

A large-scale assessment of exact lumping of quantitative models in the BioModels repository

Questa è la versione preprint della seguente opera:

Original

A large-scale assessment of exact lumping of quantitative models in the BioModels repository / Perez-Verona, I. C.; Tribastone, M.; Vandin, A.. - In: THEORETICAL COMPUTER SCIENCE. - ISSN 0304-3975. - 893:(2021), pp. 41-59. [10.1016/j.tcs.2021.06.026]

Availability:

This version is available at: 20.500.11771/20859

Publisher:

Published

DOI:10.1016/j.tcs.2021.06.026

Terms of use:

This publication is made accessible in accordance with the terms for deposit in the institutional repository, as defined by the IMT School for Advanced Studies Lucca's Open Access Policy. (https://library.imtlucca.it/sites/default/files/regolamento-policy-open-access-imtlib_0.pdf).

Si prega di consultare le pagine informative dell'editore relative alle politiche di autoarchiviazione.

(Article begins on next page)

A large-scale assessment of exact lumping of quantitative models in the BioModels repository

Isabel Cristina Perez-Verona

IMT School for Advanced Studies Lucca, Italy

Mirco Tribastone

IMT School for Advanced Studies Lucca, Italy

Andrea Vandin

*Sant'Anna School of Advanced Studies, Italy
DTU Technical University of Denmark*

Abstract

Chemical reaction networks are a popular formalism for modeling biological processes which supports both a deterministic and a stochastic interpretation based on ordinary differential equations and continuous-time Markov chains, respectively. In most cases, these models do not enjoy analytical solution, thus typically requiring expensive computational methods based on numerical solvers or stochastic simulations. Exact model reduction techniques can be used as an aid to lower the analysis cost by providing reduced networks that preserve the dynamics of interest to the modeler without incurring any approximation error. We hereby consider a family of techniques for both deterministic and stochastic networks which are based on equivalence relations over the species in the network, leading to a coarse graining which provides the exact aggregate time-course evolution for each equivalence class. We present a large-scale empirical assessment on the BioModels repository by measuring their compression capability over 579 models. Through a number of selected case studies, we also show their ability in yielding physically interpretable reductions that can reveal dynamical patterns of the bio-molecular processes under consideration, independently of the values of the kinetic parameters used in the models.

Keywords: Model reduction, Chemical Reaction Networks, Equivalence relations

1. Introduction

Computational models in systems biology combine biochemical and physiological knowledge to inform highly detailed mechanistic models of biological networks such as signaling pathways, protein-protein interaction networks, and genetic cascades. Mathematical models which equip such interaction networks with kinetic information generally lead to a dynamical-system representation in terms of a formal chemical reaction network (CRN)¹, with two main interpretations based on ordinary differential equations (ODEs) and continuous-time Markov chains (CTMCs), respectively. In either case the model tracks the time-course evolution of all biochemical species in the network. In the ODE interpretation each species is associated with a variable of a system of (typically nonlinear) ODEs, which are analyzed from an initial condition that represents the initial concentration of each species [55]. In the CTMC interpretation [28], species are tracked discretely and each state is a vector of molecular counts, one component for each

Email addresses: isabel.perez@imtlucca.it (Isabel Cristina Perez-Verona), mirco.tribastone@imtlucca.it (Mirco Tribastone), andrea.vandin@santannapisa.it (Andrea Vandin)

¹The paper uses a number of acronyms summarized in Table 1.

species. It is well known that these two representations can be formally related to each other under appropriate conditions, with the ODEs being the limit behavior when the number of molecules in the CRN is large enough [36]. Often it is useful to consider both interpretations—one would take the CTMC semantics as the ground-truth representation and the ODE as an approximation that estimates the first-order moments of the species' populations.

Unfortunately, in both cases the analysis can be expensive due to the lack of analytical solutions in general. Indeed, the modeler is typically left with computational approaches such as the numerical integration of ODEs (e.g., [3]) or stochastic simulation [28]. This is a major motivating issue for several lines of research aiming at easing the computational cost of the analysis, including efficient simulation methods (e.g., [29]) and approximation methods for stochastic chemical kinetics (e.g., [49]). Model reduction is a major line of research concerned with lowering the computational cost of the analysis of dynamical models. Under this term belong several techniques which exploit various properties such as time-scale separation (i.e., by identifying reactions evolving at considerably different rates), lumping (i.e., grouping variables into macro-variables for which it is possible to construct a coarse-grained model), and sensitivity analysis (i.e., by removing variables that have little influence on the dynamics of interest to the modeler). For a thorough review in the field of computational biology, the reader is referred to [50], and to [44] for a classical earlier review in chemical engineering. It is worth noting that the problem of model reduction transcends chemistry and biology, and is of relevance in many other computational sciences; see, for instance, the book by Antoulas for applications to control theory [1].

For the purposes of the present paper it is useful to distinguish between *exact* and *approximate* model reduction methods. The former class identifies a reduced model which preserves an aggregate dynamics of the original model precisely, i.e., without any approximation error [40, 53]. In other words, the reduction induces some loss of information, but the information which is kept remains exact. The latter class gives aggregate dynamics with some approximation error, whose degree of controllability depends on the specific technique under consideration. Exact model reduction methods can be seen as a universal pre-processing step for computational analysis because they do not add further inaccuracy beyond the inevitable disagreement between a model and the real (biological) system under study. On the other hand, they may require stricter criteria to guarantee exactness, which may limit their applicability in practice.

The main goal of this paper is to carry out a systematic evaluation on benchmark biological models of a class of methods for the exact reduction of a CRN. These methods are based on the computation of *equivalence relations* between its species. Here the original set of species is partitioned into blocks such that every original species belongs to exactly one block. The simplification consists in defining an *aggregate CRN* where each macro-species represents the overall dynamics (i.e., the sum of the original dynamics) of the species in each block. According to established terminology, such an equivalence relation identifies a so-called *proper lumping* [44], whereby every variable is guaranteed to contribute to precisely one of the macro-variables. We remark that, although methods concerned with improper lumping are available (e.g., [40, 45]), they are outside the scope of the present contribution.

We consider five different techniques [10, 11, 13, 18, 17]. These are all defined as equivalence relations which share a common property, namely the possibility of developing an algorithm for computing the coarsest lumping using partition refinement, based on the seminal algorithm in [46]. Such algorithms apply iterative refinements to a given initial partition of species with which, for instance, one can isolate the observable species to be preserved in the reduction. The techniques differ according to the dynamical interpretation that they consider (i.e., deterministic versus stochastic), the assumption on the structure of the input CRN, and the dynamical property that they preserve.

- Forward equivalence (FE) and backward equivalence (BE) apply to CRNs with deterministic semantics based on ODEs with polynomial derivatives (including mass-action kinetics as a notable instance), and lead to reduced models where the solution of each macro-species exactly corresponds to the sums of the solutions of the original variables belonging to a block [13].
 - With FE the reduced model can be used for any choice of the initial concentrations of the original model, but the time course of the concentration of one original species cannot be recovered from the reduced model in general.
 - BE aggregates species that have the same solutions at all time points, hence an individual solution may be recovered by simply dividing the solution of the corresponding aggregate by the cardinality of the equivalence class; however the dynamics of the reduced model can be related to that of the original one only if all species in the same partition block start with the same initial concentration in the original model.

<i>Acronym</i>	<i>Meaning</i>
ODE	Ordinary differential equation
CTMC	Continuous-time Markov chain
CRN	Chemical reaction network
MRN	Chemical reaction network with mass-action kinetics (either stochastic or deterministic semantics)
PRN	Chemical reaction network with deterministic semantics yielding polynomial ODEs
NRN	Chemical reaction network with deterministic semantics yielding arbitrary nonlinear ODEs
FE	Forward equivalence (applied to PRNs and MRNs with deterministic semantics)
BE	Backward equivalence (applied to PRNs and MRNs with deterministic semantics)
FDE	Forward differential equivalence (applied to NRNs with deterministic semantics)
BDE	Backward differential equivalence (applied to NRNs with deterministic semantics)
SE	Species Equivalence (applied to MRNs with stochastic semantics)
SMT	Satisfiability Modulo Theories (used to check the validity of logical formulas with real predicates)

Table 1: Acronyms for semantic models, reaction networks, equivalence relations, and other techniques used in the paper.

- Forward differential equivalence (FDE) and backward differential equivalence (BDE) are generalizations of FE and BE, respectively, that can be applied to CRNs where the underlying ODEs have nonlinearities beyond polynomials, e.g., rational expressions as in Hill kinetics [11].
- Species equivalence (SE) is an equivalence relation for mass-action CRNs with CTMC interpretation. It identifies a reduced CRN where the marginal probability distribution of each macro-species gives the marginal probability distribution of the sum of the molecular counts of all species in that block.

FDE and BDE (hence, FE and BE) are shown to be conservative generalizations of the notions of ordinary and exact lumpability for CTMCs [9], respectively, when one considers their forward equations of motion, which are a system of linear ODEs. As a consequence of the fact that ordinary and exact lumpability are not comparable—i.e., neither implies the other—FDE (resp., FE) and BDE (resp., BE) are not comparable either. This will also show in the models analyzed in this paper.

Assisted by ERODE [14], a publicly available software tool that implements the aforementioned equivalences, we carry out an assessment of the BioModels database [39], a well-known repository of quantitative models of biochemical systems. The models are available at <https://www.erode.eu/examples.html>. Our goal is to answer the following three main evaluation questions:

- Q1.** *How restrictive are the assumptions required by the considered notions of equivalence?* We answer this question by detailing how we translated the BioModels descriptions, available in the SBML format, into the input format of ERODE.
- Q2.** *What is the effectiveness of exact model reduction by the equivalence relations?* We measure effectiveness as the percentage of models that can be aggregated, as well as the compression ratio provided by the largest equivalence that preserves the observables specified in the original model.
- Q3.** *What is the robustness of the reductions with respect to the original parameter values?* For this question, we present a more detailed discussion of a selected number of models, for which we also provide a physical interpretation of the obtained reductions.

Relationship with [47]. This paper is based on and extends [47], where

- a prototypical version of our SBML importer has been presented;
- a less detailed evaluation has been performed on the BioModels repository;
- no indication on the robustness of our reduction techniques to specific parameter values has been provided;

- a preliminary, now superseded, notion of reduction for the CTMC semantics of MRN has been considered.

