



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Synchronous Balanced Analysis

Citation for published version:

Beica, A & Danos, V 2016, Synchronous Balanced Analysis. in The 5th International Workshop on Hybrid Systems Biology, HSB'16: Hybrid Systems Biology. Lecture Notes in Computer Science, vol. 9957, Springer, pp. 85-94, 5th International Workshop on Hybrid Systems Biology, Grenoble, France, 20/10/16. https://doi.org/10.1007/978-3-319-47151-8_6

Digital Object Identifier (DOI):

[10.1007/978-3-319-47151-8_6](https://doi.org/10.1007/978-3-319-47151-8_6)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

The 5th International Workshop on Hybrid Systems Biology, HSB'16

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Synchronous balanced analysis (short paper)

Andreea Beica and Vincent Danos

École Normale Supérieure, Rue d’Ulm 45, 75005 Paris, France,
beica@di.ens.fr

Abstract. When modeling Chemical Reaction Networks, a commonly used mathematical formalism is that of Petri Nets, with the usual interleaving execution semantics. We aim to substitute to a Chemical Reaction Network, especially a “growth” one (i.e., for which an exponential stationary phase exists), a piecewise synchronous approximation of the dynamics: a resource-allocation-centered Petri Net with maximal-step execution semantics. In the case of unimolecular chemical reactions, we prove the correctness of our method and show that it can be used either as an approximation of the dynamics, or as a method of constraining the reaction rate constants (an alternative to flux balance analysis, using an emergent formally defined notion of “growth rate” as the objective function), or a technique of refuting models.

Keywords: chemical reaction networks, approximation, resource allocation, max-parallel execution of Petri Nets, flux balance analysis

1 Introduction

When studying certain cellular processes, the assumption is that the remainder of the cell can either be ignored or considered constant. Despite this assumption, intracellular processes rarely work in isolation, but rather in continuous interaction with the rest of the cell. Furthermore, the cell has finite resources, so committing resources to one task reduces the amount of resources available to others. All cells experience these trade-offs, which potentially modify all cellular processes, but are often overlooked. In this paper, we propose a piecewise synchronous approximation of the dynamics of Chemical Reaction Networks (CRN) based on finite resource allocation between reactions, that puts these trade-offs front stage. One goal is to rephrase the mass action run of the system as a problem of optimization: the inter-phase between synchronous runs defines an unknown, the resource split, and we can ask for the best split (e.g., the one which minimizes parallel completion time, or maximizes growth rate). Our method allows us to define a formal notion of growth rate for our type of Petri Net execution, that can serve as an improved “biomass objective function” [11] for a constraint method similar to flux balance analysis (FBA) [11]: “Synchronous Balanced Analysis”.

Related work. While most intracellular growth processes are well characterized, the manner in which they are coordinated under the control of a scheduling

policy is not well understood. When fast replication is sought, a schedule that minimizes the completion time is naturally desirable. But when resources are scarce, in the worst case it is computationally hard to find such a schedule [1],[2]. The scheduling problem of a self-replicating bacterial cell is studied in [3]. A mathematical cell model that respects the resource trade-offs experienced by cells is built in [4]. The concept of maximally parallel execution already appears in the literature on P-systems [5], and in Levy’s family reductions [6], while in [7], the authors use it to develop a Petri Net execution semantics that resembles biology. The scheduling policy of cells is also tackled in [8], where the notion of *bounded asynchrony* is introduced. In [9], the author introduces a constraint method that generalizes FBA to the stochastic case, allowing models to be discriminated using second order moments.

This paper is organised as follows. The next section contains an overview of how CRNs can be modeled using Petri Nets (PNs), as well as our definition of max-parallel execution of a PN. Next we introduce our piecewise synchronous execution semantics and show it encompasses max-parallel execution. We then demonstrate that, at least in the case of unimolecular reactions, it recreates the usual ODE system dynamics, and that it can be used either as an approximation of the dynamics, or as an alternative to flux balance analysis. The final section concludes with a summary and outlook regarding further work on the subject.

