

Equilibrium and Termination II: the case of Petri nets

Vincent Danos[†] and Nicolas Oury

School of Informatics, University of Edinburgh

Received 3 November 2011

This paper is concerned with the asymptotic properties of a restricted class of Petri nets equipped with stochastic mass action semantics. We establish a simple algebraic criterion for the existence of an *equilibrium*, that is to say an invariant probability that satisfies the *detailed balance* condition familiar from the thermodynamics of reaction networks. We also find that when such a probability exists, it can be described by a free energy function which combines an internal energy term and an entropy one. Under strong additional conditions, we show how the entropy term can be deconstructed using the finer-grained individual-token semantics of Petri nets.

1. Introduction

Markov chains (MC) and differential equations (DE) have been the fundamental means of describing dynamical systems for a very long time. But today, as quantitative modeling efforts try to address decentralized dynamics of increasing complexity and connectedness (Easley and Kleinberg(2010)), a direct usage of such basic descriptive tools is becoming unwieldy. New modeling situations present a diversity of structures and scales where the representational challenge is insurmountable if one does not use indirect, abstract and structured syntaxes to describe the dynamics of interest (Epstein(1999)). Thus, part of the modeling activity has to morph into a sort of domain-specific programming, where MC/DEs play the role of an assembly language best not written by hand - or considered extensionally at all.

A case in point, which has drawn considerable attention in the past decade and will be the focus of this paper, is the modeling of what is arguably the complex distributed system *par excellence*, namely the combinatorial processes at work in biomolecular networks (BMN). There is a clear need for higher syntaxes such as Petri nets (Goss and Peccoud(1998)) or derivatives of pi-calculus (Regev et al.(2001)), or, the more recent rule-based propositions (Danos and Laneve(2004); Danos et al.(2007); Faeder et al.(2009)). Rule-based methods in particular introduce a whole new level of flexibility by giving the ability to express refined context-independence assumptions about a particular reaction mechanism. As a result, rules describe succinctly entire reaction classes, and they do so in

[†] Corresponding author: vdanos@inf.ed.ac.uk

a way that makes the oftentimes necessary revision of a model a lot easier (Danos(2009)). It should be noted that the purpose of creating new modeling languages, and in so doing enriching the representational apparatus of a particular modeling domain, is not just to make the modeling more agile and far-reaching. As has been long recognised in the context of programming, one can lean on syntactic structures to develop various analyses that would otherwise be impossible. For instance, in our own rule-based framework, which uses the Kappa language, we have developed and adapted methodologies for detecting dead rules (Danos et al.(2008b)), extracting long-range causal dependencies (Danos et al.(2007)), and obtaining compressed representations of the models' differential semantics (and, therefore, more accurate approximations of their stochastic behaviour) (Harmer et al.(2010); Feret et al.(2009); Danos et al.(2010)).

The more structure one can impose on the problem, the easier the analysis. Now, and unlike in the case of man-made programmed distributed systems, when dealing with natural systems, one is not at liberty to impose a structure. To make further progress in the specific modeling and analysis of BMNs, a compelling idea is to borrow structuring features from biophysics and develop a modeling language where energetic and thermodynamics constraints are put front and center.

This is not a new idea. There has been interesting attempts at developing syntaxes for reactions that would guarantee thermodynamic consistency by construction (Ederer and Gilles(2008); Ederer and Gilles(2007)). Quite recently, those have been successfully extended to restricted forms of rule-based systems (Ollivier et al.(2010)). This approach is congruent with the fact that whether a computable CTMC admits an energy function is undecidable. This is a result which we have established using an encoding of the Post correspondence problem in an earlier paper (Danos and Oury(2010)). Building to the conclusions that thermodynamic consistency should be ensured by construction rather than verified *ex post*. This parallels the way in which strongly typed languages guarantee termination, an otherwise notoriously undecidable problem.

In the present paper, we reconsider the thermodynamic consistency problem (aka the existence of a stochastic equilibrium, or, equivalently, of an energy function) for simple and symmetric Petri nets, equipped with a stochastic mass action dynamics. (Where simple means no two transitions have the same net balance; and symmetric means every reaction is reversible.) We show that in this very restricted class of computable CTMCs, and in contrast with the more expressive Kappa language used in Ref. (Danos and Oury(2010)), deciding the existence of an energy function is possible. The criterion we obtain is of a purely linear algebraic nature. In essence, any reaction invariant (a multiset of reactions with net balance zero) must have zero energy balance. This captures Kolmogorov's criterion (aka Wegscheider's condition in the chemical literature) which states that the product of the CTMC rates along any cycle must equal that of the rates on the reverse cycle. Thanks to the linear structure on the transition graph of a Petri net, the above criterion can be decided by inspecting a finite basis of reaction invariants.

The key ingredient for the algebraisation of the equilibrium problem is the thermodynamic notion of entropy, defined as $\Omega(x) = -\sum_{A \in S} \ln x(A)!$, where $x \in \mathbb{N}^S$ is a state of a Petri net - ie a multiset over a finite set S . This term ensures that the invariant measure, if any, will have finite mass. This unconditional convergence does not happen

with non mass action semantics where the jumping rate of the CTMC only depends on the reaction, not the state. Neither does it happen in the more general case of rule-based CTMCs described in the Kappa language, where one is treading a line, which, in Petri net terms, could be described as having an unbounded number of potential species, and there the entropic contribution does no longer suffice. In fact, a consequence of our result is that entropy-driven convergence can *only* be disrupted in a Kappa model by an unlimited creation of new species, as, otherwise, the dynamics can be faithfully mapped to that of a Petri net (Danos et al.(2008a)).

Unlike the literature on non-deterministic PNs where algebraic invariants have a fundamental role (Pedersen(2008); Chaouiya(2007)), the literature on stochastic Petri nets seems more oriented towards rich and scalable modeling environments than towards analysing asymptotic properties of infinite state systems (Marsan(1990)). In the finite case, the theory of continuous-time Markov chains is already largely computational (Norris(1998)), and, as the size of the state space permits, such systems are amenable to automated verification techniques (eg see Ref. (Desharnais and Panangaden(2003))). In the chemistry literature, one can find similar arguments using Lyapounov potentials for the existence of an equilibrium for reversible Petri nets equipped with a differential mass action semantics (Schuster and Schuster(1989)); and there are also numerous proofs of our result for the simple case of finite-state systems (Yang et al.(2008); Kimura et al.(2007)). None seem to cover the specific stochastic infinite-state scenario we propose here.

