

Forward and Backward Bisimulations for Chemical Reaction Networks

Luca Cardelli¹, Mirco Tribastone², Max Tschaikowski³, and Andrea Vandin⁴

¹ Microsoft Research & University of Oxford, UK

luca@microsoft.com

²⁻⁴ IMT Institute for Advanced Studies Lucca, Italy

{mirco.tribastone,max.tschaikowski,andrea.vandin}@imtlucca.it

Abstract

We present two quantitative behavioral equivalences over species of a chemical reaction network (CRN) with semantics based on ordinary differential equations. *Forward CRN bisimulation* identifies a partition where each equivalence class represents the exact sum of the concentrations of the species belonging to that class. *Backward CRN bisimulation* relates species that have identical solutions at all time points when starting from the same initial conditions. Both notions can be checked using only CRN syntactical information, i.e., by inspection of the set of reactions. We provide a unified algorithm that computes the coarsest refinement up to our bisimulations in polynomial time. Further, we give algorithms to compute quotient CRNs induced by a bisimulation. As an application, we find significant reductions in a number of models of biological processes from the literature. In two cases we allow the analysis of benchmark models which would be otherwise intractable due to their memory requirements.

1998 ACM Subject Classification D.2.4 Software/Program Verification, G.1.7 Ordinary Differential Equations, J.2 Physical Sciences and Engineering

Keywords and phrases Chemical reaction networks, ordinary differential equations, bisimulation, partition refinement

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2015.226

1 Introduction

At the interface between computer science and systems biology is the idea that biological systems can be interpreted as computational processes [23, 12], leading to a number of formal methods applied to study biomolecular systems [5, 18, 25]. In this context, chemical reaction networks (CRNs), a popular mathematical model of interaction in natural sciences, can also be seen as a kernel concurrent language for natural programming.

In this paper we present, for the first time to our knowledge, quantitative bisimulation equivalences for CRNs with the well-known interpretation based on ordinary differential equations (ODEs). (To make the paper self-contained, all background is given in Section 2.) In this semantics, each species is associated with an ODE giving the deterministic evolution of its concentration starting from an initial condition. Our bisimulations are equivalences over species that induce a reduced CRN that exactly preserves the dynamics of the original one. This is an important goal, especially in order to cope with the potentially very large number of species and reactions in many biological networks [16, 17].

We study two equivalences, developed in the Larsen-Skou style of probabilistic bisimulation [29], that are based on two distinct ideas of observable behavior. *Forward CRN*



© Luca Cardelli, Mirco Tribastone, Max Tschaikowski, and Andrea Vandin;
licensed under Creative Commons License CC-BY

26th International Conference on Concurrency Theory (CONCUR 2015).

Editors: Luca Aceto and David de Frutos Escrig; pp. 226–239



Leibniz International Proceedings in Informatics

LIPIcs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

bisimulation yields an aggregated ODE where the solution gives the exact sum of the concentrations of the species belonging to each equivalence class. In *backward CRN bisimulation*, instead, equivalent species have the same solution *at all time points*; in other words, backward CRN bisimulation relates species whose ODE solutions are equal whenever they start from identical initial conditions. The use of “forward” and “backward” has a long tradition in models of computation based on labelled transition systems [19]. In the case of quantitative variants, for instance those defined for process algebra with a continuous-time Markov chain (CTMC) semantics [26, 27, 8, 4], forward bisimulations are equivalences that induce a CTMC aggregation in the sense of ordinary lumpability [7], where the probability of an equivalence class is equal to the sum of the probabilities of the states belonging to that class. This is found by checking conditions on the *outgoing transitions* of related states in the transition diagram. A backward bisimulation induces a CTMC aggregation in the sense of weak lumpability [21], where all states in the same equivalence class have a time-invariant conditional probability distribution; exact lumpability is a special case where the conditional probability distribution is uniform, in the sense that any two states of each equivalence class have the same probability at any point in time whenever they have the same initial probabilities. It is found by relating states according to conditions on their *predecessor states* [21, 34, 7].

Despite being similar in spirit, technically our bisimulations are fundamentally different for two reasons. First, they concern a continuous-state semantics based on ODEs instead of a discrete-state CTMC. Second, they operate at the structural, syntactical level, because they are defined with quantities that can be computed by only inspecting the reactions of a CRN. Still, the repercussions of our bisimulations on the semantics are explained according to certain theories of aggregation. In particular, forward CRN bisimulation yields an aggregated system in the sense of ODE lumpability [36, 30]. This theory covers linear transformations of the original state variables in general; here we consider an instance, which we call *ordinary fluid lumpability*, where the transformation is induced by a partition of state variables. (Forward bisimulation is presented in Section 3.1.) Backward bisimulation (presented in Section 3.2) is related to *exact fluid lumpability*, introduced in the context of process algebra with fluid semantics [37], identifying process terms with the same ODE solution when initialized equally. The disadvantage of forward CRN bisimulation is that it is lossy (yet exact) because, similarly to the forward stochastic analogues, from the aggregated ODE system in general it is not possible to recover the solutions for the individual species within the same equivalence class. On the other hand, it does not impose restrictions on the initial conditions, which instead are present in our backward variant. As a further important difference, forward CRN bisimulation (again, like its stochastic analogues) turns out to be a sufficient condition for ODE lumpability. Instead, backward CRN bisimulation enjoys a full characterization, in the sense that there exists a backward CRN bisimulation between two species if and only if they have the same ODE solutions (provided that they start from equal initial conditions). More in general, by means of a number of examples we will show that the two equivalences are complementary because not comparable. In other words, there exist models that can be reduced up to forward CRN bisimulation but not by the backward variant, and vice versa; at the same time, there are models that can be reduced by both.

To enhance the usefulness of these notions, we present (in Section 5) a *template* partition-refinement algorithm that is parametric with respect to the equivalence of interest, computing the coarsest refinement up to either variant in polynomial time. To use our equivalences as an automatic model reduction tool, we further give two algorithms (in Section 4) that provide the quotient CRN induced by either bisimulation. With a prototype implementation available at <http://sysma.intlucca.it/crnreducer/>, we show (in Section 6) that we are able to

reduce a number of case studies taken from the literature. Our bisimulations yielded quotient CRNs with number of reactions and species up to four orders of magnitude smaller than the original CRNs, leading to speed-ups in the ODE solution runtimes of up to five orders of magnitude. In two cases, it was possible to analyze models that were otherwise intractable directly within our experimental environment due to excessive memory requirements.

