psi^* \\ m & \xrightarrow{\quad} & m^{(r,\psi)} \end{array} \quad (1.4)$$

m^* when (r,ψ) is clear from the context) and an embedding $\psi^* : r_R \rightarrow m^*$, where $|m^*|$ has the same agents and sites as $|m|$, i.e. $\mathcal{A}_{|m^*|} = \mathcal{A}_{|m|}$, $\mathcal{S}_{|m^*|} = \mathcal{S}_{|m|}$, $\sigma_{|m^*|} = \sigma_{|m|}$, $m_{\mathcal{A}}^* = m_{\mathcal{A}}$, $m_{\mathcal{S}}^* = m_{\mathcal{S}}$, and $\mathcal{E}_{|m^*|} = \mathcal{E}_{|m|} - \psi_{\mathcal{S}}(\mathcal{E}_{|r_L|}) + \psi_{\mathcal{S}}^*(\mathcal{E}_{|r_R|})$. Since the set of agents and sites are equal, ψ^* is given by $\psi_{\mathcal{A}}^* = \psi_{\mathcal{A}}$ and $\psi_{\mathcal{S}}^* = \psi_{\mathcal{S}}$. The inverse of r , defined as $r^\dagger := (r_R, r_L)$ is also a valid rule. By applying r^\dagger to ψ^* we recover m and ψ .

Lemma 3. *Let $r = (r_L, r_R)$ be a rule, r_L/\mathbf{rSGe}_C the coslice category under r_L , and r_R/\mathbf{rSGe}_C the coslice category under r_R . The categories r_L/\mathbf{rSGe}_C and r_R/\mathbf{rSGe}_C are isomorphic.*

Proof. We construct a functor $F : r_L/\mathbf{rSGe}_C \rightarrow r_R/\mathbf{rSGe}_C$ by mapping an embedding $\psi : r_L \rightarrow m$ to the result of applying r to it, $\psi^* : r_R \rightarrow m^{(r,\psi)}$. By definition $\psi_{\mathcal{A}}^* = \psi_{\mathcal{A}}$ and $\psi_{\mathcal{S}}^* = \psi_{\mathcal{S}}$. Hence, the mapping of embeddings induced by F is injective: whenever the application of r to two embeddings $\psi : r_L \rightarrow m$ and $\phi : r_L \rightarrow m'$ results in ϕ^* and ψ^* with $\phi^* = \psi^*$, then $\phi = \psi$. By an analogous argument, we construct a functor $G : r_R/\mathbf{rSGe}_C \rightarrow r_L/\mathbf{rSGe}_C$ using r^\dagger that maps embeddings injectively in the reverse direction. Applying r followed by r^\dagger to ψ results in ψ itself. Therefore $GF = \mathbf{1}_{r_L/\mathbf{rSGe}_C}$ and $FG = \mathbf{1}_{r_R/\mathbf{rSGe}_C}$. ■

Intuitively, this property characterises a *reversible* rule. Reversibility will be important to obtain detailed balance as every rule r must be in balance with its inverse r^\dagger .

Given a finite set of rules \mathcal{R} over C , an initial mixture m_0 and a rate map k from \mathcal{R} to $\mathbb{R}_{\geq 0}$, we construct a CTMC $(\mathcal{S}, \xi(0), Q)$ with $\mathcal{S} = \mathbf{M}(C)$ and

$$\begin{aligned} \xi(0)(x) &= \delta_{x, m_0} \quad (\text{Kronecker delta}) \\ q_{mn} &= \sum_{\substack{r \in \mathcal{R} \\ r = (r_L, r_R)}} k(r) \left| \left\{ \psi \in [r_L; m] \mid m^{(r,\psi)} \simeq n \right\} \right| \end{aligned} \quad (1.5)$$

The number of embeddings between any two contact maps x, y is finite. Hence, $q_{mm} = -\sum_{n \neq m} q_{mn} > -\infty$ for any fixed $m \in \mathcal{S}$ and Q is a well-defined q -matrix. Here we use the law of mass action for the rate of rules, where the rate is proportional to the number of embeddings of their left-hand sides. In fact, the law of mass action simply amounts to say that rewrites induced by rule r are independent processes, each one occurring at rate $k(r)$. This is clear when we derive the CTMC from a *labelled transition system*.

Definition 13. A labelled transition system \mathcal{L} is a tuple $(\mathcal{S}, \Lambda, \rightarrow)$ with \mathcal{S} a set of states, Λ a set of labels, and $\rightarrow \subseteq \mathcal{S} \times \Lambda \times \mathcal{S}$ a set of transitions (x, α, y) between states $x, y \in \mathcal{S}$ labelled by $\alpha \in \Lambda$.

Given \mathcal{R} , we define a labelled transition system $\mathcal{L}_{\mathcal{R}}$ on mixtures over C where a transition from a mixture m is labelled by an *event* (r, ψ) , as in diagram 1.4, with r in \mathcal{R} and ψ in $[r_L; m]$. The CTMC can then be equivalently constructed from $\mathcal{L}_{\mathcal{R}}$ by assigning rate $k(r)$ to an event of the form (r, ψ) . We write $\mathcal{L}_{\mathcal{R}}^k$ for the CTMC. Note that the (strongly) connected components of $\mathcal{L}_{\mathcal{R}}$ are finite as agents cannot be destroyed nor created by rules and there can only be one edge between two sites.

With all the ingredients on the table we can proceed now to formulate more precisely the main question addressed in this thesis. In the context of the dynamics of biomolecular networks, the direct problem is stated as, given a contact graph C and an energy function on mixtures over C , how do we generate a finite set of rules \mathcal{R} with a corresponding rate map $k : \mathcal{R} \rightarrow \mathbb{R}_{\geq 0}$ such that the CTMC $\mathcal{L}_{\mathcal{R}}^k$ has detailed balance with respect to the probability distribution π as defined in Eq. 1.3? To add flexibility we ask as well for an initial set of (rate-less) rules from which to derive the set with detailed balance. We write \mathcal{G} for the set of generator rules that is given as input. This set delimits which moves are possible at all. The full method is presented in §2.

On the other hand, the inverse problem is stated as, given a contact graph C , a set of rules \mathcal{R} over C and a rate map k , does the CTMC $\mathcal{L}_{\mathcal{R}}^k$ have detailed balance? If so, how do we obtain its invariant probability distribution and energy function? The former question has been proven to be undecidable by Danos and Oury (2010) using an encoding of the Post correspondence problem (Post, 1946) in Kappa. In §3 we address the inverse problem for restricted versions of the Kappa language that are decidable.

1.3 Related work

Ollivier, Shahrezaei and Swain (2010) have developed a language to describe biomolecular interaction networks that have an energy function and a tool, called the *Allosteric Network Compiler* (ANC), to generate chemical reaction systems (i.e. Petri nets) from them. The biomolecular interaction networks they introduce have ANC structures as nodes. These structures can contain hierarchical components and interaction sites. The latter can be catalytic sites, covalently modified sites or ligand-binding sites. Hierarchical components on the other hand can contain any number of interaction sites and nested hierarchical components. If a hierarchical component is marked as allosteric, it transitions between two conformational states. The transition rates are modified by the state of the sites and other components present in the same structure according to parameters given for each of them. An additional parameter is required for covalently modified and ligand-binding sites which determines the change in the ratio between the two conformational states when a ligand is bound or the site is modified.

The edges of the ANC network are connections between ligand-binding sites or between a catalytic site and a covalently modified site. Rules can be of two types, binding or enzymatic, and are only allowed to depend on the conformational state of the two participants. Binding rules specify the association and dissociation of two ligand-binding sites. Enzymatic rules follow a Michaelis-Menten mechanism in which the enzymatic site reversibly binds a covalently modified site first and then changes the state of the covalently modified site as it unbinds it (with each step following the law of mass action). Each type of rule is parameterised differently. The energy function is obtained implicitly from all the parameters of the model. The language formalises concepts that are familiar to molecular biologists and biochemists in a way that reflects the apparent complexity of these interaction networks. However, the language's many concepts and classifications make it difficult to see the big picture and obfuscate the energy function. Also it arbitrarily restricts the form and size of the rules and the form of the energy function. The work presented in §2 is a generalisation of ANC. In §3.2 we will show a language similar to ANC for which the form of the energy function is uniquely determined from the kinetic rate parameters.


Another related development is that of biomolecular interaction databases. During the last two decades databases like BindingDB (Chen, Liu and Gilson, 2001; Gilson et al., 2016), BIND (Bader et al., 2001), MINT (Licata et al., 2012; Zanzoni et al., 2002), MatrixDB (Chautard et al., 2009) and BioLiP (Yang, Roy and Zhang, 2013) have appeared with the aim to collect data pertinent to biomolecular interactions that have been described in the scientific literature. Of particular interest to us is that they gather information about thermodynamic and kinetic parameters like equilibrium and rate constants for these interactions. The proliferation of molecular interaction databases have prompted the creation of the International Molecular Exchange (IMEx) consortium (Orchard et al., 2012) and a common data format, the Human Proteome Organization Proteomics Standards Initiative Molecular Interactions (HUPO-PSI-MI) format (Hermjakob et al., 2004; Kerrien et al., 2007). The latest version of this format, which is used by all molecular interaction databases, provides a means to specify thermodynamic and kinetic parameters of an interaction. Also of interest are databases specialised in the thermodynamic properties of reactions like the Thermodynamics of Enzyme-Catalyzed Reactions Database at the National Institute of Standards and Technology (NIST) (Goldberg, Tewari and Bhat, 2004) and others (Li et al., 2011). Moreover, sometimes it is possible to predict these thermodynamic constants and how they vary in different solutions (Flamholz et al., 2012; Mavrouniotis, 1991). Together all of these tools, databases and algorithms provide a strong infrastructure that facilitates the construction of thermodynamic models of biomolecular interaction networks.

Thermodynamic models have already been used successfully in many areas of biology. For instance, in chemotaxis⁶ thermodynamic models have been put forward to explain the positioning of the chemoreceptors on the membrane (H. Wang, Wingreen and Mukhopadhyay, 2008), their cooperative adaptation mechanisms to keep a high sensitivity for different ligands in different environments (Lan et al., 2011), and the activity of the engine that bacteria use to travel by turning the filaments of their flagellum (Bai et al., 2010). By recreating the latter model we show in §2.7 how to use our method to construct such thermodynamic models. Other examples include models in metabolism (Cannon, 2014; Thomas et al., 2014), macromolecular assembly

⁶ Chemotaxis is the process by which a living organism moves away of certain molecules like poisons and chases other molecules like food.

(Saiz and Vilar, 2006), transcription regulation (Bintu et al., 2005) and more (Kiselev, Marenduzzo and Goryachev, 2011).

Understanding the relationship between the kinetics and thermodynamics of biomolecular interactions may help understand the relationship between animate and inanimate matter. It has been argued by Pascal, Pross and Sutherland (2013), in an attempt to bridge the worlds of animate and inanimate matter, that the former corresponds to the set of *persisting molecular replicators*. A molecular replicator is a molecule or set of molecules that, usually in several steps, can fulfil the following transformation,



where R is the replicator, F is a molecule consumed by R to construct a second R (mnemonically named F for food) and W (for waste) is what is left of F that was not used for the replication. This transformation is known as *autocatalysis*. Under unlimited resources a molecular replicator grows exponentially fast. On the other hand, when F is exhausted the concentration of R will converge to that dictated by the thermodynamic equilibrium. Persistent replicators are those that manage to keep themselves in a far-from-equilibrium regime and thus continue replicating to maintain its population. It is my belief that this type of processes can be investigated by introducing additional rules to the set of thermodynamically-consistent rules generated by our method.⁷ These sorts of modifications opens up a possible new line of work which we might call far-from-equilibrium graph thermodynamics. The convergence properties of thermodynamically-consistent sets of rules presented in §2.3 would not hold in a far-from-equilibrium regime but might nevertheless serve as a reference for comparison. Moreover, the rates calculated for the rules by the method would still be valid as these depend only on the chemical properties of molecules in solution (reactants, products, solvent, etc). In §2.5 a framework is proposed to systematically assign rates to rules based on some of these chemical considerations. It allows the exploration of kinetics in a thermodynamically consistent way. Hence this framework might play an important role in the study of far-from-equilibrium systems.

⁷ Note that when the additional irreversible rules do not intersect the reversible ones, one still can get an extended notion of detailed balance by an analogous argument to that put forward by Gorban, Mirkes and Yablonsky (2013).

Chapter 2

The direct problem

From energy to rules

In this chapter we show how to construct a set of reversible rules and their forward and backward rate constants from an energy function. In the spirit of rule-based modelling languages like Kappa where rules and observables are defined in terms of patterns,¹ we use a set of connected *energy patterns* \mathcal{P} for our energy function. We assign an *energy cost* $\epsilon(g)$ to each of them and build the energy function as a linear combination of their number of occurrences.

$$E(m) = \sum_{g \in \mathcal{P}} \epsilon(g) |[g; m]| \quad (2.1)$$

This is reminiscent of group contribution methods used to estimate the standard Gibbs free energy of formation of biomolecules (Mavrovouniotis, 1991).

As mentioned at the end of §1.2, we will derive the set of rules with detailed balance from a set of generator rules \mathcal{G} (without rates). We suppose that \mathcal{G} is closed under rule inversion, i.e. $\mathcal{G} = \mathcal{G}^\dagger$. Given a contact graph C , a simple option would be to include every possible minimal rule in this set, that is, include a creation and a destruction rule for each edge in the contact graph. Each of these rules is minimal in the sense that it only asks for the presence of the two participating agents and sites. The example rule in §1.2 (page 11) where agents of type 1 and 2 bind regardless of the context is one such minimal rule that can be derived from the contact graph T . We call this rule r_{12}^+ .

$$r_{12}^+ := \begin{array}{c} \text{[Diagram: Two separate nodes, one orange with site 'r' and one blue with site 'l']} \end{array} \longrightarrow \begin{array}{c} \text{[Diagram: Two nodes connected by an edge, one orange with site 'r' and one blue with site 'l']} \end{array} \quad (2.2)$$

¹ Recall that a pattern is a contact map used to find subgraphs in states.

This option is *maximally permissive* with respect to the contact graph.² Even if all transformations are possible, many of them may be unlikely due to having a high energy. Still one might prefer to forbid certain transformations in some scenarios. This is indeed the case in the example that will be presented in §2.7.

In our previous example (§1.2), we might want to favour the formation of triangles over chains and other cycles. For this we give a negative energy cost to the triangle t , i.e. $\epsilon(t) < 0$. If t is the only energy pattern, then the energy of a state m is $E(m) = \epsilon(t) |[t; m]|$. In this model one might, for instance, wonder how low the energy cost of t must be to have at least 90% of all agents in a triangle at equilibrium at least 90% of the time.

We would like to find rules that have detailed balance with respect to this energy function. Consider the rule r_{12}^+ and its inverse r_{12}^- , the unbinding of agents 1 and 2. We first ask ourselves if this pair of rules could have detailed balance for some assignment of kinetic rates. Suppose we assign kinetic rates k^+ and k^- to r_{12}^+ and r_{12}^- . Recall from §1.1 that $e^{E(n)-E(m)} = q_{nm}/q_{mn}$ for systems with detailed balance. From Eq. 1.5


$$q_{mn} = \sum_{\substack{r \in \mathcal{G} \\ r=(r_L, r_R)}} k(r) \left| \left\{ \psi \in [r_L; m] \mid m^{(r, \psi)} = n \right\} \right|$$

where $m^{(r, \psi)}$ is the outcome of rewriting m with event (r, ψ) . At most one of the two rules can bring us from state m to n , say it is r_{12}^+ . By rule reversibility (Lemma 3) r_{12}^- brings us from n back to m and the number of matches of r_{12}^- in n is equal to the number of matches of r_{12}^+ in m . Hence, $e^{E(n)-E(m)} = k^+/k^-$. In words, the change in energy produced by the rule application fixes the ratio between the kinetic rates. As a consequence, each rule application should produce the same energy change for there to be an assignment of kinetic rates with detailed balance. Whenever a rule produces the same energy change regardless of where it is applied we say that the rule has an *unambiguous energy balance* or is \mathcal{P} -balanced. More generally, we define \mathcal{P} -balance as follows.

Definition 14. Given a contact graph C and a set \mathcal{P} of contact maps over C , a rule r is \mathcal{P} -balanced if, for all mixtures m and embeddings $\psi : r_L \rightarrow m$, the number of occurrences of $p \in \mathcal{P}$ produced and consumed by r when applied to ψ is a fixed number $\Delta_r p = |[p; m^{(r, \psi)}]| - |[p; m]|$. We refer to $\Delta_r p$ as the balance of r with respect to p .

² Intuitively, this is analogous to the case of classical mechanics where, a priori, movement is not constrained along any coordinate.

The following two rule applications show that r_{12}^+ is not \mathcal{P} -balanced.




We see that, while the application on the left does not produce any change in energy ($\Delta E = 0$), the one on the right creates a triangle and thus $\Delta E = \epsilon(t)$. We must then split r_{12}^+ into subrules that check the surroundings of the rule application to make sure that, for instance, every application of such a subrule creates one triangle or none at all. It is important that the partition of the rule has certain properties. In particular, one would like that every match of the rule can be mapped to exactly one match of one of the subrules. Prior work by Murphy et al. (2010) has shown how one can obtain a partition of rules with this property and will be presented, in a slightly modified version, in §2.2.

But before diving into rule partitioning, or rule refinement as we call it, it would be good to have a more rigorous idea of when a rule is \mathcal{P} -balanced or not. In the examples shown above we see that our energy pattern, the triangle, must be fully incorporated into the left- or the right-hand side of the rule to be sure it produces or consumes it in every application. On the other hand, a rule that is incompatible with our energy pattern will also be \mathcal{P} -balanced by making it impossible for the rule to match a triangle. This is true whenever there is no glueing of the left-hand side of a rule with the energy pattern where they overlap in a site that is modified by the rule. In the next section, we introduce the concept of overlapping glueings of contact maps by means of multi-sums, a concept related to local coproducts and relative pushouts.

2.1 Minimal glueings

The category \mathbf{SG} has all pullbacks, constructed from those in \mathbf{Set} , and they indeed restrict to \mathbf{rSGe}_C .