The outline of the paper is as follows. We begin with a general reminder on continuous-time Markov chains; then we do the same for Petri nets, narrowing down to the specific class of simple and symmetric nets, and in the next section we establish a criterion for the existence of an equilibrium for this class. The conclusion discusses the limitations we had to impose on the notion of Petri net to obtain the result, as well as the potential avenues for further research.

Acknowledgments: The authors would like to thank Kousha Eteessami and Ian Stark for both inspiring and useful discussions during the preparation of this manuscript.

2. CTMCs

In this section we put in place the basics of continuous-time Markov chain that we will need. We suppose given, once and for all, a countable state space X . We write $\Delta_X = \{(x, x) \mid x \in X\} \subseteq X^2$ for the diagonal of X .

An exponential timer with parameter $\lambda > 0$ is a random variable T on $[0, +\infty)$ defined by $p(T > t) = \exp(-\lambda t)$ for $t \geq 0$. If we pick finitely many independent timers T_i with parameters λ_i , it is easy to prove that:

$$p(T_i = \min T_j, T_i < t) = \frac{\lambda_i}{\sum_j \lambda_j} \cdot (1 - e^{-(\sum_j \lambda_j)t}) \xrightarrow{t \rightarrow 0^+} \lambda_i t \quad (1)$$

Definition 1. A rate function on X is a map $q : X^2 \setminus \Delta_X \rightarrow \mathbb{R}_+$ such that for x in X , $\{y \mid q(x, y) > 0\}$ is finite.

Given a rate function q , we define a binary relation on X , written $|q|$, as $(x, y) \in |q|$

iff $q(x, y) > 0$. We call $|q|$ the *transition graph* or the support of q . We can think of $|q|$ as a directed graph on X , and of q as weighing edges of $|q|$ with positive reals. As according to our definition $q(x, x)$ is not defined, $|q|$ is irreflexive. We also assume that $|q|$ is image-finite or equivalently has finite out-degree. We say that q is *symmetric* if $|q|$ is, that is to say if $q(x, y) = 0 \Rightarrow q(y, x) = 0$.

Given a rate function one can generate a random sequence with values in $X \times [0, +\infty)$ in the following inductive way. Supposing we are at state x , we draw independently for each jump from x an exponential time $T(y)$ with parameter $q(x, y)$. Then we set:

- the next state to be the (almost surely unique) y such that $T(y) = \min T(y')$, and
- the time increment to be $T(y)$.

Image-finiteness guarantees that there is only finitely many timers to choose from. If $q(x, y) > 0$, it follows from (1) that one jumps to y , with a probability proportional to $q(x, y)$. If $q(x, y) = 0$, the chain never jumps from x to y .

This simulation protocol defines a probability to be in state y at time t , starting in state x at time 0, which we will denote by $P(t, x, y)$, that is known as a continuous-time Markov chain (CTMC).

Technically, one has to assume more from q to prevent Zenonian explosions, where a process accumulates infinitely many jumps in finite time (Norris(1998)). We do not need to get into these delicate questions as the actual CTMCs we are interested in will be sufficiently regular. Indeed, our rate functions will come from a certain class (to be defined in details in the next section) of simple and symmetric Petri nets (which is why we need a countable state space), and will therefore have a homogeneous structure - with a bounded out-degree (though rates themselves will be unbounded).

Definition 2. Given q a rate function on X , one says a probability p on X is:

- invariant if for x in X , $t \geq 0$, $p(x) = \sum_y p(y)P(t, y, x)$
- an equilibrium if for distinct x, y in X , $p(x)q(x, y) = p(y)q(y, x)$

Note that the equilibrium condition breaks naturally in two conditions of a different nature:

- (E1): $p(x)q(x, y) = p(y)q(y, x)$ and
- (E2): $\sum_x p(x) = p(X) = 1$ (convergence).

Condition (E1) is a purely algebraic condition for which there is a simple criterion (Kolmogorov's or Wegscheider's, see the introduction). It is known as *reversibility* in the probabilistic literature, and *detailed balance* in the chemical one. Condition (E2), on the other hand, is a convergence condition which only comes into play because the state space X is infinite.

An equilibrium is a particular kind of invariant probability. To see this, let us return to (1) which says that an exponential timer of parameter λ fires within $[0, \tau]$ with a probability which tends to $\lambda\tau$ as $\tau \rightarrow 0^+$. Using this first order expansion for a small time τ , we get that $P(\tau, y, x) \sim \tau q(y, x)$ if $x \neq y$, and therefore $P(\tau, x, x) \sim 1 - \sum_{y \neq x} \tau q(x, y)$.

Hence, the invariant distribution equation given above can be rewritten as:

- (S): $p(x) \sum_{y \neq x} q(x, y) = \sum_{y \neq x} p(y)q(y, x)$.

The (S) condition expresses the fact that p is a probabilistic state of the system such that the compound rate at which one leaves x (the left hand side) is the same as the rate

at which one enters x (the right hand side). By contrast, (E1) expresses that p is a state such that the rate at which one jumps from x to y is the same as the rate at which one jumps from y to x . Clearly (S) is implied by (E1).

From the classic theory of CTMCs (Norris(1998)), we know that if q 's transition graph is strongly connected (aka q is irreducible), and p is a probability invariant under q then:

- p is unique,
- q is recurrent, meaning the probability to return to any x is 1,
- q is positive-recurrent, meaning the mean return time to any x is finite,
- and importantly, for $x, y \in X$, $P(t, x, y) \rightarrow p(y)$ when $t \rightarrow +\infty$.