Related work. Behavioral equivalences have been recently proposed in [32] for comparing CRNs; however, the analysis is carried out at the qualitative level, i.e., ignoring the dynamical evolution. In [37] is introduced the notion of *label equivalence* for process algebra with fluid semantics, which captures exact fluid lumpability (processes are equivalent whenever their ODE solutions are equal at all time points). However, unlike backward CRN bisimulation, label equivalence is only a sufficient condition for ODE reduction. Indeed, it works at a coarser level of granularity as it relates *sets* of ODE variables, each corresponding to the behavior of a sequential process. Instead, backward CRN bisimulation relates individual ODE variables. Further, the conditions for equivalence, specific to the process algebra, are difficult to check automatically because of the universal quantifiers over the ODE variables. More important, no algorithm for computing the coarsest partition was developed. Similar considerations apply to the process-algebra specific ordinary fluid lumpability in [38].

Cardelli’s notion of *emulation* between two CRNs is a (structural) mapping of species and reactions that, like backward CRN bisimulation, guarantees the equality between the ODE solutions at all time points [11]. An emulation requires a source and a target CRN — the modeler 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 equivalences between species *within the same given CRN*. Thus, emulation is not useful for model reduction because a-priori information about the structure of a quotient CRN is not available. Furthermore, no algorithm is provided in [11] to find emulations automatically. Since backward CRN bisimulation fully characterizes exact fluid lumpability, it is not difficult to show that backward CRN bisimulation generalizes emulation in the sense that any emulation between two CRNs can be understood in terms of a backward CRN bisimulation over the species of a “union CRN” that contains all the reactions of the two CRNs of interest.

Model reductions have been studied in related models for biomolecular networks (e.g. [17, 22, 10]), most notably for rule-based systems such as BioNetGen [5] and the κ calculus [18]. These offer an intensional modeling approach, by providing graph-rewrite rules of interaction instead of a complete enumeration of all chemical reactions involved. *Differential fragments* for κ are self-consistent aggregates found by a static analysis on the model, identifying sums of chemical species for which an ODE system can be explicitly written [17]. In this sense, they are analogous to our CRN bisimulations, but with notable differences. First, fragmentation works directly at the rule-based level. This has the advantage that the analysis is performed on a set of rewrite rules, which is typically much more compact than the fully enumerated CRN. However, fragmentation is domain-specific, hence the model must be conveniently expressed as a biomolecular system (e.g., with complex formation or internal state modification). On the other hand, CRN bisimulations work for a generic language-independent CRN, which however must be explicitly given. Further, unlike CRN bisimulations, fragmentation is performed on a “static” view of the model, without information on the reaction rates. The ODE aggregations of both forward CRN bisimulation and fragmentation introduce loss of information (in contrast to backward CRN bisimulation). But, unlike our forward variant, in fragmentation the same species may be present in more than one fragment. Additionally, species may occur in fragments with multiplicity numbers. Thus, fragmentation can be seen

as a form of *improper lumping* that is not necessarily induced by a partition of the original state-space variables [30]. Overall, because of these differences, it is not difficult to find models that can be reduced by our CRN bisimulations and not by fragmentation, and vice versa. This is presented in detail in Section 6.

2 Background

Notation. We write $A \rightarrow B$ and B^A for the functions from A to B . When $f \in A \rightarrow B$ and $a \in A$, we set $f_a := f(a)$. Moreover, for any $X \subseteq A$ and $b \in B$, we define $f(X) := \{b \in B \mid \exists a \in X.(f(a) = b)\}$. Sets and multisets are denoted by $\{\dots\}$ and $\{\!\!\{ \dots \}\!\!\}$, respectively. Also, we shall not distinguish among an equivalence relation and the partition induced by it, and shall use the symbol $\sim_{\mathcal{H}}$ to denote the equivalence relation with $\mathcal{H} = S/\sim_{\mathcal{H}}$. Finally, given two partitions \mathcal{H}_1 and \mathcal{H}_2 of a given set S , we say that \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 $H_1 \subseteq H_2$.

2.1 Chemical Reaction Networks

Formally, a CRN (S, R) is a pair consisting of a finite set of species S taken from a countable infinite universe of all species, and a finite set of chemical reactions R . A reaction is a triple written in the form $\rho \xrightarrow{\alpha} \pi$, where ρ and π are the multisets of species *reactants* and *products*, respectively, and $\alpha > 0$ is the reaction rate. In particular, we focus on basic chemistry where only *elementary reactions* are considered, where at most two reactants (possibly of the same species) interact. No restrictions are instead imposed on products. Several models found in the literature (including those discussed in Section 6) belong to this class. Also, this is consistent with the physical considerations which stipulate that reactions with more than two reactants are very unlikely to occur in nature [24]. We denote by $\rho(X)$ the multiplicity of species X in the multiset ρ , and by $\mathcal{MS}(S)$ the set of finite multisets of species in S . To adhere to standard chemical notation, we shall use the operator $+$ to denote multiset union, e.g., $X + Y + Y$ (or just $X + 2Y$) denotes the multiset $\{X, Y, Y\}$. We may also use X to denote either the species X or the singleton $\{X\}$.

The (autonomous) ODE system $\dot{V} = F(V)$ underlying a CRN (S, R) is $F : \mathbb{R}_{\geq 0}^S \rightarrow \mathbb{R}^S$, where each component F_X , with $X \in S$ is defined as:

$$F_X(V) := \sum_{\rho \xrightarrow{\alpha} \pi \in R} (\pi(X) - \rho(X)) \cdot \alpha \cdot \prod_{Y \in S} V_Y^{\rho(Y)}.$$

This represents the well-known *mass-action* kinetics, where the reaction rate is proportional to the concentrations of the reactants involved. Since the ODE system of a CRN is given by polynomials, the vector field F is locally Lipschitz. Hence, the theorem of Picard-Lindelöf ensures that for any $V(0) \in \mathbb{R}_{\geq 0}^S$ there exists a unique non-continuable solution of $\dot{V} = F(V)$.