Lemma 4. *Given a cospan $\phi_1 : g_1 \rightarrow h \leftarrow g_2 : \phi_2$ in \mathbf{rSGe}_C there is a unique span $\psi_1 : g_1 \leftarrow p \rightarrow g_2 : \psi_2$ (up to unique isomorphism) such that any span $\omega_1 : g_1 \leftarrow q \rightarrow g_2 : \omega_2$ that forms a commuting square $\omega_1, \omega_2, \phi_1, \phi_2$ factors uniquely through it.*



Proof. We construct contact map $p : G \rightarrow C$ by taking the intersection of the agents, sites and edges in the image of ϕ_1, ϕ_2 and restricting σ accordingly. With some abuse of notation, we have


$$\begin{aligned}\mathcal{A}_G &= \phi_{1,\mathcal{A}}(\mathcal{A}_{|g_1|}) \cap \phi_{2,\mathcal{A}}(\mathcal{A}_{|g_2|}) \\ \mathcal{S}_G &= \phi_{1,\mathcal{S}}(\mathcal{S}_{|g_1|}) \cap \phi_{2,\mathcal{S}}(\mathcal{S}_{|g_2|}) \\ \mathcal{E}_G &= \phi_{1,\mathcal{E}}(\mathcal{E}_{|g_1|}) \cap \phi_{2,\mathcal{E}}(\mathcal{E}_{|g_2|})\end{aligned}$$

and $\sigma_G = \sigma|_h|_{\mathcal{S}_G}$. Functions $p_{\mathcal{A}}, p_{\mathcal{S}}$ are the restriction of $h_{\mathcal{A}}, h_{\mathcal{S}}$ to $\mathcal{A}_G, \mathcal{S}_G$, respectively. Embeddings ψ_1 and ψ_2 map agents and sites in G to their pre-images along ϕ_1 and ϕ_2 ; by construction, all agents and sites in G are guaranteed to have such a pre-image. It is easy to see that (i) ψ_1 and ψ_2 are type-preserving and thus embeddings in \mathbf{rSGe}_C ; and that (ii) the square formed by $\psi_1, \psi_2, \phi_1, \phi_2$ commutes.

Consider any span $\omega_1 : g_1 \leftarrow q \rightarrow g_2 : \omega_2$ in \mathbf{rSGe}_C . If the square formed by $\omega_1, \omega_2, \phi_1, \phi_2$ commutes, then q can have at most one copy of each agent and site in the intersection of the images of ϕ_1 and ϕ_2 because $\phi_1 \omega_1$ and $\phi_2 \omega_2$ are injective. Hence, every agent and site in the image of ω_1, ω_2 has a *unique* pre-image along ψ_1, ψ_2 , respectively, with the same type. This fixes a pair of functions $\omega_{\mathcal{A}}, \omega_{\mathcal{S}}$ that map agents and sites in q to those in p injectively and form an embedding ω in \mathbf{rSGe}_C . Since the pre-image along ψ_1, ψ_2 always exists and is unique, any embedding $\omega' : p \rightarrow q$ must be equal to ω whenever $\phi_1 \omega' = \omega_1$ and $\phi_2 \omega' = \omega_2$. ■

SG also has all pushouts and all sums, but these do not in general restrict to \mathbf{rSGe}_C , just as pushouts and sums in **Set** do not restrict to the subcategory of injective functions. However, \mathbf{rSGe}_C has *multi-sums*.

Lemma 5. *For all pairs of contact maps over C , $g_1 : G_1 \rightarrow C$ and $g_2 : G_2 \rightarrow C$, there exists a finite family of cospans $\theta_1^i : g_1 \rightarrow s_i \leftarrow g_2 : \theta_2^i$, such that any cospan $\phi_1 : g_1 \rightarrow h \leftarrow g_2 : \phi_2$ factors through exactly one of the family and does so uniquely.*



Proof. Take subsets A_i of the cartesian product of $\mathcal{A}_{|g_1|}$ and $\mathcal{A}_{|g_2|}$ that have each agent of g_1 and g_2 at most once ($((a,b) \in A_i \wedge (a,b') \in A_i \Rightarrow b = b')$) and where each pair $(a,b) \in A_i$ has the same type, i.e. $g_{1,\mathcal{A}}(a) = g_{2,\mathcal{A}}(b)$. To each A_i assign all subsets S_{ij} of $\mathcal{S}_{|g_1|} \times \mathcal{S}_{|g_2|}$ that are type-compatible and whose elements belong to agents paired in A_i , that is, if $(x,y) \in S_{ij}$ then $g_{1,\mathcal{S}}(x) = g_{2,\mathcal{S}}(y)$ and $(\sigma_{|g_1|}(x), \sigma_{|g_2|}(y)) \in A_i$. Note how this fixes a mapping σ_{ij} between elements of S_{ij} to elements of A_i defined by $\sigma_{ij}((x,y)) = (\sigma_{|g_1|}(x), \sigma_{|g_2|}(y))$. For each A_i keep only the set S_{ij} that is a superset of all other sets S_{ik} ($k \neq j$). There must be one such maximal set because if any two pairs of sites $(x_1, y_1), (x_2, y_2)$ are type-preserving and belong to the same agents, then there will be one set among the S_{ij} s that has both and thus $\{S_{ij}\}_j$ is a directed partial order for the inclusion relation. Let S_i be the maximal element of $\{S_{ij}\}_j$, which exists by directedness and finiteness of this family, and σ_i the corresponding mapping to A_i . Intuitively, the maximal set S_i is the set of all sites that are defined in both agents at the same time. Next we discard those pairs A_i, S_i whose elements do not agree on their edge structure; if $(x,y) \in S_i$ then either both sites must be free or connected to sites $(x', y') \in S_i$.

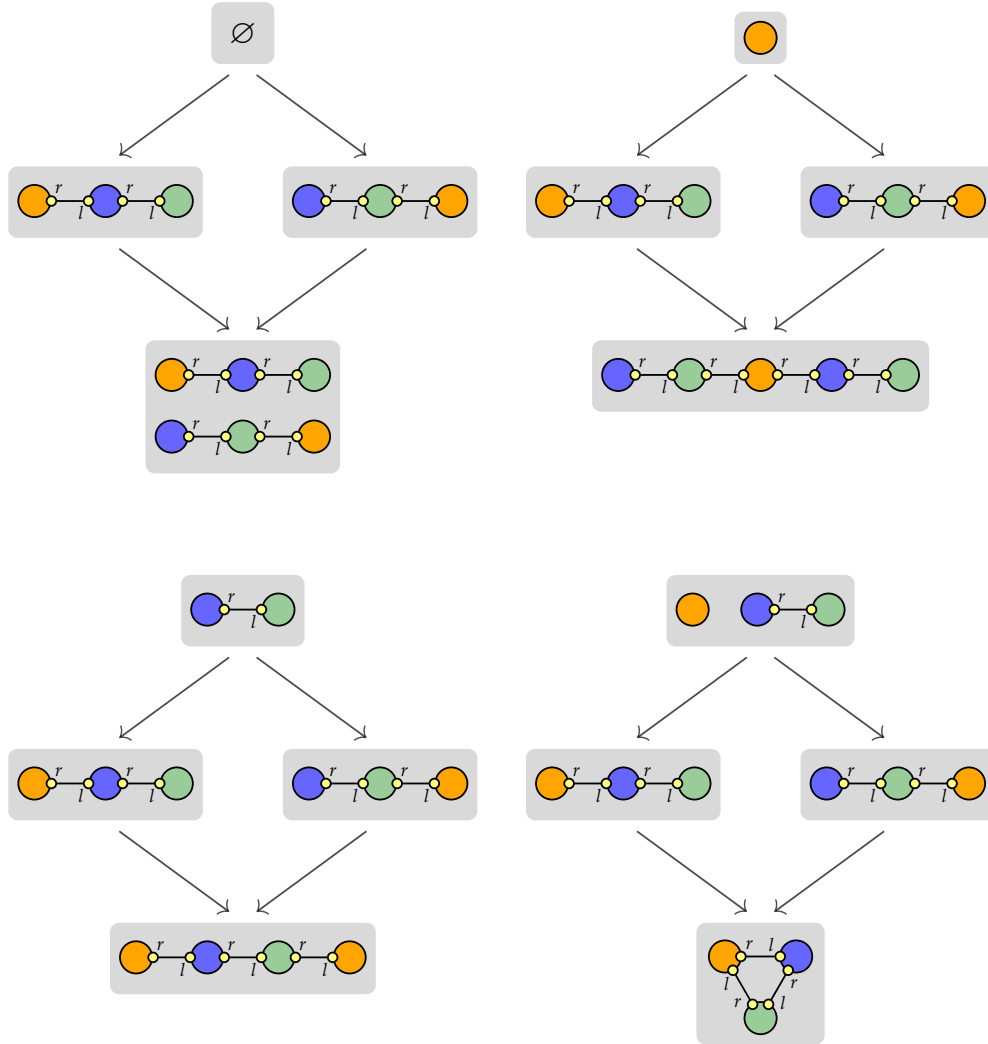
We construct a family of contact maps $p_i : P_i \rightarrow C$ using $\mathcal{A}_{P_i} = A_i$ as its agents, $\mathcal{S}_{P_i} = S_i$ as its sites, $\sigma_{P_i} = \sigma_i$ and $\mathcal{E}_{P_i} = \{((x_1, y_1), (x_2, y_2)) \in S_i \times S_i \mid x_1 \mathcal{E}_{|g_1|} x_2 \wedge y_1 \mathcal{E}_{|g_2|} y_2\}$. Functions $p_{i,\mathcal{A}}, p_{i,\mathcal{S}}$ are defined straightforwardly. Spans $\psi_1^i : g_1 \leftarrow p_i \rightarrow g_2 : \psi_2^i$ are then obtained by mapping agents (a,b) in p_i to a in g_1 and b in g_2 and similarly for sites. Multi-sums $\theta_1^i : g_1 \rightarrow s_i \leftarrow g_2 : \theta_2^i$ are pushouts of such spans: they are obtained by adding to p_i all the missing agents, sites and edges from g_1 and g_2 . Since all sites that are in g_1 but not in p_i cannot be in g_2 by maximality of S_i , there can be no conflict when adding sites or edges. The same argument holds for sites in g_2 that are not in p_i .

Note that the family A_i is finite and thus the family of multi-sums is finite as well. Also, it is easy to see that the spans ψ_1^i, ψ_2^i are pullbacks of θ_1^i, θ_2^i . Hence, (isomorphism classes of) multi-sums are in a one-to-one correspondence with (isomorphism classes of) pullbacks. This implies that there is only one multi-

sum that factors any given cospan. ■

The pairs θ_1^i, θ_2^i enumerate all minimal ways in which one can glue g_1 and g_2 . Hence, we refer to them as minimal glueings. The notion of multi-sum dates back to Diers (1978). They are very close to relative pushouts (Leifer and Milner, 2000) and will be used in the same way, to minimise rewriting contexts. Indeed, each minimal glueing i in the family of cospans θ_1^i, θ_2^i accounts for one minimal rewriting context.

To illustrate how this construction operates, consider the minimal glueings of the following two contact maps over T with their respective pullbacks.




I have implemented an online tool that computes minimal glueings available at <https://rhz.github.com/thesis/mg.html>. Its source code can be found at <https://github.com/rhz/thesis/>.

Using minimal glueings we can test whether a rule r is \mathcal{P} -balanced, that is, whether r consumes and produces the same number of instances of each


energy pattern p when applied to any mixture m . In particular, for an r -event ψ to *consume* an instance ϕ of p in a mixture m , ϕ_S and ψ_S must have images which intersect on at least one site which is modified by r (e.g. by adding an edge if it was free). This is the case iff the minimal glueing ϕ', ψ' of r_L and p that factors the cospan ϕ, ψ has the same property. Likewise, for an r -event to *produce* an instance of p , the associated minimal glueing between p and r_R must have a modified intersection. We call such minimal glueings *relevant*.


$$\begin{array}{ccccc}
 p & \xrightarrow{\phi'} & s & \xleftarrow{\psi'} & r_L \\
 & \searrow \phi & & \nearrow \psi & \\
 & & m & &
 \end{array} \quad (2.3)$$

To illustrate the idea of relevant minimal glueings, let us consider a different example. In this example, the contact graph is very simple: just one agent type with two sites, l and r , that can bind each other. Imagine we have the following reversible rule.



Take the chain of 3 agents as our energy pattern. The minimal glueings of the left-hand side of the rule with the energy pattern are shown below. On the left of each diagram is the energy pattern. The relevant minimal glueings are marked with a light green background.





Whenever $\psi' : r_L \rightarrow s$ in diagram 2.3 is an iso, then the energy pattern p is fully included in the left-hand side of rule r . This implies the rule contains all the relevant context needed to make sure that an instance of p is consumed by any r -event $\psi : r_L \rightarrow m$. We say that r is \mathcal{P} -left-balanced iff, for all $p \in \mathcal{P}$ and relevant minimal glueings $\theta_1^i : p \rightarrow s_i \leftarrow r_L : \theta_2^i$, the right leg θ_2^i is an isomorphism. Symmetrically, one says that r is \mathcal{P} -right-balanced iff r^\dagger is \mathcal{P} -left-balanced. Then r is \mathcal{P} -balanced iff it is \mathcal{P} -left- and \mathcal{P} -right-balanced.

Lemma 6. *Rule r is \mathcal{P} -balanced if and only if r is \mathcal{P} -left- and \mathcal{P} -right-balanced. Moreover, if r is \mathcal{P} -balanced then, for any mixture m , embedding $\psi : r_L \rightarrow m$, and energy pattern $p \in \mathcal{P}$,*

$$\Delta_r p = |[p; m^{(r, \psi)}]| - |[p; m]| = |[p; r_R]| - |[p; r_L]|$$

Proof. Suppose there are two mixtures m, n and embeddings $\psi : r_L \rightarrow m, \phi : r_L \rightarrow n$ such that, when r is applied to ψ and ϕ , it has a different balance with respect to a pattern $p \in \mathcal{P}$, i.e. $|[p; m^{(r, \psi)}]| - |[p; m]| \neq |[p; n^{(r, \phi)}]| - |[p; n]|$. We have

$$|[p; m]| = |\{p \rightarrow m \xleftarrow{\psi} r_L\}| = \left| \left\{ \begin{array}{ccc} p & \longrightarrow & s \xleftarrow{\psi} r_L \\ & \searrow & \uparrow \\ & m & \end{array} \right\} \right|$$

where $p \rightarrow s \xleftarrow{\psi} r_L$ is the minimal glueing that factors the cospan $p \rightarrow m \xleftarrow{\psi} r_L$. A similar equality can be obtained for $r_R, m^{(r, \psi)}$ and ψ^* . The *irrelevant* minimal glueings on each side of the rule are in bijection: the rule does not destroy


nor create them. Hence, when taking the difference $||[p; m^{(r, \psi)}]| - |[p; m]|$ they cancel each other out and we are left with a difference of *relevant* minimal glueings on each side. Since $s \simeq r_L$ for each relevant minimal glueing on the left then

$$\left| \left\{ \begin{array}{c} p \longrightarrow s \simeq r_L \\ \searrow \quad \vdots \quad \swarrow \\ \quad m \quad \psi \end{array} \right\} \right| = |[p; r_L]|$$


Again, a similar equality can be obtained for r_R , $m^{(r, \psi)}$ and ψ^* . Thus we have proved that $||[p; m^{(r, \psi)}]| - |[p; m]| = |[p; r_R]| - |[p; r_L]|$ for any m and ψ , contradicting our original assumption. ■

2.2 Refinements


A rule is refined into another rule by adding context. For example, we can add a common neighbour to the agents in r_{12}^+ to obtain a refinement.



This refinement happens to be \mathcal{P} -balanced. Another refinement of r_{12}^+ could be




Here we have added a free site to the blue node. This second refinement is also \mathcal{P} -balanced because the free r site on the blue node guarantees that (i) the rule will never create a triangle and (ii) there is no embedding from the left-hand side into a triangle and hence no triangle can be destroyed by the action of the rule. The following refinement, however, is not \mathcal{P} -balanced.




We add context to a rule $r = (r_L, r_R)$ by applying the rule to an embedding $\psi : r_L \rightarrow g$. This operation is well-defined even if the codomain of the embedding is not a mixture. The pair of contact maps $(g, g^{(r, \psi)})$ is itself a valid rule since they only differ in their edge structure. In this way, an extension of a rule is determined uniquely by an embedding.

Epis³ of \mathbf{rSGe}_C are good candidates for extensions. They are characterised as follows: an embedding $\psi : g \rightarrow h$ is an epi iff every connected component of $|h|$ contains at least one agent in the image of ψ_A . This ensures that no new connected component is added to the rule while extending it. However, for technical reasons that will become apparent in Th. 8, we use prefixes of epis instead of epis to extend rules — an embedding $\psi : g \rightarrow h$ is said to be a *prefix* of $\phi : g \rightarrow h'$ if there is some embedding $\theta : h \rightarrow h'$ that makes the composition of ψ and θ equal to ϕ (i.e. $\theta \psi = \phi$) and write $\psi \leq \phi$ for this. We refer to a prefix of an epi $\psi : g \rightarrow h$ as an *extension* of g . In the category of extensions of g , a morphism between objects $\psi : g \rightarrow h$ and $\phi : g \rightarrow h'$ is an embedding $\theta : h \rightarrow h'$ such that the triangle on the right commutes. If θ is an iso we write $\psi \cong_g \phi$.



One might wonder when the prefix of an epi is not itself an epi. The following diagram illustrates such a situation, where ψ is a prefix of epi ϕ but is not itself an epi since the connected component of the blue node in the codomain of ψ is not in the image of ψ_A .



Rule application preserves epis and in fact also prefixes of epis:

Lemma 7. *Let $r = (r_L, r_R)$ be a rule and $\psi : r_L \rightarrow g$ be an embedding with r_L, r_R, g contact maps in \mathbf{rSGe}_C . The embedding $\psi^* : r_R \rightarrow g^*$ that results from applying r to ψ is a prefix of an epi iff ψ is.*

Proof. This amounts to proving that some embedding $\phi^* \geq \psi^*$ is an epi if there is an epi $\phi \geq \psi$; the converse is true by symmetry of rules. For this it is enough to consider the case where the rule adds or deletes exactly one edge since rules that modify more than one edge at a time can be decomposed as sequences of deletions and insertions of edges; given that each deletion and insertion preserves the property, the sequence will preserve it as well.

³ Epi, mono and iso are short for epimorphism, monomorphism and isomorphism.