In words, wherever one starts in X , the dynamics will invariably lead the system to p - a strong property which reminds one of both termination and confluence. All of the above holds for an invariant probability, and *a fortiori* for one that satisfies the stronger detailed balance condition (E1). (This raises the question of whether having an invariant probability is decidable for Petri nets; a question which we do not attempt to answer here.)

We will call the set of states where an equilibrium p is not zero, its *support*. It is easy to see that the sub-graph induced by the support of p in $|q|$ is symmetric (hence strongly connected), even if $|q|$ is not, and terminal (ie has no outgoing edges). It follows that all equilibria are convex combinations of equilibria localized to such symmetric terminal components of $|q|$.

2.1. Examples

Here is a couple of examples to illustrate the definitions so far.

Take $X = \mathbb{N} \setminus \{0\}$ and $q(x, x+1) = 1$; this is a walk on X where only the times are random, the sequence of states being inflexibly determined. It is connected but clearly not symmetric, nor strongly connected. To make the transition graph strongly connected we can add 'resets', eg set $q(2^k, 1) = 1$ for powers of 2, $k > 0$. The obtained graph is still not symmetric, and one can also comment that with this specific choice, resets are so sparse that the mean return time to 1 is actually infinite, and hence the system has no invariant probability. This could not happen with a finite transition graph.

Take $X = \mathbb{N}$ and $q(x, x+1) = \alpha \cdot q(x+1, x)$ with $\alpha > 0$. This (biased) random walk on \mathbb{N} is symmetric and (hence) strongly connected. We see that (E1) has geometric solutions verifying $p(x+1)/p(x) = \alpha$, and therefore (E2) only holds if $\alpha < 1$.

2.2. A thermodynamic aside: energy

The condition (E1) of detailed balance refers to some notion of thermodynamic consistency. Indeed, the equilibrium problem can be reformulated as the existence of a function $H : X \rightarrow \mathbb{R}$ such that:

- (F1) $\ln(q(y, x)/q(x, y)) = H(y) - H(x)$, and
- (F2) $\sum_x \exp(-H(x)) \leq \infty$.

One can choose to describe an equilibrium either as a probability p or by means of a real-valued map as above. The correspondence between the two definitions is given by

Boltzmann's relation:

$$p(x) = e^{-H(x)} / \sum_x e^{-H(x)} \quad (2)$$

The map H is called an *energy* function and is defined up to an additive constant (which reflects the convention that a probability should sum up to 1).

We will return to this thermodynamic formulation to discuss our main result in the last section.

3. PNs

In this section we start with a quick reminder of what Petri nets are - with an eye on: 1) narrowing down our question to a subclass of Petri nets for which we can derive a criterion, 2) describing the reaction invariants which correspond to loops in the state space, and which we will use to decide (E1). The terminology of Petri nets witnesses the history of the subject and we will mostly use the chemical side of the dictionary, eg talking about reactions rather than transitions, and species rather than places.

3.1. Basic definition

A Petri net consists of:

- two disjoint finite sets S (species), R (reactions), and
- an input and an output function $i, o : R \rightarrow S \rightarrow \mathbb{N}$

The idea is that $i(r)(A)$ is the number of tokens of species A that reaction r consumes, while $o(r)(A)$ is the number it produces. The state space is $S \rightarrow \mathbb{N} \simeq \mathbb{N}^S$ which one can see equivalently as: multisets over S , integer vectors of size $|S|$, or integer-valued maps defined on S . States form a subset of the real vector space \mathbb{R}^S of dimension $|S|$; they are closed under linear combination with integer coefficients, and can be added, subtracted and compared. We write \leq for the pointwise order on finite maps.

The data above allows one to define a labelled transition system on the state space, that is to say a family of binary relations on \mathbb{N}^S labelled in R , written $x \rightarrow_r y$. Specifically, a transition labelled by $r \in R$ (aka a firing of r) is defined by:

$$x \rightarrow_r x + o(r) - i(r) \text{ if } i(r) \leq x$$

where the condition $i(r) \leq x$ expresses the fact that r needs a minimal amount of inputs to fire.

We will write $r \cdot x$ for $x + o(r) - i(r)$, ie the new state resulting from applying r to x .

The transition system we just defined is *monotonic* in the sense that $x \rightarrow_r y \Rightarrow x + x' \rightarrow_r y + x'$ for any $x' \in \mathbb{N}^S$. More tokens never inhibit a transition. Extensions of Petri nets including reaction where the *absence* of a species can be required for a reaction to fire, make the entire framework a lot more complex.

3.2. Linear structure

Using the linear structure of the state space, we can represent the input and output functions as matrices indexed on $S \times R$ with coefficients in \mathbb{N} . This gives the so-called *stoichiometric* matrix C as $C = o - i$.

The column vector $C_r \in \mathbb{Z}^S$ represents the net effect of applying r , or the jump/translation in the state space that results from applying r . This jump does not depend on the state one applies it to (and on the manner in which one applies it - unlike in the case of rules (Danos and Oury(2010))); what does depend on the source state is the application condition $x \geq i(r)$, the fulfillment of which will determine whether the jump is possible (or enabled) from x .

More generally, a vector $y \in \mathbb{N}^R$ represents a linear combination of reactions and $Cy \in \mathbb{Z}^S$ is its total resulting effect. In general, y will be realised by countably many trajectories in the underlying state space \mathbb{N}^S . Specifically, any trajectory whose sequence of reactions \tilde{y} enumerate the multiset y will have to have enough inputs at each step for the next reaction to proceed. By monotony, any y with integer coefficients can be realised provided one starts with a large enough state.

A *reaction invariant*, or R -invariant, is a $y \in \mathbb{N}^R$ such that $Cy = 0$. Equivalently a multiset of reactions y such that, wherever realisable, the underlying trajectory is a loop in \mathbb{N}^S . One can think of such invariants as loop schemes.

One also has the dual notion of *species invariant*, or S -invariant, namely a $u \in \mathbb{Z}^R$ in the null-space of C^t the transpose of C ; equivalently u such that $\langle u, Cy \rangle = \langle C^t u, y \rangle = 0$, that is to say the linear form $\langle u, \cdot \rangle$ is a (linear) invariant of the trajectories.

