Exact Fluid Lumpability for Chemical Reaction Networks

Luca Cardelli[1], Mirco Tribastone[2], and Max Tschaikowski[2]

[1] Microsoft Research & University of Oxford
luca@microsoft.com
[2] University of Southampton
{m.tribastone,m.tschaikowski}@soton.ac.uk

Abstract. We present a technique for the automatic reduction of chemical reaction networks (CRNs). We study *exact fluid lumpability*, a partition of species satisfying the dynamical property that species in the same block have the same output at all times if initialized equally. Inspired by analogous approaches in traditional models of computation such as labelled transition systems, we characterize this property as a behavioral equivalence over species which can be checked only using CRN structural information. We give an algorithm to construct a quotient CRN induced by exact fluid lumpability. Further, we provide an algorithm that computes the coarsest partition in polynomial time. As an application, we find significant, lossless reductions in a number of models of biological processes available in the literature. In two cases we allow the analysis of benchmark models which would be otherwise intractable due to their memory requirements.

1 Introduction

Chemical reaction networks (CRNs) are a popular mathematical model of systems in a number of disciplines including organic and inorganic chemistry, biochemistry, and systems biology. Indeed, they represent the kernel language for higher-level graph-rewriting approaches such as BioNetGen [1] and the κ-calculus [11], for modeling biochemical reaction networks, and the Systems Biology Markup Language, `http://sbml.org/`, the reference format for models of biological processes.

The dynamics of a CRN is typically expressed in terms of a system of ordinary differential equations (ODEs) or as a Markov process giving the trajectory of the species' concentrations or molecular counts, respectively, as a function of time. In either case, since no closed-form solutions are available in general, CRNs are typically analyzed using numerical solvers. Thus, finding model reductions that preserve the dynamics at a cheaper computational cost is imperative, especially for CRNs of complex biological mechanisms consisting of a large number of species and reactions [8, 7, 15, 13, 10].

In this paper we introduce a technique for the automatic reduction of CRNs according to the notion of *exact fluid lumpability* (EFL). It identifies a partition of the species of a CRN where species within the same partition block have undistinguishable ODE solutions whenever they start with equal initial conditions. Building on a species-as-process analogy, we characterize EFL in terms of behavioral relations between species.

The main consequence is that EFL can be checked *structurally*, i.e., statically by analyzing the reactions in the model, instead of considering the underlying differential system. Furthermore, we develop an efficient partition refinement algorithm [19] which computes the coarsest EFL refinement of an input partition of species in polynomial time. Since equivalent species have the same ODE trajectories, EFL will induce an aggregated ODE system yielding one equation for each partition block which incurs no loss of information with respect to the original model. We lift this aggregation to the CRN level by defining, for any given EFL partition, a quotient CRN in a canonical form characterized by the fewest number of reactions to describe the aggregated ODEs.

We studied the effectiveness of EFL in reducing realistic models of biological interest. Despite the arguable strictness imposed by EFL — the ODE trajectories must coincide at all time points — we did find significant reductions in a number of models of biological processes available in the literature, leading to speed-ups in the ODE solution runtimes of up to four orders of magnitude. In two cases, EFL automatically enabled a very rapid analysis of benchmark models presented in [21] which were intractable in our experimental set-up due to the excessive memory requirements.

EFL can be used in a *black-box* manner as an automated tool for CRN reduction. However, in our case studies it afforded an interpretation also of biological relevance. For instance, we found EFL beneficial in biochemical models involving complexes (e.g., proteins) having multiple sites that can be modelled as being independent and with equivalent capability, intended as their ability to change state (e.g., phosphorylate or methylate) or to bind other molecules with equal rates. The original CRN associates a distinct species to each possible configuration of the complex, to track which site is in which state, leading to an exponential blow-up. Instead, the EFL reduction maintains only canonical representatives that do not track the specific site but how many independent sites are in a particular configuration of local states.

Further related work. Behavioral equivalences have been recently proposed in [20] for comparing CRNs; however, the analysis is carried out at the qualitative level, i.e., ignoring rates and dynamical evolution. The study of EFL for Markovian process algebra in [24] is different for two reasons. First, checking a candidate partition is difficult due to universal quantifiers over the ODE variables present in the conditions. Second, no algorithm for computing the coarsest partition was developed.

There are complementary techniques in related approaches. The notion of *fragmentation* for the κ language presented in [10] identifies a transformation of the state space yielding a sub-space with a closed dynamics, i.e., whose ODEs depend only on the variables of that subspace. This transformation is intended in the sense of the theory ODE lumpability [23, 18], whereby in the aggregated ODE system each variable has a solution equal to a linear combination of the solutions of the original variables. This causes loss of information in general since the concentration of individual species cannot be recovered, similarly to related (but non-algorithmic) earlier approaches to model reductions in biochemical networks [2, 8, 7]. In EFL, instead, the canonical representative gives the exact trajectory of each species within the equivalence class — i.e., the reduction is lossless. In addition, it is also possible to identify the following differences.

First, fragmentation works at a higher modeling level, as it is based on a structural analysis of the rewrite rules instead of the lower-level CRN. Second, it is concerned

$mi_1 : Y_0 + Z_0 \to^{\alpha_1} Z_0 + Y_1$

$mi_2 : Y_1 + Z_0 \to^{\alpha_2} Z_0 + Y_2$

$mi_3 : Y_2 + Y_0 \to^{\alpha_3} Y_0 + Y_1$

$mi_4 : Y_1 + Y_0 \to^{\alpha_4} Y_0 + Y_0$

$mi_5 : Z_2 + Z_0 \to^{\alpha_1} Z_0 + Z_1$

$mi_6 : Z_1 + Z_0 \to^{\alpha_2} Z_0 + Z_0$

$mi_7 : Z_0 + Y_0 \to^{\alpha_3} Y_0 + Z_1$

$mi_8 : Z_1 + Y_0 \to^{\alpha_4} Y_0 + Z_2$

(a) CRN for MI

$am_1 : X_0 + X_2 \to^{\alpha_1} X_2 + X_1$

$am_2 : X_1 + X_2 \to^{\alpha_2} X_2 + X_2$

$am_3 : X_2 + X_0 \to^{\alpha_3} X_0 + X_1$

$am_4 : X_1 + X_0 \to^{\alpha_4} X_0 + X_0$

(b) CRN for AM

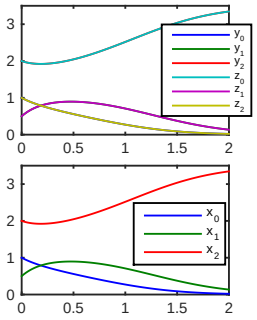

(c) ODE solutions

Fig. 1: Running example. Emulation between a mutual inhibition (MI) network and a CRN for approximate majority (AM).

with *static* information, i.e., it does not involve information about the kinetic rates of the reactions. As a result, it will yield sufficient conditions for aggregation. Instead, EFL is fully characterized, that is, two species are in the same block *if and only if* their ODE solutions are equal at all time points. On the other hand, [10] can be checked on the set of rewrite rules, which is typically much smaller than the CRN size. Third, the symmetries detected at the underlying ODE level are in general not comparable; that is, it is not difficult to find ODE models which can be aggregated by EFL and not by ODE lumpability (e.g., our running example, presented next), and vice versa. Fourth, fragmentation may introduce an additional loss of information since there may be species that cannot be observed in any fragment.

The work in [3] proposes a fragmentation approach to capture *symmetric sites* in κ, i.e., sites with equivalent capabilities in the same sense as in the aforementioned interpretation of EFL, but through ODE lumpability. It was not possible to empirically compare EFL with [3] since an implementation is currently not available. However we remark that, unlike [2, 8, 7, 10, 3], EFL is not domain-specific: we present next a case study where EFL can be used to relate the dynamics of models of biological mechanisms [4] which do not involve post-translational modifications and complex formation.

2 Running Example

In [4], Cardelli presents the notion of *emulation* between two CRNs. Essentially, it states that a mapping of species and reactions from one *source* CRN to a *target* CRN is such that related species have the same ODE trajectories whenever they start from identical initial conditions. This dynamical property makes [4] the closest approach to ours, hence we consider a case study to clarify the relationship with [4], and to introduce a running example for this paper.

Figure 1a shows the CRN for a mechanism of mutual inhibition (MI) [14], using standard notation that will be defined precisely later. There are two molecules, Y and

Z, which can evolve through *molecular states* 0, 1, and 2 (denoted by CRN species Y_0, Y_1, and Y_2). States 0 and 2, model the *activated* and *inhibited* state, respectively; state 1 is an intermediate condition that is reached at every transition between activation and inhibition. Species can influence each other. For instance, reactions mi_1 and mi_2 model that the activated state of Z inhibits Y, since they implement a two-step modification that turns Y_0 into Y_2. Dually, reactions mi_7 and mi_8 model that an active Y inhibits Z.

Figure 1b represents a network that computes the well-known approximate majority (AM) population protocol [5]. It is capable of turning all elements of either species 0 or 2 into the other (passing through the intermediate, *undecided*, state 1), depending on which of the two species had the highest initial concentration. For instance, reaction am_1 says that an encounter between X_0 and X_2 will cause X_0 going into the undecided state X_1. Instead, reactions am_2 and am_4 say that undecided elements take the state of the partner species with which they interact.

In this example, it can be shown that there is an emulation between the two CRNs via a mapping of the species and reactions given by

$$Y_0 \mapsto X_0, Y_1 \mapsto X_1, Y_2 \mapsto X_2, Z_0 \mapsto X_2, Z_1 \mapsto X_1, Z_2 \mapsto X_0 \tag{1}$$

$$mi_1 \mapsto am_1, mi_2 \mapsto am_2, mi_3 \mapsto am_3, mi_4 \mapsto am_4,$$
$$mi_5 \mapsto am_1, mi_6 \mapsto am_2, mi_7 \mapsto am_3, mi_8 \mapsto am_1 \tag{2}$$

respectively. The top plot in Figure 1c shows the ODE solution of MI when species that are mapped onto the same species in AM are initialized equally, showing trajectories overlapping in pairs. The bottom plot shows that the trajectories correspond to those of AM under analogous initial conditions. (It also illustrates the intent of AM — the concentration of X_0 vanishes as time goes by because it was lower than X_2 initially.)

Emulations can be studied when both the source and the target CRNs are available — the modeller is intended to have the suspicion that, for some given CRN, another CRN might be related to it. But emulation cannot be used when one wants to discover symmetries among species *within the same given CRN* — e.g., to detect in MI alone that the trajectories of Y_2 and Z_0 are equal. Thus emulation is not useful for model reduction because a-priori information about the structure of a quotient CRN is not available.