The case of adding an edge is easy as the image of ϕ^* has fewer connected components to intersect than ϕ . The case where r deletes an edge can introduce new connected components, however in this case both ends u, v of the deleted edge must be in r_L , so whether the deletion disconnects or not the codomain of ψ , the components of $\phi^*(u)$ and $\phi^*(v)$ will have a pre-image, namely u and v . ■

It follows that the category of extensions of r_L and r_R are isomorphic. Hence, any extension ϕ to a rule r can be mapped to an extension of its inverse rule r^+ .

A family of epis $\phi_i : g \rightarrow g_i$ *uniquely decomposes* g , or is a *refinement* of g , if, for all mixtures m and embeddings $\psi : g \rightarrow m$, there exists a unique i and ψ' such that $\psi = \psi' \phi_i$. This is the basic requirement for a reasonable notion of rule refinement: it guarantees that the left-hand side g of a given rule splits into a non-overlapping and exhaustive collection of more specific cases g_i .

A method to easily construct such decompositions was proposed by Murphy et al. (2010) which works by detailing which agents and sites should be added to g . This «extension plan» is called growth policy. A *growth policy* Γ for contact map g over C is a family of functions Γ_ϕ , indexed by all extensions $\phi : g \rightarrow h$, where Γ_ϕ maps $u \in \mathcal{A}_{|h|}$ to a subset $\Gamma_\phi(u)$ of $\sigma_C^{-1}(h_{\mathcal{A}}(u))$, i.e. each agent in $|h|$ is allocated a subset of the sites belonging to the agent type $h_{\mathcal{A}}(u)$ it is mapped to in the contact graph. An agent in $|h|$ may cover some, or all, of these sites or even completely extraneous sites:


- (i) if for all u in $\mathcal{A}_{|h|}$, $h_{\mathcal{S}}(\sigma_{|h|}^{-1}(u)) \subseteq \Gamma_\phi(u)$, we say that ϕ is *immature*;
- (ii) if for all u the inclusion is an equality and ϕ is an epi, ϕ is *mature*;
- (iii) otherwise ϕ is said to be *overgrown*.

The functions Γ_ϕ must satisfy, for all extensions ϕ and $\phi' \geq \phi$, the *faithfulness* property, $\Gamma_\phi = \Gamma_{\phi'} \psi_{\mathcal{A}}$ with ψ such that $\psi \phi = \phi'$; so a site requested by ϕ must be requested by any further extension. Additionally, this property forces Γ to eagerly ask for all sites that will be eventually requested at any given agent in the codomain of ϕ . If ϕ is not overgrown then no $\phi' \leq \phi$ is overgrown either.

Given a contact map g over C and a growth policy Γ for g , we define $\Gamma(g)$ by choosing one representative per \cong_g -isomorphism class of the set of all extensions of g which are mature according to Γ .

The following theorem guarantees that factorisations through $\Gamma(g)$ are unique when they exist, but *not* that they necessarily do exist. In section §2.3, we will construct a specific growth policy for which the exhaustivity of the decomposition can be proved by hand. As such, it fulfils our desired criteria of providing an exhaustive collection of mutually exclusive subcases.

Theorem 8. *Let g and m be contact maps over C and Γ a growth policy for g . If an embedding $\psi : g \rightarrow m$ can be decomposed in two ways as $\gamma_1\phi_1$ and $\gamma_2\phi_2$ with $\phi_i : g \rightarrow h_i$ in $\Gamma(g)$ and $\gamma_i : h_i \rightarrow m$, then $\phi_1 = \phi_2$ and $\gamma_1 = \gamma_2$.*



Proof. Suppose that $\gamma_1\phi_1 = \gamma_2\phi_2$, where ϕ_1 and ϕ_2 are mature extensions of g according to Γ and $\phi_1 \neq \phi_2$. As shown in diagram 2.6, we have an inner square formed by the pullback π_1, π_2 , and the minimal glueing θ_1, θ_2 of h_1, h_2 that factors γ_1, γ_2 . Every connected component of m has a pre-image in h_1 or h_2 , and thus also in g , since ϕ_1 and ϕ_2 are epis as mature extensions. Because every connected component of m has an image in h_1 and h_2 , then every connected component of m has a pre-image in both h_1 and h_2 . Hence θ_1 and θ_2 are epis.

The nodes in the images of θ_1 and θ_2 might be the same or differ. When they differ, some site z sitting on a node in the intersection of the images of θ_1, θ_2 is connected to a node outside the image, since θ_1, θ_2 are epis. However, z cannot be in the intersection of the images unless the site it is connected to is also part of the intersection (Lemma 5). Therefore the nodes in the images must be the same. In this case there has to be a site z that is not in the image of one of them or θ_1, θ_2 are both isos. So there must be a pair u, z , consisting of a node u in m with pre-images u_1, u_2 in h_1, h_2 and a site z of u , such that z has no pre-image in exactly one of θ_1, θ_2 . Say it is θ_2 . Since ϕ_1 is not overgrown, $z \in \Gamma_{\phi_1}(u_1)$ and, by faithfulness, $z \in \Gamma_{\phi}((u_1, u_2))$, where (u_1, u_2) is the pullback pre-image of u_1 and u_2 . So again, by faithfulness, $z \in \Gamma_{\phi_2}(u_2)$ which contradicts


our original assumption. Hence, θ_1 and θ_2 are isos. It follows that $\phi_1 = \phi_2$ as there is only one representative per \cong_g -isomorphism class in $\Gamma(g)$. Finally, $\gamma_1 = \gamma_2$ because ϕ_1 is an epi. ■

Given a rule r and an extension $\phi : r_L \rightarrow g$, we write r_ϕ for the refined rule associated to ϕ , that is, r_ϕ is the pair $(g, g^{(r, \phi)})$. Given Γ a growth policy for r_L , we write $\Gamma(r)$ for the family of rules obtained by refining r according to Γ , that is, $\Gamma(r)$ is the family of rules r_ϕ for ϕ ranging in $\Gamma(r_L)$. If ϕ is a \mathcal{P} -balanced extension of r , the refined rule r_ϕ has a *balance vector* in $\mathbb{Z}^{\mathcal{P}}$, written $\Delta\phi$, where, for each $p \in \mathcal{P}$, $\Delta\phi(p)$ is the difference in the number of copies of p produced and consumed by *any* r_ϕ -event.

An example of growth policy is the *ground* policy which assigns all possible sites to all agents. In this case, $\Gamma(g)$ is simply the set, possibly infinite, of all epis of g into mixtures, considered up to \cong_g . The ground refinement $\Gamma(r)$ contains all refinements of r along those epis. The refined rules therefore manipulate mixtures directly. It is easy to see that the ground refinement of r_{12}^+ in our example is infinite, since r_{12}^+ can trigger the extension of a chain of any length. A similar argument is true for r_{12}^- . Note that ground refinements of a rule r are trivially \mathcal{P} -balanced but, in general, the set of refined rules is impractically large or infinite as above. Instead, the growth policy that we introduce in the next section will always be finite.


2.3 Thermodynamic growth policy

An extension ϕ of a rule r is \mathcal{P} -balanced if it generates a refined rule r_ϕ that is \mathcal{P} -balanced. To find such extensions it seems natural to use minimal glueings: take as extensions the right leg θ_2^i of each relevant minimal glueing $\theta_1^i : p \rightarrow s_i \leftarrow r_L : \theta_2^i$ of $p \in \mathcal{P}$ and r_L (or r_R). For instance, the only relevant minimal glueing of the right-hand side of r_{12}^+ and the triangle is




If we use ϕ — the embedding corresponding to ϕ^* on the left-hand side — as an extension of r_{12}^+ we obtain rule 2.4. Now, having found the only extension of r_{12}^+ that produces a triangle, we are left with the problem of finding the extensions that cover the cases when r_{12}^+ can be applied without producing a triangle. Otherwise the decomposition would not be exhaustive; this is in general the case when using minimal glueings as extensions.

Whenever one of the participating agents in r_{12}^+ has a free site in addition to the two free sites that are bound by the rule, the formation of a triangle is excluded. In rule 2.5 we added a free r site to the blue node. The following extension of r_{12}^+ adds a free l site to the orange node.



Both extensions are minimally \mathcal{P} -balanced because any prefix of them that is \mathcal{P} -balanced is isomorphic to them as an extension of r_L . We call minimally \mathcal{P} -balanced extensions *primes*. Prime extensions are epis since erasing an untouched connected component in the codomain preserves balance. However, primes may overlap as shown by the following rule applications and therefore do not define in general a valid refinement.




It is thus apparent that an energy-based rule refinement has to proceed cautiously to be exhaustive and mutually exclusive. This is where our growth policy technique comes in handy to define such refinements. It divides the problem in a much simpler group of problems: each extension ϕ must declare the set of sites that it requires to be mature and \mathcal{P} -balanced. Minimal glueings play a guiding role here. They tell us whether an extension has successfully avoided or absorbed completely an energy pattern.

In our example, we extend our rule r_{12}^+ step by step to see this idea in action. First take no extension at all or, more precisely, take the identity arrow as an extension. On the left-hand side there is only one minimal glueing, the disjoint union, which, as it is always the case, is irrelevant. On the right-hand side instead we have two minimal glueings: the disjoint union and the triangle

itself, as in diagram 2.7. The latter is indeed relevant and informs us of which sites are missing in the extension, namely the l site on the orange node and the r site on the blue node. So we ask for both and set $\Gamma_{\mathbf{1}_{r_R}}(u) = \{l, r\}$ for all $u \in \mathcal{A}_{|r_R|}$. Now let us add one of them as a free site and ask again which sites each agent requires. This extension, call it ϕ_1 , has codomain the left-hand side of rule 2.8. The codomain of the corresponding extension ϕ_1^* on the right-hand side does not glue relevantly with the triangle anymore. However, $\mathbf{1}$ is a prefix of ϕ_1 and hence, due to faithfulness, Γ_{ϕ_1} should ask for the same sites that $\Gamma_{\mathbf{1}}$ does, i.e. $\Gamma_{\phi_1}(u) = \Gamma_{\mathbf{1}}(u)$ for all agents u in the image of $\mathbf{1}$. So here again caution must be exercised. The solution is to remember which sites have been asked for in the past and to keep asking for them in future extensions.

Given contact graph C and r in \mathcal{G} we define our growth policy Γ for r_L as follows. Suppose $\phi : r_L \rightarrow g$ is an extension of r_L . We set Γ_ϕ to request a site z in $\sigma_C^{-1}(g_{\mathcal{A}}(u))$ at agent u in $\mathcal{A}_{|g|}$ iff either

- (i) $u = \phi_{\mathcal{A}}(u_0)$ and $z = \phi_{\mathcal{S}}(z_0)$ for some u_0 in $\mathcal{A}_{|r_L|}$ and z_0 in $\mathcal{S}_{|r_L|}$; or
- (ii) ϕ factorises as $\phi_2 \phi_1$, where $\phi_1 : r_L \rightarrow g_1$, and there is a relevant minimal glueing $\gamma : p \rightarrow s \leftarrow g_1 : \theta$, with p in \mathcal{P} , and some u_1 in $\mathcal{A}_{|g_1|}$ and a site z_1 in $\sigma_{|s|}^{-1}(\theta_{\mathcal{A}}(u_1))$ such that $u = \phi_{2,\mathcal{A}}(u_1)$ and $z = s_{\mathcal{S}}(z_1)$; or




- (iii) $z = g_{\mathcal{S}}(z_2)$ for some z_2 in $\mathcal{S}_{|g|}$ such that $z_2 \mathcal{E}_{|g|} z_3$ and $g_{\mathcal{S}}(z_3)$ in $\Gamma_\phi(u)$.

In words, clause (i) ensures that all sites in r_L are asked for while clause (ii) adds sites z in \mathcal{S}_C corresponding to sites z_1 in $\mathcal{S}_{|s|}$ which appear by glueing with p at some point between r_L and g . Clause (iii), on the other hand, asks for sites that are bound to sites that are requested by the growth policy so that extensions that avoid minimal glueings are not overgrown. We refer to the extension $\phi_2 : g_1 \rightarrow g$ as a *rewind* of ϕ and say that the request of z at u originates from u_1 . By rewinding extensions we can remember which sites have been asked for in the past.

Symmetrically, we define a growth policy Γ^* for r_R by applying the same definition to the inverse rule r^\dagger . Finally, we define our growth policy $\Gamma^{\mathcal{P}}$ as the union of both growth policies, that is, $\Gamma^{\mathcal{P}}_\phi(u) = \Gamma_\phi(u) \cup \Gamma^*_\phi(u)$.


According to this growth policy, the extension ϕ_1^* of the right-hand side of r_{12}^+ in our example is immature (despite being \mathcal{P} -balanced) since the following rewind asks for a site that is missing in its image.




So we must add an r site to the blue node. There are two possibilities when the site is added: it can be free or it can be bound. In particular, the contact graph C tells us that an r site on a blue node can only be bound to an l site on a green node. We obtain then two new extensions, with codomains:



The first extension cannot possibly ask for any more sites. However, the second extension, call it ϕ_2^* , may ask for the r site on the green node. If this is the case there must be a rewind of ϕ_2^* which contains a pre-image of the green node and glues relevantly with the triangle.



Therefore ϕ_2^* is immature as well. We must reveal the r site on the green node and so we obtain




Finally, all extensions are mature. Note that the second extension has an l site on the rightmost orange node which would not be asked by the growth policy if it were not for clause (iii). In the absence of clause (iii) we would have moved from an immature extension to an overgrown extension in just one step, leaving us in a strange situation by allowing the growth policy to define an empty refinement. Next we prove that the growth policy that we have introduced in this section is in general well-defined and well-behaved.

Theorem 9. *The above $\Gamma^{\mathcal{P}}$ is indeed a growth policy for r_L and the induced refined family of rules $\Gamma^{\mathcal{P}}(r)$ is exhaustive, non-empty, \mathcal{P} -balanced, and finite.*

Proof. We take the same notations as in diagram 2.9.

Growth policy: Clearly, $\Gamma_{\phi_1}^{\mathcal{P}}(u_1) \subseteq \Gamma_{\phi}^{\mathcal{P}}(u)$ as every request for a site in g_1 will propagate to g by definition. To prove the other direction, we need to verify that the requests generated by rewinds do not depend on the choice of factorisation as $\Gamma^{\mathcal{P}}(\phi)(u)$ must be a subset of $\Gamma^{\mathcal{P}}(\phi_1)(u_1)$ for every ϕ_1 . So, without loss of generality, assume there are two factorisations of ϕ given by $\phi_2 \phi_1 = \phi = \psi_2 \psi_1$ and consider a site request in u originating from some u_2 in g_2 , as in the following diagram.



Consider g_0 the pull-back of the two rewinds (i.e. the lower cospan). Because $\phi_{1,\mathcal{A}}(u_1) = \psi_{1,\mathcal{A}}(u_2) = u$ the pullback must contain a pre-image for u_1 and u_2 , say (u_1, u_2) . The relevant minimal glueing of p and g_2 that makes the site request restricts to another minimal glueing of p and g_0 . This new minimal glueing is still relevant as it contains the same overlap with the original r_L . As such, the same site request is made by the pre-image agent (u_1, u_2) in g_0 which then propagates to u_1 in g_1 as required.

Exhaustive: Take any embedding ψ of r_L into a mixture m . We can restrict the codomain of ψ to be the connected closure n of the image of ψ in m , resulting in an epi $\psi_n : r_L \rightarrow n$. Let us further restrict n by removing (i) all sites not requested by the growth policy and (ii) all agents that have no sites requested by the growth policy. The result, call it g , has the same number of connected components as r_L since Γ^S only requests sites which appear by glueing and are thus (perhaps indirectly) connected to the sites that are modified by the rule. We thus obtain an epi $\phi : r_L \rightarrow g$ which is mature with respect to Γ^P since, by construction, its image contains all sites requested by Γ^P and no other foreign site. It is easy to see that ϕ factorises ψ .

Non-empty: Clause (i) guarantees that we request at least the sites in r which implies that $\mathbf{1}$ is not overgrown. Due to clause (iii) there is always an extension whose image contains exactly all sites requested by Γ^S and lies between an immature and an overgrown extension according to \leq . This extension is an epi because, as pointed out for exhaustivity, Γ^S only requests sites connected to those modified by the rule.

\mathcal{P} -balanced: If $\phi \in \Gamma^P(r)$ is not \mathcal{P} -balanced then there must be some relevant minimal glueing inducing a further site request. Hence, ϕ cannot be mature.

Finite: A request for a site a at some node in an extension $\phi : r_L \rightarrow g$, or $\phi^* : r_R \rightarrow g$, originates from a relevant minimal glueing of some p in \mathcal{P} with a prefix ϕ_1 of ϕ . Because this glueing is relevant, it must be that a is at a distance from the image of r_L in the codomain of ϕ_1 which is at most $\delta(p)$, the diameter of p (else p would not intersect the image of r_L). The same bound holds in the codomain of ϕ , as distances can only contract by further extension. Therefore any site requested in g has a distance to the image $\phi(r_L)$ which is bounded by $\max_{p \in \mathcal{P}} \delta(p)$. If ϕ is not overgrown, this sets a bound on the diameter of g . Hence there are finitely many mature extensions. ■

Therefore, given \mathcal{G} and \mathcal{P} , we obtain a finite \mathcal{P} -balanced rule set $\mathcal{G}_{\mathcal{P}}$, which refines \mathcal{G} exhaustively, by taking the disjoint sum of the refined rules $\mathcal{G}_{\mathcal{P}} = \sum_{r \in \mathcal{G}} \Gamma^P(r)$. To every refinement r_{ϕ} corresponds an inverse refinement $r_{\phi^*}^+$. Hence, $\mathcal{G}_{\mathcal{P}} = \mathcal{G}_{\mathcal{P}}^+$ is closed under inversion like \mathcal{G} .

2.4 Rates and detailed balance

To equip $\mathcal{G}_{\mathcal{P}}$ with rates we define a rate map $k : \mathcal{G}_{\mathcal{P}} \rightarrow \mathbb{R}_{>0}$. We use the real-valued vector of *energy costs* ϵ introduced at the beginning of this chapter (page 17) together with the balance vector $\Delta\phi$ of a refined rule r_{ϕ} in $\mathcal{G}_{\mathcal{P}}$ (page 29) to constrain the ratio between the forward and the backward rate:

$$\ln k(r_{\phi^*}^{\dagger}) - \ln k(r_{\phi}) = \epsilon \cdot \Delta\phi \quad (2.10)$$

The pair of rules r, r^{\dagger} is biased in the forward direction if $k_r(\phi) > k_{r^{\dagger}}(\phi^*)$ and Eq. 2.10 tells us that this happens when $\epsilon \cdot \Delta\phi < 0$, i.e. whenever the energy decreases as we go in that direction. This is the usual convention for energy functions.