► **Example 1.** We now provide a simple CRN, (S_e, R_e) , with $S_e = \{A, B, C, D, E\}$ and $R_e = \{A \xrightarrow{6} E, B \xrightarrow{6} D, A+B \xrightarrow{2} C, C+D \xrightarrow{5} 2C+D, E+D \xrightarrow{5} 2E+D\}$, which will be used as a running example throughout the paper. Its ODE system is given by

$$\begin{aligned} \dot{V}_A &= -6V_A - 2V_A V_B & \dot{V}_B &= -6V_B - 2V_A V_B & \dot{V}_C &= 2V_A V_B + 5V_C V_D \\ \dot{V}_D &= 6V_B & \dot{V}_E &= 6V_A + 5V_E V_D \end{aligned}$$

In the following, we shall assume that the universe of all species is well-ordered with respect to \sqsubseteq . We then say that a function $\mu : S \rightarrow S$ is a *choice function* of a partition \mathcal{H} of S , if $\mu(X) = \min_{\sqsubseteq} H$ for all $H \in \mathcal{H}$ and $X \in H$. Also, choice functions can be trivially lifted to multisets applying them element-wise, e.g., $\mu(X + Y) = \mu(X) + \mu(Y)$.

2.2 Fluid Lumpability

Ordinary Fluid Lumpability. We start by defining the notion of ordinary fluid lumpability, which is an instance of *ordinary lumpability* for ODEs [36] specialized to CRNs.

► **Definition 2** (Ordinary fluid lumpability). Let (S, R) be a CRN, F be its vector field, and $\mathcal{H} = \{H_1, \dots, H_m\}$ a partition of S . Then, \mathcal{H} is *ordinary fluid lumpable* if for all $H \in \mathcal{H}$ there exists a polynomial \wp_H in $|\mathcal{H}|$ variables such that $\sum_{X \in H} F_X(V) = \wp_H(\sum_{X \in H_1} V_X, \dots, \sum_{X \in H_m} V_X)$ for all $V \in \mathbb{R}_{\geq 0}^S$.

Informally, a partition \mathcal{H} is ordinary fluid lumpable if, for each $H \in \mathcal{H}$, the polynomial $\sum_{X \in H} F_X(V)$ in the variables $\{V_X \mid X \in S\}$ can be rewritten into a polynomial \wp_H in the variables $\{\sum_{X \in H} V_X \mid H \in \mathcal{H}\}$. In particular, if \mathcal{H} is known to be an ordinary fluid lumpable partition of (S, R) and V denotes the solution of $\dot{V} = F(V)$ subject to $V(0) \in \mathbb{R}_{\geq 0}^S$, the solution of the aggregated ODE system $(\dot{W}_{H_1}, \dots, \dot{W}_{H_m}) = (\wp_{H_1}(W), \dots, \wp_{H_m}(W))$ with $W_H(0) = \sum_{X \in H} V_X(0)$ is such that $W_H(t) = \sum_{X \in H} V_X(t)$ for all $t \in \text{domain}(V)$.

► **Example 3.** Consider the ODEs of (S_e, R_e) of Example 1, and let $\mathcal{H}_O = \{\{A\}, \{B\}, \{C, E\}, \{D\}\}$. By applying a variable renaming consistent with the blocks of \mathcal{H}_O , i.e., $V_{CE} = V_C + V_E$, and by exploiting the linearity of the differential operator we get

$$\dot{V}_A = -6V_A - 2V_A V_B \quad \dot{V}_B = -6V_B - 2V_A V_B \quad \dot{V}_{CE} = 2V_A V_B + 6V_A + 5V_D V_{CE} \quad \dot{V}_D = 6V_B$$

That is, we obtained an ODE system in terms of block variables only. ◀

Exact Fluid Lumpability. We extend to CRNs the notion of exact fluid lumpability in [37].

► **Definition 4** (Exact fluid lumpability). Let (S, R) be a CRN, F its vector field, and \mathcal{H} a partition of S . We call $V \in \mathbb{R}^S$ *constant on \mathcal{H}* if $V_{X_i} = V_{X_j}$ for all $H \in \mathcal{H}$, and all $X_i, X_j \in H$. Then, \mathcal{H} is *exactly fluid lumpable* if $F(V)$ is constant on \mathcal{H} whenever V is constant on \mathcal{H} .

► **Example 5.** Consider the ODEs of (S_e, R_e) of Example 1, and let $\mathcal{H}_E = \{\{A, B\}, \{C\}, \{D\}, \{E\}\}$. It is easy to see that A and B have same concentrations at all time points if initialized equally. In these cases, we can replace the ODEs of (S_e, R_e) with the ones aggregated according to \mathcal{H}_E , obtained by removing \dot{V}_B and replacing all occurrences of V_B with V_A :

$$\dot{V}_A = -6V_A - 2V_A V_A \quad \dot{V}_C = 2V_A V_A + 5V_C V_D \quad \dot{V}_D = 6V_A \quad \dot{V}_E = 6V_A + 5V_E V_D$$

That is, we obtained a (lossless) aggregated ODE system written in terms of a variable per block, chosen according to \sqsubseteq . ◀

We remark that the above definition expresses exact fluid lumpability in terms of properties of the ODE vector field of a CRN. Instead, in [37] exactly fluid lumpability was defined directly in terms of the desired dynamical property, i.e., that the ODE solutions within any equivalence class be equal at all time points. The following result is a new contribution showing that this dynamical property is fully characterized by the vector-field based definition.

► **Theorem 6.** *Let (S, R) be a CRN and F its vector field. A partition \mathcal{H} of S is exactly fluid lumpable if and only if, for any $V(0) \in \mathbb{R}_{\geq 0}^S$ that is constant on \mathcal{H} , the underlying solution of $\dot{V} = F(V)$ is such that $V(t)$ is constant on \mathcal{H} for all $t \in \text{domain}(V)$.¹*

¹ All proofs are provided in the extended technical report [13].

3 CRN Bisimulations

Both notions of fluid lumpability given in Section 2 are not convenient to be used directly because they involve a universal quantifier over the (uncountable) state space. We address this problem by providing structural conditions that concern only the reactions of a CRN.