Yet, we will use and reconcile with the theory of emulation in two ways: (i) to explain a structural relationship between a given CRN and the quotient CRN induced by an EFL partition; and (ii) to show that emulation can be interpreted as an EFL partition on the *union* CRN with species and reactions from both the source and the target CRNs.

3 Overview of Results

Using our running example, in this section we summarize the main contributions of this paper and anticipate the results of the application of EFL to the case studies. Instead, a formal development of the theory will be provided in later sections.

Behavioral characterization of EFL (Section 4.2). We understand EFL in terms of a behavioral relation between species. To do so, we introduce the notion of *fluxes*. They summarize the contribution of each species by each reaction. A flux is a triple recording

(i) the species' net stoichiometry in the reaction (if different than 0), (ii) the rate parameter of the reaction, and (iii) the reactants involved. For instance, the flux due to reaction mi_6 in Z_0 is $(1, \alpha_2, Z_0 + Z_1)$. Thus the overall behavior of a species A is the multi-set (to account for reactions with same reactants, rates, and stoichiometries) of fluxes for all reactions, $T(A)$. For instance, in MI we will have that $T(Z_0) = \{|(1, \alpha_2, Z_0 + Z_1), (-1, \alpha_3, Y_0 + Z_0)|\}$ and $T(Y_2) = \{|(1, \alpha_2, Y_1 + Z_0), (-1, \alpha_3, Y_0 + Y_2)|\}$.

Our first contribution is to characterize EFL as a partition where species within the same block have indistinguishable fluxes whenever the species in the fluxes are renamed with canonical representatives induced by the partition. For instance, $T(Z_0)$ and $T(Y_2)$ become equal to $\{|(1, \alpha_2, Z_0+Z_1), (-1, \alpha_3, Y_0+Z_0)|\}$ once the reactants are replaced by canonical representatives of the EFL partition $\tilde{\mathcal{H}} := \{\{Y_0, Z_2\}, \{Y_1, Z_1\}, \{Y_2, Z_0\}\}$.

A partition-refinement algorithm (Section 4.3). The behavioral characterization allows for checking EFL *structurally* on the reactions in the model. Using this, we develop a polynomial-time algorithm (cf. Algorithm 1) to compute the coarsest EFL refinement (cf. [19]) of a partition of species given as an input. The input partition can be chosen arbitrarily, for instance to tell apart species that cannot be EFL equivalent because already their initial conditions are known to be different. In particular, in the case studies we will use initial partitions that are *consistent with the initial conditions* specified in the original published model: two species are in the same block if and only if their initial conditions are equal. This clearly ensures that the algorithm's output is a partition where two species in the same block have the same initial conditions in the original CRN model (because the output is a refinement of the initial partition).

An algorithm to compute the quotient CRN from an EFL partition (Section 4.4). EFL will induce an aggregated ODE system yielding one equation for each equivalence class. We lift this aggregation to the CRN language level by defining, for any given EFL partition, a quotient CRN in a canonical form characterized by the fewest number of reactions to describe the aggregated ODE. If the EFL partition is consistent with the initial conditions then the aggregated ODE system will carry the same solution, in lossless way, as the original ODE system.

Case studies. We tested EFL on biological models publicly available as BioNetGen files, comparing the runtime of the ODE solution of the original model against the runtime of the EFL-reduced model. The experiments are replicable using a Java-based prototype available at `http://users.ecs.soton.ac.uk/mt2y12/EFL.zip`.

Table 1 summarizes the results. For each model, the name corresponds to the filename available in our repository; *e2–e9* are synthetic benchmarks proposed in [21] to study the scalability of the CRN solution with respect to the presence of multiple independent sites in complexes; *m1* is a model of pheromone signaling in yeast [22], *m2* models a signaling pathway through the Fcϵ receptor complex [21, 12]; *m3* is a model of EGF and insulin signaling [7, 17]; *m4* is a model of tyrosine phosphorylation and adaptor protein binding [6]; and *m5* is a MAPK pathway model [16]. Column *Int.* shows the final time point for the ODE numerical solution, taken from the referenced papers when available; in the models *e2–e9* this information was not provided, hence we solved the ODEs until they approached steady state (estimated at ca 50 time units). Labels $|S|$ and

		Original Model			*EFL Reduction*					
Id	*Int.*	$\|S\|$	$\|R\|$	*ODE*	$\|S\|$	$\|R\|$	*ODE*	*Alg. 1*	*Alg. 2*	*Speed-up*
e9	50	262146	3538944	—	220	540	9.20E-1	1.11E+4	1.25E+1	—
e8	50	65538	786432	—	167	396	1.20E-1	6.47E+2	1.22E+0	—
e7	50	16386	172032	6.64E+3	122	280	9.70E-2	3.33E+1	3.18E-1	6.84E+4 (1.97E+2)
e6	50	4098	36864	1.77E+2	86	189	2.30E-2	2.99E+0	1.50E-1	7.69E+3 (5.59E+1)
e5	50	1026	7680	7.85E+0	58	120	2.21E-2	7.83E-1	4.10E-2	3.55E+2 (9.28E+0)
e4	50	258	1536	4.64E-1	37	70	1.60E-2	1.30E-1	3.80E-2	2.90E+1 (2.52E+0)
e3	50	66	288	7.25E-2	22	36	1.51E-2	4.80E-2	1.70E-2	4.80E+0 (9.05E-1)
e2	50	18	48	5.05E-2	12	15	1.68E-2	2.70E-2	5.00E-3	3.01E+0 (1.03E+0)
m1	3600	14531	194054	3.72E+3	6634	29481	3.17E+2	9.29E+1	8.24E-1	1.17E+1 (9.06E+0)
m2	3840	10734	187468	6.36E+2	5574	28016	7.35E+1	5.80E+1	8.57E-1	8.65E+0 (4.80E+0)
m3	120	2768	38320	4.96E+1	2719	15624	7.74E+0	6.82E+0	3.19E-1	6.41E+0 (3.33E+0)
m4	1000	730	5832	2.78E+0	217	405	4.61E-2	8.64E-1	3.70E-2	6.02E+1 (2.93E+0)
m5	600	85	487	6.47E-2	56	129	2.05E-2	6.00E-2	1.30E-2	3.16E+0 (6.92E-1)

Table 1: EFL reduction for BioNetGen CRN models.

$|R|$ indicate the number of species and reactions, respectively, in the original and in the EFL reduced models. Labels *ODE* indicate the runtime of the ODE solution from 0 until the final time point, averaged over three independent executions. The ODEs were solved using Matlab's *ode15s* function for stiff ODE systems, as available in the release R2014b. Columns *Alg. 1* and *Alg. 2* show the runtimes of our implementations of the partition-refinement and the CRN-mimimization algorithms, respectively. All runtimes are rounded to two significant digits, from measurements taken on an ordinary laptop machine equipped with a 1.7 GHz Intel Core i7 processor and 8 GB RAM. Finally, column *Speed-up* shows the ratio between the measured runtimes of the original model and the reduced model. Between parentheses is the speed-up when the EFL reduction time is also accounted for. This represents a worst-case scenario when the reduced model is used only once, but we note that the same minimization can be re-used with any other choice of initial conditions that leads to the same initial partition for Algorithm 1.

It can be observed that EFL is convenient for large-scale CRNs, with speed-ups up to four orders of magnitude. For *e8* and *e9* with EFL it was possible to analyze models whose original CRNs could not open in our experimental set-up due to out-of-memory errors (consistently with [21], where a similar problem is reported for *e9*).

Biological interpretation of EFL. Let us consider model *e2* to show how EFL can exploit the aforementioned symmetries between multiple independent sites. Table 2 shows the output of the partition-refinement algorithm. According to the BioNetGen syntax, a molecule S (singleton block 5) has two binding sites, labelled $p1$ and $p2$, which can be either phosphorylated (state P) or unphosphorylated (state U). When phosphorylated, a site can bind to molecule F; when unphosphorylated, it can bind to molecule E. For instance, singleton block 4 has the CRN species where both sites are phosphorylated and bound to two distinct copies of F. EFL does not distinguish between the two sites. For instance, block 1 collects the two species where only one

Block No.	*Partition Block*
1	$\{F(s!1).S(p1{\sim}P!1, p2{\sim}P), F(s!1).S(p1{\sim}P, p2{\sim}P!1)\}$
2	$\{E(s)\}$
3	$\{F(s)\}$
4	$\{F(s!1).F(s!2).S(p1{\sim}P!1, p2{\sim}P!2)\}$
5	$\{S(p1{\sim}U, p2{\sim}U)\}$
6	$\{E(s!1).S(p1{\sim}U!1, p2{\sim}U), E(s!1).S(p1{\sim}U, p2{\sim}U!1)\}$
7	$\{E(s!1).E(s!2).S(p1{\sim}U!1, p2{\sim}U!2)\}$
8	$\{S(p1{\sim}P, p2{\sim}U), S(p1{\sim}U, p2{\sim}P)\}$
9	$\{E(s!1).S(p1{\sim}P, p2{\sim}U!1), E(s!1).S(p1{\sim}U!1, p2{\sim}P)\}$
10	$\{F(s!1).S(p1{\sim}P!1, p2{\sim}U), F(s!1).S(p1{\sim}U, p2{\sim}P!1)\}$
11	$\{E(s!1).F(s!2).S(p1{\sim}P!2, p2{\sim}U!1), E(s!1).F(s!2).S(p1{\sim}U!1, p2{\sim}P!2)\}$
12	$\{S(p1{\sim}P, p2{\sim}P)\}$

Table 2: EFL partition for model *e2*

copy of F is bound to one of the two sites, when they are both phosphorylated; block 10 is similar, but here only one binding site is unphosphorylated.

This is not the only symmetry captured by EFL, since it may not tell apart complexes formed by different molecules. For instance, in *m1* one of the equivalence classes is

$$\big\{\text{Dig2\_gene(promoter!1).Ste12(dig1,dig2,dna!1,mapk)},\\ \text{Fus3\_gene(promoter!1).Ste12(dig1,dig2,dna!1,mapk)},\\ \text{Msg5\_gene(promoter!1).Ste12(dig1,dig2,dna!1,mapk)},\\ \text{Sst2\_gene(promoter!1).Ste12(dig1,dig2,dna!1,mapk)},\\ \text{Ste12\_gene(promoter!1).Ste12(dig1,dig2,dna!1,mapk)},\\ \text{Ste2\_gene(promoter!1).Ste12(dig1,dig2,dna!1,mapk)}\big\}.$$

EFL captures that genes Dig2, Fus3, Msg5, Sst2, Ste12, and Ste2, bind to the protein Ste12 with equal rates. This yields equivalent dynamics for these Ste12-gene complexes, and all those formed by them which are equal up to the gene bound to Ste12.