Specifically, as discussed in the reworked Section 3 (especially Section 3.2), we improved the quality of our SBML importer and fully integrated it into ERODE, allowing us to automatize all the importing and analysis tasks. Despite the number of supported models has not changed considerably with respect to what reported in [47], these improvements have allowed us to identify more mass-action models (70 rather than 43), and to considerably decrease the number of models on which the reduction techniques failed with errors or timeouts (22 rather than 198). Differently from [47], the analysis is provided separately for the curated and non-curated branches of the BioModels repository. In addition, in this paper we now consider three classes of CRNs instead of two classes in [47] (i.e., with mass-action and non-mass-action kinetics) by adding the class of *polynomial reaction networks*, that is CRNs with ODEs that have polynomial derivatives but do not obey the mass-action law. Furthermore, a section on selected case studies now shows novel applications of the techniques to identify reductions that are not dependent from the actual values of the model parameters. Finally, SMB [15], the reduction technique for the stochastic semantics of elementary MRNs has been replaced by the novel SE [18]. SE generalizes SMB because it can be applied to any MRN (rather than only to elementary ones with at most two reactants for SMB) and has more reduction power.

2. Background

In order to make the paper self-contained, in this section we briefly overview the main results regarding the equivalence notions used in our assessment. We refer to the original papers for the details and further examples, while tutorial-like presentations are given in [52, 54].

2.1. Chemical reaction networks

First, we fix the notation and terminology for reaction networks. A CRN is a pair $(\mathcal{S}, \mathcal{R})$ consisting of a finite set of species \mathcal{S} and a finite set of reactions \mathcal{R} , where each reaction is in the form $\rho \xrightarrow{f} \pi$ where:

- ρ is a multiset of species which represents the *reactants*, where the multiplicity of species S is denoted by ρ_S ;
- π is a multiset of species which represents the *products* of the reaction;
- $f : \mathbb{R}^{\mathcal{S}} \rightarrow \mathbb{R}_{\geq 0}$ is the propensity function, which, roughly speaking, gives the rate at which the reaction fires.


The domain of f denotes the current state of the system: in the deterministic ODE interpretation, this is the current concentration of each species; in the stochastic CTMC representation, the state is represented by the number of molecules of each species. The *net stoichiometry* $\pi - \rho$ gives the state update upon the reaction firing. As usual, the + and - operators denote multiset union and difference, respectively, while the multiplicity of a species denotes its stoichiometric coefficient.

Deterministic semantics. According to the deterministic semantics of CRNs [55], a CRN is associated with an ODE system which tracks the time course of the concentrations of the species at time t , $X(t) = (X_S(t))_{S \in \mathcal{S}}$, as follows:

$$\frac{dX_S(t)}{dt} = \sum_{(\rho \xrightarrow{f} \pi) \in \mathcal{R}} (\pi_S - \rho_S) \cdot f(X(t)). \quad (1)$$

In a CRN with *mass-action kinetics* [55] (denoted by MRN throughout the remainder of this paper) with deterministic semantics, the form of each propensity function depends on a kinetic parameter $\lambda > 0$ and on the products of the concentrations of the reactants, considering also their multiplicities. A mass-action reaction is usually represented as $\rho \xrightarrow{\lambda} \pi$ and the propensity function, denoted by $f_{\lambda, \rho}$, is given by $f_{\lambda, \rho}(X(t)) = \lambda \prod_{S \in \mathcal{S}} X(t)^{\rho_S}$.

Example 1. Let us consider the MRN with species S_1, S_2, S_3, S_4 , and S_5 , and reactions



According to (1), the ODEs are given by:

$$\begin{aligned}
\frac{dX_{S_1}(t)}{dt} &= -3X_{S_1}(t) \\
\frac{dX_{S_2}(t)}{dt} &= -3X_{S_2}(t) \\
\frac{dX_{S_3}(t)}{dt} &= +2X_{S_1}(t) + 2X_{S_2}(t) + 3X_{S_4}(t)X_{S_5}(t) \\
\frac{dX_{S_4}(t)}{dt} &= -3X_{S_4}(t)X_{S_5}(t) \\
\frac{dX_{S_5}(t)}{dt} &= -3X_{S_3}(t)X_{S_5}(t) - 3X_{S_4}(t)X_{S_5}(t) + 2X_{S_1}(t) + 2X_{S_2}(t).
\end{aligned} \tag{2}$$

With this ODE system is associated an initial value problem by taking an initial condition $X(0) = (X_{S_1}(0), \dots, X_{S_5}(0))$ representing the initial concentration for each species in the CRN.

Stochastic semantics. According to the stochastic semantics, a CRN is represented as a Markov population process. This is a CTMC where each state is a vector $n = (n_s)_{s \in S}$ of nonnegative integers that tracks the molecular counts of each species. Each reaction $\rho \xrightarrow{f} \pi$ may fire a transition between a state n and $n + \pi - \rho^2$ according to an exponential distribution with parameter $f(n)$. The underlying CTMC can be derived by exhaustively applying the reactions on all reachable states, starting from a given one that defines the initial population count for each species. Although the stochastic semantics can be defined for propensities with arbitrary functional form, the equivalence relation studied in this paper, SE [18], applies to MRNs only. The stochastic semantics of MRNs have the well-known interpretation of exponentially distributed reaction times with rates that are proportional to the total number of possible collisions between molecules occurring in a well-stirred solution, building on a celebrated result by Gillespie [28]. Specifically, for a given kinetic parameter $\lambda > 0$, the mass-action propensity function for a reaction $\rho \xrightarrow{\lambda} \pi$ in state n is given by $f_{\lambda, \rho}(n) = \lambda \prod_{s \in S} \binom{n_s}{\rho_s}$.

Let us consider Example 1 assuming a stochastic semantics with initial state $n = (n_{S_1}, \dots, n_{S_5}) = (1, 1, 1, 0, 0)$. Then, the following are the transitions of the CTMC from this state:

$$\begin{aligned}
(1, 1, 1, 0, 0) &\xrightarrow{2} (0, 1, 1, 0, 1) \\
(1, 1, 1, 0, 0) &\xrightarrow{2} (0, 1, 3, 0, 0) \\
(1, 1, 1, 0, 0) &\xrightarrow{2} (1, 0, 2, 0, 0) \\
(1, 1, 1, 0, 0) &\xrightarrow{1} (1, 0, 1, 0, 2)
\end{aligned} \tag{3}$$

where we have used the customary notation $n_s \xrightarrow{q} n_t$ where n_s is the source state, n_t is the target state and q is the transition rate computed according to mass-action kinetics evaluated at the source state.

2.2. Equivalence relations

Forward and backward equivalence. By (1), it is clear that a deterministic MRN gives rise to a system of ODEs with polynomial derivatives. However, not every system of polynomial ODEs can be represented as an MRN. Indeed, according to (1) every monomial with negative coefficient appearing in the derivative of some variable X_S must feature X_S itself with a positive exponent. It turns out that it is possible to represent any polynomial ODE system as a CRN by further allowing kinetic parameters of mass-action reactions to be negative [12, 13]. Although negative kinetic parameters do not carry any physical interpretation, mathematically the finitary representation of a polynomial ODE system in terms of a reaction network enables the definition of FE and BE as equivalence relations over species which

²With a slight abuse of notation, here and in the rest of the paper we use multisets π and ρ to denote also vectors $n_\pi = (\pi_s)_{s \in S}$ and $n_\rho = (\rho_s)_{s \in S}$ counting the multiplicity of each species in the multiset.

can be checked by using only structural conditions on the reactions [13]. Because of this, throughout the remainder of this paper we will use the term *polynomial reaction networks* (PRN) to indicate CRNs which lead to polynomial ODE systems that do not follow the mass-action kinetics.

The next example further clarifies the relationship between polynomial ODEs and PRNs by sketching how to systematically convert the former into the latter.

Example 2. *Let us consider a CRN where species S_1 , S_2 , and S_3 are involved in the single reaction*



By (1), this leads to the following polynomial ODE system:

$$\frac{dX_{S_1}(t)}{dt} = -\lambda X_{S_1}(t) X_{S_3}(t), \quad \frac{dX_{S_2}(t)}{dt} = \lambda X_{S_1}(t) X_{S_3}(t), \quad \frac{dX_{S_3}(t)}{dt} = 0. \quad (5)$$

By [12, 13], this can be turned into a PRN by encoding each monomial m appearing in the right-hand side of the ODE of species S_i as a mass-action reaction having: all species in m with positive exponents as reagents; S_i as well as all species in m as products; the coefficient of the monomial as kinetic parameter. This transformation leads to the following PRN:



As a side remark, we note that the above systematic transformation is not unique. Indeed, reaction (4) implicitly uses species S_3 as a catalyst, i.e., a species that affects the actual propensity function of the reaction, but that is neither consumed nor produced by it. The same ODE system (5) is generated by the following single mass-action reaction:



Both notions of FE and BE adhere to the same template which can be summarized as follows.

For $\chi \in \{\text{FE}, \text{BE}\}$, given a PRN $(\mathcal{S}, \mathcal{R})$, a partition \mathcal{H} of species is χ if and only if for any two blocks $H, H' \in \mathcal{H}$ and any two species $S_i, S_j \in H$ it holds

$$\mathbf{c}_\chi(S_i, \eta, H') = \mathbf{c}_\chi(S_j, \eta, H') \quad \forall \eta. \exists (S_k + \eta \xrightarrow{\lambda} \pi) \in \mathcal{R} \text{ for } k = i, j, \quad (7)$$

where \mathbf{c}_χ maps a species (S_i, S_j) , a multiset of reagent partners (η) and a block of species (H') into a real number computed by inspecting the reactions once [13]. The appropriate definition of \mathbf{c}_χ characterizes the type of dynamical aggregation that can be obtained at the level of the underlying polynomial ODEs, as introduced in Section 1.

Specifically, if a partition of species that satisfies the FE criteria then there exists an aggregated polynomial ODE system which describes the dynamics of the sums of the solutions of the species in each partition block/equivalence class. In fact, this turns out to be a characterization result, namely if such an aggregated system can be found, then the FE conditions are satisfied [13]; in practice, this is a relevant property because it ensures that with the syntactic checks based on the aforementioned template, no dynamical aggregations are missed.