We show that the set of refined rules $\mathcal{G}_{\mathcal{P}}$ with any such rate map has detailed balance. To simplify notation, we write $\mathcal{P}(m)$ for the \mathcal{P} -indexed vector which maps p to $[p; m]$. Using vector notation, the energy $E(m)$ of a state m (as defined in Eq. 2.1) can then simply be written as $\epsilon \cdot \mathcal{P}(m)$. Moreover, we write $\mathcal{L}_{\mathcal{G}}(m)$ for the finite strongly connected component of m in $\mathcal{L}_{\mathcal{G}}$ and recall the definition of a probability distribution π_m on $\mathcal{L}_{\mathcal{G}}(m)$ from Eq. 1.3, which after substituting $E(m)$ reads

$$\pi_m(x) = \frac{e^{-\epsilon \cdot \mathcal{P}(x)}}{\sum_{y \in \mathcal{L}_{\mathcal{G}}(m)} e^{-\epsilon \cdot \mathcal{P}(y)}} \quad (2.11)$$

We can now prove the main theorem of this chapter.

Theorem 10. *Let \mathcal{G} , \mathcal{P} , $\mathcal{G}_{\mathcal{P}}$, k , and π_m be defined as above; then (i) $\mathcal{L}_{\mathcal{G}_{\mathcal{P}}}$ and $\mathcal{L}_{\mathcal{G}}$ are isomorphic as symmetric labelled transition systems; and (ii) for any mixture m , the time-homogeneous continuous-time Markov chain $\mathcal{L}_{\mathcal{G}_{\mathcal{P}}}^k$ has detailed balance for, and converges to, π_m on $\mathcal{L}_{\mathcal{G}_{\mathcal{P}}}(m)$.*

Proof. Both $\mathcal{L}_{\mathcal{G}}$ and $\mathcal{L}_{\mathcal{G}_{\mathcal{P}}}$ offer transitions from a mixture m : the former are labelled by pairs (r, ψ) with r in \mathcal{G} and ψ in $[r_L; m]$ while the latter by pairs (r_{ϕ}, γ) with r_{ϕ} the refinement of r along a mature extension $\phi : r_L \rightarrow g$ and γ in $[g; m]$. Steps in the latter can be mapped to steps in the former by transforming labels as follows: $(r_{\phi}, \gamma) \mapsto (r, \gamma\phi)$. By Th. 9, each event (r, ψ) is factored by exactly one event (r_{ϕ}, γ) and thus this correspondence is a bijection, which establishes the first claim.

Since we have multiple rules in $\mathcal{L}_{\mathcal{G}_P}$, each of which can be applied in several ways, there can be more than one transition from m to the same n — each uniquely described by a (r_ϕ, γ) label. Each such (r_ϕ, γ) has an inverse $(r_{\phi^\star}^\dagger, \gamma^\star)$ and we have a bijection between them and thus between transitions from m to n and those from n to m due to Lemmas 3 and 7.

Consider a pair t, t^\dagger of such corresponding events due to r_ϕ and $r_{\phi^\star}^\dagger$. Because t is a transition from m to n and ϕ is \mathcal{P} -balanced, we have $\mathcal{P}(n) = \mathcal{P}(m) + \Delta\phi$ and hence $\epsilon \cdot \Delta\phi = \epsilon \cdot (\mathcal{P}(n) - \mathcal{P}(m))$. So, by Eq. 2.10, the rates of t, t^\dagger are such that:

$$k(t^\dagger) e^{-\epsilon \cdot \mathcal{P}(n)} = k(t) e^{-\epsilon \cdot \mathcal{P}(m)}$$

and by summing this equation over all pairs, we obtain detailed balance for the probability local to the component $\mathcal{L}_{\mathcal{G}_P}(m) = \mathcal{L}_{\mathcal{G}_P}(n)$, defined above as $\pi_m = \pi_n$, since:

$$q_{nm} e^{-\epsilon \cdot \mathcal{P}(n)} = q_{mn} e^{-\epsilon \cdot \mathcal{P}(m)}$$

The convergence statement then follows from Lemma 2 applied to the finite irreducible continuous-time Markov chain $\mathcal{L}_{\mathcal{G}_P}^k(m)$ that is obtained by cropping all states not in $\mathcal{L}_{\mathcal{G}_P}(m)$. ■

Note that the subset of the state space which is reachable from m in $\mathcal{L}_{\mathcal{G}}$, namely $\mathcal{L}_{\mathcal{G}}(m)$, is finite. Hence, the *partition function* $Z(m) := \sum_{y \in \mathcal{L}_{\mathcal{G}}(m)} e^{-E(y)}$ which figures in the denominator of π_m is also finite. In the presence of rules which increase the number of agents, the components $\mathcal{L}_{\mathcal{G}}(m)$ can be infinite and $Z(m)$ may diverge. For mass action stochastic Petri nets (§1.1), convergence is guaranteed if detailed balance holds, but it is not true in general for Kappa (Danos and Oury, 2010, 2013).

Another point worth making is that the result holds symbolically regardless of the energy costs ϵ . Therefore ϵ can be seen as a set of parameters. This is an ideal support for machine learning techniques if one were contemplating fitting a model to data.

2.5 Linear kinetic model

The theorem in the previous section holds for any rate map that agrees with Eq. 2.10. In this section, we show how to obtain a concrete rate $k(r_\phi)$ for each refined rule r_ϕ in \mathcal{G}_P . To simplify the task, we pick rates from a tractable subset

of all possible choices by performing a log-affine expansion on the so-called «thermodynamic drive» $\Delta E = \epsilon \cdot \Delta\phi$. The expansion uses, for each generator rule r in \mathcal{G} , a constant $c_r \in \mathbb{R}$ and a real-valued matrix A_r of dimension $|\mathcal{P}| \times |\mathcal{P}|$. Then we assign rates according to the following equality

$$\ln k(r_\phi) = c_r - A_r \epsilon \cdot \Delta\phi \quad (2.12)$$

subject to the following constraints

$$\begin{aligned} c_r &= c_{r^\dagger} \\ A_r + A_{r^\dagger} &= I \end{aligned}$$

with I the $|\mathcal{P}| \times |\mathcal{P}|$ identity matrix. We verify that k satisfies Eq. 2.10 by subtracting $\ln k(r_{\phi^\star}^\dagger)$ and $\ln k(r_\phi)$, giving us

$$\ln k(r_{\phi^\star}^\dagger) - \ln k(r_\phi) = (c_{r^\dagger} - A_{r^\dagger} \epsilon \cdot \Delta\phi^\star) - (c_r - A_r \epsilon \cdot \Delta\phi)$$

We have $\Delta\phi^\star = -\Delta\phi$ by reversibility of rules and so

$$\begin{aligned} \ln k(r_{\phi^\star}^\dagger) - \ln k(r_\phi) &= c_{r^\dagger} - c_r + A_r \epsilon \cdot \Delta\phi + A_{r^\dagger} \epsilon \cdot \Delta\phi \\ &= (A_r + A_{r^\dagger}) \epsilon \cdot \Delta\phi \\ &= I \epsilon \cdot \Delta\phi = \epsilon \cdot \Delta\phi \end{aligned}$$

The kinetic model of Eq. 2.12 requires $|\mathcal{P}|^2 \times |\mathcal{G}| + |\mathcal{G}|$ parameters: one $A_{r,pq}$ for each generator rule $r \in \mathcal{G}$ and pair $p, q \in \mathcal{P}^2$, plus one c_r for each $r \in \mathcal{G}$. In practice one needs even fewer parameters as only those energy patterns that are relevant to a given generator rule r , i.e. those that have a non-zero balance for at least one rule in $\Gamma^\mathcal{P}(r)$, need to be considered when building A_r . Typically, for larger models, this will be a far smaller number than $|\mathcal{P}|$. This relative parsimony is compounded by the fact that the number of *independent* parameters will be often lower because the $\Delta\phi$ family often has low rank, meaning that, for a set of extensions ϕ , the balance vectors $\Delta\phi$ can be determined as a linear combination of a smaller basis set. By way of comparison, if we were to assign kinetic rates to each refined rule, we would need $\sum_{r \in \mathcal{G}} |\Gamma^\mathcal{P}(r)|$ parameters.

We find two special cases for the kinetic model presented here that seem appealing as a first choice for parameterisation. First, by setting $c_r = c_{r^\dagger} = 0$, $A_r = I$ and $A_{r^\dagger} = 0$, we get $k(r_\phi) = e^{-\epsilon \cdot \Delta\phi}$ and $k(r_{\phi^\star}^\dagger) = 1$. Whenever r^\dagger is

the thermodynamically favoured direction (and we can always choose it so), this choice amounts to being exponentially reluctant to climb up the energy gradient. In this way, this choice can be thought of as continuous-time version of the Metropolis algorithm introduced in §1.1.

The second special case, on the other hand, is completely symmetric and can be obtained by fixing $A_r = A_{r^\dagger} = I/2$ and $C_r = e^{c_r}$:

$$\begin{aligned} k(r_\phi) &= C_r e^{-\epsilon \cdot \Delta\phi/2} \\ k(r_{\phi^\star}^\dagger) &= C_r e^{\epsilon \cdot \Delta\phi/2} \end{aligned} \quad (2.13)$$

Note the similarity of Eq. 2.13 to the Arrhenius equation.

$$k = A e^{-E_a}$$


where E_a is the *activation energy* of the reaction (expressed in units of $1/k_B T$ as in Eq. 1.3) and A is a pre-exponential factor that defines the rate at which the molecules involved in the reaction collide in the correct orientation for the reaction to occur. Eq. 2.13 is a special case of the Arrhenius equation when we equate $A = C_r$ and $E_a = \epsilon \cdot \Delta\phi/2$. The first equality is therefore interpreted as an assumption that the rate of molecules colliding in the right orientation depends only on the molecular motifs present in the left-hand side of the generator rule r (i.e. not on the context revealed by the refined rule). Albeit an approximation, it might prove useful whenever the generator rules specify enough context to determine, for instance, the accessibility of the reaction centre. Another possible approach would be to compute A based on properties of the refinement, e.g. how big the surrounding molecular complex is.

The second equality, $E_a = \epsilon \cdot \Delta\phi/2$, tells us that the energetic barrier between the reactants and the products is determined only based on the energy patterns that are destroyed and created by the refined rule and their energy cost. Since the activation energy is allowed to depend on the context revealed by the refined rule, this assumption imposes a softer constrain than the previous one.


2.6 Example: Triangles all the way down

In this section we will complete and conclude the example on the thermodynamical control of the formation of triangles that we have used throughout this chapter. Additionally we present the model in the format used by the Kappa

simulator, KaSim version 4.0 (Boutillier et al., 2014), and run a few simulations to get an idea of how the model behaves. We use the symmetric linear kinetic model of Eq. 2.13 with $c_r = 0$ to derive the rates and add three more energy patterns (in addition to the triangle) to demonstrate how they interact in the expansion of the rates. In particular, we add one energy pattern for each type of edge, that is, for contact maps



The analysis on §2.3 unveiled five refinements for rule r_{12}^+ (Eq. 2.2) which we enumerate below.



The four subcases that do not create a triangle have $\Delta E = \epsilon \cdot \Delta\phi = \epsilon(d_{12})$ where d_{12} is the dimer of an agent of type 1 and an agent of type 2. Hence, their rate under the symmetric linear kinetic model with $c_r = 0$ would be $e^{-\epsilon(d_{12})/2}$. On the other hand, the fourth refined rule creates a triangle and thus its $\Delta E = \epsilon \cdot \Delta\phi = \epsilon(d_{12}) + \epsilon(t)$ where $\epsilon(t)$ is the energy cost of the triangle. Its rate then is $e^{-(\epsilon(d_{12}) + \epsilon(t))/2}$. The inverse generator rule r_{12}^- produces as refinements the inverse of the five subrules enumerated above. The other generator rules follow a similar pattern of refinement.

Now we present the KaSim model. The rules of the model have been manually compressed to take advantage of KaSim's extended syntax (e.g. binding types). Note also that by using KaSim's variables we can define the rates parametrically, allowing us to easily try out different values for the energy costs.

```

1  # Agent declarations
2  %agent: A(l,r)
3  %agent: B(l,r)
4  %agent: C(l,r)
5
6  # Energy costs
7  %var: 't' -10
8  %var: 'ab' 1
9  %var: 'bc' 1
10 %var: 'ca' 1
11
12 # Observable
13 %obs: 'T' |A(l!1, r!2), B(l!2, r!3), C(l!3, r!1)|
14
15 # Rules
16 # A(r), B(l) -> A(r!1), B(l!1) refines into:
17 A(l,r), B(l,r) -> A(l,r!1), B(l!1,r) @ [exp] (-1/2 * 'ab')
18 A(l!r.C,r), B(l,r) -> A(l!r.C,r!1), B(l!1,r) @ [exp] (-1/2 * 'ab')
19 A(l,r), B(l,r!1.C) -> A(l,r!1), B(l!1,r!1.C) @ [exp] (-1/2 * 'ab')
20 A(l!1,r ), B(l ,r!3), C(l!3,r!1) -> \
21 A(l!1,r!2), B(l!2,r!3), C(l!3,r!1) @ [exp] (-1/2 * ('ab' + 't'))
22 C(r!1), A(l!1,r ), B(l ,r!3), C(l!3) -> \
23 C(r!1), A(l!1,r!2), B(l!2,r!3), C(l!3) @ [exp] (-1/2 * 'ab')
24
25 # A(r!1), B(l!1) -> A(r), B(l) refines into:
26 A(l,r!1), B(l!1,r) -> A(l,r), B(l,r) @ [exp] -(-1/2 * 'ab')
27 A(l!r.C,r!1), B(l!1,r) -> A(l!r.C,r), B(l,r) @ [exp] -(-1/2 * 'ab')
28 A(l,r!1), B(l!1,r!1.C) -> A(l,r), B(l,r!1.C) @ [exp] -(-1/2 * 'ab')
29 A(l!1,r!2), B(l!2,r!3), C(l!3,r!1) -> \
30 A(l!1,r ), B(l ,r!3), C(l!3,r!1) @ [exp] -(-1/2 * ('ab' + 't'))
31 C(r!1), A(l!1,r!2), B(l!2,r!3), C(l!3) -> \
32 C(r!1), A(l!1,r ), B(l ,r!3), C(l!3) @ [exp] -(-1/2 * 'ab')
33
34 # B(r), C(l) -> B(r!1), C(l!1) refines into:
35 B(l,r), C(l,r) -> B(l,r!1), C(l!1,r) @ [exp] (-1/2 * 'bc')
36 B(l!r.A,r), C(l,r) -> B(l!r.A,r!1), C(l!1,r) @ [exp] (-1/2 * 'bc')
37 B(l,r), C(l,r!1.A) -> B(l,r!1), C(l!1,r!1.A) @ [exp] (-1/2 * 'bc')
38 B(l!1,r ), C(l ,r!3), A(l!3,r!1) -> \
39 B(l!1,r!2), C(l!2,r!3), A(l!3,r!1) @ [exp] (-1/2 * ('bc' + 't'))
40 A(r!1), B(l!1,r ), C(l ,r!3), A(l!3) -> \
41 A(r!1), B(l!1,r!2), C(l!2,r!3), A(l!3) @ [exp] (-1/2 * 'bc')
42

```

```

43 # B(r!1), C(l!1) -> B(r), C(l) refines into:
44 B(l,r!1), C(l!1,r) -> B(l,r), C(l,r) @ [exp] -(-1/2 * 'bc')
45 B(l!r.A,r!1), C(l!1,r) -> B(l!r.A,r), C(l,r) @ [exp] -(-1/2 * 'bc')
46 B(l,r!1), C(l!1,r!1.A) -> B(l,r), C(l,r!1.A) @ [exp] -(-1/2 * 'bc')
47 B(l!1,r!2), C(l!2,r!3), A(l!3,r!1) -> \
48 B(l!1,r ), C(l ,r!3), A(l!3,r!1) @ [exp] -(-1/2 * ('bc' + 't'))
49 A(r!1), B(l!1,r!2), C(l!2,r!3), A(l!3) -> \
50 A(r!1), B(l!1,r ), C(l ,r!3), A(l!3) @ [exp] -(-1/2 * 'bc')
51
52 # C(r), A(l) -> C(r!1), A(l!1) refines into:
53 C(l,r), A(l,r) -> C(l,r!1), A(l!1,r) @ [exp] (-1/2 * 'ca')
54 C(l!r.B,r), A(l,r) -> C(l!r.B,r!1), A(l!1,r) @ [exp] (-1/2 * 'ca')
55 C(l,r), A(l,r!1.B) -> C(l,r!1), A(l!1,r!1.B) @ [exp] (-1/2 * 'ca')
56 C(l!1,r ), A(l ,r!3), B(l!3,r!1) -> \
57 C(l!1,r!2), A(l!2,r!3), B(l!3,r!1) @ [exp] (-1/2 * ('ca' + 't'))
58 B(r!1), C(l!1,r ), A(l ,r!3), B(l!3) -> \
59 B(r!1), C(l!1,r!2), A(l!2,r!3), B(l!3) @ [exp] (-1/2 * 'ca')
60
61 # C(r!1), A(l!1) -> C(r), A(l) refines into:
62 C(l,r!1), A(l!1,r) -> C(l,r), A(l,r) @ [exp] -(-1/2 * 'ca')
63 C(l!r.B,r!1), A(l!1,r) -> C(l!r.B,r), A(l,r) @ [exp] -(-1/2 * 'ca')
64 C(l,r!1), A(l!1,r!1.B) -> C(l,r), A(l,r!1.B) @ [exp] -(-1/2 * 'ca')
65 C(l!1,r!2), A(l!2,r!3), B(l!3,r!1) -> \
66 C(l!1,r ), A(l ,r!3), B(l!3,r!1) @ [exp] -(-1/2 * ('ca' + 't'))
67 B(r!1), C(l!1,r!2), A(l!2,r!3), B(l!3) -> \
68 B(r!1), C(l!1,r ), A(l ,r!3), B(l!3) @ [exp] -(-1/2 * 'ca')
69
70 # Initial mixture
71 %init: 1000 (A()), B(), C())

```

The above KaSim model uses $\epsilon(d_{12}) = \epsilon(d_{23}) = \epsilon(d_{31}) = 1$ and $\epsilon(t) = -10$. Below we will change this values to see how the production of triangles is affected by them. We have set the initial mixture to contain 1000 copies of each type of agent. To run a simulation for 50 time units and take measurements (i.e. count the number of triangles in the mixture) every 0.1 time units, we issue the following command

```
$ KaSim t.ka -o t-10.tsv -d t-10 -l 50 -p 0.1
```

The input file is `t.ka` and the measurements are saved in the `t-10.tsv` in the `t-10` folder. The resulting plots are displayed in Fig. 2.1.