4 Exact Fluid Lumpability for CRNs

In this section we formally present EFL. We start in Section 4.1 with some preliminaries on CRNs with ODE semantics based on the well-known law of mass action, where the rate of interaction is proportional to the product of the reactant's concentrations. We also briefly review the notions of emulation in [4] which will be needed in the remainder. Section 4.2 presents EFL and its structural characterization via the notion of fluxes. Section 4.3 presents the algorithm for computing the coarsest EFL partition refinement. Finally, Section 4.4 discusses how to compute the quotient CRN induced by EFL.

Notation. We write $A \to B$ and B^A for the functions from A to B. When $f \in A \to B$ and $a \in A$, we set $f_a := f(a)$. Moreover, for any $X \subseteq A$ and $b \in B$, we define

$f(X) := \{b \in B \mid \exists a \in X.(f(a) = b))\}$ and $f^{-1}(b) := \{a \in A \mid f(a) = b\}$. Sets and multisets are denoted by $\{\ldots\}$ and $\{|\ldots|\}$, respectively.

4.1 Preliminaries

Chemical reaction networks. Let $\mathcal{S}$ be a countable universe of species and let $S \subseteq \mathcal{S}$ be a finite set. Either side of a reaction is a multiset of species, i.e., a function in $\mathbb{N}_0^S$ or $\mathbb{R}_0^S$. It associates each species with its multiplicity as a reactant or product. A *reaction* r over S is a triple $(\rho, \pi, \alpha) \in \mathcal{R}_S = \mathbb{N}_0^S \times \mathbb{R}_{\geq 0}^S \times \mathbb{R}_{>0}$, representing the reaction $\rho \to^\alpha \pi$. Using standard notation (cf. Figure 1), we write $\rho \to^\alpha \pi$ as

$$\rho_{X_1} X_1 + \ldots + \rho_{X_n} X_n \to^\alpha \pi_{X_1} X_1 + \ldots + \pi_{X_n} X_n$$

where ρ_{X_i} are the multiplicities (stoichiometric numbers) of the reactant species, π_{X_i} are the multiplicities of the product species, and α is a rate.

The *instantaneous net stoichiometry* $\phi(X, r)$ of a species X in a reaction $r = \rho \to^\alpha \pi$ is the difference between product and reactant multiplicity, times the coefficient α:

$$\phi(X, r) = \phi(X, \rho \to^\alpha \pi) = \alpha(\pi_X - \rho_X)$$

A *chemical reaction network* (CRN) is a pair (S, R) where $R \subseteq \mathcal{R}_S$ is a finite set of reactions over S such that $\rho \to^\alpha \pi, \rho \to^{\alpha'} \pi \in R$ implies $\alpha = \alpha'$. If $\pi \in \mathbb{N}_0^S$ for all $\rho \to^\alpha \pi \in R$, (S, R) is a *natural chemical reaction network* (nCRN). For instance, both MI and AM in Figure 1 are nCRNs, and the theory in [4] is formally provided for nCRNs. Instead, we will use the more general CRN for the quotient network by EFL.

We now give a formulation of standard mass action kinetics. A *state* $\boldsymbol{V} \in \mathbb{R}_{\geq 0}^S$ of a CRN (S, R) is a vector of concentrations for each species. For a reaction $r \in R$ over S, its *mass action* $[r] \in \mathbb{R}_{\geq 0}^S \to \mathbb{R}$ is the product of the reagent concentrations, each to the power of its stoichiometry, that is,

$$[r]_{\boldsymbol{V}} = [\rho \to^\alpha \pi]_{\boldsymbol{V}} := \prod_{X \in S} V_X^{\rho_X} = \boldsymbol{V}^\rho$$

Thus, for instance, for MI we have that $[Y_0 + Z_0 \to^{\alpha_1} Z_0 + Y_1]_{\boldsymbol{V}} = V_{Y_0} V_{Z_0}$ while for some *homoreaction* $A + A \to^\alpha B$ we have that $[A + A \to^\alpha B]_{\boldsymbol{V}} = V_A^2$.

The (autonomous) ODE system $\dot{\boldsymbol{V}} = F(\boldsymbol{V})$ underlying an CRN (S, R) is

$$F : \mathbb{R}_{\geq 0}^S \to \mathbb{R}^S, \quad \text{with components } F_X(\boldsymbol{V}) := \sum_{r \in R} \phi(X, r)[r]_{\boldsymbol{V}}, \text{ for } X \in S\,.$$

For instance, the ODE for the MI network is given by:

$$\begin{aligned}
\dot{V}_{Y_0} &= -\alpha_1 V_{Y_0} V_{Z_0} + \alpha_4 V_{Y_0} V_{Y_1} \\
\dot{V}_{Y_1} &= -\alpha_2 V_{Y_1} V_{Z_0} - \alpha_4 V_{Y_0} V_{Y_1} + \alpha_1 V_{Y_0} V_{Z_0} + \alpha_3 V_{Y_0} V_{Y_2} \\
\dot{V}_{Y_2} &= -\alpha_3 V_{Y_0} V_{Y_2} + \alpha_2 V_{Y_1} V_{Z_0} \\
\dot{V}_{Z_0} &= -\alpha_3 V_{Y_0} V_{Z_0} + \alpha_2 V_{Z_0} V_{Z_1} \\
\dot{V}_{Z_1} &= -\alpha_2 V_{Z_0} V_{Z_1} - \alpha_4 V_{Y_0} V_{Z_1} + \alpha_1 V_{Z_0} V_{Z_2} + \alpha_3 V_{Y_0} V_{Z_0} \\
\dot{V}_{Z_2} &= -\alpha_1 V_{Z_0} V_{Z_2} + \alpha_4 V_{Y_0} V_{Z_1}\,.
\end{aligned} \tag{3}$$

Since the ODE system of a CRN is given by polynomials, function F is locally Lipschitz. Hence, the theorem of Picard-Lindelöf ensures that there exists a unique non-continuable solution of $\dot{\boldsymbol{V}} = F(\boldsymbol{V})$.

Emulation. Let us denote by μ mappings of species and reactions between CRNs. Let us consider two sets of species S and $\hat{S}$ and assume that $\mu \in S \to \hat{S}$ and $\rho \in \mathbb{N}_0^S$. Then, $\mu(\rho) \in \mathbb{N}_0^{\hat{S}}$ is given by $\mu(\rho)_X = \sum_{X' \in \mu^{-1}(X)} \rho(X')$, where $X \in \hat{S}$. Informally, $\mu(\rho)$ is given by $\rho_{X_1}\mu(X_1) + \ldots + \rho_{X_n}\mu(X_n)$. For instance, given $\rho = Y_0 + Z_0$ and the mapping $\mu(Y_0) = \mu(Z_0) = X_0$, we have that $\mu(\rho) = 2X_0$.

Definition 1. *Let (S, R) and $(\hat{S}, \hat{R})$ denote two CRNs with underlying ODE systems F and $\hat{F}$, respectively. A CRN* morphism *from (S, R) to $(\hat{S}, \hat{R})$ is a pair of functions $\mu_S \in S \to \hat{S}$ and $\mu_R \in R \to \hat{R}$.*

- *(μ_S, μ_R) is a* reactant morphism *if, for all $\rho \to^\alpha \pi \in R$ there exist $\hat{\alpha}$ and $\hat{\pi}$ with $\mu_R(\rho \to^\alpha \pi) = \mu_S(\rho) \to^{\hat{\alpha}} \hat{\pi} \in \hat{R}$.*
- *(μ_S, μ_R) is a* stoichiomorphism *whenever $\sum_{r \in \mu_R^{-1}(\hat{r})} \phi(s, r) = \phi(\mu_S(s), \hat{r})$ for all $s \in S$ and $\hat{r} \in \hat{R}$.*
- *(μ_S, μ_R) is an* emulation *if it holds that $F(\hat{\boldsymbol{V}} \circ \mu_S) = \hat{F}(\hat{\boldsymbol{V}}) \circ \mu_S$ for all $\hat{\boldsymbol{V}} \in \mathbb{R}_{\geq 0}^{\hat{S}}$. We then say that $\hat{F}$ emulates F.*

A morphism relates the *structure* of two CRNs, while an emulation relates their dynamics, since it relates the derivatives of the ODE systems underlying the two CRNs. The following result states that reactant morphism and stoichiomorphism allow for a structural characterization of emulations.

Theorem 1. *Let $\mu \in (S, R) \to (\hat{S}, \hat{R})$ be a reactant morphism and a stoichiomorphism. Then, μ is an emulation.*

This theorem is proven in [4] for nCRNs only, but it carries over to CRNs as well.

The mappings (1) and (2) are a reactant morphism and stoichiomorphism. By Theorem 1 there is an emulation between the species' concentrations, so $\dot{V}_{Y_0} = \dot{V}_{Z_2} = \dot{V}_{X_0}$, $\dot{V}_{Y_1} = \dot{V}_{Z_1} = \dot{V}_{X_1}$, and $\dot{V}_{Y_2} = \dot{V}_{Z_0} = \dot{V}_{X_2}$ in our running example. Thus, the ODE solutions will have equal trajectories if the initial conditions satisfy $V_{Y_0}(0) = V_{Z_2}(0) = V_{X_0}(0)$, $V_{Y_1}(0) = V_{Z_1}(0) = V_{X_1}(0)$ and $V_{Y_2}(0) = V_{Z_0}(0) = V_{X_2}(0)$. This formally explains the equivalence between the plots for MI and AM in Figure 1c.

4.2 Exact Fluid Lumpability

We now introduce EFL as a partition of species within a CRN, and then crucially state that species in the same block have the same ODE solutions when initialized equally.