As an example, it is possible to show that the partition \mathcal{H} given by

$$\mathcal{H} = \{\{S_1, S_2\}, \{S_3, S_4\}, \{S_5\}\} \quad (8)$$

is an FE for the MRN of Example 1. This can be verified by simple algebraic manipulations of its ODE system, showing that one can write a self-consistent ODE system in terms of the sums of variables $X_{S_1} + X_{S_2}$ and $X_{S_3} + X_{S_4}$, as well as X_{S_5} on its own:

$$\begin{aligned} \frac{d(X_{S_1}(t) + X_{S_2}(t))}{dt} &= -3(X_{S_1}(t) + X_{S_2}(t)) \\ \frac{d(X_{S_3}(t) + X_{S_4}(t))}{dt} &= +2(X_{S_1}(t) + X_{S_2}(t)) \\ \frac{dX_{S_5}(t)}{dt} &= -3X_{S_5}(t)(X_{S_3}(t) + X_{S_4}(t)) + 2(X_{S_1}(t) + X_{S_2}(t)). \end{aligned}$$

That is, by introducing the change of variables $Y_{1,2} = X_{S_1} + X_{S_2}$, $Y_{3,4} = X_{S_3} + X_{S_4}$, and for consistency $Y_5 = X_{S_5}$, one gets the aggregated ODE system

$$\begin{aligned}\frac{dY_{1,2}(t)}{dt} &= -3Y_{1,2}(t) \\ \frac{dY_{3,4}(t)}{dt} &= +2Y_{1,2}(t) \\ \frac{dY_5(t)}{dt} &= -3Y_5(t)Y_{3,4}(t) + 2Y_{1,2}(t).\end{aligned}\tag{9}$$


This means that the solution to this three-variable ODE system satisfies $Y_{1,2}(t) = X_{S_1}(t) + X_{S_2}(t)$, $Y_{3,4}(t) = X_{S_3}(t) + X_{S_4}(t)$ and $Y_5(t) = X_{S_5}(t)$ at all time points, thus yielding an exact, albeit aggregate, relationship with the solution of the original system (1). In particular we note that, in general, from the solution of the aggregated system it is not possible to recover all the individual solutions of the original variables.

Importantly, this FE relationship holds for any initial concentration of species. BE, instead, aggregates species that have the same solutions at all time points, but only when initialized equally. For example it is easy to see that the partition \mathcal{H}' given by

$$\mathcal{H}' = \{S_1, S_2\}, \{S_3\}, \{S_4\}, \{S_5\}$$

is a BE for the MRN of Example 1. This descends from the fact that the equations for X_{S_1} and X_{S_2} are not coupled; indeed their solutions are $X_{S_i}(t) = X_{S_i}(0)e^{-3t}$, for $i = 1, 2$, from which it is apparent that they have equal trajectories if $X_{S_1}(0) = X_{S_2}(0)$. Similarly to FE, also the BE criteria characterize this dynamical property. Finally, we observe that the BE partition \mathcal{H}' is a refinement of the FE partition \mathcal{H} in this specific example. In general, however, FE and BE are not comparable [10, 12].

Given an FE or a BE, the reduction algorithm presented in [13] constructs a *quotient CRN* from the original CRN such that the analysis of the former gives the solution of the aggregated ODE system. For either equivalence, each *macro-species* in the quotient CRN represents the dynamics of each equivalence class. For example, let us consider again the FE in Equation (8). Denoting by $S_{1,2}$ and $S_{3,4}$ the macro-species for the blocks $\{S_1, S_2\}$, $\{S_3, S_4\}$, respectively, the quotient MRN is given by:



It is easy to see that applying the deterministic mass-action semantics to this MRN yields the ODE system in (9).

Using the same reduction algorithm, with the BE \mathcal{H}' one gets the following quotient MRN:



This yields the following aggregated ODE system:

$$\begin{aligned}\frac{dX_{S_{1,2}}(t)}{dt} &= -3X_{S_{1,2}}(t) \\ \frac{dX_{S_3}(t)}{dt} &= +2X_{S_{1,2}}(t) + 3X_{S_4}(t)X_{S_5}(t) \\ \frac{dX_{S_4}(t)}{dt} &= -3X_{S_4}(t)X_{S_5}(t) \\ \frac{dX_{S_5}(t)}{dt} &= -3X_{S_3}(t)X_{S_5}(t) - 3X_{S_4}(t)X_{S_5}(t) + 2X_{S_{1,2}}(t).\end{aligned}\tag{10}$$

However, since this is related to BE, the relationship $S_{1,2}(t) = S_1(t) + S_2(t)$ only holds when $S_1(0) = S_2(0)$. Indeed, since $S_1(t) = S_2(t)$ under this constraint, one gets that $S_i(t) = S_{1,2}(t)/2$ for $i = 1, 2$. More in general, BE allows one to recover the individual ODE solutions from the original model by dividing the aggregated solution by the size of each corresponding block.


Forward and backward differential equivalence. FDE and BDE are generalizations of FE and BE, respectively, for CRNs with deterministic semantics that lead to ODEs with nonlinearities beyond polynomial [11, 17]. For applications in computational systems biology, these include, e.g., derivatives with rational expressions as used in Michaelis-Menten or Hill kinetics. Throughout the remainder of this paper, we shall refer to this class of networks with the term *nonlinear reaction networks* (NRNs). The greater generality of FDE/BDE comes at the cost of more computationally expensive checks based on encodings in satisfiability modulo theories (SMT) [11, 17].

For example, checking BDE can be done by encoding the property that, for species to be in the same block, their corresponding ODE variables must have the same derivative whenever they are equal. For instance, the following formula ψ encodes the check whether partition \mathcal{H}' introduced above is a BDE:

$$\psi := (X_{S_1} = X_{S_2}) \implies (-3X_{S_1} = -3X_{S_2}) \quad (11)$$

By using an SMT solver, we can check if \mathcal{H}' is a BDE by checking for the satisfiability of $\neg\psi$. If there exists an assignment for X_{S_1} and X_{S_2} that makes $\neg\psi$ true, then \mathcal{H}' is not a BDE. This is not the case, and hence it is a BDE (as expected from it being a BE). Otherwise, the SMT solver returns a counterexample, i.e. an evaluation for the variables that falsifies ψ . This means that the solver identified an assignment for the variables such that variables in same block get assigned same value, but the updated functions of variables in the same block evaluate to different values. Such counterexample is used in the partition-refinement algorithm to compute the maximal BDE. In ERODE, this computation is done using the well-known solver Z3 [24].

Species equivalence. Among the notions of equivalence considered in this paper, SE is the only one developed for stochastic semantics and is defined MRNs [18]. It is defined in a similar fashion as FE and BE by identifying criteria at the level of the reaction network. These induce a notion of aggregation for the underlying CTMC that, essentially, enables the description of the stochastic dynamics of macro-species representing the sum of the original species in each block. For instance, it can be shown that the partition $\mathcal{H}'' = \{\{S_1\}, \{S_2\}, \{S_3, S_4\}, \{S_5\}\}$ is an SE for the MRN of Example 1. In [18] it is shown that the same reduction algorithm as for FE and BE can be used to obtain the quotient network, which reads:



In order to exemplify the relationship with the original CTMC, let us consider the initial state $n = (n_1, n_2, n_{3,4}, n_5) = (1, 1, 1, 0)$. Notice that this corresponds to the original state $n = (n_1, n_2, n_3, n_4, n_5) = (1, 1, 1, 0, 0)$ from which we obtained the sample transitions (3), by summing the initial population counts for species S_3 and S_4 , as identified by SE. These two CTMCs are related by the property that the probability distribution of species $S_{3,4}$ corresponds to the probability distribution of the overall population for species S_3 and S_4 in the original CTMC. Technically, SE proves that the original CTMC enjoys an aggregation according to ordinary lumpability [9], where all original states that have the same overall counts for species S_3 and S_4 can be collapsed into a single macro-state of the CTMC, and this aggregated CTMC corresponds to that generated by the above quotient network.

In this example \mathcal{H}'' is a refinement of the FE \mathcal{H} . This is a specific case of the general result that SE implies FE, but not vice versa [15]. In Section 4 we will show that the different reduction power of these two equivalences can be observed in practice in the BioModels instances.

Partition refinement. Thus far, the overviewed criteria for equivalence require a candidate partition to be provided. For all equivalence notions there exists the *maximal* equivalence, i.e., the largest one leading to the coarsest aggregation. In all cases, this can be computed by means of a partition refinement algorithm which refines an initial partition of species (splitting its blocks) until a fixed point, corresponding to the partition that is the coarsest refinement of the input one. The maximal equivalence is thus computed in the case where the input partition is the trivial one where all species are in the same block. Other interesting equivalences can be computed by choosing different initial partitions. For example, the modeler may choose to preserve the full dynamics of a given species of interest when using FE, by isolating such species in a singleton initial block. With BE, some domain knowledge that two species cannot have the same initial condition may be realized by placing such species in distinct initial block.

For FE, BE, and SE, which are based on the template criterion in (7), the iterations are based on computing those quantities and splitting the blocks of the current candidate partition into sub-blocks with equal values. For FDE and



Figure 1: Workflow overview. Models were downloaded from the BioModels repository in the SBML format. We implemented in ERODE an importer to translate the SBML description into the CRN-like input (.ode format) of ERODE. The output of ERODE is a reduced CRN with reactions involving *macro-species*, each representing the sum within an equivalence class of original species. We manually inspected the ERODE output to provide a physical interpretation of the obtained equivalences.

BDE, which are based on SMT checks, the approach is similar but it is based on the provided counterexample: the current candidate partition is refined such that the values of the counterexample within each sub-block are equal.

3. Experimental set-up

3.1. Overview of the BioModels repository

The BioModels Database is a repository of computational models of biological processes [39]. It hosts dynamical quantitative models described in peer-reviewed scientific literature as well as models generated automatically from pathway resources such as KEGG [34], BioCarta [43], MetaCyc [19], PID [48] and SABIO-RK [57]. BioModels covers a wide range of models from several biological categories such as biochemical reaction systems, kinetic models, metabolic networks, steady-state models and signaling pathways. Models are available in the Systems Biology Markup Language (SBML) [31], a well-known machine-readable format based on XML for representing quantitative models of biological systems.

The BioModels repository is divided into two sections: the *curated branch* and the *non-curated branch*. The former contains models that have been manually checked to match the outputs presented in the scientific articles where they are presented. Furthermore, the model components are annotated using unambiguous identifiers [33] that refer to external biological databases [27, 51, 22] or ontologies (such as Gene Ontology [4], SBO [23] or ChEBI [25]). Models are curated following the Minimum Information Required in the Annotation of Models guidelines (MIRIAM) [37]. Models that are not MIRIAM-compliant are stored in the non-curated branch, which also contains non-kinetic models such as flux balance analysis models. A more detailed description of BioModels is available at [21].

In our experiments we used the BioModels repository snapshot 26 July 2017, the latest available at the time of writing this paper.³ The repository consists of 640 models in the curated branch (from id *BIOMD0000000001* to *BIOMD0000000640*) and 1000 models in the non-curated branch (with ids ranging from *MODEL0072364382* to *MODEL9811206584*).

3.2. Model conversion

We implemented an importer within ERODE for translating SBML models into ERODE’s format, using the workflow in Figure 1. SBML files were read using the *jsbml* library (version 1.2) [38, 26]. Here we briefly explain the main phases of the conversion process, which, unlike the prototype first presented in [47], is now fully automated.