3.1 Forward CRN Bisimulation

We now introduce forward CRN bisimulation, an equivalence on species that will turn out to induce ordinary fluid lumpability. We start with the notions of *reaction* and *production rate*. The former collects the rates at which the concentration of a species X decreases when reacting with a given partner. The latter collects the positive contribution that X exerts to the concentration of a species Y , again when reacting with a certain partner.

► **Definition 7** (Reaction and production rates). Let (S, R) be a CRN, $X, Y \in S$, and $\rho \in \mathcal{MS}(S)$. The ρ -*reaction rate* of X , and the ρ -*production rate* of Y -elements by X are defined respectively as

$$\mathbf{crr}[X, \rho] := (\rho(X) + 1) \sum_{X+\rho \xrightarrow{\alpha} \pi \in R} \alpha, \quad \mathbf{pr}(X, \rho, Y) := (\rho(X) + 1) \sum_{X+\rho \xrightarrow{\alpha} \pi \in R} \alpha \cdot \pi(Y)$$

Finally, for $H \subseteq S$ we define $\mathbf{pr}[X, \rho, H] := \sum_{Y \in H} \mathbf{pr}(X, \rho, Y)$.

► **Definition 8** (Forward CRN Bisimulation). Let (S, R) be a CRN, \mathcal{R} an equivalence relation over S and $\mathcal{H} = S/\mathcal{R}$. Then, \mathcal{R} is a forward CRN bisimulation (abbreviated FB) if for all $(X, Y) \in \mathcal{R}$, all $\rho \in \mathcal{MS}(S)$, and all $H \in \mathcal{H}$ it holds that

$$\mathbf{crr}[X, \rho] = \mathbf{crr}[Y, \rho] \quad \text{and} \quad \mathbf{pr}[X, \rho, H] = \mathbf{pr}[Y, \rho, H] \quad (1)$$

► **Example 9.** Consider $\mathcal{H}_O = \{\{A\}, \{B\}, \{C, E\}, \{D\}\}$ of Example 3. It can be shown that \mathcal{H}_O is an FB, as, e.g., $\mathbf{crr}[C, D] = \mathbf{crr}[E, D] = 5$, and $\mathbf{pr}[C, D, \{C, E\}] = \mathbf{pr}[E, D, \{C, E\}] = 10$. ◀

We are interested in the coarsest FB, as well as in the coarsest one refining a given initial partition of species.

► **Proposition 10.** Let (S, R) be a CRN, I a set of indices, and \mathcal{R}_i an FB for (S, R) , for all $i \in I$. The transitive closure of their union $\mathcal{R} = (\bigcup_{i \in I} \mathcal{R}_i)^*$ is an FB for (S, R) . In particular, if each \mathcal{R}_i is such that S/\mathcal{R}_i refines some partition \mathcal{G} of S , then so does S/\mathcal{R} .

► **Theorem 11** (Forward bisimulation implies ordinary fluid lumpability). Let (S, R) be a CRN. Then, \mathcal{H} is an ordinarily fluid lumpable partition of S if \mathcal{H} is an FB of S .

FB is only a sufficient condition for lumpability, as discussed in the next example. (However, Section 6 shows that FB can be effectively applied to interesting existing models.)

► **Example 12.** Consider the CRN $(\{F, G\}, \{F \xrightarrow{\alpha_1} G, G \xrightarrow{\alpha_2} F\})$, having ODEs

$$\dot{V}_F = -\alpha_1 V_F + \alpha_2 V_G \quad \dot{V}_G = -\alpha_2 V_G + \alpha_1 V_F$$

If $\alpha_1 \neq \alpha_2$, $\mathcal{H}_c = \{\{G, F\}\}$ is not an FB, as $\mathbf{crr}[F, \emptyset] = \alpha_1$ and $\mathbf{crr}[G, \emptyset] = \alpha_2$. Nevertheless, the above ODE system is lumpable. Indeed, by applying the variable renaming consistent with \mathcal{H}_c , i.e., $V_{FG} = V_F + V_G$, we get a single ODE for V_{FG} , i.e., $\dot{V}_{FG} = 0$. ◀

3.2 Backward CRN Bisimulation

We now introduce backward CRN bisimulation, an equivalence on species that will turn out to characterize exact fluid lumpability. We start with the notion of cumulative flux rate, which collects the overall contribution that reactions with a given multiset of reactants ρ exert to the concentration of a species X .

► **Definition 13** (Cumulative flux rate). Let (S, R) be a CRN, $X \in S$, $\rho \in \mathcal{MS}(S)$, and $\mathcal{M} \subseteq \mathcal{MS}(S)$. Then, we define

$$\mathbf{fr}(X, \rho) := \sum_{\rho \xrightarrow{\alpha} \pi \in R} (\pi(X) - \rho(X)) \cdot \alpha, \quad \mathbf{fr}[X, \mathcal{M}] := \sum_{\rho \in \mathcal{M}} \mathbf{fr}(X, \rho).$$

We call $\mathbf{fr}(X, \rho)$ and $\mathbf{fr}[X, \mathcal{M}]$ ρ -flux rate and cumulative \mathcal{M} -flux rate of X , respectively.

► **Definition 14** (Backward CRN bisimulation). Let (S, R) be a CRN, \mathcal{R} an equivalence relation over S , $\mathcal{H} = S/\mathcal{R}$ and μ the choice function of \mathcal{H} . Then, \mathcal{R} is a backward CRN bisimulation (BB) if for any $(X, Y) \in \mathcal{R}$ it holds that

$$\mathbf{fr}[X, \mathcal{M}] = \mathbf{fr}[Y, \mathcal{M}] \quad \text{for all } \mathcal{M} \in \{\rho \mid \rho \xrightarrow{\alpha} \pi \in R\} / \approx_{\mathcal{H}}, \quad (2)$$

where any two $\rho, \sigma \in \mathcal{MS}(S)$ satisfy $\rho \approx_{\mathcal{H}} \sigma$ if $\mu(\rho) = \mu(\sigma)$.