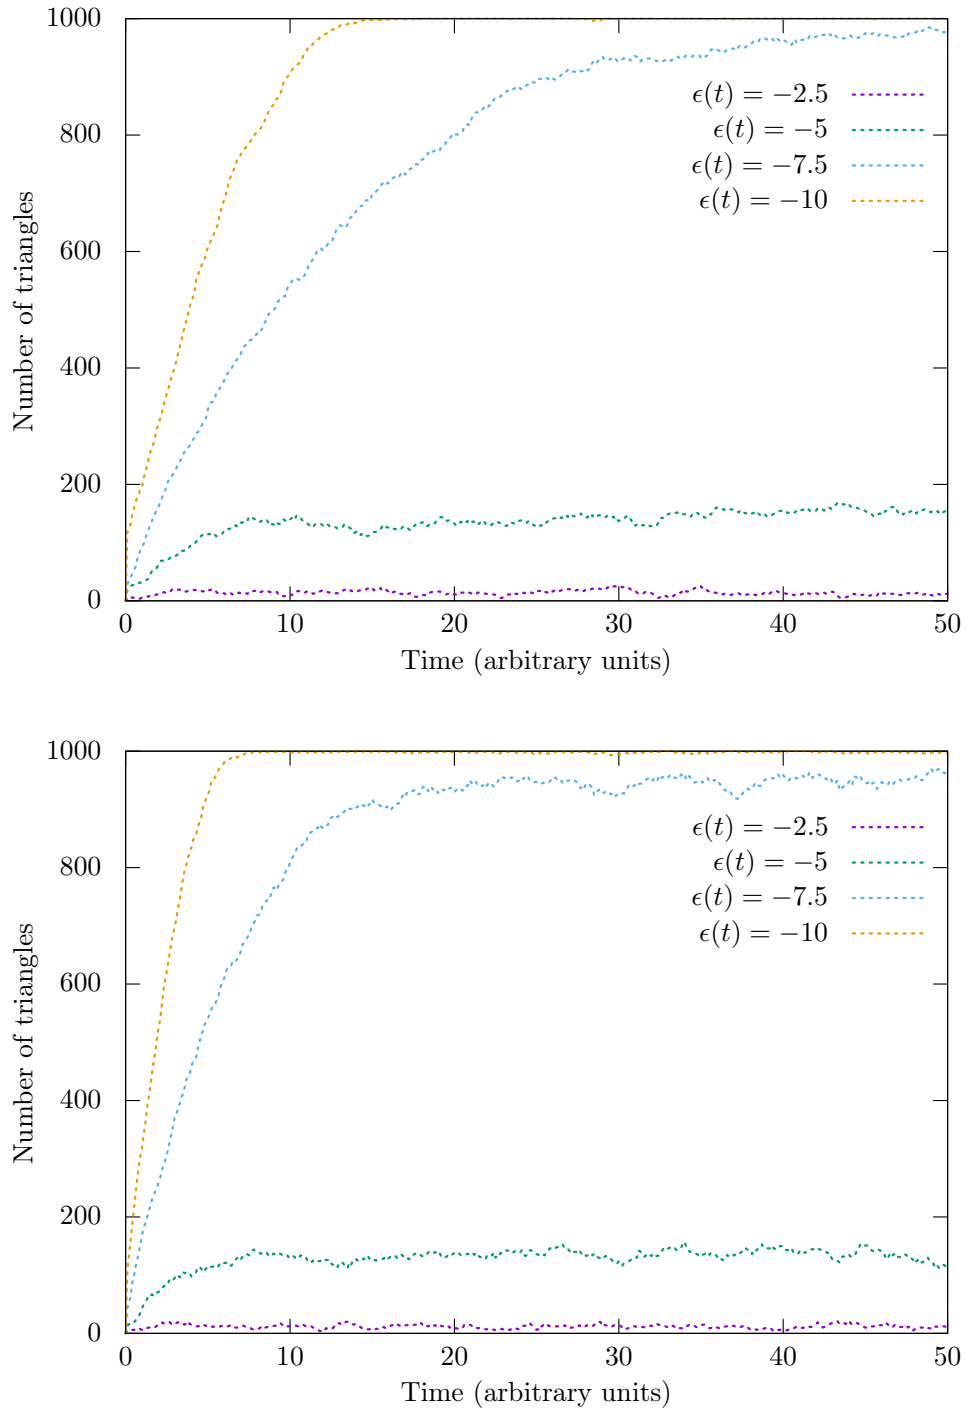


Figure 2.1: Trajectories for the number of triangles when $\epsilon(t)$ varies. In the plot above the energy cost of the dimers is 0 whereas in the plot below they are set to 1.

First, we notice that the moderate energy penalty we impose on dimers in the second plot does not change much the number of triangles at equilibrium. It does, however, have an impact on the speed at which the triangles form. This effect is perhaps counter-intuitive.

Second, notice that when $\epsilon(t) = -10$ all agents are used to build triangles. In contrast, when $\epsilon(t) = -5$ less than 20% of the agents of each type are used. In both cases the set of states with a globally minimum energy is the same, namely those states that maximise the amount of triangles. So then why is it that in the latter case there are so few triangles? The reason is entropic: although the probability of being in a state with few triangles is small, there are many such states and together they outweigh the probability of being in the few states where the energy is minimal. By further decreasing the energy of those few states we compensate for this mass effect, until at $\epsilon(t) = -10$, order wins, and the effect is not noticeable anymore.

2.7 Example: Flagellum's engine

In this section we present another model. This model is inspired in a classical object of study in molecular biology: the bacterial flagellar engine. The flagellar engine can rotate clockwise or anti-clockwise at high angular velocities. When it rotates clockwise the filaments of the flagellum move chaotically in all directions, making the bacterium tumble and thus randomly change the axis of its body and engine. When it rotates anti-clockwise the filaments of the flagellum align and move synchronously, propulsing the bacterium in the direction the engine is pointing to. In the latter regime the bacterium thus swims forward. When the bacterium detects that the levels of food are decreasing or the amount of poisonous substances is increasing, it tumbles to change the direction in which it is swimming. In this way it implements a basic chemotactic system.

A simple model of the switch between the two modes has been proposed by Bai et al. (2010). In this model the engine is seen as a ring of n identical components, called protomers or P for short, with two possible conformations, 0 and 1. Here we take $n = 34$ for simulations and diagrams but the analysis does not depend on the specific value of n . A ring homogeneously in state 0 (1) rotates (anti-)clockwise and induces tumbling (straight motion). Importantly,





Figure 2.2: Ring of protomers with some Y s bound. Since only a few Y s are bound to the ring, the majority of protomers is likely to be in state 0 (visually represented as grey nodes) and a minority in state 1 (green).

neighbouring P s on the ring prefer to have matching conformations. States of the ring with many mismatches thus incur high penalties. A small diffusible protein named CheY, which we call Y for short, binds P when it is activated. When Y is binding P , P favours state 1. Conversely, in the absence of a Y molecule binding P , P favours state 0. CheY, in turn, is activated by the system of chemoreceptors in the presence of food and absence of poisons.⁴ The configuration of the chemoreceptor cluster and its activity have also been modelled thermodynamically (Lan et al., 2011).

As each of the P s can be in four states, a ring of size 34 has on the order of 10^{18} non-isomorphic configurations. This precludes a Petri net approach to the dynamics where each state of the whole ring is considered as one chemical species. We thus use the rule-based approach pioneered in Kappa that allows us to specify events based only on a partial and local context around each protomer and derive the set of rules by applying the method of §2.3.


⁴ Here we assume that every Y is an activated CheY.

We define the contact graph of the model as



where P has 4 sites a, b, c, d . The first two form the backbone of the ring while c can bind Y s. Site d encodes the conformation state of P : we say P is in state 0 when site d is bound to an A agent and is in state 1 when bound to a B agent (A and B agents are not displayed in the contact graph above). We will never mention this site (nor A and B agents) explicitly but instead will colour the agent of type P accordingly.⁵ Also, we will draw sites a, b, c always on the left, right and top of P , respectively, and thus forgo annotating the name of the site.

The informal statements about the favoured states of P in the different configurations discussed above are captured in the definition of the energy patterns and associated energy costs. Note that the various patterns overlap.



We abuse notation by referring to both the pattern and its energy cost as ϵ_{ij} . The following constraints are imposed on the energy costs:

$$\epsilon_{00}, \epsilon_{11} < \epsilon_{10}, \epsilon_{01} \quad (2.14)$$

$$\epsilon_0 < \epsilon_1 \quad (2.15)$$


$$\epsilon_0^Y > \epsilon_1^Y \quad (2.16)$$

These inequalities enact the considerations in the discussion above. The role of Eq. 2.14 is to align the states of neighbours on the ring — essentially an Ising term which spreads conformation. Eq. 2.15 makes 0 the favoured state, while Eq. 2.16 inverts the situation in the presence of Y .


⁵ The naïve encoding where i) A and B have a free site that can bind site d of P , ii) whenever P changes from state 0 to state 1 we detach an A from P and attach a B to it, and thus iii) we have a pool of free A s and B s in the mixture, will have a problem with kinetics due to mass action: when we attach a B to P we make it less likely for the next P to bind a B since there are less B s free in the mixture. To solve this issue every P is either bound to an A that is in turn bound to a B or a B that is bound to an A . Whenever we want to change state we only need to exchange the order of the A and B .

Following §2.4 we associate to each ring configuration x the occurrence vector $\mathcal{P}(x)$ and total energy $\epsilon \cdot \mathcal{P}(x)$. For example, a ring of size n uniformly in state 0 with no bound Y s has total energy $n(\epsilon_{00} + \epsilon_0)$.


The next step is to define the set of generator rules \mathcal{G} . The first pair of rules that we include in this set are r_Y^+ the binding of P and Y and its inverse r_Y^- .



An uncoloured P means it can bind a Y regardless of the state it is in. The nature of the method presented in §2.3 allows us to refine each rule individually, so we proceed to refine r_Y^+, r_Y^- immediately. We first give the rationale for the refinements informally. The pair of rules has an ambiguous energy balance because applying the forward rule r_Y^+ to a P in state 0 will create an ϵ_0^Y pattern while applying it to a P in state 1 will create an ϵ_1^Y pattern. Hence, we cannot assign rates to these rules that satisfy detailed balance — unless $\epsilon_0^Y = \epsilon_1^Y$, which contradicts Eq. 2.16. To get \mathcal{P} -balanced rules one needs to refine r_Y^+, r_Y^- into

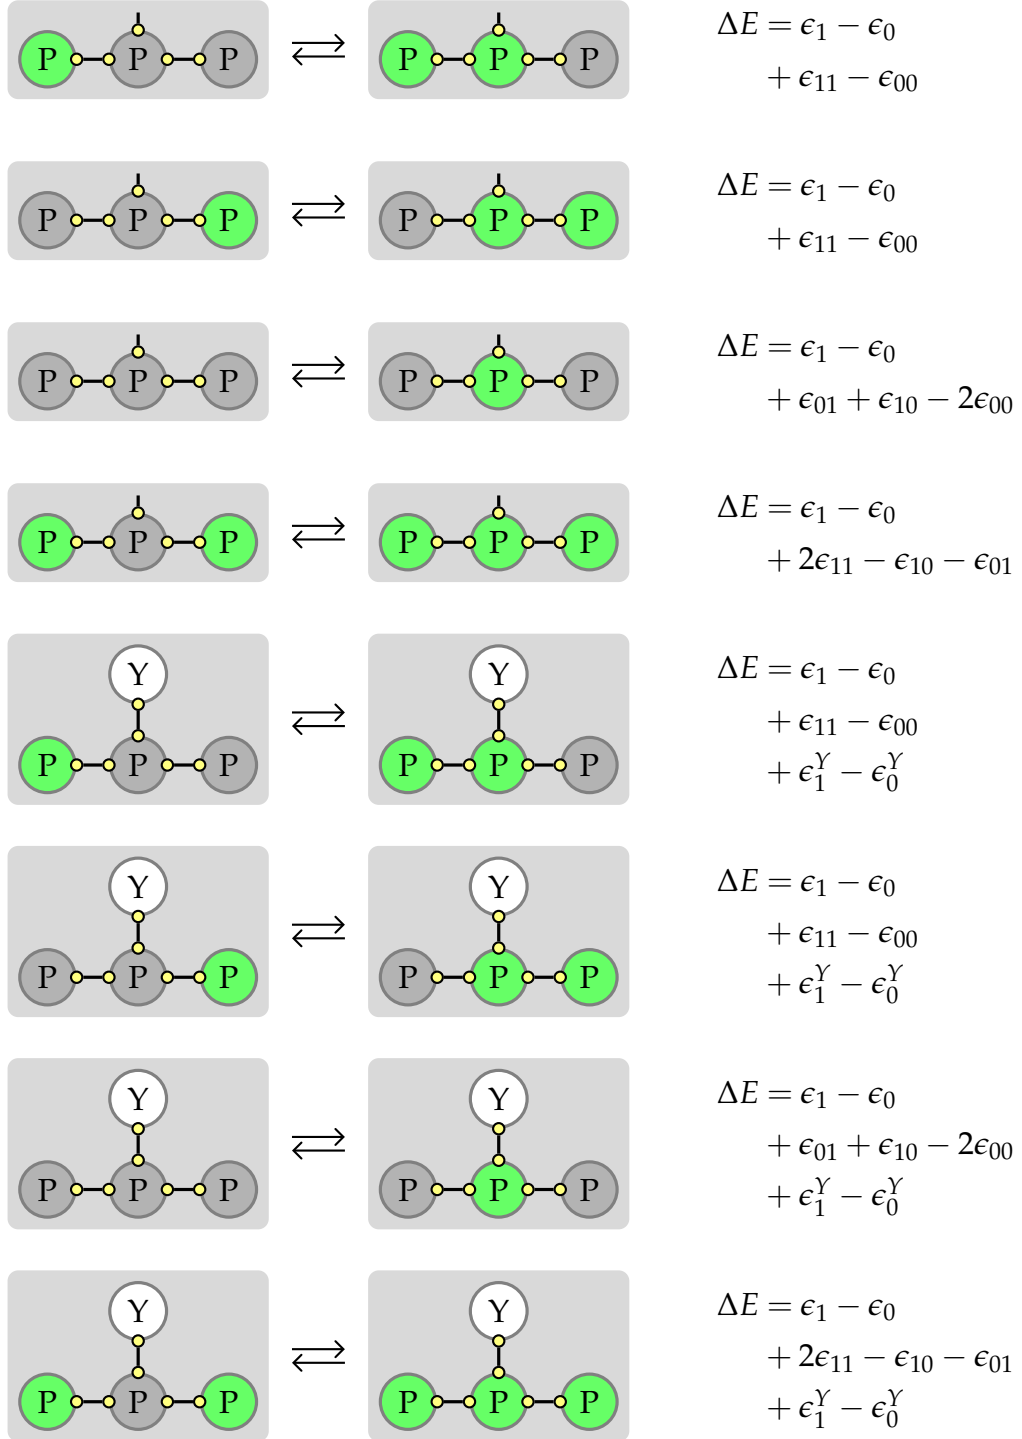


We call the refined rules r_{Y0}^+, r_{Y0}^- and r_{Y1}^+, r_{Y1}^- . Each rule r_{Yi}^+ ($i \in \{0, 1\}$) specifies enough of the context in which it applies to have a definite energy balance $\Delta E = \epsilon_i^Y$. The second pair of rules in \mathcal{G} flip the state of P :



This pair of rules generates many more refinements as changing the state of P will create and destroy matches $\epsilon_{00}, \epsilon_{11}$ and mismatches $\epsilon_{10}, \epsilon_{01}$ between P and its neighbours in the ring. The refinements must then reveal a larger context that includes at least the neighbourhood of P and therefore account for all combinations of neighbours' states. Since the state of the neighbours is not changed when the rule is applied, we do not need to reveal the state of the neighbours' neighbours, which saves us from an infinite recursion of

revelations.⁶ We must also know whether the P that is subject to the action of the rule is bound to a Y as when it is patterns ϵ_0^Y and ϵ_1^Y would be consumed and produced. Hence, the refinements of this second pair of rules are



⁶ Indeed Th. 9 guarantees that such infinite recursions never occur.




Figure 2.3: The simulation steps up the amount of Y (green curve) at $t = 100$ and down again at $t = 200$. This sends the vast majority of the ring into state 1 (orange curve) and then back to state 0 (blue curve). The number of mismatches (purple curve) stays low even during transitions. The parameters for the simulation are $\epsilon_0 = \epsilon_{00} = \epsilon_{11} = -1$, $\epsilon_1 = \epsilon_{01} = \epsilon_{10} = 1$, $\epsilon_0^Y = 2$ and $\epsilon_1^Y = -2$.

In general, if we write i for the state of the left neighbour and j for that of the right neighbour, we have that the energy balance for the first 4 refinements is $\epsilon_{i1} + \epsilon_{1j} - \epsilon_{i0} - \epsilon_{0j} + \epsilon_1 - \epsilon_0$ and for the last 4 is $\epsilon_{i1} + \epsilon_{1j} - \epsilon_{i0} - \epsilon_{0j} + \epsilon_1 - \epsilon_0 + \epsilon_1^Y - \epsilon_0^Y$. As there are 10 pairs of refined rules in total (2 + 8) and only 8 energy patterns, there must be linear dependencies between the various balances. Indeed, the family of vector balances has rank six given by basis vectors ϵ_1^Y , ϵ_0^Y , ϵ_{00} , ϵ_{11} , $\epsilon_{01} + \epsilon_{10}$ and $\epsilon_1 - \epsilon_0$. This example portrays how thermodynamic consistency (i.e. detailed balance) induces relationships between the rates of the refined rules.