³Downloaded from <ftp://ftp.ebi.ac.uk/pub/databases/biomodels/releases/>.

```

1 <reaction id="reaction_0000001" metaid="metaid_0000046" name="binding MAPKK on Tyr site of MAPK">
2   <listOfReactants>
3     <speciesReference metaid="_063184" species="M"/>
4     <speciesReference metaid="_063196" species="MAPKK"/>
5   </listOfReactants>
6   <listOfProducts>
7     <speciesReference metaid="_063208" species="M_MAPKK_Y"/>
8   </listOfProducts>
9   <kineticLaw metaid="_063220">
10     <math xmlns="http://www.w3.org/1998/Math/MathML">
11       <apply>
12         <times/>
13         <ci>cell</ci>
14         <apply>
15           <minus/>
16           <apply>
17             <times/><ci>k1</ci><ci>M</ci><ci>MAPKK</ci>
18           </apply>
19           <apply>
20             <times/><ci>k_1</ci><ci>M_MAPKK_Y</ci>
21           </apply>
22         </apply>
23       </math>
24     </kineticLaw>
25 </reaction>
26

```

Figure 2: Sample SBML reaction adapted from *BIOMD0000000030*.

3.2.1. From SBML to ERODE format: parameters and species

The input format of an ERODE specification contains three lists: parameters (to be used, e.g., in kinetic rates); species (with corresponding initial conditions); and reactions. These are followed by a list of commands for analysis, reduction, and export.

The following SBML snippet, from *BIOMD0000000030*, specifies a parameter

```
<parameter id="k1" metaid="metaid_0000019" name="k1" value="0.02"/>
```

This is translated into a simple assignment $k1 = 0.02$ within the parameter list in the ERODE specification.

The next SBML snippet, adapted from the same model by removing the annotation tag containing links to external databases, defines the species M:

```
<species compartment="cell" id="M" initialConcentration="800" metaid="metaid_0000005" name="MAPK"/>
```

It describes the compartment in which the species is located, the initial concentration and an identifier. This is translated this into the assignment $M = 800$ within ERODE's species list.

3.2.2. From SBML to ERODE format: reactions

The conversion of the reactions is less straightforward, particularly to identify MRNs, to which the specialized FE, BE, and SE can be applied, or PRNs, to which FE and BE can be applied. Indeed, SBML allows the direct specification of mass-action reactions by means of appropriate SBO labels in the `kineticLaw` tag (other labels identify different kinetics such as Michaelis-Menten and Hill). However, many situations involved reactions that, although not tagged with mass-action labels, were clearly so upon inspection of the reactions. One such example is given in Figure 2. It shows the SBML specification of a reaction containing a list of reactants, products (as well as modifiers, not used in this reaction, to model, e.g., catalysts or intermediates in the reaction). Such SBML description is translated into a formal reaction as follows:



where the coefficient 1 denotes the volume of the compartment `cell` in the SBML specification.

<i>Dataset</i>	<i>MRN</i>	<i>PRN</i>	<i>NRN</i>	<i>Total</i>
Curated branch	50	77	267	394
Non-curated branch	20	76	84	180
BioModels snapshot 26 July 2017	70	153	351	574

Table 2: Preprocessing results of the BioModels repository.

A manual inspection of the form of the propensity function reveals that it subsumes a pair of mass-action reactions (the forward reaction and its corresponding reversed one) with kinetic parameters k_1 and k_{-1} , respectively:



While a straightforward translation of (12) into ERODE would yield an NRN, the identification of an MRN (13) would enable the faster equivalences FE and BE for deterministic semantics, as well as SE in the case of the stochastic semantics.

ERODE supports a command to automatically detect such implicit mass-action reactions. Specifically, the propensity function of each reaction is parsed, expecting a polynomial expression over the species defined in the model. The command separates the positive and negative monomials, and attempts to generate a mass-action reaction for each of them. For the positive monomials, it checks that they contain the product of the reagents in agreement with the ODE semantics of MRNs given in Section 2. Instead, negative monomials are used to create reverse reactions, meaning that the reagents and the products get swapped, therefore the check is done on the products. If any of these checks fails, the command gives a detailed error message pointing to the problem found. Another ERODE command, implemented in a similar fashion, can be used to attempt a conversion into a PRN (following the encoding sketched in Example 2).

3.3. Repository Preprocessing

We performed an initial preprocessing step to filter out models that could not be used for the analysis (cf. evaluation question **Q1** in Section 1). In the non-curated branch only 428 models are kinetic models described as ODE systems, while the others are described in formalisms, such as logical or flux balance analysis models, that are outside the scope of applicability of the model-reduction techniques. Overall, we could process 394 models from the curated branch and 180 from the non-curated one, for an overall sanitized dataset of 574 models. The most frequent reasons for discarding a model were (within parenthesis we give the frequency in the curated branch, which we assume to be more stable):

- syntactic limitations in our importer, including the lack of support for models without explicit reactions where the dynamics is given by rate rules over a set of parameters, e.g., as in *BIOMD0000000020* (81);
- models with unsupported propensity functions such as tanh and exp, which make the satisfiability problem underlying FDE/BDE undecidable (82);
- models with non integer stoichiometries (multiplicities) in the reagents or products of reactions (16).

In particular, we ensured to discard all models which, after importing, resulted in having no dynamics (e.g. due to the presence of not supported features like those listed above).

On the sanitized dataset, the SBML models were first imported as NRNs, as, for instance, shown in (12). Subsequently, using the aforementioned ERODE commands, they were categorized into more specific classes by attempting a translation into an MRN, and, when this failed, into a PRN. The results of this classification are summarized in Table 2. All MRN models turned out to be elementary. Throughout the remainder of this article, when reporting model sizes the number of reactions will refer to those obtained after translation into MRNs and PRNs, when applicable. In order to test the correctness of the translation from the SBML description, on a large random selection of models we verified that the ODE solutions obtained by the ERODE’s and BioModel’s Matlab exports coincided.

It is worth noticing that the proportion of PRNs is significantly larger in the non-curated branch (ca 42%) than in the curated branch (ca 28%), while the proportion of NRNs is much smaller (about 46% and 68% in the non-curated and curated branch, respectively). Although a systematic manual investigation of the possible causes of such differences is outside the scope of this paper, the following example shows that the higher proportion of PRNs in the non-curated branch may be due to imprecisions in the original model submission that are fixed after curation. Let us take the non-curated model *MODEL1411130000*, taken from [41]. Every reaction of this model indicates the presence of a *modifier*—a species that, according to the SBML specification [32],

“appears in the kinetic rate formula of a reaction but is itself neither created nor destroyed in that reaction (for example, because it acts as a catalyst or inhibitor).”

The presence of modifiers does not allow the reaction to be recognized as mass-action (where, e.g., a catalyst can be encoded by indicating the species both in the reactants and in the products). However such a reaction may still be recognized as a PRN. Model *MODEL1411130000* is one such instance, although the use of modifiers may not be appropriate. For instance, reaction id *r4* has species *s1* as both reactant and modifier, and species *s3* as product, with a kinetic law defined to be equal to $\text{gamma_s1m} \cdot s1$, where *gamma_s1m* is a model parameter. An inspection of the generated ODEs, however, reveals that the kinetics is consistent with a mass-action reaction $s1 \xrightarrow{\text{gamma_s1m}} s3$, i.e., without the explicit indication of the modifier. This is also confirmed by an examination of the original article from which the model was taken, i.e., ref. [42]. Indeed, we found that the propensity function defined in the article is consistent with the SBML specification, but the article does not explicitly mention the presence of modifiers in the reaction. This kind of imprecision misclassifies the model, but does not affect the correctness of the translation because the underlying ODEs correspond.

4. Reduction results

We hereby report the results of our experiments. Section 4.1 describes the experimental configuration, while Section 4.2 and 4.3 discuss the reduction power of our techniques in terms of percentage of reduced models and obtained reduction ratios, respectively. Section 4.4 concludes the section by comparing the results obtained for the *forward* and *backward* reduction notions. Considerations on the actual analysis speedups obtainable thanks to our techniques are out of the scope of this paper, because the considered SBML models do not contain information related to this. However, we remark that the speedups obtainable with our techniques have been analyzed in [18] for CTMC semantics, and in [10] for the ODE one.

4.1. Experiments configuration

Here we report the configurations used for our experiments. MRNs were analyzed using FE, BE, and SE; PRNs were analyzed using FE and BE; NRNs were analyzed using FDE and BDE. For each equivalence we computed the reduced model; we used the resulting number of species and reactions as a measure of the effectiveness of the exact reduction techniques (cf. **Q2** in Section 1).

As discussed in Section 2, all the techniques considered in this paper are based on a partition-refinement algorithm. Hence, the reduction power depends on the choice of the initial partition to be refined. This information, however, cannot be directly inferred from the original SBML specification. To allow for a uniform, model-independent treatment across the whole repository, the initial partitions were systematically chosen as follows.

- For BE and BDE reductions we fixed initial partitions where each block contains species with same initial concentration; this guarantees that from the solution of the quotient reaction network one can recover the solutions of the original model when starting from the initial conditions specified in the SBML specification.
- For FE, FDE, and SE we computed the *maximal* reductions obtained by using initial partitions with one block only that contains all the species in the model. This gives an indication of the best achievable reductions—different initial partitions will yield finer quotients. The main limitation is that it may collapse individual species which are outputs of interest to the modeler. However, this information is not stored in the SBML file, hindering a systematic identification of model-dependent initial partitions across the whole repository.

In the next sections we present the reduction results for the curated and the non-curated branches separately.




Figure 3: Pie charts illustrating how many models can be reduced by which techniques. Label *N.R.* indicates that the models could not be reduced by any of our techniques as implemented in ERODE within a time limit of 30 minutes.

4.2. Reduction capability

Figure 3 shows how many techniques could reduce a model regardless of the achieved reduction sizes. Overall, 136 models (ca 23%) could not be reduced, while 224 models (ca 64%) could be reduced by at least one technique. For the remaining 67 models (ca 13%), either FDE or BDE failed to terminate. This was due either to errors raised by the Z3 solver, or after an imposed timeout expiring when a single SMT check exceeded 30 minutes.⁴ FDE failed on 16 more models where BDE terminated successfully; this is consistent with the fact that more (and more complex) SMT checks are required by FDE than BDE [11].

For MRNs and PRNs, Figures 3a, 3b, 3d, and 3e confirm that BE and FE are not comparable (only a subset of the models that can be reduced by either is also reducible by both). For the mass-action models, none of the models in Figures 3a, 3d were able to distinguish FE and SE, meaning that all models reducible by one, could be reduced also by the other. Figures 3c and 3f similarly confirm that FDE and BDE are not comparable.