► **Example 15.** Consider $\mathcal{H}_E = \{\{A, B\}, \{C\}, \{D\}, \{E\}\}$ of Example 5. We first note that $\{A\} \approx_{\mathcal{H}_E} \{B\}$, as $\approx_{\mathcal{H}_E}$ relates multisets with same number of \mathcal{H}_E -equivalent species. Also, it can be shown that \mathcal{H}_E is a BB, as, e.g., $\mathbf{fr}[A, \mathcal{M}] = \mathbf{fr}[B, \mathcal{M}] = -6$ for $\mathcal{M} = \{\{A\}, \{B\}\}$. ◀

As for FB, there exists a coarsest BB (that refines a given partition of S).

► **Proposition 16.** Let (S, R) be a CRN, I a set of indices, and \mathcal{R}_i a BB for (S, R) , for all $i \in I$. The transitive closure of their union $\mathcal{R} = (\bigcup_{i \in I} \mathcal{R}_i)^*$ is a BB for (S, R) . In particular, if each \mathcal{R}_i is such that S/\mathcal{R}_i refines some partition \mathcal{G} of S , then so does S/\mathcal{R} .

We now state the mentioned characterization of exact fluid lumpability in terms of BB.


► **Theorem 17** (Backward bisimulation characterizes exact fluid lumpability). Let (S, R) be a CRN. Then, \mathcal{H} is an exactly fluid lumpable partition of S if and only if \mathcal{H} is a BB of S .

► **Remark.** We wish to stress that FB and BB are not comparable: First, \mathcal{H}_O is not a BB, as $\mathbf{fr}[C, \{A+B\}] = 2$ and $\mathbf{fr}[E, \{A+B\}] = 0$; Second, \mathcal{H}_E is not an FB, as $\mathbf{crr}(A, B) = 2$ and $\mathbf{crr}(B, B) = 0$; Third, for the same reasons, $\{\{A, B\}, \{C, E\}, \{D\}\}$ is neither an FB nor a BB. Similar examples on models of biological relevance are provided in Section 6. ◀

4 Reduced Chemical Reaction Networks up to CRN Bisimulations

We have shown that, given a CRN and a CRN bisimulation \mathcal{R} , we can analyze the aggregated ODE system according to \mathcal{R} . We now provide the notion of reduced CRN from which the aggregated ODEs can be directly generated, as depicted in Figure 1.

► **Definition 18** (Forward reduction). Let (S, R) be a CRN, \mathcal{H} an FB, and μ its choice function. The (\mathcal{H}, F) -reduction of (S, R) is given by $(S, R)^{(\mathcal{H}, F)} = (S^{(\mathcal{H}, F)}, R^{(\mathcal{H}, F)})$, where $S^{(\mathcal{H}, F)} = \mu(S)$ and $R^{(\mathcal{H}, F)}$ is defined as follows: (F1) Discard all reactions $\rho \xrightarrow{\alpha} \pi$ such that $\rho \neq \mu(\rho)$; (F2) Replace all remaining reactions $\rho \xrightarrow{\alpha} \pi$ with $\rho \xrightarrow{\alpha} \mu(\pi)$; (F3) Fuse all reactions that have the same reactants and products by summing their rates.



■ **Figure 1** Relation among (\mathcal{H} -reduced) CRNs and (\mathcal{H} -lumped) semantics, with \mathcal{H} a bisimulation.

The idea underlying forward reduction is to discard all reactions having non-representative reagents, and to replace the products of the remaining reactions with their representatives. This can be seen as a special case of Theorem 4.4 of [10].

► **Example 19.** Consider the FB $\mathcal{H}_O = \{\{A\}, \{B\}, \{C, E\}, \{D\}\}$ used in Example 3. The (\mathcal{H}_O, F) -reduction of (S_e, R_e) is (with C being the representative of its block) $S_e^{(\mathcal{H}_O, F)} = \{A, B, C, D\}$, $R_e^{(\mathcal{H}_O, F)} = \{A \xrightarrow{6} C, B \xrightarrow{6} D, A+B \xrightarrow{2} C, C+D \xrightarrow{5} 2C+D\}$. Note that the reaction $E+D \xrightarrow{5} 2E+D$ is discarded, as E is not a representative species. ◀

We now state that the (\mathcal{H}, F) -reduction of an FB \mathcal{H} induces the ODEs aggregated according to \mathcal{H} . For example, the (\mathcal{H}_O, F) -reduction of (S_e, R_e) induces the ODEs shown in Example 3, if applying the renaming $V_C = V_{CE}$.

► **Theorem 20** (Forward reduction induces aggregation). *Let (S, R) be a CRN, \mathcal{H} an FB and μ its choice function. Then, $(S, R)^{(\mathcal{H}, F)}$ is computed in at most $\mathcal{O}(|R| \cdot |S| \cdot (\log(|R|) + \log(|S|)))$ steps. Crucially, if F is the vector field of (S, R) and \hat{F} the one of $(S, R)^{(\mathcal{H}, F)}$, then $\sum_{X \in H} F_X(V) = \hat{F}_{\mu(Y)}(\sum_{X \in H_1} V_X, \dots, \sum_{X \in H_m} V_X)$ for all $V \in \mathbb{R}_{\geq 0}^S$, $H \in \mathcal{H}$ and $Y \in H$.*

For the backward reduction, the underlying idea is to keep track only of differential contributions that affect the representative species $\mu(S)$.

► **Definition 21** (Backward reduction). Let (S, R) be a CRN, \mathcal{H} a BB, and μ its choice function. The (\mathcal{H}, B) -reduction of (S, R) is given by $(S, R)^{(\mathcal{H}, B)} = (S^{(\mathcal{H}, B)}, R^{(\mathcal{H}, B)})$, where $S^{(\mathcal{H}, B)} = \mu(S)$ and $R^{(\mathcal{H}, B)}$ is obtained as follows: (B1) Replace all reactions $\rho \xrightarrow{\alpha} \pi$ with $\rho \xrightarrow{\alpha} \tilde{\pi}$ where $\tilde{\pi}(X_i) := \pi(X_i)$ if $X_i \in \mu(S)$ and $\tilde{\pi}(X_i) := \rho(X_i)$ otherwise; (B2) Replace all such obtained reactions $\rho \xrightarrow{\alpha} \pi$ with $\mu(\rho) \xrightarrow{\alpha} \mu(\pi)$; (B3) Fuse all reactions that have the same reactants and products by summing their rates.