Definition 2 (Exact Fluid Lumpability). *Let (S, R) be a CRN and $\mathcal{H}$ a partition of S. Further, let us call $\boldsymbol{V} \in \mathbb{R}^S$* constant on $\mathcal{H}$ *whenever $V_{X_i} = V_{X_j}$ for all $H \in \mathcal{H}$ and $X_i, X_j \in H$. A partition $\mathcal{H}$ of S is called* exactly fluid lumpable *if $F(\boldsymbol{V})$ is constant on $\mathcal{H}$ whenever $\boldsymbol{V}$ is constant on $\mathcal{H}$. We say that $f : S \to S$ is a* choice function of $\mathcal{H}$ *when $f(X_i) = f(X_j) \in H$ for all $H \in \mathcal{H}$ and $X_i, X_j \in H$.*

Theorem 2. *Let (S, R) be a CRN and $\mathcal{H}$ an EFL partition of S. Then, whenever the initial condition $\boldsymbol{V}(0)$ of $\dot{\boldsymbol{V}} = F(\boldsymbol{V})$ is constant on $\mathcal{H}$, it holds that $\boldsymbol{V}(t)$ is constant on $\mathcal{H}$ for all $t \in I \cap \mathbb{R}_{\geq 0}$, with I being the time domain of $\boldsymbol{V}$.*[3]

Remark 1. As introduced in Section 2, there is a close connection between emulation and EFL, in that emulation implicitly induces an EFL partition. More precisely, let us assume that $(\mu_S, \mu_R) \in (S, R) \to (\hat{S}, \hat{R})$ is an emulation. Then, $\{\mu_S^{-1}(X) \mid X \in \hat{S}\}$ is an EFL partition of (S, R). Indeed, in our running example $\tilde{\mathcal{H}}$ (introduced in Section 3) is induced by the species mapping (1). Additionally, an emulation between two CRNs can be explained as an EFL over the union CRN. For this, let us assume without loss of generality that $S \cap \hat{S} = \emptyset$ — otherwise it is always possible to rename species of either CRN with fresh variables. Then, $\left\{\mu_S^{-1}(X) \cup \{\mu_S(X)\} \mid X \in \hat{S}\right\}$ is an EFL partition of the CRN $(S \cup \hat{S}, R \cup \hat{R})$. For instance, this leads to viewing the mappings $Y_0 \mapsto X_0$ and $Z_2 \mapsto X_0$ of (1) as inducing the EFL block $\{Y_0, Z_2, X_0\}$ on the union CRN.

Given an EFL partition, an aggregated ODE system can be obtained by associating an ODE with each equivalence class. This can be defined by picking only the ODEs of the representatives, and replacing each species in the ODEs with its representative. For instance, using the EFL partition $\tilde{\mathcal{H}}$ and the choice function $\tilde{f}$ defined as

$$\tilde{f}(Y_0) = \tilde{f}(Z_2) = Y_0, \qquad \tilde{f}(Y_1) = \tilde{f}(Z_1) = Y_1, \qquad \tilde{f}(Y_2) = \tilde{f}(Z_0) = Y_2,$$

we get the aggregated ODEs

$$\begin{aligned}
\dot{V}_{Y_0} &= -\alpha_1 V_{Y_0} V_{Y_2} + \alpha_4 V_{Y_0} V_{Y_1} \\
\dot{V}_{Y_1} &= -\alpha_2 V_{Y_1} V_{Y_2} - \alpha_4 V_{Y_0} V_{Y_1} + \alpha_1 V_{Y_0} V_{Y_2} + \alpha_3 V_{Y_0} V_{Y_2} \\
\dot{V}_{Y_2} &= -\alpha_3 V_{Y_0} V_{Y_2} + \alpha_2 V_{Y_1} V_{Y_2}
\end{aligned}$$

Unfortunately, Definition 2 is not convenient to be used directly because it involves a universal quantifier over the whole (uncountable) state space, similarly to [24], as discussed. Thus, we consider structural conditions which only concern the reactions of a CRN. For this, we first formalize the notion of *flux* anticipated in Section 3, and (novelly, to our knowledge) defined as a projection of reactions to species.

Definition 3 (Flux). *For a given CRN (S, R), let $\mathcal{F} = (\mathbb{Z} \setminus \{0\}) \times \mathbb{R}_{>0} \times \mathbb{N}_0^S$ be the set of flux terms. For a given species $X \in S$, the flux of X in (S, R), denoted by $T(X)$, is defined as the multiset of flux terms where the multiplicity of $(z, \alpha, \rho) \in \mathcal{F}$ in $T(X)$ is given by the cardinality $|\{\rho \to^\alpha \pi \in R \mid \pi(X) - \rho(X) = z\}|$.*

A multiset is required because multiple reactions can contribute the same flux term. For instance, in the CRN $A + B \to^\alpha C$, $A + B \to^\alpha D$, we have $T(A) = \{|(-1, \alpha, A + B), (-1, \alpha, A + B)|\}$. It is easy to see that $T(X)$ contains all the information to construct the ODE for species X. In this form, the species' behavior with respect to a partition is obtained by renaming the reactants of fluxes with their representatives.

[3] *Note to reviewers.* Proofs of all statements are given in the appendix, which will be made publicly available in case of acceptance.

Definition 4. *Let (S, R) be a CRN and let $\mathcal{H}$ denote a partition of S. Further, let f be a choice function of $\mathcal{H}$. Then, $T_{\mathcal{H}}(X)$ is defined as the multiset of flux terms where the multiplicity of $(z, \alpha, \rho) \in \mathcal{F}$ in $T_{\mathcal{H}}(X)$ is given by $|\{|(z, \alpha, \rho_0) \in T(X) \mid f(\rho_0) = \rho|\}|$. If no choice function is explicitly given we assume an arbitrary one, e.g., based on a lexicographical order over species.*

Informally, the multiset $T_{\mathcal{H}}(X)$ arises from $T(X)$ by replacing every term (z, α, ρ) in $T(X)$ with $(z, \alpha, f(\rho))$. For instance, using the choice function $\tilde{f}$ for $\tilde{\mathcal{H}}$ we get $T_{\tilde{\mathcal{H}}}(Z_0) = \{|(1, \alpha_2, Y_2 + Y_1), (-1, \alpha_3, Y_0 + Y_2)|\} = T_{\tilde{\mathcal{H}}}(Y_2)$. That is, with respect to partition $\tilde{\mathcal{H}}$, the behavior of Z_0 and Y_2 is characterized by equivalent fluxes.

The *cumulative flux* of X with respect to $\mathcal{H}$, $T^{\Sigma}_{\mathcal{H}}(X)$, sums all elements of $T_{\mathcal{H}}(X)$ which have the same reactants (an idea that appears also in Chemical Reaction Network Theory, see [9] and references therein). Informally, $T^{\Sigma}_{\mathcal{H}}(X)$ is obtained from $T_{\mathcal{H}}(X)$ by replacing all terms $(z_1, s_1, \rho), \ldots, (z_k, s_k, \rho) \in T_{\mathcal{H}}(X)$ with $(z_1 s_1 + \ldots + z_k s_k, \rho)$. In particular, $T^{\Sigma}_{\mathcal{H}}(X)$ is, unlike $T_{\mathcal{H}}(X)$ and $T(X)$, an ordinary set.

Definition 5 (Cumulative Flux). *Let (S, R) be a CRN. Then, for any partition $\mathcal{H}$ of S and $X \in S$, the cumulative flux of X with respect to $\mathcal{H}$ is given by the set*

$$T^{\Sigma}_{\mathcal{H}}(X) = \big\{(s, \rho) \mid \rho \in \mathfrak{R}\big(T_{\mathcal{H}}(X)\big) \ \wedge \ s = \sum_{(z,\alpha,\rho) \in T_{\mathcal{H}}(X)} z\alpha \ \wedge \ s \neq 0\big\},$$

where $\mathfrak{R}\big(T_{\mathcal{H}}(X)\big) = \{\rho \mid \exists z \in \mathbb{Z}.\ \exists \alpha \in \mathbb{R}.\ ((z, \alpha, \rho) \in T_{\mathcal{H}}(X))\}$.

The cumulative fluxes finally allow for a structural characterization of EFL.

Theorem 3. *A partition $\mathcal{H}$ of S is exactly fluid lumpable if and only if $T^{\Sigma}_{\mathcal{H}}(X_i) = T^{\Sigma}_{\mathcal{H}}(X_j)$ for all $H \in \mathcal{H}$ and $X_i, X_j \in H$. The computation of $\{T^{\Sigma}_{\mathcal{H}}(X) \mid X \in S\}$ can be done in $\mathcal{O}\big(|S|^2|R|\log(|R|)\big)$ steps. Similarly, deciding whether a given partition $\mathcal{H}$ is EFL can be done in $\mathcal{O}\big(|S|^2|R|\log(|R|)\big)$ steps.*

4.3 Partition Refinement Algorithm for EFL

Now we turn to automatically finding EFL partitions. We first recall the notion of refinement and introduce afterwards two equivalence relations over the set of species.

Definition 6. *Let $\mathcal{H}_1$ and $\mathcal{H}_2$ denote two partitions of S. Then, $\mathcal{H}_1$ is a* refinement *of $\mathcal{H}_2$ if for any $H_1 \in \mathcal{H}_1$ there exists a (unique) $H_2 \in \mathcal{H}_2$ such that $H_1 \subseteq H_2$.*

Definition 7. *Let (S, R) be a CRN. For a partition $\mathcal{H}$ of S and $X_i, X_j \in S$, set $X_i \sim_{T^{\Sigma}_{\mathcal{H}}} X_j :\Leftrightarrow T^{\Sigma}_{\mathcal{H}}(X_i) = T^{\Sigma}_{\mathcal{H}}(X_j)$. Moreover, set $X_i \sim_{\mathcal{H}} X_j :\Leftrightarrow \exists H \in \mathcal{H}.\ (X_i, X_j \in H)$.*

Next, we show that, for any given partition $\mathcal{G}$ of S, there exists a unique coarsest EFL partition that refines $\mathcal{G}$.

Proposition 1. *Let (S, R) be a CRN. Then, if $\mathcal{H}_1$ and $\mathcal{H}_2$ are EFL partitions of (S, R), the partition $S/(\sim_{\mathcal{H}_1} \cup \sim_{\mathcal{H}_2})^*$, where the asterisk denotes the transitive closure, is EFL as well. Moreover, if $\mathcal{H}_1, \mathcal{H}_2$ are refinements of $\mathcal{G}$, then so is $S/(\sim_{\mathcal{H}_1} \cup \sim_{\mathcal{H}_2})^*$. Crucially, there exists a unique coarsest EFL partition that refines $\mathcal{G}$.*

Algorithm 1 Construction of the coarsest EFL partition via partition refinement.

Require: A CRN $(S, R) = (\{X_1, \ldots, X_n\}, R)$.
Require: A partition $\mathcal{G}$ of S.
  $\mathcal{H} \longleftarrow \mathcal{G}$
  while true do
    Compute $T^{\Sigma}_{\mathcal{H}}(X_1), \ldots, T^{\Sigma}_{\mathcal{H}}(X_n)$
    $\mathcal{H}' \longleftarrow S/(\sim_{T^{\Sigma}_{\mathcal{H}}} \cap \sim_{\mathcal{H}})$
    if $\mathcal{H}' = \mathcal{H}$ **then**
      return $\mathcal{H}$
    else
      $\mathcal{H} \longleftarrow \mathcal{H}'$
    end if
  end while

Algorithm 1 computes the coarsest EFL partition that refines a given input partition. The following theorem proves its correctness.