2 Modeling Chemical Reaction Networks

Definition 1. *A chemical reaction network (CRN) is a tuple $\langle \mathcal{S}, \nabla^-, \nabla^+, \mathcal{R}, \kappa \rangle$, where $\mathcal{S} = \{S_1, \dots, S_s\}$ is a finite set of species, ∇^- and ∇^+ are $r \times s$ consumption, respectively production stoichiometry matrices, $\mathcal{R} = \{r_1, \dots, r_r\}$ is a finite set of reactions, and $\kappa : \mathcal{R} \rightarrow \mathbb{R}_{>0}$ associates a (positive) rate constant to each reaction.*

Each reaction $r_i \in \mathcal{R}$ is of the form $\sum_{j=1}^s \nabla_{ij}^- S_j \xrightarrow{k_i} \sum_{j=1}^s \nabla_{ij}^+ S_j$, and the reaction network can be written compactly in matrix-vector form as $\nabla^- S \xrightarrow{k} \nabla^+ S$, with S the species column vector, and k the rates column vector.

The state of a system can be represented as a multiset of the concentrations of all the chemical species in the network, denoted by $\mathbf{x} = (\mathbf{x}_{S_1}, \dots, \mathbf{x}_{S_s}) \in \mathbb{N}^s$. Applying the law of mass action, the dynamics of the reaction network assumed to be in state \mathbf{x} are given by the kinetic equations:

$$\frac{d\mathbf{x}}{dt} = (\nabla^+ - \nabla^-)^T \cdot K \cdot \mathbf{x}^{\nabla^-} \quad (1)$$

with $K = \text{diag}(\kappa_1, \dots, \kappa_r)$, and \mathbf{x}^A the “vector-matrix” exponentiation: for $x = [x_1 \dots x_q]^T \in \mathbb{R}^q$ and non-negative $A = [A_{ij}] \in \mathbb{R}^{p \times q}$, x^A denotes the element of \mathbb{R}^p whose i^{th} component is $\prod_{j=1}^q x_j^{A_{ij}}$ (see Appendix A for an example).

This describes the continuous, deterministic model of a chemical reaction network, which is a limit of the stochastic model when all species are highly abundant [10]. One way to model CRNs is by using Petri Nets, as we recall below.

Definition 2. A Petri Net is a tuple $N = \langle S, T, W, m_0 \rangle$, where S is a finite set of places, T is a finite set of transitions, $W : ((S \times T) \cup (T \times S)) \rightarrow \mathbb{N}$ is the arc weight mapping and $m_0 : P \rightarrow \mathbb{N}$ is the marking representing the initial distribution of tokens.

A transition is enabled when all of its requirements are met (in the current marking, every place that has an incoming arc to the transition has at least as many tokens as the weight of its incoming arc), and it is fired by consuming all required tokens and producing new tokens.

Representing a CRN using Petri Nets is straightforward: places represent species (genes, proteins, complexes), and transitions represent reactions.

The most commonly used execution semantics of Petri Nets is the interleaving execution semantics: in each step, select one enabled transition non-deterministically, fire it, then repeat. This semantics describes totally asynchronous behaviour, which does not capture the concurrent nature of cellular behavior, where all reactions can happen in parallel. A better suited semantics, proposed in [7], and which we adapt in this paper, is presented below.

2.1 Max-parallel execution semantics of Petri Nets


The max-parallel execution semantics can be informally described as “execute greedily as many transitions as possible in one step”[7]; we formalize this description in Def.3.

The markings of a PN can be regarded as non-negative integer S-vectors. Its transition relation can be then described as a pair of $|S| \times |T|$ incidence matrices: ∇^- defined by $\forall s, t : \nabla^-(s, t) = W(s, t)$ and ∇^+ defined by $\forall s, t : \nabla^+(s, t) = W(t, s)$. Then their difference $\nabla = \nabla^+ - \nabla^-$, the composite change matrix, can be used to describe the reachable markings: for each sequence of transitions, w , $o(w)$ will denote T-the vector that maps every transition to its number of occurrences in w . Then, we have $reach(m_0) = \{m \mid \exists w : m = m_0 + \nabla \cdot o(w) \wedge w \text{ is a firing sequence}^1 \text{ of } m_0\}$.

Definition 3. A max-parallel execution step in a PN at state m is a positive T-vector v such that:

1. v is **compatible** with m (i.e., there are enough tokens to do everything, in any order): $0 \leq m - \nabla^- v$
2. v is **exhaustive** (i.e., no reaction is enabled after firing): $\forall j \in T, m - \nabla^- v \not\geq r_j$, where r_j is the j^{th} column of ∇^- .