► **Example 22.** Considering the CRN (S_e, R_e) and the BB \mathcal{H}_E , (B1) changes $B \xrightarrow{6} D$ in $B \xrightarrow{6} D+B$, and $A+B \xrightarrow{2} C$ in $A+B \xrightarrow{2} C+B$, while (B2) yields $\{A \xrightarrow{6} E, A \xrightarrow{6} D+A, A+A \xrightarrow{2} C+A, C+D \xrightarrow{5} 2C+D, E+D \xrightarrow{5} 2E+D\}$. Finally, (B3) does not introduce any change. ◀

► **Theorem 23** (Backward reduction induces aggregation). *Let (S, R) be a CRN, \mathcal{H} a BB and μ its choice function. Then, $(S, R)^{(\mathcal{H}, B)}$ is computed in at most $\mathcal{O}(|R| \cdot |S| \cdot (\log(|R|) + \log(|S|)))$ steps. Crucially, if \hat{F} denotes the vector field induced by $(S, R)^{(\mathcal{H}, B)}$, it holds that $F_X(V) = \hat{F}_X(V)$ for all $X \in \mu(S)$ and $V \in \mathbb{R}_{\geq 0}^S$ that are constant on \mathcal{H} .*

5 Partition Refinement Algorithms for CRN Bisimulations

We study a polynomial-time algorithm for the computation of the coarsest bisimulations that refine an arbitrary input partition. We start introducing two auxiliary equivalence relations.

► **Definition 24** (Splitter equivalences). Let (S, R) be a CRN and \mathcal{H} a partition over S . Then, we write $X \sim_{\mathcal{H}}^F Y$ if (1) is fulfilled by (X, Y) . Similarly, write $X \sim_{\mathcal{H}}^B Y$ if (X, Y) satisfies (2).

Algorithm 1 Template partition refinement algorithm for the construction of the coarsest CRN bisimulations that refine some given initial partition \mathcal{G} .

Require: A CRN (S, R) , a partition \mathcal{G} of S and $\chi \in \{F, B\}$.

```

 $\mathcal{H} \leftarrow \mathcal{G}$ 
while true do
   $\mathcal{H}' \leftarrow S / (\sim_{\mathcal{H}}^{\chi} \cap \sim_{\mathcal{H}})$ 
  if  $\mathcal{H}' = \mathcal{H}$  then
    return  $\mathcal{H}$ 
  else
     $\mathcal{H} \leftarrow \mathcal{H}'$ 
  end if
end while

```

Algorithm 1 iteratively computes the coarsest forward or backward bisimulation (when $\chi = F$ or $\chi = B$, respectively) that refines a given input partition of species of a CRN. Note that, contrary to CRN reduction algorithms, one (parametric) algorithm suffices for both bisimulations. Using the above splitter equivalences, at each iteration the blocks of the current partition $S/\sim_{\mathcal{H}}$ are split in sub-blocks of $\sim_{\mathcal{H}}^{\chi}$ -equivalent species $S/(\sim_{\mathcal{H}}^{\chi} \cap \sim_{\mathcal{H}})$. The algorithm terminates when no refinement is performed.

The freedom in choosing the initial partition \mathcal{G} is useful in both bisimulations. For FB it allows to single out species that are the “observables” of the CRN. These are the species for which the modeler is interested in obtaining distinct ODE solutions, information which would otherwise be lost if such species are found in larger equivalence classes. BB is lossless, hence this issue does not arise. However BB requires the same initial conditions for equivalent species. In this case, an appropriate input partition may tell apart species for which it is known that the initial conditions are different.

► **Theorem 25 (Correctness).** *Given a CRN (S, R) and a partition \mathcal{G} of S , Algorithm 1 calculates the coarsest forward and backward bisimulation that refines \mathcal{G} . In both cases, the number of steps needed is polynomial in the number of species and reactions.*

Note that, due to space constraints, we only focussed on the existence of a polynomial-time algorithm, and in the next section we provide numerical evidence of its scalability. The proof of this theorem gives a bound of $\mathcal{O}(|R|^2 \cdot |S|^5)$ on the number of steps. Tighter bounds could be obtained by extending classical partition refinement approaches available for labeled transitions systems [31, 1] to CRNs, which is however the subject of future work.

6 Evaluation

We now evaluate FB and BB. We first study their effectiveness in reducing the ODEs of a number of biochemical models from the literature given in the `.net` format of BioNetGen [5], version 2.2.5-stable. Using selected models we discuss how FB and BB relate with each other, and provide a biological interpretation of the aggregations. Finally, we compare them against κ 's fragmentation. All experiments are replicable using a prototype available at <http://sysma.imtlucca.it/crnreducer/>.

Numerical results. Table 1 lists our case studies: four synthetic benchmarks to obtain combinatorially larger CRNs by varying the number of phosphorylation sites (M1–M4) [33]; a model of pheromone signaling (M5, [35]); two signaling pathways through the Fc ϵ complex

■ **Table 1** Forward and backward reductions and corresponding speed-ups in ODE analysis. Speed-up entries “—” indicate that the original model could not be solved; entries “x” indicate that the coarsest bisimulation did not reduce the original model.

<i>Id</i>	<i>Original model</i>		<i>Forward reduction</i>				<i>Backward reduction</i>			
	<i> R </i>	<i> S </i>	<i>Red.(s)</i>	<i> R </i>	<i> S </i>	<i>Speed-up</i>	<i>Red.(s)</i>	<i> R </i>	<i> S </i>	<i>Speed-up</i>
M1	3538944	262146	4.61E+4	990	222	—	7.65E+4	2708	222	—
M2	786432	65538	1.92E+3	720	167	—	3.68E+3	1950	167	—
M3	172032	16386	8.15E+1	504	122	1.16E+3	1.77E+2	1348	122	5.34E+2
M4	48	18	1.00E-3	24	12	1.00E+0	2.00E-3	45	12	1.00E+0
M5	194054	14531	3.72E+1	142165	10855	1.03E+0	1.32E+3	93033	6634	1.03E+0
M6	187468	10734	3.07E+1	57508	3744	1.92E+1	2.71E+2	144473	5575	3.53E+0
M7	32776	2506	1.26E+0	16481	1281	6.23E+0	1.66E+1	32776	2506	x
M8	41233	2562	1.12E+0	33075	1897	1.12E+0	1.89E+1	41233	2562	x
M9	5033	471	1.91E-1	4068	345	1.04E+0	4.35E-1	5033	471	x
M10	5797	796	1.61E-1	4210	503	1.47E+0	7.37E-1	5797	796	x
M11	5832	730	3.89E-1	1296	217	1.32E+1	6.00E-1	2434	217	7.55E+0
M12	487	85	2.00E-3	264	56	1.88E+0	6.00E-3	426	56	1.31E+0
M13	24	18	1.20E-2	24	18	x	7.00E-3	6	3	1.00E+0