Theorem 4. *Given a CRN (S, R) and a partition $\mathcal{G}$ of S, Algorithm 1 calculates the coarsest EFL partition that refines $\mathcal{G}$ in at most $\mathcal{O}\big(|S|^3|R|\big(\log(|S|)+\log(|R|)\big)\big)$ steps.*

The idea is to study the sequence $\mathcal{H}_n := S/(\sim_{T^{\Sigma}_{\mathcal{H}_{n-1}}} \cap \sim_{\mathcal{H}_{n-1}})$, where $\mathcal{H}_0$ is the initial input partition $\mathcal{G}$, and $n \geq 1$. The theorem establishes that $(\mathcal{H}_n)_n$ converges to $\mathcal{H}$, the coarsest EFL partition that refines $\mathcal{G}$, by proving that $\mathcal{H}$ is a refinement of $\mathcal{H}_n$ and that $\mathcal{H}_n$ is a refinement of $\mathcal{H}_{n-1}$ for all $n \geq 1$. Obviously, by intersecting with $\sim_{\mathcal{H}_{n-1}}$ in the definition of $\mathcal{H}_n$, one ensures that $\mathcal{H}_n$ is a refinement of $\mathcal{H}_{n-1}$. To see that the intersection with $\sim_{\mathcal{H}_{n-1}}$ is also necessary, consider the CRN $X \to^1 Y$, $X \to^1 Z$. Then, the sequence given by $\mathbb{H}_n := S/\sim_{T^{\Sigma}_{\mathbb{H}_{n-1}}}$ and $\mathbb{H}_0 := \mathcal{G}$ yields $\mathbb{H}_1 = \{\{X\}, \{Y, Z\}\}$ if $\mathcal{G} = \{\{X\}, \{Y\}, \{Z\}\}$. That is, $\mathbb{H}_1$ does not refine $\mathbb{H}_0$ in general. However, in the special case where $\mathcal{G} = \{\{S\}\}$, $(\mathbb{H}_n)_n$ can be shown to converge to the coarsest EFL partition refining $\mathcal{G}$.

4.4 Computation and Characterization of the Quotient CRN by EFL

We now present an algorithm which provides a CRN that induces the aggregated ODE system underlying an EFL partition. As will be seen later, in general there exist several CRNs that have the same ODE system. Therefore, here we consider a canonical CRN representation with two notable properties. (P1) It preserves the structural properties of the original CRN. In particular, the reduced CRN can be related to the original one by means of a reactant morphism which is also a stoichiomorphism. (P2) The CRN is of minimal size, in the sense that any other CRN that yields the same ODEs and is related to the original CRN by means of a reactant morphism which is also a stoichiomorphism has at least the same number of reactions as the canonical representation.

Our algorithm will make suitable transformations to the original CRN in order to build the quotient one. A reactant morphism is induced naturally by mapping all species to their representatives. The original ODE system and the aggregated one are related by

an emulation. Thus, for studying (P1) we need to identify conditions under which reactant morphism and emulation imply a stoichiomorphism. The following result from [4] carries over from nCRNs to CRNs and is a first step towards this goal.

Theorem 5. *Let us assume that $\mu : (S, R) \to (\hat{S}, \hat{R})$ is a reactant morphism and emulation. Assume further that the reactants of $\hat{R}$ are pairwise different, i.e. $\hat{\rho} \to^{\hat{\alpha}} \hat{\pi}, \hat{\rho}' \to^{\hat{\alpha}'} \hat{\pi}' \in \hat{R}$ implies $\hat{\rho} \neq \hat{\rho}'$. Then, μ is a stoichiomorphism.*

Fortunately, unlike for nCRNs, Theorem 5 can be improved in the case of CRNs, removing the assumption of pairwise different reactants. For this, let us first introduce the notion of the ρ-normalization morphism.

Definition 8. *Let (S, R) be a CRN and $\rho \in \mathbb{N}_0^S$.*

- *The ρ-filter R_ρ is given by $R_\rho = \{\tilde{\rho} \to^{\tilde{\alpha}} \tilde{\pi} \in R \mid \tilde{\rho} = \rho\}$.*
- *For $I_\rho \neq \emptyset$ such that $R_\rho = \{\rho \to^{\alpha_i} \pi_i \mid i \in I_\rho\}$, we define the ρ-merge reaction r_ρ by $r_\rho := \rho \to^{\alpha_\rho} \pi_\rho$, where $\alpha_\rho = \sum_{i \in I_\rho} \alpha_i$ and $\pi_\rho = \sum_{i \in I_\rho} \frac{\alpha_i}{\alpha_\rho} \pi_i$.*
- *For $R_\rho \neq \emptyset$, the ρ-normalization morphism $\mu \in (S, R) \to \big(S, (R \setminus R_\rho) \cup \{r_\rho\}\big)$ is given by $\mu_S := id_S$ and $\mu_R(r) := \begin{cases} r & , r \in R \setminus R_\rho \\ r_\rho & , r \in R_\rho \end{cases}$*

For instance, if (S, R) is the AM network in Figure 1, we have that

$$R_{X_0+X_2} = \{X_0 + X_2 \to^{\alpha_1} X_2 + X_1, X_2 + X_0 \to^{\alpha_2} X_0 + X_1\} \quad \text{and}$$
$$r_\rho : X_0 + X_2 \to^{\alpha_1+\alpha_2} \frac{\alpha_2}{\alpha_1+\alpha_2} X_0 + \frac{\alpha_1+\alpha_2}{\alpha_1+\alpha_2} X_1 + \frac{\alpha_1}{\alpha_1+\alpha_2} X_2$$

The next proposition shows that replacing the set of reactions R_ρ by the single reaction r_ρ does not change the underlying ODE system and leads to a CRN which is structurally related to the original one.

Proposition 2. *Let $\mu \in (S, R) \to (\hat{S}, \hat{R})$ be the CRN ρ-normalization morphism, then μ is a reactant morphism, a stoichiomorphism and an emulation.*

Proposition 3 tells us that Proposition 2 can be applied several times without changing the structure or the underlying ODE system.

Proposition 3. *Given a CRN (S, R), by repeated application of Proposition 2 to any $|R_\rho| > 1$, we obtain a CRN $(S, \hat{R})$ with pairwise different reactants. The composition of all ρ-normalization morphisms $(\mu_S, \mu_R) \in (S, R) \to (S, \hat{R})$ is a reactant morphism, a stoichiomorphism and emulation with $\mu_S = id_S$.*

By combining Theorem 1 and Proposition 3, we infer our desired result — an enhanced version of Theorem 5 where the assumption on pairwise different reactants is dropped owing to Proposition 3.

Theorem 6. *Let us assume that $(\mu_S, \mu_R) : (S, R) \to (\hat{S}, \hat{R})$ is an emulation and reactant morphism. Moreover, let $(id_{\hat{S}}, \mu_{\hat{R}}) : (\hat{S}, \hat{R}) \to (\hat{S}, \underline{\hat{R}})$ denote the composition of ρ-normalization morphisms from Proposition 3 which ensures that $\underline{\hat{R}}$ has pairwise different reactants. Then, $(\mu_S, \mu_{\hat{R}} \circ \mu_R) : (S, R) \to (\hat{S}, \underline{\hat{R}})$ is an emulation, reactant morphism and stoichiomorphism.*

Algorithm 2 Calculation of the CRN in canonical form which induces the aggregated ODE system underlying a given EFL partition.

Require: A CRN $(S, R) = \big(\{X_1, \ldots, X_n\}, R\big)$.
Require: An EFL partition $\mathcal{H}$ underlying (S, R).
Require: A choice function μ_S of $\mathcal{H}$, with $\hat{S} := \mu_S(S)$.

(O1) Replace any $\rho \to^{\alpha} \pi$ with $\rho \to^{\alpha} \tilde{\pi}$ where $\tilde{\pi}(X_i) := \begin{cases} \pi(X_i) & , X_i \in \hat{S}, \\ \rho(X_i) & , X_i \notin \hat{S}. \end{cases}$

(O2) Replace any $\rho \to^{\alpha} \pi$ with $\tilde{\rho} \to^{\alpha} \tilde{\pi}$ where $\tilde{\rho} = \mu_S(\rho)$ and $\tilde{\pi} = \mu_S(\pi)$.

(O3) For any $\rho \in \mathbb{N}_0^S$ with $R_\rho \neq \emptyset$, replace the set of reactions R_ρ by the ρ-merge reaction r_ρ as discussed in Definition 8.

return $(\hat{S}, \hat{R})$, where $\hat{R}$ denotes the as modified set of reactions.

We are ready to present Algorithm 2, which yields the CRN in canonical form underlying the quotient ODE system by an EFL partition. Let us explain it using the MI network with the EFL partition $\tilde{\mathcal{H}}$ and the choice function $\mu_S = \tilde{f}$.

The steps (O1)–(O3) can be explained with the observation that the aggregated ODE system contains only the ODEs of the representatives, denoted by $\hat{S}$. Thus we have only to keep track of the fluxes which affect them. For this, in (O1) we first remove all products which are not elements of $\hat{S}$. For instance, reaction mi_1 becomes $Y_0 + Z_0 \to^{\alpha_1} Y_1$. Then, we complete (O1) by introducing *additional* products that make sure that the net stoichiometry of elements outside $\hat{S}$ is zero. In our example, this yields

$$\begin{aligned} Y_0 + Z_0 &\to^{\alpha_1} Y_1 + Z_0 & \qquad Z_2 + Z_0 &\to^{\alpha_1} Z_2 + Z_0 \\ Y_1 + Z_0 &\to^{\alpha_2} Y_2 + Z_0 & \qquad Z_1 + Z_0 &\to^{\alpha_2} Z_1 + Z_0 \\ Y_2 + Y_0 &\to^{\alpha_3} Y_0 + Y_1 & \qquad Z_0 + Y_0 &\to^{\alpha_3} Y_0 + Z_0 \\ Y_1 + Y_0 &\to^{\alpha_4} Y_0 + Y_0 & \qquad Z_1 + Y_0 &\to^{\alpha_4} Y_0 + Z_1 \end{aligned}$$

In (O2), we replace any species with the representative of its partition block, yielding,

$$\begin{aligned} Y_0 + Y_2 &\to^{\alpha_1} Y_1 + Y_2 & \qquad Y_0 + Y_2 &\to^{\alpha_1} Y_0 + Y_2 \\ Y_1 + Y_2 &\to^{\alpha_2} Y_2 + Y_2 & \qquad Y_1 + Y_2 &\to^{\alpha_2} Y_1 + Y_2 \\ Y_2 + Y_0 &\to^{\alpha_3} Y_0 + Y_1 & \qquad Y_2 + Y_0 &\to^{\alpha_3} Y_0 + Y_2 \\ Y_1 + Y_0 &\to^{\alpha_4} Y_0 + Y_0 & \qquad Y_1 + Y_0 &\to^{\alpha_4} Y_0 + Y_1 \end{aligned}$$

In (O3), instead, each nonempty reaction set $R_\rho \neq \emptyset$ is merged to the ρ-merge reaction as discussed in Definition 8. In the case of our example, we get

$$\begin{aligned} Y_0 + Y_2 &\to_{2\alpha_1 + 2\alpha_3} \frac{\alpha_1 + \alpha_3}{2\alpha_1 + 2\alpha_3} Y_1 + \frac{2\alpha_1 + \alpha_3}{2\alpha_1 + 2\alpha_3} Y_2 + \frac{\alpha_1 + 2\alpha_3}{2\alpha_1 + 2\alpha_3} Y_0 \\ Y_1 + Y_2 &\to_{2\alpha_2} \frac{3\alpha_2}{2\alpha_2} Y_2 + \frac{\alpha_2}{2\alpha_2} Y_1 \\ Y_1 + Y_0 &\to_{2\alpha_4} \frac{3\alpha_4}{2\alpha_4} Y_0 + \frac{\alpha_4}{2\alpha_4} Y_1 \end{aligned}$$

The next theorem states that the sketched algorithm is correct and provides us with a minimal number of chemical reactions in polynomial time.