Figure 1 depicts the Petri Net of the CRN (we ignore the products):



with initial marking $m_0 = (9, 9, 9)$, and its possible max-parallel strategies.

¹ a sequence of transitions that can be fired consecutively starting from a marking




Fig. 1. A network with exactly 2 possible maximally parallel steps: $\{t_0 \times 3, t_2 \times 9\}$ and $\{t_0 \times 2, t_1, t_2 \times 6\}$

3 Piecewise synchronous execution semantics of CRNs

In order to deal with resource allocation in CRNs, we construct an execution semantics of PNs that is piecewise synchronous (and includes the maximally parallel strategies): among all traces of execution of a PN, we single out a subset of semi-synchronous ones.

Execution proceeds in an alternation of resource allocation (“split”) and depletion (“burst”) (plus a phase of collection of the products between depletion and the next allocation). Allocation of tokens to their possible transitions is done via a $|T| \times |S|$ matrix α , where α_{ij} denotes the fraction of resource j being allocated to reaction i , meaning that:

$$\forall j \in S, \sum_{i \in T} \alpha_{ij} \leq 1 \quad (2)$$

The “burst” phase consists of the execution of all transitions in parallel until the available input are reduced to a small constant fraction of the initial amount (for reasons explained below). This small constant remainder we impose on our semantics is both the reason for the inequality sign in (2), and the reason our executions will be in the spirit of max-parallel executions, rather than max-parallel in the strict sense: whereas the max-parallel execution seeks to deplete all available resources, ours consumes them up to a fixed level (noted ϵ in the following).

3.1 Resource allocation: relation to max-parallel execution

Suppose a CRN and assume $\alpha \in \mathbb{R}_+^{|T| \times |S|}$ a resource allocation matrix defined as above, $m \in \mathbb{R}^{|S| \times 1}$ a marking of the PN (i.e., a resource array) and v a (potentially max-parallel) reaction vector. We note that zero-order reactions ($\emptyset \rightarrow \dots$) are not taken into account, as the question of resource allocation does not apply to them.

We define the operation \star as:

Definition 4. $(\alpha \star m)_j \stackrel{\text{def}}{=} \min_{i \in S} \left(\frac{\alpha_{ji}}{\sqrt{v_{ij}}} \cdot m_i \right)$

Then Theorem 1 states that our execution semantics encompasses the max-parallel strategy (each max-parallel strategy is associated with a resource allocation matrix α , see Appendix B for proof).

Theorem 1. $\forall v$ compatible with a resource array m (and potentially max-parallel), $\exists \alpha$ resource allocation matrix s.t. $v = \alpha \star m$. Furthermore, if the CRN is unary, there is unicity of α .

For bimolecular reactions, α defined as in Appendix B is no longer the unique solution of $v = \alpha \star m$; intuitively, for a bimolecular reaction $r_k : A + B \rightarrow \dots$, a different resource allocation matrix α' can be created by allocating to r_k whatever amount of species A is not allocated in α . The use of \min in Def. 4 ensures that $v = \alpha' \star m$ (see Appendix C for an example).

3.2 Unary CRNs and growth rate

Consider a CRN comprised exclusively of unimolecular reactions. Then, for m an initial marking, ∇ the composite change matrix, and α a resource allocation matrix, the state of the system after one execution with the α split is given by the matrix $(I + \nabla \cdot \alpha) \cdot m$, with I the identity matrix. More generally, after k iterations of the “split-burst” execution with the same split α , the state of the system is:

$$D_\alpha^k \cdot m, \text{ with } D_\alpha = I + \nabla \cdot \alpha \quad (3)$$

Let $\lambda_1 > \lambda_2 > \dots$, the eigenvalues of D_α , and $E(\lambda_i)$ the eigenspace associated to each λ_i . If the initial marking vector can be decomposed as $m = \sum_i m_i$, with $m_i \in E(\lambda_i)$, then we can rewrite:

$$D_\alpha^k \cdot m = \lambda_1^k \cdot [m_1 + \sum_{i \geq 2} (\frac{\lambda_i}{\lambda_1})^k \cdot m_i] \quad (4)$$

If $\lambda_1 < 1$, given (4), the system will eventually go extinct. Also, if $m_1 \in E(\lambda_1)$ is not unique, one has redundancy of growth (i.e., growth on multiple species/sources). We thus assume that $\lambda_1 > 1$, alongside uniqueness of m_1 and unidimensionality of eigenspaces.