For MRNs and NRNs, a comparison of the curated and the non-curated branches highlights that the percentage of models that could not be reduced by any technique is significantly lower in the latter branch, while it is almost the same for PRNs. In particular, the percentage of models that can be reduced by BE and BDE is higher in the non-curated branch (ca 30%) than in the curated one (ca 46%). It is outside the scope of this paper to investigate whether

⁴It is worth to note that in [47] we reported a much higher number of models in which FDE failed. The better performance reported here is due to improvements in the implementation of FDE in ERODE and in the importer, and due to the fact that the new class of PRNs is treated with FE/BE rather than FDE/BDE.




Figure 4: Distribution of size of original and reduced models (log-log scale). Pairs of dots for an original/reduced model is connected by a line

Table 3: Average reduction ratios 'reduced over original' for species and reactions

		<i>FE</i>	<i>FDE</i>	<i>SE</i>	<i>BE</i>	<i>BDE</i>
<i>Curated branch</i>	Species	60%	63%	46%	81%	83%
	Reactions	66%	69%	48%	89%	93%
<i>Non-curated branch</i>	Species	70%	50%	49%	85%	75%
	Reactions	76%	60%	46%	93%	88%

such statistical differences are due to model-specific properties or by the presence of imprecisions in the models in the non-curated branch; indeed, this is the very aim of the manual curation process conducted by the BioModels team. However, for the sake of illustration, the following example shows that more care must be exercised for the analysis of non-curated models: model *MODEL4821294342* is a PRN with 12 species, of which 4 with no dynamics, i.e., their corresponding equations have derivative equal to zero. Because of their initial conditions, these species are placed into two distinct blocks in the initial partition, which are never refined. These kinds of species may actually represent model parameters rather than biochemical compounds and could be effectively removed from the description without affecting the validity of the results.

4.3. Reduction sizes

Figure 4 shows scatter plots to summarize the distribution of the size of the original (red) and reduced (blue) models from the two branches. For each model we considered the size corresponding to the best reduction, avoiding though too aggressive reductions with 1 species. This can happen for FE, SE, and FDE, because, as discussed in Section 4.1, we considered initial partitions leading to maximal reductions. Each pair of blue/red dots referring to the same models is connected by a line. We can see that most of such lines are diagonal, indicating that a reduction in the number of species leads most of the times also to a reduction in the number of reactions. In the curated branch the original models are mostly concentrated in the region corresponding to the interval $[10^1, 10^2]$ for the species axis, and to $[10^1, 10^3]$ for the reactions axis. Visually, the sizes of the reduced models are translated to the bottom-left of the figures, indicating the effectiveness in practice of the reductions. Focusing on the top-right areas of the two figures we can see both cases of large models that could be effectively reduced (red dots *far away* from the corresponding blue ones), and cases of limited reduction (red dots *close* to the corresponding blue ones).

Table 3 shows the average reduction ratios for each technique in the respective sets of models where they have been applied (see Section 4.1) from the two branches, measured as the number of species or reactions in the reduced model over the corresponding number in the original model. We note that FDE on the non-curated branch has the




Figure 5: Reduction ratios for species and reactions for all techniques. The y-axis counts the number of models, while the x-axis gives the reduction ratio of the reduced species (reactions) over the original species (reactions). The number within the rightmost bar refers to the number of models whose number of reactions could not be reduced.

best performance, followed by the other FE and FDE cases. SE, applied only to MRNs, is the technique that performs best, in particular better than FE and FDE. This might surprise at first, because SE is stricter than FE (and therefore of FDE). However, this can be explained by the fact that FE has been applied to both MRNs and PRNs. Indeed, the average reduction ratios of FE for MRNs only, not shown in the table, coincide with the ones of SE. Therefore, the models from this repository do not allow us to distinguish FE and SE neither in terms of percentage of reduced models (Fig 3), nor in terms of actual reduction ratios.

Figure 5 provides more information on the magnitude of the obtained reductions. For each technique, the models are grouped in 5 bins (0%-20%, ..., 80%-100%), for which we show the reduction ratios in the number of species and reactions as separate series. In both branches, the figures for BE and FE present distributions similar to the corresponding techniques for non-mass-action non-polynomial models (FDE and BDE, respectively). The only exception is that FDE in the non-curated branch has a higher percentage of models in the bins 0%-20%. A similar pattern can be appreciated comparing each reduction technique in the two branches: all figures are pairwise similar, apart for the 0%-20% bins of FDE.

It is worth to note that the bins 80%-100% of reactions are always the largest one in the corresponding figure. In fact, these bins also include the cases with models showing no reductions in the number of reactions. This can be due to an equivalence among species with no dynamical role in the network, as they can be interpreted as distinct auxiliary species that are used to model zero-order reactions, such as I in reaction $I \rightarrow I + A$, a purely catalytic species C in a reaction like $A + C \rightarrow B + C$, or $SINK$ in a degradation reaction such as $A \rightarrow SINK$. In the first two cases, these species are associated with zero-derivative variable, while in the last case the variable for $SINK$ does not appear in any ODE in the system. Each figure provides the number of such models in bold text in the middle of the last reactions bin.

4.4. Comparison between forward and backward reductions

The results provided in this section show that a high number of models from the curated branch could be reduced by at least one of our techniques: 241 (ca 61%) out of the 394 that could be imported. This percentage is even higher in the non-curated branch, where 130 (ca 70%) models could be reduced by at least one technique. When considering networks with deterministic semantics, it is possible to compare *forward reductions* (i.e., FE and FDE) against *backward reductions* (i.e., BE and BDE). Forward reductions were applied to 212 curated models (54%), and 118 non-curated ones (66%). Backward reductions, instead, were successful on fewer models: 124 curated models (31%) and 84 non-curated ones (46%). As shown in Table 3, this trend is confirmed also when looking at the actual reduction power of the techniques in terms of number of obtained reduced reactions and species. This pattern is less

evident in Figure 5, however we note that the bins from 0% to 60% for FE and FDE have higher values than the corresponding ones for BE and BDE.

5. Robustness of the reduction techniques to parameters values on selected case studies

A CRN with deterministic semantics can be extended into a new one with dynamically equivalent ODEs where the reductions are independent from the actual values of the model parameters. This enables the analysis of the robustness of the reduction techniques, a property that is particularly relevant in computational systems biology, where the parameters are not precisely known in general. The main idea behind this extension is to promote each parameter to a species with no dynamics (i.e., which appears only as catalyst in mass-action reactions, or that appears only in propensity functions in non-mass-action reactions). For instance, given a mass-action reaction



the parameter λ is translated into a new species S_λ ; the reaction in Equation (14) is replaced with one where S_λ is added to both the reagents and products, with kinetic constant set equal to 1:



According to Section 2, the two reactions describe the same underlying ODEs, apart for the fact that in the latter case we get an additional ODE for S_λ :


$$\frac{dX_{S_\lambda}(t)}{dt} = 0.$$

The value of S_λ does not change and will be equal to its initial value for all time points. Therefore, if the initial condition is $X_{S_\lambda}(0) = \lambda$, we will get ODEs with same solution in the original and extended models. Intuitively, S_λ can be seen as a *symbolic* version of parameter λ , allowing to abstract from the actual value of the latter. For non-mass-action CRNs the transformation is even simpler since it does not involve the modification of the reactants but only the replacement of all occurrences of the kinetic parameters with their corresponding species in every propensity function. This CRN extension is further exemplified in Example 3.

Example 3. *The CRN from Example 2 can be extended by adding a new species S_λ for the parameter λ , changing the reaction in Equation (4) into $S_1 \xrightarrow{X_{S_\lambda}(t) \cdot X_{S_1}(t) \cdot X_{S_3}(t)} S_2$. The ODEs underlying such CRN are the same as those in Equation (5), apart for the fact that λ is replaced by X_{S_λ} , which also gets assigned an ODE:*

$$\frac{dX_{S_1}(t)}{dt} = -X_{S_\lambda}(t) \cdot X_{S_1}(t) \cdot X_{S_3}(t) \quad \frac{dX_{S_2}(t)}{dt} = X_{S_\lambda}(t) \cdot X_{S_1}(t) \cdot X_{S_3}(t) \quad \frac{dX_{S_3}(t)}{dt} = 0 \quad \frac{dX_{S_\lambda}(t)}{dt} = 0 \quad (16)$$


Provided we set $X_{S_\lambda}(0) = \lambda$, then these ODEs have same solution as those in Equation (5). Let us now consider instead the MRN version of the same system consisting of the reaction in Equation (6). Here we can apply as well the extension promoting λ to the species S_λ . However, this requires to change the reaction in Equation (6) as follows:



Such reaction will produce the ODEs in Equation (16).

The advantage in considering an extended CRN rather than the original one is that the considered reduction techniques ignore the values of the initial concentrations. Indeed, FDE and FE do not depend on the initial conditions. Thus, any reduction obtained on such expanded CRN will hold for any value of the original parameters. Instead for BDE and BE it is convenient to initialize the partition-refinement algorithm with an initial partition where every species corresponding to original model parameters is isolated in a singleton blocks. This would guarantee meaningful reductions for all initial values of all variables, therefore for all possible values assigned to the parameters.

Example 4. Here we consider an extended version of the MRN in Example 1 where kinetic constants 1, 2, and 3 have been first replaced by kinetic parameters λ_1 , λ_2 , and λ_3 , respectively. The extended MRN has species $S_1, S_2, S_3, S_4, S_5, S_{\lambda_1}, S_{\lambda_2}$, and S_{λ_3} , and reactions



According to (1), the ODEs are given by:

$$\begin{aligned} \frac{dX_{S_1}(t)}{dt} &= -(X_{S_{\lambda_1}}(t) + X_{S_{\lambda_2}}(t))X_{S_1}(t) \\ \frac{dX_{S_2}(t)}{dt} &= -(X_{S_{\lambda_1}}(t) + X_{S_{\lambda_2}}(t))X_{S_2}(t), \\ \frac{dX_{S_3}(t)}{dt} &= +2X_{S_{\lambda_1}}(t)X_{S_1}(t) + X_{S_{\lambda_2}}(t)X_{S_2}(t) + X_{S_{\lambda_3}}(t)X_{S_4}(t)X_{S_5}(t) \\ \frac{dX_{S_4}(t)}{dt} &= -X_{S_{\lambda_3}}(t)X_{S_4}(t)X_{S_5}(t) \\ \frac{dX_{S_5}(t)}{dt} &= -X_{S_{\lambda_3}}(t)X_{S_5}(t)(X_{S_3}(t) + X_{S_4}(t)) + X_{S_{\lambda_2}}(t)X_{S_1}(t) + 2X_{S_{\lambda_1}}(t)X_{S_2}(t) \\ \frac{dX_{S_{\lambda_1}}(t)}{dt} &= 0 \\ \frac{dX_{S_{\lambda_2}}(t)}{dt} &= 0 \\ \frac{dX_{S_{\lambda_3}}(t)}{dt} &= 0. \end{aligned}$$

Let us consider again the FE for the original CRN in Equation (8), extended with singleton blocks for the promoted species: $\{\{S_1, S_2\}, \{S_3, S_4\}, \{S_5\}, \{\lambda_1\}, \{\lambda_2\}, \{\lambda_3\}\}$. This is not an FE on the extended CRN because it is not possible anymore to express the ODE of S_5 in terms of $X_{S_1} + X_{S_2}$. In order to obtain an FE for the extended network we need to split the first block, obtaining the FE $\{\{S_1\}, \{S_2\}, \{S_3, S_4\}, \{S_5\}, \{\lambda_1\}, \{\lambda_2\}, \{\lambda_3\}\}$.