It is important to note that the refined rules shown above are those that assume the P s lie on a ring and the ring is fixed, i.e. it does not break. This is true in our model as long as no rule able to form or break bonds between the P s is included in the generator rules and we make sure the initial mixture contains




Figure 2.4: Snapshots of the ring configuration taken at times 50, 150, and 250. At 50 and 250 no Y is bound (because they have not been yet injected into the system or already removed) and the ring is globally in state 0, up to tiny fluctuations. At time 150, it is globally in state 1 as a consequence of the binding of Y s.

no open P -chains. The method of §2.3, which makes no such assumptions, generates many more rules as it takes into account the cases where, for instance, the P that changes state is an end of the chain of protomers.

The final step is to choose concrete rates for our refined rules. We do so by using the symmetric linear kinetic model of Eq. 2.13. In Fig. 2.3 one can see the result of a simulation when Y is stepped up and down again. The model behaves as the one-dimensional cyclic Ising model where the role of the magnetic field is played by Y . Fig. 2.4 shows the state of the simulation before, during and after the injection of Y s. The simulations were run using KaSim with the following model file.

```

1 # agent signatures
2 %agent: P(a,b,c,d~0~1)
3 %agent: Y(p)
4
5 # energy costs
6 %var: '0' -1
7 %var: '1' 1
8 %var: 'Y1' -2
9 %var: 'Y0' 2
10 %var: '00' -1
11 %var: '11' -1
12 %var: '01' 1
13 %var: '10' 1
14

```

```

15 # 2 reversible binding rules
16 'bind_0'    P(d~0,c), Y(p) -> P(d~0,c!1), Y(p!1) \
17             @ [exp] (-1/2 * 'Y0')
18 'unbind_0'  P(d~0,c!1), Y(p!1) -> P(d~0,c), Y(p) \
19             @ [exp] ( 1/2 * 'Y0')
20
21 'bind_1'    P(d~1,c), Y(p) -> P(d~1,c!1), Y(p!1) \
22             @ [exp] (-1/2 * 'Y1')
23 'unbind_1'  P(d~1,c!1), Y(p!1) -> P(d~1,c), Y(p) \
24             @ [exp] ( 1/2 * 'Y1')
25
26 # 8 reversible flipping rules
27 'flip_000'  P(d~0,b!1), P(a!1,d~0,b!2,c), P(a!2,d~0) -> \
28             P(d~0,b!1), P(a!1,d~1,b!2,c), P(a!2,d~0) \
29             @ [exp] (-1/2 * ('1' - '0' + '01' + '10' - 2 * '00'))
30 'flip_010'  P(d~0,b!1), P(a!1,d~1,b!2,c), P(a!2,d~0) -> \
31             P(d~0,b!1), P(a!1,d~0,b!2,c), P(a!2,d~0) \
32             @ [exp] ( 1/2 * ('1' - '0' + '01' + '10' - 2 * '00'))
33
34 'flip_100'  P(d~1,b!1), P(a!1,d~0,b!2,c), P(a!2,d~0) -> \
35             P(d~1,b!1), P(a!1,d~1,b!2,c), P(a!2,d~0) \
36             @ [exp] (-1/2 * ('1' - '0' + '11' - '00'))
37 'flip_110'  P(d~1,b!1), P(a!1,d~1,b!2,c), P(a!2,d~0) -> \
38             P(d~1,b!1), P(a!1,d~0,b!2,c), P(a!2,d~0) \
39             @ [exp] ( 1/2 * ('1' - '0' + '11' - '00'))
40
41 'flip_001'  P(d~0,b!1), P(a!1,d~0,b!2,c), P(a!2,d~1) -> \
42             P(d~0,b!1), P(a!1,d~1,b!2,c), P(a!2,d~1) \
43             @ [exp] (-1/2 * ('1' - '0' + '11' - '00'))
44 'flip_011'  P(d~0,b!1), P(a!1,d~1,b!2,c), P(a!2,d~1) -> \
45             P(d~0,b!1), P(a!1,d~0,b!2,c), P(a!2,d~1) \
46             @ [exp] ( 1/2 * ('1' - '0' + '11' - '00'))
47
48 'flip_101'  P(d~1,b!1), P(a!1,d~0,b!2,c), P(a!2,d~1) -> \
49             P(d~1,b!1), P(a!1,d~1,b!2,c), P(a!2,d~1) \
50             @ [exp] (-1/2 * ('1' - '0' + 2 * '11' - '10' - '01'))
51 'flip_111'  P(d~1,b!1), P(a!1,d~1,b!2,c), P(a!2,d~1) -> \
52             P(d~1,b!1), P(a!1,d~0,b!2,c), P(a!2,d~1) \
53             @ [exp] ( 1/2 * ('1' - '0' + 2 * '11' - '10' - '01'))
54
55 'flip_000_Y' P(d~0,b!1), P(a!1,d~0,b!2,c!_), P(a!2,d~0) -> \
56             P(d~0,b!1), P(a!1,d~1,b!2,c!_), P(a!2,d~0) \

```



```

57      @ [exp] (-1/2 * ('1' - '0' + '01' + '10' - 2 * '00' \
58      + 'Y1' - 'Y0'))
59      'flip_010_Y' P(d~0,b!1), P(a!1,d~1,b!2,c!_), P(a!2,d~0) -> \
60      P(d~0,b!1), P(a!1,d~0,b!2,c!_), P(a!2,d~0) \
61      @ [exp] ( 1/2 * ('1' - '0' + '01' + '10' - 2 * '00' \
62      + 'Y1' - 'Y0'))
63
64      'flip_100_Y' P(d~1,b!1), P(a!1,d~0,b!2,c!_), P(a!2,d~0) -> \
65      P(d~1,b!1), P(a!1,d~1,b!2,c!_), P(a!2,d~0) \
66      @ [exp] (-1/2 * ('1' - '0' + '11' - '00' \
67      + 'Y1' - 'Y0'))
68      'flip_110_Y' P(d~1,b!1), P(a!1,d~1,b!2,c!_), P(a!2,d~0) -> \
69      P(d~1,b!1), P(a!1,d~0,b!2,c!_), P(a!2,d~0) \
70      @ [exp] ( 1/2 * ('1' - '0' + '11' - '00' \
71      + 'Y1' - 'Y0'))
72
73      'flip_001_Y' P(d~0,b!1), P(a!1,d~0,b!2,c!_), P(a!2,d~1) -> \
74      P(d~0,b!1), P(a!1,d~1,b!2,c!_), P(a!2,d~1) \
75      @ [exp] (-1/2 * ('1' - '0' + '11' - '00' \
76      + 'Y1' - 'Y0'))
77      'flip_011_Y' P(d~0,b!1), P(a!1,d~1,b!2,c!_), P(a!2,d~1) -> \
78      P(d~0,b!1), P(a!1,d~0,b!2,c!_), P(a!2,d~1) \
79      @ [exp] ( 1/2 * ('1' - '0' + '11' - '00' \
80      + 'Y1' - 'Y0'))
81
82      'flip_101_Y' P(d~1,b!1), P(a!1,d~0,b!2,c!_), P(a!2,d~1) -> \
83      P(d~1,b!1), P(a!1,d~1,b!2,c!_), P(a!2,d~1) \
84      @ [exp] (-1/2 * ('1' - '0' + 2 * '11' - '10' - '01' \
85      + 'Y1' - 'Y0'))
86      'flip_111_Y' P(d~1,b!1), P(a!1,d~1,b!2,c!_), P(a!2,d~1) -> \
87      P(d~1,b!1), P(a!1,d~0,b!2,c!_), P(a!2,d~1) \
88      @ [exp] ( 1/2 * ('1' - '0' + 2 * '11' - '10' - '01' \
89      + 'Y1' - 'Y0'))
90
91      # P ring
92      %init: 1 (P(a!0 , b!1 ), P(a!1 , b!2 ), P(a!2 , b!3 ), \
93      P(a!3 , b!4 ), P(a!4 , b!5 ), P(a!5 , b!6 ), \
94      P(a!6 , b!7 ), P(a!7 , b!8 ), P(a!8 , b!9 ), \
95      P(a!9 , b!10), P(a!10, b!11), P(a!11, b!12), \
96      P(a!12, b!13), P(a!13, b!14), P(a!14, b!15), \
97      P(a!15, b!16), P(a!16, b!17), P(a!17, b!18), \
98      P(a!18, b!19), P(a!19, b!20), P(a!20, b!21), \

```

```

99          P(a!21, b!22), P(a!22, b!23), P(a!23, b!24), \
100          P(a!24, b!25), P(a!25, b!26), P(a!26, b!27), \
101          P(a!27, b!28), P(a!28, b!29), P(a!29, b!30), \
102          P(a!30, b!31), P(a!31, b!32), P(a!32, b!33), P(a!33, b!0))
103
104 # observables
105 %obs: 'P01' | P(d~0, b!1), P(a!1, d~1) | # 'P10' = 'P01'
106 %obs: 'Y' | Y() |
107 %obs: 'P0' | P(d~0) |
108 %obs: 'P1' | P(d~1) |
109
110 # injection and removal of Ys
111 %var: 'nY' 34
112 %mod: [T] > 100 do $ADD 'nY' Y()
113 %mod: [T] > 200 do $DEL 'nY' Y()
114
115 # snapshots
116 %mod: [T] > 50 do $SNAPSHOT "t50"
117 %mod: [T] > 150 do $SNAPSHOT "t150"
118 %mod: [T] > 250 do $SNAPSHOT "t250"

```

Now we briefly show how the growth policy of §2.3 generates the refinements introduced informally above. First, consider the extensions of the binding rule: only patterns ϵ_i^Y can glue relevantly on it, so the corresponding (unique) site request is for P to reveal its site c and its state (i.e. site d). This gives us the two refinements presented earlier.

Regarding the more interesting extensions of the flipping rule we see that:

- (i) Patterns ϵ_i glue relevantly but do not generate any site request.
- (ii) Patterns ϵ_i^Y asks P to reveal its site c , resulting in two possible extensions: one in which P is bound to a Y and one in which it is free.
- (iii) Patterns ϵ_{ij} can be glued on both sides of P , inducing a request to reveal sites a and b . This results in four possible extensions: a free or bound to a P and the similarly for b .
- (iv) Once a neighbour P has been revealed, patterns ϵ_{ij} induce a further site request, this time on the neighbour P , to reveal its state.

2.8 Non-linear energy functions

At the beginning of this chapter, we made the key assumption in Eq. 2.1 that the energy function is linear in the cost and number of occurrences of energy patterns. Here we consider a more general situation in which the energy function E is no longer asked to be linear. Instead we assume the much weaker property that E can be factored as $v \circ \mathcal{P}(_)$, where $\mathcal{P}(_)$ is the function that counts the number of occurrences of energy patterns in some finite set \mathcal{P} . That is, the energy function is computed by an arbitrary function v on the number of occurrences of energy patterns, not the graph itself. Schematically we have

$$\mathbf{rSGe}_C \xrightarrow{\mathcal{P}(_)} \mathbb{N}^{\mathcal{P}} \xrightarrow{v} \mathbb{R}, \quad (2.17)$$

We can reconstruct Eq. 2.1 by using the linear function $v(x) = \epsilon \cdot x$. As an example of a non-linear energy function, consider the contact graph



and a pair of generator rules r^+, r^- that create/delete the unique edge type. The successive application of these rules can form chains and cycles of arbitrary length. Let us write c_3 for a cycle of length 3 (a triangle) and t_3 for an open chain with 3 nodes. We define a quadratic energy function $E(m) = |[c_3; m]|^2$. In terms of diagram 2.17, we factor E using $\mathcal{P} = \{c_3\}$ and $v(x) = x(c_3)^2$. Applying r^+ to t_3 in a mixture m will create a new copy of c_3 and give the following energy balance:

$$\Delta E = (|[c_3; m]| + 1)^2 - (|[c_3; m]|)^2 = 2|[c_3; m]| + 1 \quad (2.18)$$

Note that the refinement r_ϕ^+ of r^+ that extends the left-hand side of r^+ into t_3 is \mathcal{P} -balanced. As we have seen at the end of §2.2, whenever a rule is \mathcal{P} -balanced the \mathcal{P} -vector $\Delta\phi$ associated to r_ϕ^+ — where each component $\Delta\phi(p)$ is defined as the difference $|[p; n]| - |[p; m]|$ for an r_ϕ^+ -transition from m to n — is the same for all m, n . In the example $\Delta\phi$ has only one component, $\Delta\phi(c_3) = 1$. Despite being \mathcal{P} -balanced and having a constant $\Delta\phi$, Eq. 2.18 shows us that its ΔE is not constant and so detailed balance forces the log-ratio of the backward and forward rates of an edge creation, $\ln(k(r_\phi^-)/k(r_\phi^+))$, to depend on m . This is unlike the case of linear energy functions examined before where the log-ratio is independent of m .

More generally, whenever the refined rule r_ϕ is \mathcal{P} -balanced one can visualise the situation as follows.

$$\begin{array}{ccc}
 m & \xrightarrow{r_\phi} & n \\
 \mathcal{P} \downarrow & & \downarrow \mathcal{P} \\
 \mathcal{P}(m) & \xrightarrow{+\Delta\phi} & \mathcal{P}(n) \\
 v \downarrow & & \downarrow v \\
 \mathbb{R} & \xrightarrow{+\Delta E} & \mathbb{R}
 \end{array}$$

In this setting, detailed balance amounts to asking for

$$K_r := \ln \frac{k(r_{\phi^\star}^\dagger)}{k(r_\phi)} = v(\mathcal{P}(m) + \Delta\phi) - v(\mathcal{P}(m)) \quad (2.19)$$

If v happens to be linear then this is the usual condition $K_r = v(\Delta\phi)$. If v is not linear, detailed balance does not seem very helpful as *a priori* one has to know m to compute the right-hand side. However, since $\Delta\phi$ only depends on r_ϕ , we see that K_r factors through $\mathcal{P}(_)$ just like E and thus does not depend on a full knowledge of m , but only on $\mathcal{P}(m)$. In the example, $K_r = 2|[c_3; x]| + 1$ and $\psi_r(x) = 2x(c_3) + 1$. This is good enough to define rates for r_ϕ . For example, by analogy with the linear kinetic model of §2.5, we can choose log-rates (seen as real-valued functions on $\mathbb{N}^{\mathcal{P}}$) as follows:

$$\ln k(r_\phi) = \alpha_r - \beta_r w_\phi \quad (2.20)$$

with $w_\phi(x) = v(x + \Delta\phi) - v(x)$ and α_r, β_r real-valued functions on $\mathbb{N}^{\mathcal{P}}$ such that $\alpha_{r^\dagger} = \alpha_r$ and $\beta_{r^\dagger} + \beta_r = 1$. This assignment solves the constraint imposed by Eq. 2.19 as $w_{\phi^\star} + w_\phi = 0$.

From the *simulation* point of view, this added generality requires two things: (i) that rates can be made to depend explicitly on observables; (ii) that the internal state of the simulation be extended to incorporate $\mathcal{P}(m)$. Both possibilities are already generically available in the current version of KaSim. A modification of the engine could obtain direct updates to $\mathcal{P}(m)$ as, by assumption, applying r_ϕ leads to a constant $+\Delta\phi$ update; and the same holds for propagating these updates to the rates of the rules which depend on them, e.g. as in Eq. 2.20. Thus, the complexity properties of the simulation algorithm established by Danos, Feret et al. (2007) would be preserved.

We have explained how we can deal with non-linear energy functions that depend on local energy patterns. An interesting extension would be to deal with non-local forms of energies expressing long-range interactions, where the energy is a function of the graph itself. Non-local energy functions, however, would generate many more refined rules, making the simulation of such systems unfeasible unless the simulation algorithm is improved, e.g. by partitioning rules according to energy balances for faster selection. Interesting examples of non-local energy functions include electrostatic interactions like shielded potentials (Kiselev, Marenduzzo and Goryachev, 2011).

Chapter 3

The inverse problem

From rules to energy

In this chapter we would like to explore restricted versions of Kappa for which it is possible to infer the energy function from the rewriting rules and their associated rates. Recall from the introduction that in Kappa itself this problem is undecidable (Danos and Oury, 2010). One such restriction is when agents do not have sites and thus cannot bind. This is Petri nets. We briefly present here the result obtained by Danos and Oury (2013) and show the construction of the energy function for *simple* and *symmetric* Petri nets with *mass action* semantics (sisma Petri net for short).

Definition 15. A sisma Petri net is a Petri net on species Σ and reactions \mathfrak{R} for which:

- (i) (*simple*) there are no two reactions that have the same stoichiometry (net change in species).
- (ii) (*symmetric*) for each reaction $r \in \mathfrak{R}$, there is a reaction $r^\dagger \in \mathfrak{R}$ that has the reverse direction, i.e. the inputs of one are the outputs of the other.
- (iii) (*mass action*) the jumping rate $q_{xy,r}$ of going from a state x to y by a reaction r is proportional to the number of occurrences of its left-hand side in x . In particular, given a rate constant $k(r)$ for reaction r , we have

$$q_{xy,r} = k(r) \prod_{A \in \Sigma} \frac{x(A)!}{(x(A) - \Delta_r(A))!}$$

where $x(A)$ is the number of A s in x and $\Delta_r(A)$ is the net change of A in reaction r .

Note that simple implies no two distinct reactions can be applied to a state x to obtain state y . A sisma Petri net has an energy function

$$E(x) = \sum_{A \in \Sigma} \epsilon(A)x(A) + \ln(x(A)!) \quad (3.1)$$

for some function $\epsilon : \Sigma \rightarrow \mathbb{R}$ such that, for all $r \in \mathfrak{R}$,


$$\sum_{A \in \Sigma} \Delta_r(A)\epsilon(A) = \ln(k(r^*)) - \ln(k(r)).$$

If there is no such function ϵ , the Petri net does not have detailed balance and an energy function. In the rest of the chapter we introduce two other restrictions of Kappa and show how to construct their energy function.

3.1 Cooperative assembly systems

The first restriction is when rules can only create or destroy one edge at a time and their rates can only depend on how many bound sites the endpoints of the edge have. Therefore sites are treated as *indistinguishable*. In addition, agents of the same type cannot bind. Danos, Koepl and Wilson-Kanamori (2011) have proposed these restrictions and a simple formalism incorporating them to study the thermodynamics of polymer formation when there are two types of monomers. Here we extend their result to any number of monomer types.