We will use later the following basic linear algebraic fact. If C is a (stoichiometric) matrix, then $\ker(C)^\perp = \text{Im}(C^t)$. To see this, suppose $\kappa = C^t \epsilon$ for some $\epsilon \in \mathbb{R}^S$, and pick ϕ in $\ker(C)$. We have $\langle \kappa, \phi \rangle = \langle C^t \epsilon, \phi \rangle = \langle \epsilon, C\phi \rangle = 0$. Hence $\text{Im}(C^t) \subseteq \ker(C)^\perp$, and since both subspaces have the same dimension this inclusion is in fact an equality.

Note that C^t might not be injective (equivalently the associated system might have non-trivial S -invariants). But if we pick ϵ_1, ϵ_2 such that $C^t \epsilon_1 = C^t \epsilon_2$, then the linear forms $\langle \epsilon_1, \cdot \rangle, \langle \epsilon_2, \cdot \rangle$ are equal when restricted to $\text{Im}(C)$; in other words, $\epsilon_1 - \epsilon_2$ is an S -invariant. In the sequel, this will imply that any choice of ϵ amounts to the same.

We will restrict ourselves to Petri nets that are simple and symmetric, as defined below.

Definition 3. Let N be a net, with reactions R and stoichiometric matrix C ; N is said to be:

- simple if there are no identical columns in C
- symmetric if every reaction $r \in R$ has an inverse reaction $r^* \in R$ with $i(r^*) = o(r)$ and $i(r) = o(r^*)$.

The first condition says that no two reactions have identical jumps; the second one says that every reaction is reversible. We remark that if N is symmetric and simple, then there is a unique r^* inverse to r , so in this case one can unambiguously refer to *the* reaction inverse to r . We can also remark that r can be inverse to itself when $i(r) = o(r)$, that is to say when r has no effect on the state.

One can always symmetrize a net by adding an inverse where one lacks one. It is unclear however how one can naturally ‘simplicize’ a given net; one could select among identical columns but that would be arbitrary.

3.3. Examples

In practice, we present a Petri net as a list of reactions, here are 3 examples:



Ex. (5) and (4) are simple, but ex. (3) is not as both transitions have the same net balance. Ex. (5) is symmetric; ex. (4) would have a symmetric underlying transition graph if we were to add $S \rightarrow I$, but it would no longer be simple, neither would it be symmetric in the intensional sense we have used here. This illustrates the difference between the extensional notion of a symmetric transition graph and the more restrictive and intensional one which we have defined and will be using for Petri nets.

Continuing on the same example (4), we see that there is a reaction invariant $y^t = (1\ 1)$, which can be realised by a loop provided one starts from a state x where $x(I) > 0$. It also has a species invariant, $x^t = (1\ 1)$, expressing the fact that $x(S) + x(I)$ is invariant under any transition (the actual value depending on the initial state).

3.4. Mass action semantics

We turn to the quantitative aspects of Petri nets.

PNs have a countably infinite state space and finitely many reactions - so they will generate countably infinite transition graphs with finite degree which fits the definition of rate function of §2. It remains to assign to each transition a rate.

Suppose x, z are multisets over S , we define: $x! := \prod_{A \in S} x(A)!$, the number of symmetries of any enumeration of x ; $|x| := \sum_{A \in S} x(A)$ the total number of elements of x ; and, when $z \leq x$, the number of matches for z in x :

$$[z; x] := \prod_{A \in S} [z(A); x(A)] := \prod_{A \in S} x(A)! / (x(A) - z(A))! = x! / (x - z)! \quad (6)$$

Definition 4. Given a simple PN on S, R , and a *rate constant* map $k : R \rightarrow \mathbb{R}^+ \setminus \{0\}$, we define the jumping rate as:

$$x \rightarrow_r x + o(r) - i(r) \text{ with rate } \tau(x, r) = k(r) \cdot [i(r); x]$$

By the application condition, $i(r) \leq x$ and hence $[i(r); x]$ is well-defined.

Note that if the PN were not simple the above definition would induce a compound rate function:

$$q(x, y) = \sum_{\{r | r \cdot x = y\}} \tau(x, r)$$

As we will see, this would create a problem when it comes to constructing equilibria -

which is why we restrict to simple PNs. As a consequence, the set of reactions leading from x to y is either empty or a singleton $\{r\}$, and we can write $q(x, r \cdot x) = \tau(x, r)$ and $q(x, y) = 0$ otherwise.

Another point worth making is that we have chosen to use $[a; b] = b!/(b-a)!$ the number of injections of a in b to count matches. Sometimes, one takes instead as a counting principle $\binom{b}{a} = [a; b]/a!$, ie the number of subsets. The difference between the two conventions is independent of the state (ie it is static), and can (therefore) be entirely hidden in the rate constant $k(r)$. The convention we follow is noticeably more natural when one considers the rule-based extension of PNs and their refinement theory (Danos et al.(2008a)).

3.5. Comparison with constant rate semantics

The semantics above will be called the *mass action* semantics.

It is important to contrast it with another semantics which associates to a jump $x \rightarrow_r x + o(r) - i(r)$ a flat rate $\tau(x, r) = k(r)$. These nets are often referred to as stochastic Petri nets with marking-independent rates (Marsan(1990)), as indeed the rate of a jump only depends on the reaction r and not on the state x to which the reaction applies.

The difference between the two semantics has drastic consequences on their long-term behaviours. To see this, let us reconsider example (5), a birth-and-death process $\rightarrow_k A \rightarrow_d$ where rate constants are indicated as subscripts (which we recall are both > 0). If we follow the non mass action definition, condition (E1) becomes $p(x)k = dp(x+1)$, so $p(x)$ is a geometric sequence, and will verify (E2) iff $\mu := k/d < 1$. With the mass action definition the (E1) condition reads $p(x)k = d(x+1)p(x+1)$, and one can verify easily that the solution is a Poisson distribution:

$$p(x) = \exp(-\mu) \mu^x / x!$$

which therefore converges independently of the value of μ . This unconditional convergence is a general phenomenon and so is the occurrence of the Poisson distribution; somewhat paradoxically, mass action will turn out to be a mathematically simpler semantics when it comes to deciding equilibria.