Under these assumptions, as $\frac{\lambda_i}{\lambda_1} < 1$, for a big-enough k , the state of the system will converge to $\lambda_1^k \cdot m_1$, meaning that the growth rate of the system is given by λ_1 , the biggest eigenvalue of D_α .


4 Depletion time of unary CRNs

Consider a unary reaction $S_i \rightarrow \dots @k$; the time evolution of the concentration of species S_i is given by the ordinary differential equation $\frac{dS_i}{dt} = -k \cdot [S_i]$. Then at time t , the concentration of species S_i is: $S_i(t) = S_i(0) \cdot e^{-kt}$. Equivalently, the mean time of depletion of the reaction (i.e., bringing the level of species S_i to a specified amount $0 < s_i \leq S_i(0)$) is

$$\tau = k^{-1} \log\left(\frac{S_i(0)}{s_i}\right) \quad (5)$$

We note that s_i is a convention (the remainder of the reaction), or rather $\frac{S_i(0)}{s_i}$ is the relative amount we consume off the input (e.g., $s_i = 1\% S_i(0)$); the point being that we cannot deplete the whole amount of S_i , as that would take time $\tau = \infty$. Herein lies the main difference between our method and max-parallel execution, as mentioned in the beginning of Section 3.

Now suppose n unary reactions with the same input :



We then allocate in parallel the S_i 's between the n reactions, according to our execution semantics: $\forall j \in N_i$, reaction j receives $\alpha_{ji} \cdot S_i(0)$ input, and has a remainder of $s_{i,j}$. Then the depletion time of reaction j is :

$$\tau_j = k_j^{-1} \cdot \log(\alpha_{ji} \cdot \frac{S_i(0)}{s_{i,j}}) \quad (7)$$

4.1 Isochronicity and iso-remainder assumptions

In order to have a synchronous execution, we fix the same depletion time, τ for all reactions of the unary CRN. Then, from (7):

$$\alpha_{ji} = \beta^{k_j} \cdot \epsilon_{i,j} \quad (8)$$

where $\beta = e^\tau$ and $\epsilon_{i,j} = \frac{s_{i,j}}{S_i(0)}$. In this notation, $\epsilon_{i,j}$ is the remaining *percentage* (relative amount) of the total amount of S_i available in the beginning of the split round (i.e., $S_i(0)$), after reaction j is executed.

Furthermore, if we assume the same relative amount, s_i , remains after executing all n reactions, we have:

$$\forall j \in N_i, \alpha_{ji} = \epsilon_i \cdot \beta^{k_j}, \quad (9)$$

with $\epsilon_i = \frac{s_i}{S_i(0)}$.

Under these assumptions, the dynamics of the system in state m , for species S_i , is given by:

$$\Delta m(\tau) = \nabla \cdot (\alpha_{\cdot i} - \epsilon_{\cdot i}) \cdot m, \quad (10)$$

where $\alpha_{\cdot i}$ denotes the i^{th} column of the resource allocation matrix α :

$$\forall j \in T, \alpha_{ji} = \begin{cases} \epsilon_i \cdot \beta^{k_j}, & \text{if } j \in N_i \\ 0, & \text{otherwise} \end{cases}, \text{ and } \forall j \in T, \epsilon_{ji} = \begin{cases} \epsilon_i, & \text{if } j \in N_i \\ 0, & \text{otherwise} \end{cases}.$$

Then, from (2)² and (9), we have $\epsilon_i = \frac{1}{\sum_{j \in N_i} 1 + \beta^{k_j}}$ and:

$$\hat{\alpha}_{S_i}(\tau, k_{S_i}) \stackrel{\text{def}}{=} [\alpha_{\cdot i} - \epsilon_{\cdot i}] = \left[\frac{e^{\tau \cdot k_j} - 1}{\sum_j 1 + e^{\tau \cdot k_j}} \right] \quad (11)$$