(M6–M7, [20, 33]); two models of enzyme activation (M8–M9, [2]); a model of a tumor suppressor protein (M10, [3]); a model of tyrosine phosphorylation and adaptor protein binding (M11, [14, 15]); a MAPK model (M12, [28]); and an *influence network* (M13, [11]).

Headers $|R|$ and $|S|$ give the number of reactions and species of the CRN (and of its reductions), respectively. The reduction times (*Red.*) account also for the computation of the quotient CRNs. The speed-up is the ratio between the time to solve the ODEs of the original CRN and that of the reduced one including the time to reduce the CRN. Measurements were taken on a 2.6 GHz Intel Core i5 with 4 GB of RAM. The time interval of the ODE solution was taken from the original papers; for M1–M4, where this data was not available, time point 50.0 was used as an estimate of steady state. The initial conditions for the ODEs were also taken from the original papers. The initial partition for FB was chosen to be the trivial one containing the singleton block $\{S\}$ (i.e., no species was singled out). Instead, the initial partition for BB was chosen consistently with the ODE initial conditions; that is, two species may be equivalent only if they have the same initial conditions in the original CRN. This ensured that the backward reduced CRN was a lossless aggregation of the original CRN.


We make three main observations: (i) FB and BB can reduce a significant number of models. In the two largest models of our case studies, M1 and M2, the bisimulations were able to provide a compact aggregated ODE system which could be straightforwardly analyzed, while the solutions of the original models did not terminate due to out-of-memory errors, consistently with [33]. (ii) FB and BB are not comparable in general. For instance, both reduce M5 to 10855 and 6634 species, respectively, while M6 is reduced to 3744 species by FB, and to 5574 by BB. Also, FB was able to reduce M7–M10, while BB did not aggregate. The influence network M13 shows the opposite; in fact, none of the influence networks presented in [11] can be reduced up to FB (here we showed M13, which is the largest one from [11]). (iii) Models M1–M4 and M12 show that the intersection between FB and BB is nonempty.

Biological interpretation. Models M1 and M2 enjoy significant reductions and ODE analysis speed-ups. Here we use them to explain that FB and BB are effective at aggregating species representing symmetric sites in a complex. For this, let us consider M4, chosen for space reasons. A typical equivalence class is for instance $\{E(s!1).S(p1\sim P, p2\sim U!1), E(s!1).S(p1\sim U!1, p2\sim P)\}$. According to the syntax of the BioNetGen language, the CRN species are

formed from basic *molecules* S and E . Molecule S has two binding sites ($p1$, and $p2$) which can be either in *phosphorylated* state (P) or not (U); E has one stateless binding site (s) which can bind to those of S to form a complex. The two sites of S have equivalent capabilities in terms of binding with other species or changing state. For instance, the above equivalence class contains two species composed by S and E , with E bound to the unphosphorylated site of S (here the exclamation mark links the binding sites used to form the species). Models M1 and M2 exhibit a fast growth of the number of species due to a larger number of symmetric sites, requiring distinct species to track exactly which site exhibits a particular phosphorylation state. This form of symmetry has also been studied in [9] where the authors propose an approach to detect it directly at the κ level. However, an experimental comparison could not be performed because [9] is not yet implemented. Although both bisimulations give the same equivalence classes in these cases, the reduced CRNs have different reactions, since FB provides the dynamics of the sums of equivalent species, while BB considers the distinct dynamics of representative species. Instead, aggregation of identical binding sites is supported by BioNetGen. This can be seen in models M6 and M7, since they both have $Lig(l, l)$, a ligand with two copies of site l . Intuitively, the rule




gives rise to only one chemical complex in the underlying CRN, $Rec(a!1).Lig(l!1, l)$. This represents the (forward and backward) canonical representative of a ligand bound to a single receptor $Rec(a)$. To see this, let us rename the two sites and *expand* the rule appropriately:




Then, this underlying CRN will distinguish the two sites. However, applying either of our CRN bisimulations leads to the CRN for Equation (3).

We remark that the original CRN sizes of M6 and M7 already account for the aggregations obtained with BioNetGen. Nevertheless, our CRN bisimulations allow for further (significant) reductions. For instance, part of the reductions for M6 are due to the presence of $Rec(a, b, g_1, g_2)$, a molecule with symmetric sites g_1 and g_2 , similarly to those of M4.

Symmetric sites are not the only property captured by our bisimulations. For instance in both M8 and M9 one of the FB equivalence classes is given by:



A biological interpretation is that a species containing the molecule J behaves in the same way as long as it is bound to a molecule R having binding site i in state “on”. This is independent of whether R is further complexed with other molecules via its binding site l ; For instance, the first species models that R is only bound to J , while in the second and third species it is also bound to L . Finally, in M5, one of the BB equivalence classes is



It 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.