This extension can be used also for SE reductions. However, given that the CTMC semantics associates integer populations to species, it is necessary that each kinetic constant is a rational number (See Section 2 of [18] for a deeper discussion on this aspect. This is the case for the model in Section 5.2 used to study the robustness of both FE and SE.

In the remainder of this section we apply this method to selected case studies to state the robustness of our reduction techniques to specific parameters values. Furthermore, we also discuss in detail the obtained equivalences among species, giving insights on their physical meaning.

5.1. BE case study: STAT1 homodimers binding to DNA

Model *BIOMOD000000501* is a PRN describing a mechanism of binding of STAT1 homodimers to DNA, taken from [5]. The model considers three main kinds of reactions: reversible binding of STAT1 homodimers to non-GAS sites; reversible binding of STAT1 homodimers to GAS sites; and reversible polymerization of adjacent bound STAT1 homodimers. Furthermore, for STAT1 to unbind from DNA, polymers need to depolymerize first. The model considers DNA with 4 sites, of which the middle ones are GAS sites and the external ones are non-GAS sites; the kinetic parameters of the reaction differ according to the kind of site where the reaction occurs. The state of system can thus be represented through formal species indicating whether sites are bound or unbound and if a polymer is formed at adjacent bound site. Each species can thus be represented as DNA_{xxxx} where x can be either 0 (unbound) or 1 (bound); the presence of a polymer is denoted by a dash between two site identifiers, e.g. DNA_{001-1} indicates two bound STAT1 homodimers (one at a GAS site and one at a non-GAS site) that have polymerized.

The two GAS sites are intentionally modeled to be symmetric [5, Supplementary Note 1]; as a consequence, the authors state that the proportion of DNA with STAT1 occupying the GAS site can be measured at either site. BE mathematically confirms this claim (which however holds only if the initial concentrations of symmetric DNA configurations are symmetric, as happens in the original model). Moreover, it proves that also the non-GAS sites are symmetric.




Figure 6: (a) Adaptation of the SPOC dynamical model from [20]. The SPB compartment is depicted in the yellow-circle background. Reactions crossing the compartment boundary represent the intrinsic Tem1 (blue rectangle) GTPase-cycle and reversible SPB association in terminal T . (b) Reduced mechanism where both FE and SE equate all Tem1 molecules up to their GTP (green)- or GDP (red)-bound state (indicated by the green/red ellipsis).

That is, BE places in the same block species that have equal DNA configuration up to the position in symmetric sites; for example, it yields the partition blocks $\{\text{DNA}_{1000}, \text{DNA}_{0001}\}$ and $\{\text{DNA}_{1011}, \text{DNA}_{1101}\}$. A similar aggregation occurs also independently of the polymerization state, e.g., BE yields the block $\{\text{DNA}_{01-11}, \text{DNA}_{011-1}, \text{DNA}_{1-110}, \text{DNA}_{11-10}\}$. Overall, BE reduces the original model with 35 species to one with 19 species, independently of the values of the 6 kinetic parameters used in the model.

5.2. FE and SE case study: Spindle Position Checkpoint model

Model *MODEL1202090002* is an MRN of the Spindle Position Checkpoint (SPOC) [20]. SPOC intervenes in the process of cell division by verifying all requirements to pass to the next phase in the cell cycle. In particular, it prevents the separation of the duplicated chromosomes until each chromosome is properly attached to the spindle apparatus. The most upstream event of the pathway involves GTPase Tem1. Tem1 binds to the yeast centrosomes (called spindle pole bodies, SPBs) via GAP-dependent and GAP-independent sites (Figure 6a). The intrinsic GTPase switching cycle of Tem1 is modeled as a reversible first-order reaction that converts Tem1^{GTP} into Tem1^{GDP} and vice versa. The model consists of 24 species and 71 irreversible reactions. Both FE and SE equate the two forms of the GTPase Tem1 (Figure 6), moreover this equivalence extends to all Tem1 molecular complexes, yielding a reduced model with 16 species and 36 reactions, independently of the values of all 34 kinetic parameters in the model.

5.3. FDE case study: Signaling cascade

Model *BIOMOD0000000033* is an NRN modeling a signaling pathway concluding in ERK activation [8]. Its most upstream event (Figure 7) starts with the binding of EGF and NGF to their respective receptors (EGFR, NGFR). Once bound, both receptors can activate molecular SOS and trigger the Ras cascade. Here, molecular components are modeled representing the species active and inactive state, i.e. mSOS^* and mSOS , yielding a model with 32 species and 26 reactions. Here FDE collapses the active and inactive form of those species. Moreover, the dynamics of the active and inactive species sum up to zero if aggregated. Furthermore, also the phosphatases P1, P2, and P3 are aggregated in the same class. Instead, the free and bound EGF, NGF and EGFR receptors are left in singleton blocks. This results in the FDE reduction in Figure 7b, with 7 species and 4 reactions, independently of the values of all 48 kinetic parameters in the model.

5.4. BDE case study: model of cell under limited resources

Model *MODEL1502270000*, presented in [56], is a mechanistic model of the behaviour of a cell. The model has been specifically designed to keep into account an aspect often overlooked when studying cells' models: cells have finite resources (e.g. finite cellular energy, mass, or levels of ribosomes), and therefore using resources for a task (e.g. growth or protein expression), reduces resources available for other tasks. The model is recognized as an NRN because it features Michaelis-Menten kinetics for enzymatically catalyzed reactions; as such it is only analyzable by means of FDE and BDE. The model consists of 22 species and 61 reactions. In particular, it contains the five proteins e_t , e_m , r , q , and p , each represented using a distinct formal species. Species e_t is a transporter enzyme for

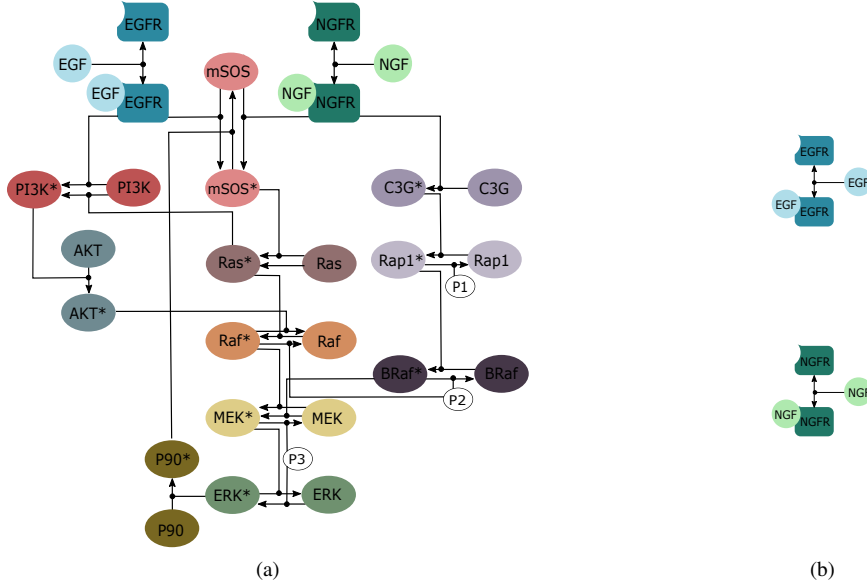


Figure 7: (a) Adaptation of the signaling network in [8]. The activation of the molecular SOS by either of the receptors triggers the Ras cascade, concluding in ERK activation. EGF can also use the left branch involving PI3K to modulate Erk activity through Raf1 downregulation, and NGF can upregulate Mek using the right branch containing Rap1. P1,P2 and P3 represent unregulated phosphatases. (b) FDE reduction.

nutrients (denoted by species s_i) into the cell, while e_m is an enzyme which metabolizes the internalized nutrients s_i , producing energy (denoted by species a). The remaining proteins denote ribosomes (r), house-keeping proteins (q) and gratuitous proteins that do not contribute to the cell growth (p). For each of such proteins, the model contains three mRNAs species representing their free (m_x), ribosome-bound (c_x) and bound with ribosome and chloramphenicol (zm_x) mRNAs, with $x \in \{t, m, r, q, p\}$. The amount of chloramphenicol is constant and therefore it is modeled as a parameter rather than as a species. Each protein and related mRNA species undergo a number of reactions which involve the processes of *transcription*, *degradation*, *ribosome binding/unbinding*, *translation*, *chloramphenicol binding*, or *dilution*. The reactions involving the enzymes e_t and e_m and corresponding mRNAs are given in Table 4.

On this model, BDE is able to capture a type of *chains of aggregations* typical of bisimulation-based reduction techniques like ours. Indeed, BDE is able to relate e_t with e_m and zm_t with zm_e . This is due to the fact that BDE relates m_t with m_m , which in turn allows to relate c_t with c_m , which finally allows to relate zm_t with zm_m , and e_t with e_m . Actually, zm_t and zm_m belong to a larger BDE block with five species, including all chloramphenicol-bound mRNAs. Overall, we obtain a BDE-reduced model with 15 species and 49 reactions.