² the inequality of (2) is here explicitly expressed via the remainder $\epsilon : \sum_{i \in T} \alpha_{ij} \leq 1$ is the same as $\sum_{i \in T} (\alpha_{ij} + \epsilon_j) = 1$

(NB: $\left[\frac{e^{\tau \cdot k_j} - 1}{\sum_j 1 + e^{\tau \cdot k_j}}\right]$ actually denotes a T-vector that has 0 in the components representing reactions $j \notin N_i$)

Then, by injecting (11) into (10), one can easily observe that when $\tau \rightarrow 0$, $\Delta m(\tau)$ recreates the usual ODE system dynamics (as in formula 1):

$$\frac{\Delta m}{\tau} \approx \nabla \cdot [k_j] \cdot m \quad (12)$$

(NB: when $x \rightarrow 0$, $e^x \approx 1 + x$)

5 Applications

Based on formulae (4) and (11), our method can be interpreted either in a concrete way, as an approximation of the real system’s dynamics (and be used for simulation purposes), or in an abstract way, as an alternative to flux balance analysis (“Synchronous Balanced Analysis”).

5.1 Concrete interpretation: approximation of system dynamics

As an approximation of the dynamics, under the unary/isochronous/iso-remainder assumptions, ours is a temporised discrete execution dynamics, that, when $\tau \rightarrow 0$, recreates the usual ODE dynamics. If we fix τ the execution time-step, and \hat{k} the reaction rate vector, we can determine α , the resource allocation matrix, and ϵ , the remainder percentage (cf. (11)). The “iso” assumption represents a way of decoupling production and consumption in the chemical network, in the spirit of Karr’s modular systems [12]; intuitively, it can be interpreted as: “in a parallel execution of a reaction set, there is no waiting for the slowest reaction to complete”.

As a simulation method, it can be viewed as a big-step approximation of an integrator, resembling a deterministic τ -leaping [13].

5.2 Abstract interpretation: Synchronous Balanced Analysis

Conversely, if the resource allocation matrix α is fixed, our execution semantics can be interpreted as an alternative to Flux Balance Analysis [11], in order to determine the limitations of a metabolic system; this is maybe the most important application of our method. FBA is a constraint-based approach that creates a solution space based on stoichiometric information, which impose constraints on the flow of metabolites through the network. The (flux) solution space can be further reduced via optimization with respect to a mathematical “objective function” representing a biological objective that is relevant to the problem being studied. In the case of predicting growth, the objective is biomass production, which is mathematically represented by an “objective function” Z that indicates how much each reaction contributes to the phenotype.

Our method uses λ , the growth rate, as the objective function. Cf. (4) maximizing λ means finding the resource allocation matrix α with the maximal

biggest eigenvalue. Once this α is fixed, (11) can be used in order to constrain the reaction rates, \hat{k} , as well as the time-step τ , and the remainder ϵ . Hence, our method can be used as a technique of refuting models.

The advantages, when compared to FBA, are twofold: firstly, our method is applicable to growth systems (the implicit assumption of FBA is that the system has reached steady state), thus taking into account the real system kinetics; secondly, the idea of maximizing the biomass is preserved, but the invention of an “objective biomass function” is no longer needed, as it emerges directly from our method: *the growth rate*. The downside of our “Synchronous Balanced analysis” resides in the difficulty of maximizing the biggest eigenvalue of matrix D_α .

6 Conclusions and future work

In this paper, we propose a piecewise synchronous approximation of the dynamics of a (growth) chemical reaction network: a parallel execution semantics of Petri Nets, based on resource allocation. Our method can be interpreted either as an approximation of the real dynamics of the system, or as a constraint method similar to flux balance analysis, and has the advantage of being able to characterize the behavior of a cell using only one construction: the resource allocation matrix α . Consequently, one can eliminate the mechanistic details that deal with resource allocation, and replace them by an abstract vector (α). Furthermore, when compared to flux balance analysis, our method is applicable to growth systems.