In the case of two monomers, rules are of the form



where the three grey dots on the sides of each graph are an ellipsis to mean that monomers can be bound to an arbitrary number of monomers of the other type as long as each site is bound only once and there is a finite number of sites per monomer fixed by the monomer type. Hence, the rule schema represents a family of rules indexed by the number of bound sites in the two monomers.

We formalise the ideas in the first paragraph as follows. Let \mathcal{T} be the set of monomer types and $\nu : \mathcal{T} \rightarrow \mathbb{N}$ the map that assigns to each type the number of sites a monomer of that type has, which we refer to as their valence.

A monomer u has type $\tau(u) \in \mathcal{T}$ and degree $d_x(u) \in \mathbb{N}$ in state x . We simply write $v(u)$ for $v(\tau(u))$. The rate constant of a rule that binds a monomer of type $t \in \mathcal{T}$ and degree i with a monomer of type $t' \in \mathcal{T}$ and degree j is $\gamma_{t,i,t',j}^+$. The rate constant of the reverse rule (unbinding) is $\gamma_{t,i,t',j}^-$. The jumping rate q_{xy} from state x to y is then linearly determined by the number of occurrences of the left-hand side of the rule and the rate constant (mass action semantics). We assume that any two agents can be bound only once.

The binding or unbinding of any two nodes u, v in x can only be carried out by one rule, namely the one that operates on degrees $d_x(u), d_x(v)$. The binding rule has rate constant $\alpha(u, v) := \gamma_{\tau(u), d_x(u), \tau(v), d_x(v)}^+$ while the unbinding rate constant is $\beta(u, v) := \gamma_{\tau(u), d_x(u), \tau(v), d_x(v)}^-$. When binding we are free to choose one site among the $v(u) - d_x(u)$ free sites of u and one among the $v(v) - d_x(v)$ free sites of v in order to apply the binding rule. On the other hand, when unbinding we have only one choice, namely removing the only edge between u and v . Hence, q_{xy} is equal to $\alpha(u, v) (v(u) - d_x(u)) (v(v) - d_x(v))$ when the binding rule is applied to x to obtain y and to $\beta(u, v)$ when unbinding.

The theorem below shows under which conditions the type of systems presented in this section have an energy function.

Proposition 1. *Let \mathcal{T} be a finite set of monomer types and $\gamma_{t,i,t',j}^-, \gamma_{t,i,t',j}^+$ families of real values indexed by types $t, t' \in \mathcal{T}$, $0 \leq i < v(t)$ and $0 \leq j < v(t')$ as above. Given a family $\Gamma_{t,i}$ of non-zero real values the following two statements are equivalent*

- (i) *The q -matrix Q as defined above by q_{xy} has detailed balance with respect to the probability distribution π determined by the energy function*

$$E(x) = \sum_{t \in \mathcal{T}} \sum_{0 \leq i \leq v(t)} \epsilon_t(i) x(t_i)$$

where $x(t_i)$ is the number of nodes of type t with degree i in x and

$$\epsilon_t(i) = \sum_{0 \leq j < i} \ln \frac{\Gamma_{t,j}}{v(t) - j}$$

- (ii) *For all $t, t' \in \mathcal{T}$, $0 \leq i < v(t)$ and $0 \leq j < v(t')$ we have*

$$\frac{\gamma_{t,i,t',j}^-}{\gamma_{t,i,t',j}^+} = \Gamma_{t,i} \Gamma_{t',j} \tag{3.2}$$

Proof. (\Rightarrow): Recall the detailed balance condition from Def. 3 which says that for all states x, y

$$\pi_x q_{xy} = \pi_y q_{yx}$$

By substituting π_x and π_y as in Eq. 1.3 we obtain

$$\exp \left[- \sum_{t \in \mathcal{T}} \sum_{0 < i \leq v(t)} \epsilon_t(i) x(t_i) \right] q_{xy} = \exp \left[- \sum_{t \in \mathcal{T}} \sum_{0 < i \leq v(t)} \epsilon_t(i) y(t_i) \right] q_{yx}$$

and by rearranging

$$\prod_{t \in \mathcal{T}} \prod_{0 < i \leq v(t)} e^{\epsilon_t(i) (y(t_i) - x(t_i))} = \frac{q_{yx}}{q_{xy}}$$

When y is obtained from x by binding nodes u, v , the difference $y(t_i) - x(t_i)$ is equal to 0 for all pairs t, i except i) when $t = \tau(u), i = d_x(u)$ or $t = \tau(v), i = d_x(v)$, then $y(t_i) - x(t_i) = -1$; and ii) when $t = \tau(u), i = d_y(u) = d_x(u) + 1$ or $t = \tau(v), i = d_y(v) = d_x(v) + 1$, then $y(t_i) - x(t_i) = 1$. Let $t_u = \tau(u), t_v = \tau(v), d_u = d_x(u)$ and $d_v = d_x(v)$. It follows that the last equation can be rewritten as

$$\exp [\epsilon_{t_u}(d_u + 1) + \epsilon_{t_v}(d_v + 1) - \epsilon_{t_u}(d_u) - \epsilon_{t_v}(d_v)] = \frac{q_{yx}}{q_{xy}}$$

By substituting ϵ we get

$$\frac{\prod_{0 \leq i < d_u + 1} \frac{\Gamma_{t_u, i}}{v(u) - i} \prod_{0 \leq i < d_v + 1} \frac{\Gamma_{t_v, i}}{v(v) - i}}{\prod_{0 \leq i < d_u} \frac{\Gamma_{t_u, i}}{v(u) - i} \prod_{0 \leq i < d_v} \frac{\Gamma_{t_v, i}}{v(v) - i}} = \frac{q_{yx}}{q_{xy}}$$

Products on the left cancel out and yield, after substituting q on the right,

$$\frac{\Gamma_{t_u, d_u} \Gamma_{t_v, d_v}}{(v(u) - d_u) (v(v) - d_v)} = \frac{\beta(u, v)}{\alpha(u, v) (v(u) - d_u) (v(v) - d_v)}$$

which then simplifies to

$$\Gamma_{t_u, d_u} \Gamma_{t_v, d_v} = \frac{\gamma_{t_u, d_u, t_v, d_v}^-}{\gamma_{t_u, d_u, t_v, d_v}^+}$$


This equality holds in general for nodes of any degree and type.

(\Leftarrow): We prove that, whenever (ii) holds, π verifies the detailed balance condition. For all x, y such that $q_{xy} = 0$ the equality $\pi_x q_{xy} = \pi_y q_{yx}$ holds as rules are reversible and (ii) dictates that a rate constant is zero if the reverse rate constant is. When $q_{xy} > 0$ then y can be obtained from x by binding or unbinding some nodes u, v . By substituting t for $\tau(u)$, t' for $\tau(v)$, i for $d_x(u)$ and j for $d_x(v)$ in Eq. 3.2 we obtain the last equation in the first part of the proof. We can replay the transformations backwards to obtain $\pi_x q_{xy} = \pi_y q_{yx}$ when y is obtained by binding. The case of unbinding follows a similar argument. ■

3.2 Flipping and binding

Now we look at systems whose nodes have sites that possess an internal state. This internal state is used to decide when to bind other nodes or change the internal state of other sites. For simplicity, internal states can take one of only two possible values. Unlike cooperative assembly systems, here sites are *distinguishable* as in Kappa. Hence, we extend contact maps g as defined in §1.2 with a map δ_g that assigns an internal state vector $\delta_g(u)$ in $\{0,1\}^{\sigma_C^{-1}(g_{\mathcal{A}}(u))}$ to agents u in $\mathcal{A}_{|g|}$. The internal state vector is indexed by the sites of the agent type $g_{\mathcal{A}}(u)$ of u . We use these extended contact maps as graphs in this section.


A site's internal state can be changed by rules we call *flips*. The rate at which we flip a site may depend on the type of the site and the node it belongs to, the internal state of sites on the same node and the type of the neighbours. Note that it cannot depend on the internal states of the neighbours' sites or the nodes they are bound to. Graphically, flips are of the form



where the dotted lines denote an optional node or edge and site x changes state from white to black. As usual, we write r_L for the left-hand side contact map of rule r and r_R for that of the right-hand side, both contact maps over some fixed contact graph C .

In a flip we have a complete view over the internal state of sites in u and those of the neighbours, which we characterise as vectors indexed by site types in $I := \sigma_C^{-1}(r_{L,\mathcal{A}}(u))$. The rate constants of the forward and backward flip rules are then parametrised by the agent type $a := r_{L,\mathcal{A}}(u) = r_{R,\mathcal{A}}(u)$ of u in C , the site type $i := r_{L,\mathcal{S}}(x) = r_{R,\mathcal{S}}(x)$ of x , the internal state vector $\mathbf{s} \in \{0,1\}^I$ of u , and the binding state vector $\mathbf{n} \in ((\mathcal{A}_C \times \mathcal{S}_C) \cup \{\star\})^I$ where \star is used to denote a free site. We write $\lambda_{a,i,\mathbf{s},\mathbf{n}}^+$ for the rate constant of the forward rule, $\lambda_{a,i,\mathbf{s},\mathbf{n}}^-$ for that of the backward rule, and $\Lambda_{a,i,\mathbf{s},\mathbf{n}} = \lambda_{a,i,\mathbf{s},\mathbf{n}}^- / \lambda_{a,i,\mathbf{s},\mathbf{n}}^+$ for their ratio.

The second type of rules we allow are *binds*. They are of the form




where the three grey dots are an ellipsis meaning that nodes u and v must declare an internal state to all sites they have. The rate constant of the binding and unbinding rules depends on the internal state of the sites on the participating nodes \mathbf{s} and \mathbf{s}' as well as the type of the nodes a, b and the bound sites i, j . We write $\Gamma_{a,i,\mathbf{s},b,j,\mathbf{s}'} = \gamma_{a,i,\mathbf{s},b,j,\mathbf{s}'}^- / \gamma_{a,i,\mathbf{s},b,j,\mathbf{s}'}^+$ for the ratio between the backward and forward rate constants of binds.

We refer to systems composed by flips and binds as *flip-bind systems* or FB-systems for short. We will show how the energy function of an FB-system looks like. But first, we prove two lemmas that show us how detailed balance fixes the value of some ratios of rate constants, reducing so the number of free parameters in the system. To simplify notation in the following lemmas, let $\mathbf{s} + i$ be the vector \mathbf{s} with site i flipped.

Lemma 11. *Let C be a contact graph. For all agent types $a \in \mathcal{A}_C$, let $I = \sigma_C^{-1}(a)$, and for all site types $i \in I$, internal state vectors $\mathbf{s} \in \{0,1\}^I$ and binding state vectors $\mathbf{n} \in ((\mathcal{A}_C \times \mathcal{S}_C) \cup \{\star\})^I$, an FB-system with detailed balance verifies*

$$\Lambda_{a,i,\mathbf{s},\mathbf{n}} \Lambda_{a,j,\mathbf{s}+i,\mathbf{n}} = \Lambda_{a,j,\mathbf{s},\mathbf{n}} \Lambda_{a,i,\mathbf{s}+j,\mathbf{n}} \quad (3.3)$$

Proof. Pick a state x and a node u in x . We can find a square in the transition graph with 4 flips starting from x : flip site i first then j and flip j first then i .



By detailed balance we have that the product of rates along a cycle must be equal to 1. Hence, starting from x and going through the cycle in one direction and the reverse we get

$$\lambda_{a,i,\mathbf{s},\mathbf{n}}^+ \lambda_{a,j,\mathbf{s}+i,\mathbf{n}}^+ \lambda_{a,i,\mathbf{s}+j,\mathbf{n}}^- \lambda_{a,j,\mathbf{s},\mathbf{n}}^- = \lambda_{a,j,\mathbf{s},\mathbf{n}}^+ \lambda_{a,i,\mathbf{s}+j,\mathbf{n}}^+ \lambda_{a,j,\mathbf{s}+i,\mathbf{n}}^- \lambda_{a,i,\mathbf{s},\mathbf{n}}^-$$

By rearranging we obtain Eq. 3.3. ■

When we consider all flips for a node with n sites, we find an n -hypercube in the transition graph. This hypercube has $2^{n-1}n$ edges where an edge corresponds to a pair of forward and backward flips. It follows that there are the same number of Λ ratios. In addition, each face of the hypercube generates an equation by Lemma 11 and there are $2^{n-3}(n-1)n$ faces in an n -hypercube. This is a severe constraint on the number of parameters that can be freely set in an FB-system with detailed balance.

A further constrain to the values of the rate constant ratios of flips and bindings can be obtained and its proved in the following lemma. For this lemma we will use a fixed total order $<_a$ on the sites of an agent type a in the contact graph C .

Lemma 12. *Let C be a contact graph. For all agent types $a \in \mathcal{A}_C$, let $I = \sigma_C^{-1}(a)$, and for all site types $i \in I$, binding state vectors $\mathbf{n} \in ((\mathcal{A}_C \times \mathcal{S}_C) \cup \{\star\})^I$ and internal state vectors $\mathbf{s} \in \{0,1\}^I$, $\mathbf{s}_j \in \{0,1\}^{\sigma_C^{-1}(b)}$ with $j \in I$ and b the agent type in $\mathbf{n}(j)$, an FB-system with detailed balance verifies*

$$\frac{\Lambda_{a,i,\mathbf{s},\mathbf{n}}}{\Lambda_{a,i,\mathbf{s},\emptyset}} = \prod_{\substack{i' \in I \\ \mathbf{n}(i') \neq \star \\ (b,j) := \mathbf{n}(i')}} \frac{\Gamma_{a,i',\mathbf{s},b,j,\mathbf{s}_{i'}}}{\Gamma_{a,i',\mathbf{s}+i,b,j,\mathbf{s}_{i'}}} \quad (3.4)$$

where \emptyset a vector of free sites, i.e. $\emptyset(i) = \star$ for all $i \in I$.

Proof. We construct a series of squares with 2 flips and 2 bindings each. Pick a state x and a node u in x . Strip u of all its neighbours and call that state x_0 . Choose a site i of u . The first square starts from x_0 , then i) flips site i and ii) binds the smallest site i_1 according to the order $<_{x_A(u)}$ to site j of an agent of type b that has internal vector state \mathbf{s}_{i_1} , where $(j,b) = \mathbf{n}(i_1)$. After performing (i) then (ii) or (ii) then (i) we reach state x_1 .

$$\begin{array}{ccccccc} x_0 & \xrightleftharpoons[\gamma_{a,i_1,\mathbf{s},b,j,\mathbf{s}_{i_1}}^-]{\gamma_{a,i_1,\mathbf{s},b,j,\mathbf{s}_{i_1}}^+} & \hat{x}_1 & \xrightleftharpoons[\dots]{\dots} & x_n & \xrightleftharpoons[\gamma_{a,i_n,\mathbf{s},b',j',\mathbf{s}_{i_n}}^-]{\gamma_{a,i_n,\mathbf{s},b',j',\mathbf{s}_{i_n}}^+} & \hat{x}_{n+1} \\ \uparrow \lambda_{a,i,\mathbf{s},\emptyset}^- & \lambda_{a,i,\mathbf{s},\emptyset}^+ & \uparrow \lambda_{a,i,\mathbf{s},\mathbf{n}_1}^- & \lambda_{a,i,\mathbf{s},\mathbf{n}_1}^+ & \uparrow \lambda_{a,i,\mathbf{s},\mathbf{n}_n}^- & \lambda_{a,i,\mathbf{s},\mathbf{n}_{n+1}}^+ & \uparrow \lambda_{a,i,\mathbf{s},\mathbf{n}_{n+1}}^+ \\ \hat{x}_0 & \xrightleftharpoons[\gamma_{a,i_1,\mathbf{s}+i,b,j,\mathbf{s}_{i_1}}^-]{\gamma_{a,i_1,\mathbf{s}+i,b,j,\mathbf{s}_{i_1}}^+} & x_1 & \xrightleftharpoons[\dots]{\dots} & \hat{x}_n & \xrightleftharpoons[\gamma_{a,i_n,\mathbf{s}+i,b',j',\mathbf{s}_{i_n}}^-]{\gamma_{a,i_n,\mathbf{s}+i,b',j',\mathbf{s}_{i_n}}^+} & x_{n+1} \end{array}$$

with $\mathbf{n}_1(i) = \star$ for all i except i_1 , where $\mathbf{n}_1(i_1) = \mathbf{n}(i_1)$. In general, $\mathbf{n}_n(i)$ is equal to $\mathbf{n}(i)$ if $i \leq_{x_A(u)} i_n$ and \star otherwise.

By detailed balance we obtain the following relation for the first square

$$\Lambda_{a,i,s,\emptyset} \Gamma_{a,i_1,s+i,b,j,s_{i_1}} = \Gamma_{a,i_1,s,b,j,s_{i_1}} \Lambda_{a,i,s,n_1} \quad (3.5)$$

We construct the n th square starting from x_n by flipping site i and binding site i_n to site j' in an agent of type b' as indicated by $\mathbf{n}(i_n)$. This neighbour has an internal state vector \mathbf{s}_{i_n} . Again, by detailed balance we get

$$\Lambda_{a,i,s,n_n} \Gamma_{a,i,s+i,b',j',s_{i_n}} = \Gamma_{a,i,s,b',j',s_{i_n}} \Lambda_{a,i,s,n_{n+1}} \quad (3.6)$$

Eq. 3.5 can be rewritten as

$$\frac{\Lambda_{a,i,s,n_1}}{\Lambda_{a,i,s,\emptyset}} = \frac{\Gamma_{a,i,s+i,b,j,s_{i_1}}}{\Gamma_{a,i,s,b,j,s_{i_1}}}$$

Then we substitute Λ_{a,i,s,n_1} according to Eq. 3.6 and obtain

$$\frac{\Lambda_{a,i,s,n_2}}{\Lambda_{a,i,s,\emptyset}} = \frac{\Gamma_{a,i,s+i,b,j,s_{i_1}}}{\Gamma_{a,i,s,b,j,s_{i_1}}} \frac{\Gamma_{a,i,s+i,c,k,s_{i_2}}}{\Gamma_{a,i,s,c,k,s_{i_2}}}$$

We repeat until we recover Eq. 3.4. ■

Now we proceed to compute the energy function.