3.6. A lemma

We will say that a Petri net N with a rate constant map $k : T \rightarrow \mathbb{R}_+$ is *sisma*, if it is simple, symmetric (as in Def. 3), and has mass-action semantics.

Lemma 1. Let N be *sisma*, for x in \mathbb{N}^S , $r \in R$, $y = r \cdot x$, one has:

$$q(y, x)/q(x, y) = k(r^*)/k(r) \cdot y!/x!$$

Proof. Because N is simple, r is the only reaction that brings x to y , so by definition of mass action, the rate of the jump from x to y is:

$$q(x, y) = k(r) \cdot x!/(x - i(r))!$$

Because N is symmetric r^* brings y back to x , so $q(y, x) > 0$, and again, because N is simple, r^* is the only reaction that does this, so the rate of the reverse jump is:

$$q(y, x) = k(r^*) \cdot y! / (y - o(r))!$$

where we have used $i(r^*) = o(r)$ which holds by definition of r^* .

The conclusion follows as $y - o(r) = x - i(r)$. \square

4. Equilibrium

Now all is in place for our main result.

Theorem 1. Let N be sisma with stoichiometric matrix C , and define N 's transitional energy vector $\kappa \in \mathbb{R}^R$ as $\kappa(r) = \ln(k(r^*)/k(r))$ for $r \in R$; N has a nowhere zero equilibrium iff $\kappa \in \ker(C)^\perp$.

Proof. \Rightarrow : (easy direction) Pick ϕ in $\ker(C)$; as C has values in \mathbb{Z} , without loss of generality we can suppose that ϕ has coefficients in \mathbb{Z} . In fact, we can suppose ϕ has coefficients in \mathbb{N} , as N being symmetric, one can always replace a negative coefficient on r (say) with a positive one on r^* . Next, pick a state x_0 sufficiently large for ϕ to be realisable as a cycle γ_ϕ at x_0 . This is always possible by monotony. By (E1) for all jumps (x, y) in γ_ϕ , $p(x)q(x, y) = p(y)q(y, x)$. Since p is nowhere zero, this can be rewritten $p(x)/p(y) = q(y, x)/q(x, y)$, which implies:

$$\prod_{\gamma_\phi} q(y, x)/q(x, y) = \prod_{\gamma_\phi} p(x)/p(y) = 1$$

or, equivalently (by taking a log):

$$0 = \sum_{(x, y) \in \gamma_\phi} \ln(p(x)/p(y)) = \sum_{(x, y) \in \gamma_\phi} \ln(q(y, x)/q(x, y))$$

We can now evaluate $q(y, x)/q(x, y)$ using Lemma 1 -and obtain:

$$\begin{aligned} 0 &= \sum_r \kappa(r) \phi(r) + \sum_{(x, y) \in \gamma_\phi} \ln(y!/x!) \\ &= \langle \kappa, \phi \rangle + \sum_{(x, y) \in \gamma_\phi} \ln(y!/x!) \\ &= \langle \kappa, \phi \rangle \end{aligned}$$

as the terms $\ln(y!/x!)$ add up to zero along any cycle.

\Leftarrow :

Because N is symmetric, the transition graph of q is a disjoint sum of symmetric components.

Pick a component D , together with a distinguished state z_0 , and for each z in D choose a path γ_z from z_0 to z .

Up to the choice of $p(z_0)$, the following defines uniquely a function p on D :

$$\begin{aligned} \ln(p(z_0)/p(z)) &:= \sum_{(x, y) \in \gamma_z} \ln(q(y, x)/q(x, y)) \\ &= \langle \kappa, \tilde{\gamma}_z \rangle + \ln(z!/z_0!) \end{aligned}$$

where $\tilde{\gamma}_z$ is the R -vector to which γ_z projects.

We can rewrite the above as:

$$\ln p(z) + \ln z! = \ln p(z_0) + \ln z_0! - \langle \kappa, \tilde{\gamma}_z \rangle$$

This expression for $p(z)$ does not depend on the choice of γ_z - since $\langle \kappa, \tilde{\gamma}_z \rangle$ is constant over all paths from z_0 to z . This is because we assume that $\langle \kappa, \tilde{\gamma} \rangle = 0$ for any cycle γ . For the same reason, this assignment verifies (E1).

We also have to take care of (E2).

Using an earlier remark in §3.2, we know there is an $\epsilon \in \mathbb{R}^S$ such that $\langle \kappa, \tilde{\gamma}_z \rangle = \langle \epsilon, C\tilde{\gamma}_z \rangle$, and since $C\tilde{\gamma}_z = z - z_0$, we can rewrite the above as:

$$\ln p(z) + \ln z! + \langle \epsilon, z \rangle = \ln p(z_0) + \ln z_0! + \langle \epsilon, z_0 \rangle \quad (7)$$

From Eq. (7) follows that:

$$\sum_{z \in D} p(z) \propto \sum_{z \in D} e^{-\langle \epsilon, z \rangle} / z! \leq \sum_{z \in \mathbb{N}^S} e^{-\langle \epsilon, z \rangle} / z!$$

Let us write $\mu_A = e^{-\epsilon(A)}$. We have:

$$\sum_{z \in \mathbb{N}^S} e^{-\langle \epsilon, z \rangle} / z! = \sum_{z \in \mathbb{N}^S} \prod_{A \in S} \mu_A^{z(A)} / z(A)! = \prod_{A \in S} e^{\mu_A} < \infty$$

Hence (E2) holds as well. As we can repeat the construction for each component, we can define a nowhere zero equilibrium. \square

Note that in the (E2) part of the argument, the choice of ϵ such that $C^t \epsilon = \kappa$ does not matter. Pick another ϵ' , it follows from (7) that $\ln(p'(z)/p(z)) = \langle \epsilon' - \epsilon, z_0 - z \rangle = 0$, because $\epsilon' - \epsilon$ is an S-invariant (§3.2) and z_0, z are connected in the transition graph.