Theorem 7. *Let $\mathcal{H}$ be an EFL partition of (S, R) and μ_S be a choice function of $\mathcal{H}$. Further, set $\hat{S} := \mu_S(S)$.*

- *Algorithm 2 returns, after at most $\mathcal{O}\big(|S||R|\log(|R|)\big)$ steps, a CRN $(\hat{S}, \hat{R})$ which can be related by means of a reactant morphism, stoichiomorphism and emulation $(\mu_S, \mu_R) : (S, R) \to (\hat{S}, \hat{R})$, where $\mu_R(\rho \to^{\alpha} \pi') := \mu_S(\rho) \to^{\beta} \pi$ and β, π are uniquely determined by the reactant $\mu_S(\rho)$.*
- *$|\hat{R}|$ is minimal, i.e., if $(\hat{S}, \underline{\hat{R}})$ is also a CRN such that there exists a reactant morphism and stoichiomorphism $(\mu_S, \underline{\mu}_R) : (S, R) \to (\hat{S}, \underline{\hat{R}})$, then $|\hat{R}| \leq |\underline{\hat{R}}|$.*
- *The CRN $(\underline{\hat{S}}, \underline{\hat{R}})$ returned by Algorithm 2 using $\underline{\mu}_S$, another choice function of $\mathcal{H}$, is such that $|\hat{R}| = |\underline{\hat{R}}|$. That is, the size of the CRN is invariant under the choice of the representatives of $\mathcal{H}$.*

The lower time complexity of Algorithm 2 with respect to Algorithm 1 explains the differences in their respective runtimes measured in the case studies of Table 1.

Remark 2. Let (S, R) denote the CRN of the MI network and let $\underline{\hat{R}}$ be

$$\begin{aligned} &Y_0 + Y_2 \to^{\alpha_1+\alpha_3} Y_1 + \frac{\alpha_3}{\alpha_1+\alpha_3} Y_0 + \frac{\alpha_1}{\alpha_1+\alpha_3} Y_2 \\ &Y_1 + Y_2 \to^{\alpha_2} Y_2 + Y_2 \\ &Y_1 + Y_0 \to^{\alpha_4} Y_0 + Y_0 \end{aligned}$$

Then, there exists a unique $\underline{\mu}_R : R \to \underline{\hat{R}}$ such that $(\mu_S, \underline{\mu}_R) : (S, R) \to (\hat{S}, \underline{\hat{R}})$ is an emulation, reactant morphism and stoichiomorphism. Thus, as anticipated, there are in general more than one CRNs which induce the same aggregated ODE system and which are structurally related to the original CRN.

5 Conclusion

Exact fluid lumpability (EFL) is a property of a system of ordinary differential equations that guarantees equal solutions for variables that are initialized equally. We have characterized EFL for chemical reaction networks (CRNs) in terms of structural conditions that only concern the set of reactions. This has allowed an efficient way of checking candidate EFL partitions and of computing the coarsest one. We provided an algorithm to construct a quotient CRN induced by EFL in a canonical form. This is related by means of appropriate mappings to the original CRN, thus fully reconciling our theory with the morphisms recently developed in [4] to study relations between CRNs. Subject of ongoing work is the implementation of a mature version of our prototype.

Acknowledgement. This work is partially supported by the EU project QUANTICOL, 600708, and by the DFG project FEMPA.

References

1. Blinov, M.L., Faeder, J.R., Goldstein, B., Hlavacek, W.S.: BioNetGen: software for rule-based modeling of signal transduction based on the interactions of molecular domains. Bioinformatics 20(17), 3289–3291 (2004)
2. Borisov, N.M., Markevich, N.I., Hoek, J.B., Kholodenko, B.N.: Signaling through receptors and scaffolds: Independent interactions reduce combinatorial complexity. Biophysical Journal 89(2), 951–966 (2005)
3. Camporesi, F., Feret, J.: Formal reduction for rule-based models. Electronic Notes in Theoretical Computer Science 276, 29–59 (2011), MFPS XXVII
4. Cardelli, L.: Morphisms of reaction networks that couple structure to function. BMC Systems Biology 8(1), 84 (2014)
5. Cardelli, L., Csikász-Nagy, A.: The cell cycle switch computes approximate majority. Sci. Rep. 2 (2012)
6. Colvin, J., Monine, M.I., Faeder, J.R., Hlavacek, W.S., Hoff, D.D.V., Posner, R.G.: Simulation of large-scale rule-based models. Bioinformatics 25(7), 910–917 (2009)
7. Conzelmann, H., Fey, D., Gilles, E.D.: Exact model reduction of combinatorial reaction networks. BMC Syst Biol 2, 78 (2008)
8. Conzelmann, H., Saez-Rodriguez, J., Sauter, T., Kholodenko, B., Gilles, E.: A domain-oriented approach to the reduction of combinatorial complexity in signal transduction networks. BMC Bioinformatics 7(1), 34 (2006)
9. Craciun, G., Tang, Y., Feinberg, M.: Understanding bistability in complex enzyme-driven reaction networks. Proceedings of the National Academy of Sciences 103(23), 8697–8702 (2006)
10. Danos, V., Feret, J., Fontana, W., Harmer, R., Krivine, J.: Abstracting the differential semantics of rule-based models: Exact and automated model reduction. In: LICS. pp. 362–381 (2010)
11. Danos, V., Laneve, C.: Formal molecular biology. Theoretical Computer Science 325(1), 69–110 (2004)
12. Faeder, J.R., Hlavacek, W.S., Reischl, I., Blinov, M.L., Metzger, H., Redondo, A., Wofsy, C., Goldstein, B.: Investigation of early events in FcεRI-mediated signaling using a detailed mathematical model. The Journal of Immunology 170(7), 3769–3781 (2003)
13. Feret, J., Danos, V., Krivine, J., Harmer, R., Fontana, W.: Internal coarse-graining of molecular systems. Proceedings of the National Academy of Sciences 106(16), 6453–6458 (2009)
14. Gardner, T.S., Cantor, C.R., Collins, J.J.: Construction of a genetic toggle switch in Escherichia coli. Nature 403(6767), 339–342 (01 2000)
15. Heath, J., Kwiatkowska, M., Norman, G., Parker, D., Tymchyshyn, O.: Probabilistic model checking of complex biological pathways. Theoretical Computer Science 391(3), 239–257 (2008), converging Sciences: Informatics and Biology
16. Kocieniewski, P., Faeder, J.R., Lipniacki, T.: The interplay of double phosphorylation and scaffolding in MAPK pathways. Journal of Theoretical Biology 295, 116–124 (2012)
17. Kolczyk, K., Samaga, R., Conzelmann, H., Mirschel, S., Conradi, C.: The process-interaction-model: a common representation of rule-based and logical models allows studying signal transduction on different levels of detail. BMC Bioinformatics 13(1), 251 (2012)
18. Okino, M.S., Mavrovouniotis, M.L.: Simplification of mathematical models of chemical reaction systems. Chemical Reviews 2(98), 391–408 (1998)
19. Paige, R., Tarjan, R.: Three partition refinement algorithms. SIAM Journal on Computing 16(6), 973–989 (1987)
20. Shin, S.W., Thachuk, C., Winfree, E.: Verifying chemical reaction network implementations: A pathway decomposition approach. In: Verification of Engineered Molecular Devices and Programs. Vienna Summer of Logic (2014)

21. Sneddon, M.W., Faeder, J.R., Emonet, T.: Efficient modeling, simulation and coarse-graining of biological complexity with NFsim. Nature Methods 8(2), 177–183 (Feb 2011)
22. Suderman, R., Deeds, E.J.: Machines vs. ensembles: Effective MAPK signaling through heterogeneous sets of protein complexes. PLoS Comput Biol 9(10), e1003278 (10 2013)
23. Toth, J., Li, G., Rabitz, H., Tomlin, A.S.: The effect of lumping and expanding on kinetic differential equations. SIAM Journal on Applied Mathematics 57(6), 1531–1556 (1997)
24. Tschaikowski, M., Tribastone, M.: Exact fluid lumpability for Markovian process algebra. In: CONCUR. pp. 380–394. LNCS (2012)

A Proofs

Proof (Theorem 2). Let $\mu_S : S \to S$ be some choice function of $\mathcal{H}$, set $\hat{S} := \mu_S(S)$ and define $G_{\hat{X}}(\hat{\boldsymbol{V}}) := F_{\hat{X}}(\hat{\boldsymbol{V}} \circ \mu_S)$ for any $\hat{\boldsymbol{V}} \in \mathbb{R}^{\hat{S}}$ and $\hat{X} \in \hat{S}$. Further, let $\hat{\boldsymbol{V}}$ denote the unique ODE solution of $\frac{d}{dt}\hat{\boldsymbol{V}}(t) = G\big(\hat{\boldsymbol{V}}(t)\big)$ subject to some given initial condition $\hat{\boldsymbol{V}}(0)$. Then, for all $X \in S$, it holds that

$$\frac{d}{dt}(\hat{\boldsymbol{V}}(t) \circ \mu_S)_X = \Big(\frac{d}{dt}\hat{\boldsymbol{V}}(t)\Big)_{\mu_S(X)} = \\ = G_{\mu_S(X)}\big(\hat{\boldsymbol{V}}(t)\big) = F_{\mu_S(X)}(\hat{\boldsymbol{V}}(t) \circ \mu_S) = F_X(\hat{\boldsymbol{V}}(t) \circ \mu_S)$$

Thus, $t \mapsto \hat{\boldsymbol{V}}(t) \circ \mu_S$ is the unique solution of the ODE system $\frac{d}{dt}\boldsymbol{V}(t) = F(\boldsymbol{V}(t))$ subject to $\hat{\boldsymbol{V}}(0) \circ \mu_S$. Since $t \mapsto \hat{\boldsymbol{V}}(t) \circ \mu_S$ is constant on $\mathcal{H}$, the proof is complete.