Proposition 2. *Given an FB-system with detailed balance, its energy function is*

$$E(x) = \sum_{s \in \mathcal{S}_{|x|}} \ln \Lambda_{a,i,s_u(i),\emptyset} + \sum_{(s,t) \in \mathcal{E}_{|x|}} \ln \Gamma_{a,i,\delta_x(u),b,j,\delta_x(v)} \quad (3.7)$$

where $u = \sigma_{|x|}(s)$, $v = \sigma_{|x|}(t)$, $a = x_{\mathcal{A}}(u)$, $b = x_{\mathcal{A}}(v)$, $i = x_{\mathcal{S}}(s)$, $j = x_{\mathcal{S}}(t)$, and $\mathbf{s}_w : \mathcal{S}_C \rightarrow \{0,1\}^I$ a family of functions indexed by $w \in \mathcal{A}_{|x|}$, with $I = \sigma_C^{-1}(x_{\mathcal{A}}(w))$, defined as

$$\mathbf{s}_u(i)(j) = \begin{cases} \delta_x(u)(j) & \text{if } j <_{x_{\mathcal{A}}(u)} i \\ 0 & \text{otherwise} \end{cases}$$

Proof. Let x_0 be the state where all nodes are disconnected and the internal state of their sites is set to 0. We set $E(x_0) = 0$ and construct a canonical path from x_0 to x to assign an energy $E(x)$ to a state x as the sum of the energy contributions of each step along the path. The canonical path starts by performing all flips first (in the order set by $<_a$) and then all bindings. Using $E(y) - E(x) = \ln q(y, x) - \ln q(x, y)$ to compute the energy contribution of each flip and bind, we obtain Eq. 3.7. ■

Finally, we would like to know when an FB-system has detailed balance.

Proposition 3. *An FB-system has detailed balance if and only if Eq. 3.3 and 3.4 hold for all states, nodes and sites in the system.*

Proof. By Lemmas 11 and 12, the forward direction has been already proved. The backward direction amounts to proving that the definition of Eq. 3.7 is independent of the choice of path, i.e. that any alternative energy function E' that uses a non-canonical path for its definition is equivalent to E as defined in Prop. 2. In a non-canonical path we might have two types of non-canonical flips: those that occur after a bind and those that happen in an order different to the one specified by $<_a$. The terms contributed by the latter can be transformed into canonical ones using Eq. 3.3 and those contributed by the former using Eq. 3.4. Non-canonical binds that occur before a flip can be transformed into canonical binds by a similar argument. ■

Bibliography

- Anderson, William J. (1991). *Continuous-Time Markov Chains. An Applications-Oriented Approach*. Springer Series in Statistics. Springer-Verlag New York. ISBN: 978-1-4612-7772-9 978-1-4612-3038-0. DOI: 10.1007/978-1-4612-3038-0.
- Bader, Gary D., Ian Donaldson, Cheryl Wolting, B. F. Francis Ouellette, Tony Pawson and Christopher W. V. Hogue (2001). 'BIND—The Biomolecular Interaction Network Database'. In: *Nucleic Acids Research* 29.1, pp. 242–245. ISSN: 0305-1048, 1362-4962. DOI: 10.1093/nar/29.1.242.
- Bai, Fan et al. (2010). 'Conformational Spread as a Mechanism for Cooperativity in the Bacterial Flagellar Switch'. In: *Science* 327.5966, pp. 685–689. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1182105.
- Bintu, Lacramioara et al. (2005). 'Transcriptional regulation by the numbers: applications'. In: *Current opinion in genetics & development* 15.2, pp. 125–135. ISSN: 0959-437X. DOI: 10.1016/j.gde.2005.02.006.
- Boutillier, Pierre, Jérôme Feret, Jean Krivine and Lý Kim Quyên (2014). *KaSim and KaSa Reference Manual*. URL: https://kappa-dev.github.io/docs/KaSim-manual-master/KaSim_manual.htm.
- Cannon, William R. (2014). 'Simulating Metabolism with Statistical Thermodynamics'. In: *PLoS One* 9.8, e103582. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0103582.
- Chautard, Emilie, Lionel Ballut, Nicolas Thierry-Mieg and Sylvie Ricard-Blum (2009). 'MatrixDB, a database focused on extracellular protein–protein and protein–carbohydrate interactions'. In: *Bioinformatics* 25.5, pp. 690–691. ISSN: 1367-4803. DOI: 10.1093/bioinformatics/btp025.
- Chen, X., M. Liu and Michael K. Gilson (2001). 'BindingDB: A Web-Accessible Molecular Recognition Database'. In: *Combinatorial Chemistry & High Throughput Screening* 4.8, pp. 719–725. DOI: 10.2174/1386207013330670.
- Danos, Vincent, Jérôme Feret, Walter Fontana and Jean Krivine (2007). 'Scalable simulation of cellular signaling networks'. In: *Asian Symposium on Programming Languages and Systems*, pp. 139–157.
- Danos, Vincent, Russ Harmer and Glynn Winskel (2013). 'Constraining rule-based dynamics with types'. In: *Mathematical Structures in Computer Science* 23.2, pp. 272–289.
- Danos, Vincent, Heinz Koeppl and John Wilson-Kanamori (2011). 'Cooperative assembly systems'. In: *DNA Computing and Molecular Programming*, pp. 1–20.

- Danos, Vincent and Cosimo Laneve (2004). 'Formal molecular biology'. In: *Theoretical Computer Science. Computational Systems Biology* 325.1, pp. 69–110. ISSN: 0304-3975. DOI: 10.1016/j.tcs.2004.03.065.
- Danos, Vincent and Nicolas Oury (2010). 'Equilibrium and Termination'. In: *Proceedings Sixth Workshop on Developments in Computational Models: Causality, Computation, and Physics*. Ed. by S. Barry Cooper, Prakash Panangaden and Elham Kashefi. Vol. 26. EPTCS, pp. 75–84. DOI: 10.4204/EPTCS.26.7.
- (2013). 'Equilibrium and Termination II: the case of Petri Nets'. In: *Mathematical Structures in Computer Science* 23.2, pp. 290–307.
- Delbrück, Max (1940). 'Statistical Fluctuations in Autocatalytic Reactions'. In: *The Journal of Chemical Physics* 8.1, pp. 120–124. ISSN: 0021-9606, 1089-7690. DOI: 10.1063/1.1750549.
- Diaconis, Persi and Laurent Miclo (2009). 'On characterizations of Metropolis type algorithms in continuous time'. In: *ALEA : Latin American Journal of Probability and Mathematical Statistics* 6, pp. 199–238. URL: <https://hal.archives-ouvertes.fr/hal-00353845> (visited on 08/11/2016).
- Diers, Y. (1978). 'Familles universelles de morphismes'. In: *Tech. report, Université des Sciences et Techniques de Lille I*.
- Doob, J. L. (1937). 'Stochastic processes depending on a continuous parameter'. In: *Transactions of the American Mathematical Society* 42.1, pp. 107–140. ISSN: 0002-9947, 1088-6850. DOI: 10.1090/S0002-9947-1937-1501916-1.
- Einstein, Albert (1905). 'Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen'. In: *Annalen der Physik* 322, pp. 549–560. DOI: 10.1002/andp.19053220806.
- Evans, Denis J. and Debra J. Searles (2002). 'The Fluctuation Theorem'. In: *Advances in Physics* 51.7, pp. 1529–1585. ISSN: 0001-8732. DOI: 10.1080/00018730210155133.
- Feller, Willy (1940). 'On the integro-differential equations of purely discontinuous Markoff processes'. In: *Transactions of the American Mathematical Society* 48.3, pp. 488–515. ISSN: 0002-9947, 1088-6850. DOI: 10.1090/S0002-9947-1940-0002697-3.
- Flamholz, Avi, Elad Noor, Arren Bar-Even and Ron Milo (2012). 'eQuilibrator—the biochemical thermodynamics calculator'. In: *Nucleic Acids Research* 40 (D1), pp. D770–D775. ISSN: 0305-1048, 1362-4962. DOI: 10.1093/nar/gkr874.
- Gillespie, Daniel T (1976). 'A general method for numerically simulating the stochastic time evolution of coupled chemical reactions'. In: *Journal of Computational Physics* 22.4, pp. 403–434. ISSN: 0021-9991. DOI: 10.1016/0021-9991(76)90041-3.
- Gilson, Michael K., Tiqing Liu, Michael Baitaluk, George Nicola, Linda Hwang and Jenny Chong (2016). 'BindingDB in 2015: A public database for medicinal chemistry, computational chemistry and systems pharmacology'. In: *Nucleic Acids Research* 44 (D1), pp. D1045–D1053. ISSN: 0305-1048, 1362-4962. DOI: 10.1093/nar/gkv1072.
- Goldberg, Robert N., Yadu B. Tewari and Talapady N. Bhat (2004). 'Thermodynamics of enzyme-catalyzed reactions—a database for quantitative

- biochemistry'. In: *Bioinformatics* 20.16, pp. 2874–2877. ISSN: 1367-4803. DOI: 10.1093/bioinformatics/bth314.
- Gorban, A. N., E. M. Mirkes and G. S. Yablonsky (2013). 'Thermodynamics in the limit of irreversible reactions'. In: *Physica A: Statistical Mechanics and its Applications* 392.6, pp. 1318–1335. ISSN: 0378-4371. DOI: 10.1016/j.physa.2012.10.009.
- Hamilton, Sir William Rowan (1834). 'On a General Method in Dynamics; by which the Study of the Motions of all free Systems of attracting or repelling Points is reduced to the Search and Differentiation of one central Relation, or characteristic Function'. In: *Philosophical Transactions of the Royal Society* (part II for 1834), pp. 247–308.
- Hastings, W. K. (1970). 'Monte Carlo sampling methods using Markov chains and their applications'. In: *Biometrika* 57.1, pp. 97–109. ISSN: 0006-3444, 1464-3510. DOI: 10.1093/biomet/57.1.97.
- Helmholtz, Hermann von (1887). 'Über die physikalische Bedeutung des Prinzips der kleinsten Wirkung'. In: *Journal für die reine und angewandte Mathematik*, pp. 137–166.
- Hermjakob, Henning et al. (2004). 'The HUPO PSI's molecular interaction format—a community standard for the representation of protein interaction data'. In: *Nature Biotechnology* 22.2, pp. 177–183. ISSN: 1087-0156. DOI: 10.1038/nbt926.
- Kerrien, Samuel et al. (2007). 'Broadening the horizon – level 2.5 of the HUPO-PSI format for molecular interactions'. In: *BMC Biology* 5, p. 44. ISSN: 1741-7007. DOI: 10.1186/1741-7007-5-44.
- Kiselev, Vladimir Yu., Davide Marenduzzo and Andrew B. Goryachev (2011). 'Lateral Dynamics of Proteins with Polybasic Domain on Anionic Membranes: A Dynamic Monte-Carlo Study'. In: *Biophysical Journal* 100.5, pp. 1261–1270. ISSN: 0006-3495. DOI: 10.1016/j.bpj.2011.01.025.
- Kolmogorov, Andrey (1933). *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Berlin: Julius Springer.
- Lagrange, Joseph-Louis (1811). *Mécanique analytique*. II. Courcier, Paris.
- Lan, Ganhui, Sonja Schulmeister, Victor Sourjik and Yuhai Tu (2011). 'Adapt locally and act globally: strategy to maintain high chemoreceptor sensitivity in complex environments'. In: *Molecular Systems Biology* 7.1, p. 475. ISSN: 1744-4292. DOI: 10.1038/msb.2011.8.
- Leifer, James J. and Robin Milner (2000). 'Deriving Bisimulation Congruences for Reactive Systems'. In: *CONCUR*. Ed. by Catuscia Palamidessi. Vol. 1877. Lecture Notes in Computer Science. Springer, pp. 243–258. ISBN: 3-540-67897-2. DOI: 10.1007/3-540-44618-4_19.
- Lewis, Gilbert N. (1925). 'A new principle of equilibrium'. In: *Proceedings of the National Academy of Sciences* 11.3, pp. 179–183. ISSN: 0027-8424.
- Li, Xin, Fan Wu, Feng Qi and Daniel A. Beard (2011). 'A database of thermodynamic properties of the reactions of glycolysis, the tricarboxylic acid cycle, and the pentose phosphate pathway'. In: *Database: The Journal of Biological Databases and Curation* 2011, bar005. ISSN: 1758-0463. DOI: 10.1093/database/bar005.

- Licata, Luana et al. (2012). 'MINT, the molecular interaction database: 2012 update'. In: *Nucleic Acids Research* 40 (Database issue), pp. D857–861. ISSN: 1362-4962. DOI: 10.1093/nar/gkr930.
- Mavrovouniotis, M. L. (1991). 'Estimation of standard Gibbs energy changes of biotransformations.' In: *Journal of Biological Chemistry* 266.22, pp. 14440–14445. ISSN: 0021-9258, 1083-351X. URL: <http://www.jbc.org/content/266/22/14440> (visited on 08/01/2017).
- McQuarrie, Donald A. (1967). 'Stochastic Approach to Chemical Kinetics'. In: *Journal of Applied Probability* 4.3, pp. 413–478. ISSN: 0021-9002. DOI: 10.2307/3212214.
- (2000). *Statistical mechanics*. OCLC: 43370175. Sausalito, Calif.: University Science Books. ISBN: 978-1-891389-15-3.
- Metropolis, Nicholas, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller and Edward Teller (1953). 'Equation of State Calculations by Fast Computing Machines'. In: *The Journal of Chemical Physics* 21.6, pp. 1087–1092. ISSN: 0021-9606, 1089-7690. DOI: 10.1063/1.1699114.
- Murphy, Elaine, Vincent Danos, Jérôme F  ret, Jean Krivine and Russell Harmer (2010). 'Rule-Based Modeling and Model Refinement'. In: *Elements of Computational Systems Biology*. Ed. by Huma M. Lodhi and Stephen H. Muggleton. John Wiley & Sons, Inc., pp. 83–114. ISBN: 978-0-470-55675-7.
- Newton, Isaac (1687). *Philosophi   naturalis principia mathematica*. Londini, Jussu Societatis Regi   ac Typis Josephi Streater. Prostat apud plures Bibliopolas.
- Ollivier, Julien F., Vahid Shahrezaei and Peter S. Swain (2010). 'Scalable rule-based modelling of allosteric proteins and biochemical networks'. In: *PLoS Computational Biology* 6.11, e1000975. ISSN: 1553-7358. DOI: 10.1371/journal.pcbi.1000975.
- Orchard, Sandra et al. (2012). 'Protein interaction data curation: the International Molecular Exchange (IMEx) consortium'. In: *Nature Methods* 9.4, pp. 345–350. ISSN: 1548-7091. DOI: 10.1038/nmeth.1931.
- Pascal, Robert, Addy Pross and John D. Sutherland (2013). 'Towards an evolutionary theory of the origin of life based on kinetics and thermodynamics'. In: *Open Biology* 3.11, p. 130156. ISSN: 2046-2441. DOI: 10.1098/rsob.130156.
- Perrin, Jean (1908). 'L'agitation mol  culaire et le mouvement brownien'. In: *Comptes rendus hebdomadaires des s  ances de l'acad  mie des sciences* 146, pp. 967–970.
- Petri, Carl Adam (1962). 'Kommunikation mit Automaten'. PhD thesis. University of Bonn.
- Post, Emil L. (1946). 'A variant of a recursively unsolvable problem'. In: *Bulletin of the American Mathematical Society* 52.4, pp. 264–268. ISSN: 0002-9904, 1936-881X. DOI: 10.1090/S0002-9904-1946-08555-9.
- Saiz, Leonor and Jose M. G. Vilar (2006). 'Stochastic dynamics of macromolecular-assembly networks'. In: *Molecular Systems Biology* 2, 2006.0024. ISSN: 1744-4292. DOI: 10.1038/msb4100061.
- Santilli, Ruggero Maria (1978). *Foundations of Theoretical Mechanics I: The Inverse Problem in Newtonian Mechanics*. New York: Springer-Verlag.

- Thomas, Dennis G., Sebastian Jaramillo-Riveri, Douglas J. Baxter and William R. Cannon (2014). 'Comparison of Optimal Thermodynamic Models of the Tricarboxylic Acid Cycle from Heterotrophs, Cyanobacteria, and Green Sulfur Bacteria'. In: *The Journal of Physical Chemistry B* 118.51, pp. 14745–14760. ISSN: 1520-6106. DOI: 10.1021/jp5075913.
- Tolman, Richard C. (1925). 'The Principle of Microscopic Reversibility'. In: *Proceedings of the National Academy of Sciences of the United States of America* 11.7, pp. 436–439. ISSN: 0027-8424.
- Wang, G. M., E. M. Sevick, Emil Mittag, Debra J. Searles and Denis J. Evans (2002). 'Experimental Demonstration of Violations of the Second Law of Thermodynamics for Small Systems and Short Time Scales'. In: *Physical Review Letters* 89.5, p. 050601. DOI: 10.1103/PhysRevLett.89.050601.
- Wang, Hui, Ned S. Wingreen and Ranjan Mukhopadhyay (2008). 'Self-Organized Periodicity of Protein Clusters in Growing Bacteria'. In: *Physical Review Letters* 101.21, p. 218101. DOI: 10.1103/PhysRevLett.101.218101.
- Wegscheider, Rudolf (1901). 'Über simultane Gleichgewichte und die Beziehungen zwischen Thermodynamik und Reaktionskinetik homogener Systeme'. In: *Monatshefte für Chemie und verwandte Teile anderer Wissenschaften* 22.8, pp. 849–906. ISSN: 0343-7329, 1434-4475. DOI: 10.1007/BF01517498.
- Yang, Jianyi, Ambrish Roy and Yang Zhang (2013). 'BioLiP: a semi-manually curated database for biologically relevant ligand–protein interactions'. In: *Nucleic Acids Research* 41 (D1), pp. D1096–D1103. ISSN: 0305-1048, 1362-4962. DOI: 10.1093/nar/gks966.
- Zanzoni, Andreas, Luisa Montecchi-Palazzi, Michele Quondam, Gabriele Ausiello, Manuela Helmer-Citterich and Gianni Cesareni (2002). 'MINT: a Molecular INTeraction database'. In: *FEBS Letters* 513.1, pp. 135–140. ISSN: 1873-3468. DOI: 10.1016/S0014-5793(01)03293-8.