We conclude immediately that whether a sisma Petri net N has a nowhere zero equilibrium is decidable. One can equally use the theorem to choose a κ in the solution space $\ker(C)^\perp$. This solution space has dimension $d = |R| - \dim(\ker(C))$, which means one can fix d of the $|R|$ transitional energies to arbitrary values and fill in the rest uniquely according to the constraint. In particular, there is always the trivial choice $\kappa = \mathbf{0}$, corresponding to $k(r) = k(r^*)$ for all r in R .

Ref. (Yang et al.(2008)) explores in the case of a finite state space, how the solution space can be further constrained by experimental evidence. Our criterion shows that this is possible in general for a countable state space.

It is clear from the proof that, provided κ is the solution space, one can define a unique equilibrium with support any of the components of q ; all other equilibria can then be obtained by convex combinations of such minimal ones.

It is also clear that the convergence part of the proof (E2) makes great use of the specific mass action shape of the rates. This does not come as a surprise as we have seen earlier with the birth-and-death example (§3.5) that the constant rate semantics does not always yield a convergent solution to (E1). Moreover, as foreseen in §3.5, the equilibrium solution must be a Poisson distribution when q is irreducible.

Corollary 1. Let N be sisma and irreducible, its invariant probability if it exists is unique and is a product of Poisson distributions.

Proof. To see this, observe that from eq. (7), one gets for $z \in \mathbb{N}^S$:

$$p(z) \propto \prod_{A \in S} \mu_A^{z(A)} / z(A)! \quad (8)$$

with $\mu_A = e^{-\epsilon(A)}$. Now if N 's transition graph is (strongly) connected, this is exactly saying that $p(z)$ is an S -indexed product of Poisson distributions with parameters μ_A . \square

Hence a thermodynamically consistent and irreducible sisma net N is equivalent in the long term to a juxtaposition of independent birth-and-death processes with parameter ratios μ_A . All correlations between A s and B s are transient. In the case N is not irreducible, the formula above is still correct, but the normalisation factor will for $p(z)$ will depend on the the component of z , so while there can now be long-term correlations, they will have to be entirely encoded in the qualitative reachability properties of N . This intriguing lack of expressivity in the asymptotic behaviour of (irreducible) sisma nets - which one might call their 'normal forms' to borrow from the terminology of rewriting systems - can be traced back to the fact that the transitional energies κ can be 'tokenized' as $C^t\epsilon$ as soon as N satisfies (E1).

4.1. A stronger version?

The result only decides the existence of an equilibrium which is everywhere non-zero. (Or equivalently, whether it is possible, for every connected component of N 's transition graph to define an equilibrium with support this component.) Now, what if one is interested in deciding whether N has an equilibrium with support the component D of some specific initial state z_0 ? Admittedly, this is a rather gratuitous question, as in such situations which are not already covered by our result, the 'physics' of our Petri net is consistent for some specific inputs only. Nevertheless, it seems fairly subtle.

Returning to the second part of the proof, one sees that to build an equilibrium on D , one needs κ to be orthogonal only to those reaction invariants ϕ that are realized by a loop γ_ϕ in the distinguished component D . So the refined statement becomes that N has a solution to (E1) with support D iff κ is orthogonal to those. The problem is that we no longer know how to conclude to (E2), and even if we could prove it always holds, it is unclear how to decide the partial orthogonality property above.

True, it is easy to decide whether any given ϕ is realized in D . First, ϕ is realisable in D iff one of the $|\phi|!$ orderings ϕ_σ of ϕ is. For each ϕ_σ , and $A \in S$, we can compute the minimal amount of A s needed to complete $\phi_\sigma(A)$, say $x(A)$ - and ϕ_σ is realizable in D iff there exists a z reachable from z_0 such that $z \geq x$ (in the PN speak, whether x is *coverable* starting from z_0). This is decidable - eg using coverability trees à la Karp-Miller (Karp and Miller(1969)).

But checking realisability within D on a generating family of R -invariants (in \mathbb{N}^R), is no longer enough, as non-realisation is *not* stable under linear combination (with integer coefficients).


Consider the following two 'cooperating' (ordered) R -invariants:

$$\begin{aligned} \phi_1 &= A_1 \rightarrow B_1, \mathbf{B}_2 + B_1 \rightarrow \mathbf{B}_2 + C_1, C_1 \rightarrow A_1 \\ \phi_2 &= A_2 \rightarrow B_2, \mathbf{C}_1 + B_2 \rightarrow \mathbf{C}_1 + C_2, C_2 \rightarrow A_2 \end{aligned}$$

ϕ_1 needs an A_1 and a B_2 , ϕ_2 needs an A_2 and a C_1 .

If we start from $z_0 = A_1 + A_2$, none of the loops go through; nevertheless, one can

realize $\phi_1 + \phi_2$, because ϕ_1, ϕ_2 can exchange intermediates \mathbf{B}_2 , and \mathbf{C}_1 :



Whether one can work around this additional complication, and find a finite description of the loops realisable in an arbitrary component of a Petri net remains to be seen.

5. A thermodynamic aside: entropy

It is instructive to rephrase our main result in terms of energy, and to build a different intuition thereof.

We have proved that a sisma Petri net N , with transitional energies κ has an equilibrium iff $\kappa \in \ker(C)^\perp$ essentially using Eq. (7) to define the equilibrium. Equivalently, using the correspondence explained in eq. (2) in §2.2, the equilibrium can be described as the following (free) energy assignment:

$$F(x) = \langle \epsilon, x \rangle + \ln x! \quad (9)$$

where ϵ is such that $\kappa = C^t \epsilon$. (Recall we write $x!$ for $\prod_A x(A)!$). Note that the choice of a particular representative of ϵ will only change the above definition by an additive constant per connected component.

One sees that $F(x)$ decomposes as $E(x) - \Omega(x)$ with:

- $E(x) = \langle \epsilon, x \rangle$ which one can interpret as the internal energy of x , or the energy of creating x
- $\Omega(x) = -\ln x!$ a combinatorial symmetry discount term which one can interpret as the entropy of creating x (see the discussion below).