Proof (Theorem 3). We first show the first claim. Fix two arbitrary $X_i, X_j \in S$. Further, let $\mu_S : S \to S$ be some choice function of a given partition $\mathcal{H}$. Then, together with $\hat{S} := \mu_S(S)$, it suffices to show that $T^{\Sigma}_{\mathcal{H}}(X_i) = T^{\Sigma}_{\mathcal{H}}(X_j)$ if and only if $F_{X_i}(\hat{V} \circ \mu_S) = F_{X_j}(\hat{V} \circ \mu_S)$ for all $\hat{V} \in \mathbb{R}^{\hat{S}}$. The if direction follows because $T^{\Sigma}_{\mathcal{H}}(X_i) = T^{\Sigma}_{\mathcal{H}}(X_j)$ implies that the coefficients of the polynomials $\hat{V} \mapsto F_{X_i}(\hat{V} \circ \mu_S)$ and $\hat{V} \mapsto F_{X_j}(\hat{V} \circ \mu_S)$ are the same. To see that the only-if direction holds true, instead, recall that the multivariate version of Taylor's theorem ensures that two real polynomials p, q that satisfy $p \equiv q$ have necessarily the same coefficients.

To see the second claim, we first observe that computing $\{T(X) \mid X \in S\}$ takes at most $|S||R|$ steps and that $|T(X)| \leq |R|$ for all $X \in S$. Thus, rewriting $T(X)$ into $T_{\mathcal{H}}(X)$ takes at most $|S||R|$ steps. To infer $T^{\Sigma}_{\mathcal{H}}(X)$ from $T_{\mathcal{H}}(X)$, we first sort the entries of $T_{\mathcal{H}}(X)$ with respect to reactants. This can be done in $\mathcal{O}\big(|S||R|\log(|R|)\big)$ steps, because $|T_{\mathcal{H}}(X)| = |T(X)| \leq |R|$. Afterwards, we sum entries with the same reactants together. The number of summations is bounded by $|T_{\mathcal{H}}(X)| \leq |R|$, meaning that $T^{\Sigma}_{\mathcal{H}}(X)$ can be computed from $T_{\mathcal{H}}(X)$ in at most $\mathcal{O}\big(|S||R|\log(|R|) + |S||R|\big)$ steps. From this we infer that $\{T^{\Sigma}_{\mathcal{H}}(X) \mid X \in S\}$ can be computed in $\mathcal{O}\big(|S|^2|R|\log(|R|)\big)$ steps.

We next prove the third statement. If each $T^{\Sigma}_{\mathcal{H}}(X)$ is sorted with respect to the reactants, deciding whether $T^{\Sigma}_{\mathcal{H}}(X_i)$ is equal to $T^{\Sigma}_{\mathcal{H}}(X_j)$ can be done in $|S|\max\{|T^{\Sigma}_{\mathcal{H}}(X_i)|, |T^{\Sigma}_{\mathcal{H}}(X_j)|\} \leq |S||R|$ steps. Thus, if $\{T^{\Sigma}_{\mathcal{H}}(X) \mid X \in S\}$ is given as above, checking whether a given partition is EFL or not takes at most $\mathcal{O}(|S|^2|R|)$ steps.

Proof (Proposition 1). Note that $T^{\Sigma}_{\mathcal{G}}(X_i) = T^{\Sigma}_{\mathcal{G}}(X_j)$ implies $T^{\Sigma}_{\mathcal{G}'}(X_i) = T^{\Sigma}_{\mathcal{G}'}(X_j)$ for any $X_i, X_j \in S$, whenever $\mathcal{G}$ is a refinement of $\mathcal{G}'$. Thus, since $\mathcal{H}_1$ and $\mathcal{H}_2$ are EFL partitions, we infer using Theorem 3 that $T^{\Sigma}_{\mathcal{G}'}(X_i) = T^{\Sigma}_{\mathcal{G}'}(X_j)$ for any $X_i, X_j \in H$ and $H \in S/(\sim_{\mathcal{H}_1} \cup \sim_{\mathcal{H}_2})^*$. Applying Theorem 3 again shows that $S/(\sim_{\mathcal{H}_1} \cup \sim_{\mathcal{H}_2})^*$ is an EFL partition. Let us now assume that $\mathcal{H}_1, \mathcal{H}_2$ are refinements of some other partition $\mathcal{G}$. For the sake of brevity, define $\sim_i := \sim_{\mathcal{H}_i}$. We next show that $S/(\sim_1 \cup \sim_2)^*$ is a refinement of $\mathcal{G}$. To see this, let us assume that $X_0 \sim_{i_1} X_1 \sim_{i_2} \ldots \sim_{i_k} X_k$, where $i_j \in \{1, 2\}$ for all $1 \leq j \leq k$. Then, is suffices to prove that there exists a (unique) $G \in \mathcal{G}$ such that $X_1, \ldots, X_k \in G$. We show this by induction. Since the base case

$k = 1$ is trivial, let us consider the induction step $k \to k+1$. Then, $X_k \sim_{i_{k+1}} X_{k+1}$ implies the existence of some $H \in \mathcal{H}_{i_{k+1}}$ such that $X_k, X_{k+1} \in H$. Let $G_k \in \mathcal{G}$ be such that $H \subseteq G_k$. Since $X_k \in G$ by induction hypothesis and $X_k \in G_k$, it holds that $G \cap G_k \neq \emptyset$. Since $\mathcal{G}$ is a partition, this implies that $G = G_k$ and the proof is complete.

Proof (Theorem 4). For the proof of correctness, let us assume that $\mathcal{H}$ denotes the coarsest EFL partition that refines $\mathcal{H}_0 := \mathcal{G}$ and define $\mathcal{H}_{k+1} := S/(\sim_{T^\Sigma_{\mathcal{H}_k}} \cap \sim_{\mathcal{H}_k})$ for all $k \geq 0$. Then, the sequence $\mathcal{H}_0, \mathcal{H}_1, \mathcal{H}_2, \ldots$ is such that

- $\mathcal{H}$ is a refinement of $\mathcal{H}_k$
- $\mathcal{H}_k$ is a refinement of $\mathcal{H}_{k-1}$

for all $k \geq 1$. We prove this by induction on k.

- $k = 1$: Since $\mathcal{H}$ is a refinement of $\mathcal{H}_0$, $T^\Sigma_{\mathcal{H}}(X_i) = T^\Sigma_{\mathcal{H}}(X_j)$ implies $T^\Sigma_{\mathcal{H}_0}(X_i) = T^\Sigma_{\mathcal{H}_0}(X_j)$ for any $X_i, X_j \in S$. Noting that this yields $\sim_{\mathcal{H}} \subseteq \sim_{T^\Sigma_{\mathcal{H}_0}}$ we readily infer that $\sim_{\mathcal{H}} \subseteq \sim_{T^\Sigma_{\mathcal{H}_0}} \cap \sim_{\mathcal{H}_0}$. The second claim is trivial.
- $k \to k+1$: Thanks to the fact that $\mathcal{H}$ is a refinement of $\mathcal{H}_k$ by induction, $T^\Sigma_{\mathcal{H}}(X_i) = T^\Sigma_{\mathcal{H}}(X_j)$ implies $T^\Sigma_{\mathcal{H}_k}(X_i) = T^\Sigma_{\mathcal{H}_k}(X_j)$ for any $X_i, X_j \in S$. Hence, $\sim_{\mathcal{H}} \subseteq \sim_{T^\Sigma_{\mathcal{H}_k}}$ and the induction hypothesis yields $\sim_{\mathcal{H}} \subseteq \sim_{T^\Sigma_{\mathcal{H}_k}} \cap \sim_{\mathcal{H}_k}$. The second claim is trivial.

Since $\mathcal{H}$ is a refinement of any $\mathcal{H}_k$, it holds that $\mathcal{H} = \mathcal{H}_k$ whenever $\mathcal{H}_k$ is an EFL partition. Moreover, thanks to the fact that $\mathcal{H}_k$ is a refinement of $\mathcal{H}_{k-1}$ for all $k \geq 1$, we can fix the smallest $k \geq 1$ such that $\mathcal{H}_k = \mathcal{H}_{k-1}$. Let us pick arbitrary $H \in \mathcal{H}_k$ and $X_i, X_j \in H$. Since $\mathcal{H}_k = S/(\sim_{T^\Sigma_{\mathcal{H}_{k-1}}} \cap \sim_{\mathcal{H}_{k-1}})$, it holds that $T^\Sigma_{\mathcal{H}_{k-1}}(X_i) = T^\Sigma_{\mathcal{H}_{k-1}}(X_j)$. Recalling that $\mathcal{H}_k = \mathcal{H}_{k-1}$, this yields $T^\Sigma_{\mathcal{H}_k}(X_i) = T^\Sigma_{\mathcal{H}_k}(X_j)$. Since this shows that $\mathcal{H}_k$ is an EFL partition, we infer the correctness of the algorithm.

We now bound the number of steps. In our implementation, each species is represented by a structure that embodies a species integer `sid`, a list of fluxes `fluxes` and a pointer to a species structure `repr`. Moreover, `species` denotes an array of size $|S|$ that stores the pointers to all species structures. A partition $\mathcal{H}$ of S that is given by a choice function $\mu_S : S \to S$ is encoded via the `repr` fields: if `X` is a species structure with `X.sid == i` and $\mu_S(X_i) = X_j$, then `X.repr` points to the structure of X_j. We discuss next how to calculate $S/(\sim_{T^\Sigma_{\mathcal{H}}} \cap \sim_{\mathcal{H}})$.

In the first step, we exchange the entries of the array `species` in such a way that whenever `*(species[i]).repr != *(species[i+1]).repr`, then, for all $j > i$, it holds that `*(species[i]).repr != *(species[j]).repr`. In other words, our goal is to permute the entries of `species` in such a way that each equivalence class of $\mathcal{H}$ builds a connected interval in $\{1, \ldots, |S|\}$. Since comparison and exchange of pointers needs constant time, this can be done in $\mathcal{O}(|S|^2)$ steps.