Future work. Since the method presented in this paper constitutes work in progress, we plan to extend it into several directions. Firstly, we are interested in the extension of our method to binary reactions: the depletion time cannot be derived in the same way as for unary reactions. A way to potentially (but not completely) relieve this issue is by assuming that no significant change in the concentration of one of the two reactants is being caused by any other reaction during one execution step. We plan on looking into Michelis-Menten like reduction schemes to alleviate this problem. Once the issue of bimolecular reactions is solved, we can construct untrivial examples based on real-life CRNs, that will allow us to test the quality of our method.

Secondly, we plan to use our method to determine possible correlations between growth rate and different model parameters (such as reaction rate constants).

Last but not least, we would compare our method to the τ -leaping simulation, as well as the allocation method proposed by Karr [12], and further study the quality of our method.


References

1. Garey, M.R., Johnson, D.S., Sethi R.: The complexity of flowshop and jobshop scheduling. *Math Oper Res* 1, 117 – 129 (1976)

2. Sinnen, O.: Task Scheduling for Parallel Systems. Wiley-Interscience, Hoboken, NJ (2007)
3. Pugatch, Rami: Greedy scheduling of cellular self-replication leads to optimal doubling times with a log-Frechet distribution. PNAS 112/8, 2611–2616 (2015)
4. Weiße, A.Y., Oyarzún, D.A., Danos, V., Swain, P.S.: Mechanistic links between cellular trade-offs, gene expression, and growth. PNAS 112/9, E1038–E1047 (2015)
5. Păun, G., Rozenberg, G.: A guide to membrane computing. Theoretical Computer Science 287 (1), 73–100 (2002) doi:10.1016/S0304-3975(02)00136-6
6. Lévy, J.-J.: Réductions correctes et optimales dans le lambda-calcul. PhD Thesis, Université Paris 7, January 1978
7. Krepska, E., Bonzanni, N., Feenstra, A., Fokkink, W., Kielmann, T., Bal, H., Heringa, J.: Design Issues for Qualitative Modelling of Biological Cells with Petri Nets. Formal Methods in Systems Biology: First International Workshop, FMSB 2008, Cambridge, UK, June 4-5, 2008. Proceedings, 48–62 (2008)
8. Fisher, J., Henzinger, T.A., Mateescu, M., Piterman, N.: Bounded Asynchrony: Concurrency for Modeling Cell-Cell Interactions. Formal Methods in Systems Biology: First International Workshop, FMSB 2008, Cambridge, UK, June 4-5, 2008. Proceedings, 17–32 (2008)
9. Picard, V.: Réseaux de réactions : de l’analyse probabiliste à la réfutation. PhD Thesis, Université de Rennes 1, December 2015
10. Kurtz, T.G.: Limit theorems for sequences of jump Markov processes approximating ordinary differential processes. Journal of Applied Probability, 8(2),344–355 (1971)
11. Orth, J. D., Thiele, I., Palsson, B. Ø.: What is flux balance analysis? Nature Biotechnology, 28(3), 245–248. <http://doi.org/10.1038/nbt.1614>
12. Karr, J.R, et al: A Whole-Cell Computational Model Predicts Phenotype from Genotype. Cell, 150(2), 389 – 401 (2012). <http://dx.doi.org/10.1016/j.cell.2012.05.044>
13. Gillespie, D.T.: Approximate accelerated stochastic simulation of chemically reacting systems. The Journal of Chemical Physics, 115(4), 1716 (2001). doi:10.1063/1.1378322

Appendix A CRN mass-action kinetic equations

Consider the following chemical reaction network:



Then $\mathcal{S} = \{S_1, S_2\}$, $\nabla^- = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$, $\nabla^+ = \begin{bmatrix} 0 & 2 \\ 1 & 0 \end{bmatrix}$, $\mathbf{x}^{\nabla^-} = \begin{bmatrix} x_1 x_2 \\ x_2 \end{bmatrix}$ and $\frac{d\mathbf{x}}{dt} = \begin{bmatrix} -1 & 1 \\ 1 & -1 \end{bmatrix} \cdot \begin{bmatrix} \kappa_1 & 0 \\ 0 & \kappa_2 \end{bmatrix} \cdot \begin{bmatrix} x_1 x_2 \\ x_2 \end{bmatrix} = \begin{bmatrix} -\kappa_1 x_1 x_2 + \kappa_2 x_2 \\ \kappa_1 x_1 x_2 - \kappa_2 x_2 \end{bmatrix}$