Since $p(x) \propto e^{-F(x)}$, those states x that minimise $E(x)$ and maximise $\Omega(x)$ will be favoured by the invariant probability. Minimising the E -term is easy, if A has the lowest ϵ_A , then x should have only A s; maximising the Ω -term is also easy, x should have a low $x!$, which means x should be as uniformly distributed among species as possible. To see this, observe that $|x|!/x!$ is a multinomial coefficient, which for a fixed value of $|x|$ gets highest values when x partitions $|x|$ as evenly as possible. This second term, unlike the first one, favours disorder. To minimise both at once is complicated as the two goals are clearly in contention.

Consider that one can rescale κ by any factor λ , without leaving $\ker(C)^\perp$, and obtain a rescaled equilibrium:

$$H(x) = \lambda E(x) - \Omega(x) = \lambda(E(x) - \lambda^{-1}\Omega(x))$$

which leads to interpret $\|\epsilon\|^{-1}$ as a formal notion of temperature arbitrating the competition between the internal energy term and the entropic one.

To return to the entropy term, supposing $|x| := \sum_A x(A)$ is a constant n (ie the total number of tokens per component is constant), we can explain away this term in the following way.

Consider the set of words S^n , equivalently the set of maps $n \rightarrow S$. Permutations of

n have an action on this set. Write Π for the (canonical) projection Π from words to multisets.

By the orbit-stabilizer relation, for any word w :

$$n! = \text{sym}(w) \times \text{orb}(w)$$

On the other hand, $\text{sym}(w) = \prod_{A \in S} \Pi(w)(A)!$, since $\text{sym}(w)$ is the set of name-preserving permutations of w .

Hence the cardinality of the inverse image of x is:

$$|\Pi^{-1}(x)| = \left(\sum_A x(A)! \right) / \prod_A x(A)! = |x|! / \prod_A x(A)! = |x|! / x! \quad (10)$$

Using logs we get $\ln |\Pi^{-1}(x)| = \ln |x|! + \Omega(x)$. And we see that the equilibrium F is the image of E under the quotient induced by Π , that is to say

$$F(x) = E(w) + \ln |\Pi^{-1}(x)|$$

up to an additive constant $\ln n!$, for any $w \in \Pi^{-1}(x)$.

This looks like a hint at the existence of a concrete counterpart of N , acting on words - reminiscent of the *individual-token semantics* of Petri nets (Bruni et al.(1999)) - for which Π would be a functional stochastic bisimulation. It is not too difficult to explicitly carry out this programme for the simple case where one has a constant number of tokens. In so doing, the entropic term is seen as a side-effect of changing one's semantics and switching to the more abstract collective one, as described by the projection Π . It is however unclear at the time of writing whether this microscopic rendition of entropy can be freed from the simplifying assumption of an invariant number of tokens.

6. Conclusion

We have established a computable criterion for a simple and symmetric Petri net equipped with mass action semantics to admit an equilibrium probability. This result put alongside our earlier result (Danos and Oury(2010)) where we proved that the same problem is undecidable for general computable Markov chains, sheds some light on the decidability boundary. There are many questions one might want to pursue from there.

We have mentioned already the decision of the weaker property of having an invariant probability, as well as the more foundational question of the extent to which one can deconstruct entropy terms by constructing covers as sketched at the end of the preceding section. This should be done preferably in some axiomatic framework to have maximal generality -perhaps following the leads in Ref. (Bruni et al.(1999)), or the more encompassing notion of adhesive categories (Lack and Sobociński(2005); Lack and Sobociński(2004); Ehrig et al.(2004)).

Another question is how our result articulates with the MC/DE divide. If a Petri net has an invariant probability, then one would expect its deterministic DE approximation to be defined at all times (Darling and Norris(2008)). For instance, the reaction $2A \rightarrow 3A$ will give the mass action differential equation $x' = x^2$, a Riccati equation which has explosive solutions $x(t) = 1/(x(0) - kt)$. As soon as one adds the reverse reaction $2A \leftarrow 3A$

the new differential equation $x' = x^2 - x^3$ has positive solutions defined at all times. The extent to which convergence in the MC world carries over to differential limits should be investigated. A related question is the relation with Feinberg's chemical reaction theory, which is also partly based on algebraic conditions (the so-called deficiency space of a reaction system) to ensure multistationarity of the differential semantics (Feinberg(1987)). Certainly a probabilistic equilibrium is adverse to multistationarity and the link should be made precise. Ref. (Anderson et al.(2010)) might help shed some light on this issue.

Another direction, which prompted the investigation we have presented in the first place, already mentioned in a thread of discourse in the introduction, is to bring the experience gained here to bear on the study of the thermodynamical consistency of the larger class of rule-based models of BMNs. As said, there is no hope to have an all-encompassing criterion of computational significance because the question is undecidable. However, it is possible to envision the synthesis of specific classes of consistent rule sets by using local energy terms to constrain the allowed rules. This is an exciting question which we have already started investigating in trying to extend the energy-based modeling techniques of Ref. (Ollivier et al.(2010)).

Finally, we have dealt here with what one might call a quantitative termination question. It is pleasing, if unsurprising, to see thermodynamics inviting itself in the conceptual apparatus. It seems that a thermodynamics-based quantitative rewriting theory could renew interestingly the classical questions of reachability (for a recent example of such an investigation in Kappa-related formalisms, see eg Ref. (Delzanno et al.(2009))), confluence, and termination that have been its traditional concerns (for a recent example in concurrency theory, see eg Ref. (Bacci et al.(2011))).