This model features 25 parameters of which one, n_q (originally set to 4), appears as exponent in the ODE of m_q ; it is related to a translation reaction of protein q . Promoting such parameter to a species would move the ODEs outside the decidable theory of reals [11], preventing us from applying BDE to it. For this reason, we promoted only all other 24 parameters, obtaining again the equivalence of e_t with e_m coming from the equivalences among their corresponding mRNAs species. In the extended CRN the block containing all five chloramphenicol-bound mRNAs (zm_{\cdot}) is split into a block containing zm_t and zm_m , and three more singleton blocks. This is due to the term $(cl \cdot k_c m) \cdot rm_x$, with cl and $k_c m$ two parameters, appearing in the ODE of species zm_x , with $x \in \{t, m, r, p, q\}$. This monomial should distinguish zm_r , zm_p , and zm_q and relate zm_t with zm_m , as BDE relates rm_t with rm_m , and puts the other ribosome-bound species in singleton blocks. In the original model this does not happen because parameter cl has value zero.

6. Conclusion

The empirical assessment of exact model reduction on the BioModels repository has provided a number of findings along the main evaluation questions **Q1–Q3** introduced in Section 1, which can be summarized as follows.

<i>Transcription</i>	<i>Degradation</i>		
$\emptyset \xrightarrow{f_1(a)} m_x$	$m_x \xrightarrow{f_2(m_x)} \emptyset$		
	<i>Rib. (un)binding</i>	<i>Translation</i>	<i>Dilution</i>
	$r + m_x \xrightarrow{f_3(r,m_x)} c_x$	$n \cdot a + c_x \xrightarrow{f_6(a,c_x)} r + m_x + e_x$	$e_x \xrightarrow{f_5(e_x)} \emptyset$
	$c_x \xrightarrow{f_4(c_x)} r + m_x$	<i>Chlor. binding</i>	<i>Dilution</i>
		$c_x \xrightarrow{f_7(c_x)} zm_x$	$zm_x \xrightarrow{f_5(zm_x)} \emptyset$
		<i>Dilution</i>	
		$c_x \xrightarrow{f_5(c_x)} \emptyset$	

Table 4: Processes involving the enzymatic proteins and related mRNAs in the cell model from [56], for $x \in \{t, m\}$. The propensity functions and the stoichiometries n are specified in [56]. Here we provide a different name f_i to each different propensity function, and we explicitly provide the species appearing in them. The free mRNA m_x gets transcribed as long as the cell has energy. Such mRNA can then either degrade, or reversibly bind with ribosome. In the latter case, a ribosome-bound mRNA can be involved into three alternative tasks: (i) as long as the cell has energy, it can undergo translation unbinding from the ribosome to synthesize an enzymatic protein which can later dilute; (ii) it can bind with chloramphenicol and later dilute; iii) it can dilute.

Q1 Assumptions for applicability of model reductions. In the preprocessing phase (Section 3.3), we found 494 models not supported by ERODE. Among the reasons for incompatibility it is worth commenting on the models which included exponential expressions in rate functions. This is not accepted by FDE/BDE because the underlying theory is not decidable. A workaround has been sketched in [13, 17] and builds on a systematic technique which transforms an initial value problem for an ODE system with derivatives containing rational and exponential expressions into an equivalent problem with polynomial derivatives [30], to which BE and FE can be applied. In future work we plan to implement such a transformation in order to extend the range of applicability of species equivalences. In Section 3.3 we discussed that not all models in BioModels are described as ODE systems. Some are logical models described as Boolean networks [35]. ERODE has been recently extended with support and reduction techniques for such formalism [2]. In future work we plan to extend the large-scale analysis of this paper also for such formalism.

Q2 Effectiveness of the reductions. Overall, we found exact model reductions effective in terms of the number of cases in which a CRN could be reduced by at least one technique (64% of the imported ones). With regards to the compression ratios obtained, SE turned out to be as effective as FE in the considered models, despite they consider different interpretations for CRNs and the existence of results showing that SE is stricter than FE [18]. The reduction techniques FE, BE and SE terminated (possibly without reduction) on all considered models. Instead, BDE and FDE failed on a low percentage (19%) of the NRN models, including instances where such failures were due to timeouts. For this reason, it might be worth investigating alternative approaches to computing them, especially FDE, for example by parallelizing the computations.

Q3 Robustness to parameters values. In the selected case studies herein presented, the exact model reductions have revealed that symmetries in certain signaling pathways carry over to equivalences at the level of the underlying quantitative semantics. Given their moderate size, the considered models would be computationally tractable even without reduction. However, the equivalences can be used as an aid in developing more complex models where such symmetries are present in some components. In addition, we remark that exact model reduction can still be useful when the complexity is due to the many repetitions that are required (e.g., for sensitivity analysis or for simulation with tight confidence intervals) or for particularly difficult analyses such as parametric inference [49]. In these cases, and whenever the knowledge of the model parameters are subject to uncertainty, it is particularly relevant to what extent the reductions are robust with respect to the actual values of such parameters. The robustness analysis conducted in this paper on the selected case studies indicates that the patterns of aggregations uncovered by the reductions are mostly insensitive to the parameter values. A systematic analysis to the whole repository was not possible due to the lack of automated assistance to implement the model translation discussed in Section 5, which is left for future work.

For the future, it might be interesting to consider extensions of our importer or of the ERODE input language so to be able to import a higher percentage of models from the BioModels repository. The availability of ready-to-use

model conversions in a simple CRN format such as ERODE’s might stimulate similar assessments with other model reduction techniques (e.g., [7, 16]), or with techniques for bifurcation analysis based on temporal logic as in [6].

Acknowledgement. The authors are grateful to Andreas Dräger (Institut für Informatik Zentrum für Bioinformatik Tübingen) for his support with JSBML. Partially supported by the Independent Research Fund Denmark DFF Research Project 9040-00224B REDUCTO, and the Italian PRIN project “SEDUCE” no. 2017TWR CNB.

References

- [1] A. Antoulas. *Approximation of Large-Scale Dynamical Systems*. Advances in Design and Control. SIAM, 2005.
- [2] G. Argyris, A. L. Lafuente, M. Tribastone, M. Tschaikowski, and A. Vandin. Reducing Boolean Networks with Backward Boolean Equivalence. In *Computational Methods in Systems Biology - 19th International Conference, CMSB 2021, Bordeaux, France, September 22-24, 2021, Proceedings*, 2021.
- [3] U. M. Ascher and L. R. Petzold. *Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations*. SIAM, 1998.
- [4] M. Ashburner, C. A. Ball, J. A. Blake, D. Botstein, H. Butler, J. M. Cherry, A. P. Davis, K. Dolinski, S. S. Dwight, J. T. Eppig, et al. Gene ontology: tool for the unification of biology. *Nature genetics*, 25(1):25, 2000.
- [5] A. Begitt, M. Droscher, T. Meyer, C. D. Schmid, M. Baker, F. Antunes, K.-P. Knobeloch, M. R. Owen, R. Naumann, T. Decker, and U. Vinkemeier. Stat1-cooperative dna binding distinguishes type 1 from type 2 interferon signaling. *Nature Immunology*, 15(2):168–176, 2014. doi: 10.1038/ni.2794. URL <https://doi.org/10.1038/ni.2794>.
- [6] N. Benes, L. Brim, M. Demko, S. Pastva, and D. Safránek. A model checking approach to discrete bifurcation analysis. In J. S. Fitzgerald, C. L. Heitmeyer, S. Gnesi, and A. Philippou, editors, *FM 2016: Formal Methods - 21st International Symposium, Limassol, Cyprus, November 9-11, 2016, Proceedings*, volume 9995 of *Lecture Notes in Computer Science*, pages 85–101, 2016. doi: 10.1007/978-3-319-48989-6_6. URL https://doi.org/10.1007/978-3-319-48989-6_6.
- [7] M. Boreale. Algebra, coalgebra, and minimization in polynomial differential equations. In *20th International Conference on Foundations of Software Science and Computation Structures (FOSSACS)*, pages 71–87, 2017.
- [8] K. S. Brown, C. C. Hill, G. A. Calero, C. R. Myers, K. H. Lee, J. P. Sethna, and R. A. Cerione. The statistical mechanics of complex signaling networks: nerve growth factor signaling. *Physical biology*, 1(3):184, 2004.
- [9] P. Buchholz. Exact and Ordinary Lumpability in Finite Markov Chains. *Journal of Applied Probability*, 31(1):59–75, 1994. ISSN 00219002.
- [10] L. Cardelli, M. Tribastone, M. Tschaikowski, and A. Vandin. Forward and backward bisimulations for chemical reaction networks. In *26th International Conference on Concurrency Theory, CONCUR*, pages 226–239, 2015. doi: 10.4230/LIPIcs.CONCUR.2015.226.
- [11] L. Cardelli, M. Tribastone, M. Tschaikowski, and A. Vandin. Symbolic computation of differential equivalences. In *POPL*, pages 137–150, 2016. doi: 10.1145/2837614.2837649.
- [12] L. Cardelli, M. Tribastone, M. Tschaikowski, and A. Vandin. Efficient syntax-driven lumping of differential equations. In *TACAS*, pages 93–111, 2016.
- [13] L. Cardelli, M. Tribastone, M. Tschaikowski, and A. Vandin. Maximal aggregation of polynomial dynamical systems. *Proceedings of the National Academy of Sciences*, 114(38):10029–10034, 2017.
- [14] L. Cardelli, M. Tribastone, M. Tschaikowski, and A. Vandin. Erode: A tool for the evaluation and reduction of ordinary differential equations. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 310–328. Springer, 2017.
- [15] L. Cardelli, M. Tribastone, M. Tschaikowski, and A. Vandin. Syntactic Markovian bisimulation for chemical reaction networks. In *Models, Algorithms, Logics and Tools*, pages 466–483. Springer, 2017.
- [16] L. Cardelli, M. Tribastone, M. Tschaikowski, and A. Vandin. Guaranteed error bounds on approximate model abstractions through reachability analysis. In *15th International Conference on Quantitative Evaluation of Systems (QEST)*, 2018.
- [17] L. Cardelli, M. Tribastone, M. Tschaikowski, and A. Vandin. Symbolic computation of differential equivalences. *Theoretical Computer Science*, 2019.
- [18] L. Cardelli, I. C. Perez-Verona, M. Tribastone, M. Tschaikowski, A. Vandin, and T. Waizmann. Exact maximal reduction of stochastic reaction networks by species lumping. *Bioinformatics*, 2021. ISSN 1367-4803. doi: 10.1093/bioinformatics/btab081. URL <https://doi.org/10.1093/bioinformatics/btab081>. eprint: <https://academic.oup.com/bioinformatics/advance-article-pdf/doi/10.1093/bioinformatics/btab081/36287939/btab081.pdf>.
- [19] R. Caspi, T. Altman, R. Billington, K. Dreher, H. Foerster, C. A. Fulcher, T. A. Holland, I. M. Keseler, A. Kothari, A. Kubo, et al. The metacyc database of metabolic pathways and enzymes and the biocyc collection of pathway/genome databases. *Nucleic acids research*, 42(D1):D459–D471, 2013.
- [20] A. K. Caydasi, M. Lohel, G. Grünert, P. Dittrich, G. Pereira, and B. Ibrahim. A dynamical model of the spindle position checkpoint. *Molecular systems biology*, 8(1):582, 2012.
- [21] V. Chelliah, C. Laibe, and N. L. Novère. Biomodels database: a repository of mathematical models of biological processes. *Encyclopedia of Systems Biology*, pages 134–138, 2013.
- [22] U. Consortium. Uniprot: a hub for protein information. *Nucleic acids research*, 43(D1):D204–D212, 2014.
- [23] M. Courtot, N. Juty, C. Knüpfner, D. Waltemath, A. Zhukova, A. Dräger, M. Dumontier, A. Finney, M. Golebiewski, J. Hastings, et al. Controlled vocabularies and semantics in systems biology. *Molecular systems biology*, 7(1):543, 2011.
- [24] L. De Moura and N. Bjørner. Z3: An efficient SMT solver. In *TACAS*, pages 337–340, 2008.
- [25] K. Degtyarenko, P. De Matos, M. Ennis, J. Hastings, M. Zbinden, A. McNaught, R. Alcántara, M. Darsow, M. Guedj, and M. Ashburner. Chebi: a database and ontology for chemical entities of biological interest. *Nucleic acids research*, 36(suppl.1):D344–D350, 2007.
- [26] A. Dräger, N. Rodriguez, M. Dumousseau, A. Dörr, C. Wrzodek, N. Le Novère, A. Zell, and M. Hucka. JSBML: a flexible Java library for working with SBML. *Bioinformatics*, 27(15):2167–2168, 06 2011. ISSN 1367-4803. doi: 10.1093/bioinformatics/btr361.