Appendix B Proof of Theorem 1

Proof. Given a resource array m , consider v a (potentially max-parallel) vector compatible at m :

$$\nabla^- v \leq m \tag{13}$$

Then we can construct $\alpha \in \mathbb{R}_+^{|T| \times |S|}$:

$$\alpha_{ji} = \frac{\nabla_{ij}^- \cdot v_j}{m_i} \quad (14)$$

s.t.

$$(\alpha \star m)_j = \min_{i \in S} \left\{ \frac{\alpha_{ji}}{\nabla_{ij}^-} \cdot m_i \right\} = \min_{i \in S} \left\{ \frac{\nabla_{ij}^- \cdot v_j}{m_i} \cdot \frac{m_i}{\nabla_{ij}^-} \mid \nabla_{ij}^- \neq 0 \right\} = v_j \quad (15)$$

Furthermore,

$$\forall j \in S, \sum_{i \in T} \alpha_{ij} = \frac{\sum_{i \in T} \nabla_{ji}^- \cdot v_i}{m_j} \stackrel{(13)}{\implies} \forall j \in S, \sum_{i \in T} \alpha_{ij} \leq 1 \quad (16)$$

i.e. α is indeed a resource-allocation matrix.

If all reactions of the CRN are unimolecular, then :

$$\forall j \in T, \exists ! i_j \in S \text{ s.t. } \nabla_{i_j j}^- \neq 0 \implies \forall j \in T, (\alpha \star m)_j = \frac{\alpha_{j i_j}}{\nabla_{i_j j}^-} \cdot m_{i_j} \quad (17)$$

hence the uniqueness of α . □

Appendix C Non-uniqueness of α for bimolecular reactions

Example 1. (Based on Figure 1.) $m = \begin{bmatrix} 9 \\ 9 \\ 9 \end{bmatrix}$, $\nabla^- = \begin{bmatrix} 3 & 0 & 0 \\ 2 & 5 & 0 \\ 0 & 3 & 1 \end{bmatrix}$, and $v = \begin{bmatrix} 2 \\ 1 \\ 6 \end{bmatrix}$, one of the 2 possible maximally parallel steps ($\{t_0 \times 2, t_1, t_2 \times 6\}$).

$$\text{Then } \exists \alpha = \begin{bmatrix} \frac{6}{9} & \frac{4}{9} & 0 \\ 0 & \frac{5}{9} & \frac{3}{9} \\ 0 & 0 & \frac{6}{9} \end{bmatrix}, \text{ defined as in (14), s.t. } \alpha \star m = \begin{bmatrix} \frac{6}{9} \cdot 9 \cdot \frac{1}{3} \wedge \frac{4}{9} \cdot 9 \cdot \frac{1}{2} \wedge \infty \\ \infty \wedge \frac{5}{9} \cdot 9 \cdot \frac{1}{5} \wedge \frac{3}{9} \cdot 9 \cdot \frac{1}{3} \\ \infty \wedge \infty \wedge \frac{6}{9} \cdot 9 \cdot \frac{1}{1} \end{bmatrix} =$$

$$\begin{bmatrix} 2 \\ 1 \\ 6 \end{bmatrix} = v.$$

By re-allocating the excess of species A to the first reaction, we get $\alpha' =$

$$\begin{bmatrix} 1 & \frac{4}{9} & 0 \\ 0 & \frac{5}{9} & \frac{3}{9} \\ 0 & 0 & \frac{6}{9} \end{bmatrix}, \text{ a resource-allocation matrix that also verifies } \alpha' \star m = \begin{bmatrix} 1 \cdot 9 \cdot \frac{1}{3} \wedge \frac{4}{9} \cdot 9 \cdot \frac{1}{2} \wedge \infty \\ \infty \wedge \frac{5}{9} \cdot 9 \cdot \frac{1}{5} \wedge \frac{3}{9} \cdot 9 \cdot \frac{1}{3} \\ \infty \wedge \infty \wedge \frac{6}{9} \cdot 9 \cdot \frac{1}{1} \end{bmatrix} =$$

$$\begin{bmatrix} 2 \\ 1 \\ 6 \end{bmatrix} = v \text{ (non-uniqueness of } \alpha \text{ in the bimolecular case).}$$