References

- D.F. Anderson, G. Craciun, and T.G. Kurtz. Product-form stationary distributions for deficiency zero chemical reaction networks. *Bulletin of Mathematical Biology*, 72(8):1947–1970, 2010.
- G. Bacci, V. Danos, and O. Kammar. On the statistical thermodynamics of reversible communicating processes. In Andrea Corradini, Bartek Klin, and Corina Cîrstea, editors, *4th International Conference, CALCO 2011*, volume 6859 of *Lecture Notes in Computer Science*, pages 1–18, 2011.
- R. Bruni, J. Meseguer, U. Montanari, and V. Sassone. Functorial semantics for Petri nets under the individual token philosophy. In *Proceedings of CTCS*, volume 99. Citeseer, 1999.
- C. Chaouiya. Petri net modelling of biological networks. *Briefings in Bioinformatics*, 8(4):210, 2007.
- V. Danos. Agile modelling of cellular signalling. *Electr. Notes Theor. Comput. Sci.*, 229(4): 3–10, 2009.
- V. Danos and C. Laneve. Formal molecular biology. *Theor. Comput. Sci.*, 325(1):69–110, 2004.
- V. Danos and N. Oury. Equilibrium and termination. *CoRR*, abs/1006.1430, 2010.
- V. Danos, J. Feret, W. Fontana, R. Harmer, and J. Krivine. Rule-based modelling of cellular signalling. In L. Caires and V. Vasconcelos, editors, *CONCUR*, volume 4703 of *LNCS*, pages 17–41. Springer, 2007.
- V. Danos, J. Feret, W. Fontana, R. Harmer, and J. Krivine. Rule-based modelling, symmetries, refinements. In J. Fisher, editor, *FMSB*, volume 5054 of *Lecture Notes in Computer Science*, pages 103–122. Springer, 2008a.

- V. Danos, J. Feret, W. Fontana, and J. Krivine. Abstract interpretation of cellular signalling networks. In Francesco Logozzo, Doron Peled, and Lenore D. Zuck, editors, *VMCAI*, volume 4905 of *LNCS*, pages 83–97. Springer, 2008b.
- V. Danos, J. Feret, W. Fontana, R. Harmer, and J. Krivine. Abstracting the differential semantics of rule-based models: Exact and automated model reduction. In *LICS*, pages 362–381. IEEE Computer Society, 2010.
- R. Darling and J.R. Norris. Differential equation approximations for Markov chains. *Probability surveys*, 5:37–79, 2008.
- G. Delzanno, C. Di Giusto, M. Gabbriellini, C. Laneve, and G. Zavattaro. The Kappa-Lattice: Decidability Boundaries for Qualitative Analysis in Biological Languages. In *Computational Methods in Systems Biology*, pages 158–172. Springer, 2009.
- J. Desharnais and P. Panangaden. Continuous stochastic logic characterizes bisimulation of continuous-time Markov processes. *J. Log. Algebr. Program.*, 56(1-2):99–115, 2003.
- D. Easley and J. Kleinberg. *Networks, crowds, and markets: Reasoning about a highly connected world*. Cambridge Univ Press, 2010.
- M. Ederer and E.-D. Gilles. Thermodynamic constraints in kinetic modeling: thermodynamic-kinetic modeling in comparison to other approaches. *Engineering in Life Sciences*, 8(5):467–476, 2008.
- M. Ederer and E.-D. Gilles. Thermodynamically feasible kinetic models of reaction networks. *Biophysical Journal*, 92(6):1846–1857, 2007.
- H. Ehrig, A. Habel, J. Padberg, and U. Prange. Adhesive high-level replacement categories and systems. *Graph Transformations*, pages 144–160, 2004.
- J.M. Epstein. Agent-based computational models and generative social science. *Complexity*, 4(5):41–60, 1999.
- J.R. Faeder, M.L. Blinov, and W.S. Hlavacek. Rule-based modeling of biochemical systems with BioNetGen. *Methods Mol. Biol.*, 500:113–167, 2009.
- M. Feinberg. Chemical reaction network structure and the stability of complex isothermal reactors—I. The deficiency zero and deficiency one theorems. *Chemical Engineering Science*, 42(10):2229–2268, 1987.
- J. Feret, V. Danos, J. Krivine, R. Harmer, and W. Fontana. Internal coarse-graining of molecular systems. *Proceedings of the National Academy of Sciences*, 106(16):6453, 2009.
- P.J.E. Goss and J. Peccoud. Quantitative modeling of stochastic systems in molecular biology by using stochastic Petri nets. *Proceedings of the National Academy of Sciences*, 95(12):6750, 1998.
- R. Harmer, V. Danos, J. Feret, J. Krivine, and W. Fontana. Intrinsic Information carriers in combinatorial dynamical systems. *Chaos*, 20(3):037108, 2010.
- R.M. Karp and R.E. Miller. Parallel Program Schemata. *Journal of Computer and System Sciences*, 3:147–195, 1969.
- H. Kimura, H. Okano, and R. J. Tanaka. Stochastic approach to molecular interactions and computational theory of metabolic and genetic regulations. *Journal of Theoretical Biology*, 248(4):590–607, 2007.
- S. Lack and P. Sobociński. Adhesive categories. In *Foundations of Software Science and Computation Structures*, pages 273–288. Springer, 2004.
- S. Lack and P. Sobociński. Adhesive and quasiadhesive categories. *Theoretical Informatics and Applications*, 39(3):511–545, 2005.
- M. Marsan. Stochastic Petri nets: an elementary introduction. *Advances in Petri Nets 1989*, pages 1–29, 1990.
- J.R. Norris. *Markov chains*. Cambridge University Press, 1998.

- J.F. Ollivier, V. Shahrezaei, and P.S. Swain. Scalable Rule-Based Modelling of Allosteric Proteins and Biochemical Networks. *PLOS Computational Biology*, 6(11):6750, 2010.
- M. Pedersen. Compositional definitions of minimal flows in Petri nets. In Monika Heiner and Adelinde M. Uhrmacher, editors, *CMSB*, volume 5307 of *LNCS*, pages 288–307. Springer, 2008.
- A. Regev, W. Silverman, and E. Shapiro. Representation and simulation of biochemical processes using the-calculus process algebra. In *Pacific symposium on biocomputing*, volume 6, pages 459–470, 2001.
- S. Schuster and R. Schuster. A generalization of Wegscheider’s condition. Implications for properties of steady states and for quasi-steady-state approximation. *Journal of Mathematical Chemistry*, 3(1):25–42, 1989.
- J. Yang, W.J. Bruno, W.S. Hlavacek, and J.E. Pearson. On imposing detailed balance in complex reaction mechanisms. *Biophysical Journal*, 91(3):1136–1141, 2008.