In the second step, we update the fields `fluxes` by calculating $\{T^\Sigma_{\mathcal{H}}(X) \mid X \in S\}$ where the entries of each $T^\Sigma_{\mathcal{H}}(X)$ are sorted with respect to the reactants. Theorem 3 ensures that this can be done in $\mathcal{O}\big(|S|^2|R|\log(|R|)\big)$ steps.

For the third step, we note that each entry of $T^\Sigma_{\mathcal{H}}(X)$ refers to a monomial, meaning that $T^\Sigma_{\mathcal{H}}(X)$ itself represents a polynomial. Due to the fact that $T^\Sigma_{\mathcal{H}}(X)$ is sorted with

respect to monomials, it is possible to define a linear ordering $\sqsubseteq$ on the set of polynomials that allows to check in $\mathcal{O}(|S||R|)$ steps whether $T^{\Sigma}_{\mathcal{H}}(X_i) \sqsubseteq T^{\Sigma}_{\mathcal{H}}(X_j)$ or not. Armed with this notion, we now sort each single equivalence class of $\mathcal{H}$. Thanks to the first step, the latter are present as intervals in `species`. Thus, if $H_1, \ldots, H_m$ are equivalence classes of $\mathcal{H}$ and $|S|_i := |H_i|$, there exists a constant $C > 0$ such that the sorting of $H_1, \ldots, H_m$ takes up to $C|S||R| \sum_{i=1}^{m} |S|_i \log(|S|_i) \leq C|S||R| \sum_{i=1}^{m} |S|_i \log(|S|) = C|S|^2|R| \log(|S|)$ steps. Note that after the sorting, each H_i will still occupy the same interval of indices in `species`.

In the forth step, all we need to do is to update the `repr` fields in `species`. Clearly, $(X_i, X_j) \in \sim_{T^{\Sigma}_{\mathcal{H}}} \cap \sim_{\mathcal{H}}$ if and only if $T^{\Sigma}_{\mathcal{H}}(X_i) = T^{\Sigma}_{\mathcal{H}}(X_j)$ and $X_i, X_j \in H$ for some $H \in \mathcal{H}$. Thus, after taking into account that comparing and exchanging of $T^{\Sigma}_{\mathcal{H}}(X_i)$ and $T^{\Sigma}_{\mathcal{H}}(X_j)$ needs $\mathcal{O}(|S||R|)$ steps, we infer that the third step can be done in $\mathcal{O}(|S|^2|R|)$ steps.

After the forth step, `species` encodes $\mathcal{H}' := S/(\sim_{T^{\Sigma}_{\mathcal{H}}} \cap \sim_{\mathcal{H}})$ in a way where the equivalence classes of $\mathcal{H}'$ build intervals in `species`. In particular, the first step is only needed when $\mathcal{H}$ is given by the initial partition $\mathcal{G}$. In order to check whether $\mathcal{H}'$ is EFL or not, we perform the second step for $\mathcal{H}'$, i.e. we update the fields `fluxes` in `species` by calculating $\{T^{\Sigma}_{\mathcal{H}'}(X) \mid X \in S\}$. Then, since comparing $T^{\Sigma}_{\mathcal{H}'}(X_i)$ and $T^{\Sigma}_{\mathcal{H}'}(X_j)$ needs up to $\mathcal{O}(|S||R|)$ steps, determining whether $\mathcal{H}'$ is EFL or not can be done in $\mathcal{O}(|S|^2|R|)$ steps. This is because $\mathcal{H}$ is an EFL partition if and only if $\mathcal{H} = S/\sim_{T^{\Sigma}_{\mathcal{H}}}$. In the case $\mathcal{H}'$ is an EFL partition, we can terminate. Otherwise, we continue with step three.

From the above discussion, it becomes apparent that the number of steps we need is at most $|S|\big[\mathcal{O}\big(|S|^2|R|\log(|R|)\big) + \mathcal{O}(|S|^2|R|\log(|S|))\big] = \mathcal{O}\big(|S|^3|R|(\log(|S|) + \log(|R|))\big)$.

Proof (Proposition 2). We first show that μ is a reactant morphism:

- For $r = \rho \to^{\alpha} \pi \in R \setminus R_\rho$ we get $\mu_R(\rho \to^{\alpha} \pi) = \rho \to^{\alpha} \pi = \mu_S(\rho) \to^{\alpha} \pi$.
- For $r = \rho \to^{\alpha} \pi \in R_\rho$ we get $\mu_R(\rho \to^{\alpha} \pi) = r_\rho = \rho \to^{\alpha_\rho} \pi = \mu_S(\rho) \to^{\alpha_\rho} \pi_\rho$.

In order to see that μ is also a stoichiomorphism, we fix an arbitrary $X \in S$ and proceed by case distinction.

- If $\hat{r} \in R \setminus R_\rho$, then $\mu_R^{-1}(\hat{r}) = \{\hat{r}\}$. Hence $\sum_{r \in \mu_R^{-1}(\hat{r})} \phi(X, r) = \phi(\mu_S(X), \hat{r})$
- If $\hat{r} = r_\rho$, then $\mu_R^{-1}(\hat{r}) = R_\rho = \{\rho \to^{\alpha_i} \pi_i \mid i \in I_\rho\}$, yielding

$$
\begin{aligned}
\sum_{r \in \mu_R^{-1}(\hat{r})} \phi(X, r) &= \sum_{i \in I_\rho} \phi(X, \rho \to^{\alpha_i} \pi_i) = \sum_{i \in I_\rho} \alpha_i\big((\pi_i)_X - \rho_X\big) \\
&= \sum_{i \in I_\rho} \big(\alpha_i \cdot (\pi_i)_X\big) - \big(\sum_{i \in I_\rho} \alpha_i\big)\rho_X = \alpha_\rho \sum_{i \in I_\rho} \frac{\alpha_i}{\alpha_\rho}(\pi_i)_X - \alpha_\rho \rho_X \\
&= \alpha_\rho\Big(\big(\sum_{i \in I_\rho} \frac{\alpha_i}{\alpha_\rho}(\pi_i)_X\big) - \rho_X\Big) = \alpha_\rho\big(\pi_{\rho_X} - \rho_X\big) \\
&= \phi(X, r_\rho) = \phi(\mu_S(X), r_\rho)
\end{aligned}
$$

Since this shows that $\sum_{r\in\mu_R^{-1}(\hat{r})}\phi(X,r) = \phi(\mu_S(X),\hat{r})$ for any reaction $\hat{r} \in \hat{R} = (R \setminus R_\rho) \cup \{r_\rho\}$, we infer that μ is also a stoichiomorphism. Moreover, Theorem 1 ensures that μ is an emulation.

Proof (Proposition 3). It suffices to observe that the results on compositions made in [4] for nCRNs carry over to CRNs.

Proof (Theorem 6). Follows as a direct consequence of Theorem 1 and Proposition 3.

Proof (Theorem 7). Let us first consider the time complexity. By encoding $\mu_S : S \to S$, ρ and π as arrays of length $|S|$, it is obvious that (O1) needs at most $|R||S|$ steps. Similarly, it can be easily seen that (O2) needs at most $2|R||S|$ steps. We perform (O3) in two steps. First, we sort the set of reactions with respect to the reactants. Since a comparison and exchange of two reactions takes at most $|S|+6|S|$ time steps, the sorting of R does not need more than $\mathcal{O}\big(|S||R|\log(|R|)\big)$ steps. After R has been sorted with respect to the reactants, we merge the underlying reaction groups together. Here, we note that all mergings can be done in $|R||S|$ steps, because the sizes of all groups sum up to $|R|$. We now turn to the remaining part of the first statement. Let $(X, \boldsymbol{V}) \mapsto G_X(\boldsymbol{V})$ denote the drift induced by the reactions under study. We next prove that $G_{X_k}(\hat{\boldsymbol{V}} \circ \mu_S) = F_{X_k}(\hat{\boldsymbol{V}} \circ \mu_S)$ for all $\hat{\boldsymbol{V}} \in \mathbb{R}^{\hat{S}}_{\geq 0}$ and $X_k \in \hat{S}$ at any step of the algorithm. For this, let us consider a reaction change $\rho \to^\alpha \pi \mapsto \rho \to^\alpha \tilde{\pi}$ as applied in (O1). Then, the value of $G_{X_k}(\hat{\boldsymbol{V}} \circ \mu_S)$ is not changed because of $\pi(X_k) - \rho(X_k) = \tilde{\pi}(X_k) - \rho(X_k)$. Let us now consider a reaction change $\rho \to^\alpha \pi \mapsto \tilde{\rho} \to^\alpha \tilde{\pi}$ as in (O2). Such a change does not affect $G_{X_k}(\hat{\boldsymbol{V}} \circ \mu_S)$ for two reasons. First, since $\mathcal{H}$ is an EFL partition, $\mathcal{H} = \{\mu_S^{-1}(X_i) \mid X_i \in \hat{S}\}$ implies that $X_i(t) = X_j(t)$ for all $X_j \in \mu_S^{-1}(X_i)$, $X_i \in \hat{S}$ and $t \in I \cap \mathbb{R}_{\geq 0}$. Consequently, the $(\hat{\boldsymbol{V}} \circ \mu_S)^\rho = (\hat{\boldsymbol{V}} \circ \mu_S)^{\tilde{\rho}}$. Second, it holds that $\pi(X_k) - \rho(X_k) = \tilde{\pi}(X_k) - \tilde{\rho}(X_k)$. Since Proposition 3 ensures that (O3) does not change the value of $G_{X_k}(\hat{\boldsymbol{V}} \circ \mu_S)$ either, we infer that (μ_S, μ_R) is an emulation. Moreover, by construction, reactions of $\hat{R}$ have pairwise different reactants and each $\hat{\rho} \to^\beta \hat{\pi} \in \hat{R}$ is such that there exists some $\rho \to^\alpha \pi \in R$ with $\hat{\rho} = \mu_S(\rho)$. Hence, (μ_S, μ_R) is a reactant morphism. Using Theorem 5, we infer that (μ_S, μ_R) is also a stoichiomorphism. This yields the first claim. The second claim, instead, holds because $(\mu_S, \underline{\mu}_R) : (S, R) \to (\hat{S}, \underline{\hat{R}})$ is a reactant morphism and $\underline{\hat{R}}$ has pairwise different reactants. Similarly, the third claim holds true because $(\underline{\mu}_S, \underline{\mu}_R) : (S, R) \to (\underline{\hat{S}}, \underline{\hat{R}})$ is a reactant morphism, $\underline{\hat{R}}$ has pairwise different reactants and $|\hat{S}| = |\underline{\hat{S}}|$.