- [27] S. Federhen. The ncbi taxonomy database. *Nucleic acids research*, 40(D1):D136–D143, 2011.
- [28] D. Gillespie. Exact stochastic simulation of coupled chemical reactions. *Journal of Physical Chemistry*, 81(25):2340–2361, December 1977.
- [29] D. T. Gillespie. Stochastic simulation of chemical kinetics. *Annual Review of Physical Chemistry*, 58(1):35–55, 2007.
- [30] C. Gu. QLMOR: A Projection-Based Nonlinear Model Order Reduction Approach Using Quadratic-Linear Representation of Nonlinear Systems. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 30(9):1307–1320, 2011. doi: 10.1109/TCAD.2011.2142184.
- [31] M. Hucka, A. Finney, H. M. Sauro, H. Bolouri, J. C. Doyle, H. Kitano, A. P. Arkin, B. J. Bornstein, D. Bray, A. Cornish-Bowden, et al. The systems biology markup language (sbml): a medium for representation and exchange of biochemical network models. *Bioinformatics*, 19(4): 524–531, 2003.
- [32] M. Hucka, F. T. Bergmann, A. Dräger, S. Hoops, S. M. Keating, N. Le Novère, C. J. Myers, B. G. Olivier, S. Sahle, J. C. Schaff, L. P. Smith, D. Waltemath, and D. J. Wilkinson. Systems biology markup language (sbml) level 2 version 5: Structures and facilities for model definitions. *Journal of integrative bioinformatics*, 12(2):271–271, 09 2015. doi: 10.2390/biecoll-jib-2015-271. URL <https://pubmed.ncbi.nlm.nih.gov/26528569>.
- [33] N. Juty, N. Le Novère, and C. Laibe. Identifiers. org and miriam registry: community resources to provide persistent identification. *Nucleic acids research*, 40(D1):D580–D586, 2011.
- [34] M. Kanehisa and S. Goto. Kegg: kyoto encyclopedia of genes and genomes. *Nucleic acids research*, 28(1):27–30, 2000.
- [35] S. Kauffman. Metabolic stability and epigenesis in randomly constructed genetic nets. *Journal of Theoretical Biology*, 22(3):437–467, 1969. ISSN 0022-5193.
- [36] T. G. Kurtz. The Relationship between Stochastic and Deterministic Models for Chemical Reactions. *Journal of Chemical Physics*, 57(7), 1972.
- [37] N. Le Novère, A. Finney, M. Hucka, U. S. Bhalla, F. Campagne, J. Collado-Vides, E. J. Crampin, M. Halstead, E. Klipp, P. Mendes, et al. Minimum information requested in the annotation of biochemical models (miriam). *Nature biotechnology*, 23(12):1509, 2005.
- [38] N. Le Novère, N. Rodriguez, F. Wrzodek, F. Mittag, S. Fröhlich, M. Hucka, A. Thomas, B. Ø. Palsson, N. E. Lewis, A. Dräger, C. J. Myers, L. Watanabe, I. Y. Vazirabad, V. Kofia, H. F. Gómez, A. Diamantikos, E. Netz, J. Matthes, J. Eichner, R. Keller, J. Rudolph, and C. Wrzodek. JSBML 1.0: providing a smorgasbord of options to encode systems biology models. *Bioinformatics*, 31(20):3383–3386, 06 2015. ISSN 1367-4803. doi: 10.1093/bioinformatics/btv341.
- [39] C. Li, M. Donizelli, N. Rodriguez, H. Dharuri, L. Endler, V. Chelliah, L. Li, E. He, A. Henry, M. I. Stefan, J. L. Snoep, M. Hucka, N. Le Novère, and C. Laibe. BioModels Database: An enhanced, curated and annotated resource for published quantitative kinetic models. *BMC Systems Biology*, 4:92, Jun 2010.
- [40] G. Li and H. Rabitz. A general analysis of exact lumping in chemical kinetics. *Chemical Engineering Science*, 44(6):1413–1430, 1989. ISSN 0009-2509. doi: [http://dx.doi.org/10.1016/0009-2509\(89\)85014-6](http://dx.doi.org/10.1016/0009-2509(89)85014-6).
- [41] H. Miao, M. Y. Sangster, A. M. Livingstone, S. P. Hilchey, L. Zhang, D. J. Topham, T. R. Mosmann, J. Holden-Wiltse, A. S. Perelson, H. Wu, and M. S. Zand. Modeling the dynamics and migratory pathways of virus-specific antibody-secreting cell populations in primary influenza infection. *PLoS one*, 9(8):e104781–e104781, 08 2014. doi: 10.1371/journal.pone.0104781. URL <https://pubmed.ncbi.nlm.nih.gov/25171166>.
- [42] H. Miao, M. Y. Sangster, A. M. Livingstone, S. P. Hilchey, L. Zhang, D. J. Topham, T. R. Mosmann, J. Holden-Wiltse, A. S. Perelson, H. Wu, et al. Modeling the dynamics and migratory pathways of virus-specific antibody-secreting cell populations in primary influenza infection. *PLOS One*, 9(8):e104781, 2014.
- [43] D. Nishimura. Biocarta. *Biotech Software & Internet Report: The Computer Software Journal for Scientist*, 2(3):117–120, 2001.
- [44] M. S. Okino and M. L. Mavrouniotis. Simplification of mathematical models of chemical reaction systems. *Chemical Reviews*, 2(98): 391–408, 1998.
- [45] A. Ovchinnikov, I. Pérez Verona, G. Pogudin, and M. Tribastone. CLUE: exact maximal reduction of kinetic models by constrained lumping of differential equations. *Bioinformatics*, 02 2021. ISSN 1367-4803. doi: 10.1093/bioinformatics/btab010. URL <https://doi.org/10.1093/bioinformatics/btab010>. btab010.
- [46] R. Paige and R. Tarjan. Three partition refinement algorithms. *SIAM Journal on Computing*, 16(6):973–989, 1987. doi: 10.1137/0216062.
- [47] I. C. Pérez-Verona, M. Tribastone, and A. Vandin. A large-scale assessment of exact model reduction in the biomodels repository. In *Computational Methods in Systems Biology - 17th International Conference, CMSB 2019, Trieste, Italy, September 18-20, 2019, Proceedings*, pages 248–265, 2019. doi: 10.1007/978-3-030-31304-3_13. URL https://doi.org/10.1007/978-3-030-31304-3_13.
- [48] D. Pratt, J. Chen, D. Welker, R. Rivas, R. Pillich, V. Rynkov, K. Ono, C. Miello, L. Hicks, S. Szalma, et al. Ndx, the network data exchange. *Cell systems*, 1(4):302–305, 2015.
- [49] D. Schnoerr, G. Sanguinetti, and R. Grima. Approximation and inference methods for stochastic biochemical kinetics—a tutorial review. *Journal of Physics A: Mathematical and Theoretical*, 50(9):093001, 2017.
- [50] T. J. Snowden, P. H. van der Graaf, and M. J. Tindall. Methods of model reduction for large-scale biological systems: A survey of current methods and trends. *Bulletin of Mathematical Biology*, 79(7):1449–1486, 2017.
- [51] G. Stoesser, W. Baker, A. van den Broek, E. Camon, M. Garcia-Pastor, C. Kanz, T. Kulikova, R. Leinonen, Q. Lin, V. Lombard, et al. The embl nucleotide sequence database. *Nucleic acids research*, 30(1):21–26, 2002.
- [52] M. Tribastone and A. Vandin. Speeding up stochastic and deterministic simulation by aggregation: an advanced tutorial. In *2018 Winter Simulation Conference, WSC 2018, Gothenburg, Sweden, December 9-12, 2018*, pages 336–350, 2018. doi: 10.1109/WSC.2018.8632364.
- [53] T. Turanyi and A. S. Tomlin. *Analysis of Kinetic Reaction Mechanisms*. Springer, 2014.
- [54] A. Vandin and M. Tribastone. Quantitative abstractions for collective adaptive systems. In *SFM 2016, Bertinoro Summer School*, pages 202–232, 2016. doi: 10.1007/978-3-319-34096-8_7.
- [55] E. O. Voit. Biochemical systems theory: A review. *ISRN Biomathematics*, 2013:53, 2013. URL <http://dx.doi.org/10.1155/2013/897658>.
- [56] A. Y. Weiße, D. A. Oyarzún, V. Danos, and P. S. Swain. Mechanistic links between cellular trade-offs, gene expression, and growth. *Proceedings of the National Academy of Sciences*, 112(9):E1038–E1047, 2015. ISSN 0027-8424. doi: 10.1073/pnas.1416533112. URL

<https://www.pnas.org/content/112/9/E1038>.

- [57] U. Wittig, R. Kania, M. Golebiewski, M. Rey, L. Shi, L. Jong, E. Algae, A. Weidemann, H. Sauer-Danzwith, S. Mir, et al. Sabio-rk—database for biochemical reaction kinetics. *Nucleic acids research*, 40(D1):D790–D796, 2011.