Experimental comparison with κ -based reduction techniques. We now experimentally compare our CRN bisimulations and fragmentation in the case of rule-based biochemical models for which the underlying CRN can be fully enumerated. All models in Table 1 belong to this class; however, none of them was originally available in κ , the only language that supports fragmentation. Thus, we performed a manual translation of a selection of the case studies from the BioNetGen language into κ .²

We found:

- *Models that can be reduced by CRN bisimulations but not by fragmentation.* The κ encoding of M12 (a case where only cosmetic syntactical changes are required) returned 85 fragments, equal to the size of the CRN, while both FB and BB reduced to 56 species. The encodings of M6 and M7 necessitated expansions analogous to Equation (4) because κ does not currently support distinct sites with the same name. This led to bigger initial CRNs, for which fragmentation returned 58040 fragments for M6 and 10930 for M7.
- *Models that can be reduced by fragmentation but not by our bisimulations.* The κ model of early events of the EGF pathway in [6] is reduced from 356 species to 38 fragments [17], while no aggregation is obtained with either FB or BB.
- *Models that can be reduced by both our bisimulations and fragmentation.* The κ encodings of models M1–M4 present different reductions than using either bisimulation, specifically 38, 34, 30 and 10 fragments (versus 222, 167, 122, and 12 FB and BB equivalence classes, respectively). It can be shown that, in the latter examples, the reductions are complementary, in the sense that no two bisimilar species are included in the same fragment. While our bisimulations captured symmetric sites, fragments explain that the sites of S are *independent*, i.e., the state of a site does not affect the dynamics of the other. For instance, one of the fragments for model M4 is

$$\{S(p1\sim P, p2\sim P), S(p1\sim P, p2\sim U), E(s!1).S(p1\sim P, p2\sim U!1), F(s!1).S(p1\sim P, p2\sim P!1)\}$$

which essentially collects all species where the $p1$ site of molecule S is phosphorylated.

7 Conclusion

Forward and backward bisimulations are equivalence relations over the species of a chemical reaction network inducing a partition of the underlying mass-action system of ordinary differential equations. An experimental evaluation has demonstrated their usefulness by showing their complementarity as well as significant model reductions in a number of biochemical models available in the literature. This has been supported by a prototype, which currently allows a ready-to-use tool-chain with BioNetGen, a state-of-the-art tool.

Ongoing work is studying stochastic counterparts of both forward and backward bisimulations, to obtain model reductions when the semantics of chemical reaction networks based on continuous-time Markov chains is considered. Also, we plan to investigate the applicability of our bisimulations in other model repositories, e.g., those using the well-known SBML interchange format (<http://sbml.org>).

Acknowledgement. The authors thank J. Krivine and J. Feret for helpful discussions, and the anonymous referees for suggestions that improved the paper.

This work was partially supported by the EU project QUANTICOL, 600708. L. Cardelli is partially funded by a Royal Society Research Professorship. Part of this research has

² All discussed κ -encodings are provided in the technical report [13] and are available for download.

been carried out while M. Tribastone, M. Tschaikowski, and A. Vandin were at University of Southampton, UK.

References

- 1 C. Baier, B. Engelen, and M. E. Majster-Cederbaum. Deciding bisimilarity and similarity for probabilistic processes. *J. Comput. Syst. Sci.*, 60(1):187–231, 2000.
- 2 D. Barua, J. R. Faeder, and J. M. Haugh. A bipolar clamp mechanism for activation of jak-family protein tyrosine kinases. *PLoS Computational Biology*, 5(4), 2009.
- 3 D. Barua and W. S. Hlavacek. Modeling the effect of apc truncation on destruction complex function in colorectal cancer cells. *PLoS Comput Biol*, 9(9):e1003217, 09 2013.
- 4 M. Bernardo. A survey of Markovian behavioral equivalences. In *Formal Methods for Perf. Eval.*, volume 4486 of *LNCS*, pages 180–219. Springer Berlin Heidelberg, 2007.
- 5 M. L. Blinov, J. R. Faeder, B. Goldstein, and W. S. Hlavacek. BioNetGen: software for rule-based modeling of signal transduction based on the interactions of molecular domains. *Bioinformatics*, 20(17):3289–3291, 2004.
- 6 M. L. Blinov, J. R. Faeder, B. Goldstein, and W. S. Hlavacek. A network model of early events in epidermal growth factor receptor signaling that accounts for combinatorial complexity. *Biosystems*, 83:136–151, 2006.
- 7 P. Buchholz. Exact and Ordinary Lumpability in Finite Markov Chains. *Journal of Applied Probability*, 31(1):59–75, 1994.
- 8 P. Buchholz. Markovian Process Algebra: Composition and Equivalence. In *Proc. 2nd Workshop on Process Algebra and Performance Modelling*, Erlangen, Germany, 1994.
- 9 F. Camporesi and J. Feret. Formal reduction for rule-based models. *Electronic Notes in Theoretical Computer Science*, 276:29–59, 2011. MFPS XXVII.
- 10 Ferdinanda Camporesi, Jérôme Feret, Heinz Koepl, and Tatjana Petrov. Combining model reductions. *Electr. Notes Theor. Comput. Sci.*, 265:73–96, 2010.
- 11 L. Cardelli. Morphisms of reaction networks that couple structure to function. *BMC Systems Biology*, 8(1):84, 2014.
- 12 L. Cardelli and A. Csikász-Nagy. The cell cycle switch computes approximate majority. *Sci. Rep.*, 2, 2012.
- 13 L. Cardelli, M. Tribastone, M. Tschaikowski, and A. Vandin. Forward and Backward Bisimulations for Chemical Reaction Networks. Extended Version. <http://arxiv.org/abs/1507.00163>. <http://arxiv.org/abs/1507.00163>, 2015.
- 14 J. Colvin, M. I. Monine, J. R. Faeder, W. S. Hlavacek, D. D. Von Hoff, and R. G. Posner. Simulation of large-scale rule-based models. *Bioinformatics*, 25(7):910–917, 2009.
- 15 J. Colvin, M. I. Monine, R. N. Gutenkunst, W. S. Hlavacek, D. D. Von Hoff, and R. G. Posner. Rulemonkey: software for stochastic simulation of rule-based models. *BMC Bioinformatics*, 11:404, 2010.
- 16 H. Conzelmann, J. Saez-Rodriguez, T. Sauter, B. Kholodenko, and E. Gilles. A domain-oriented approach to the reduction of combinatorial complexity in signal transduction networks. *BMC Bioinformatics*, 7(1):34, 2006.
- 17 V. Danos, J. Feret, W. Fontana, R. Harmer, and J. Krivine. Abstracting the differential semantics of rule-based models: Exact and automated model reduction. In *LICS*, pages 362–381, 2010.
- 18 V. Danos and C. Laneve. Formal molecular biology. *TCS*, 325(1):69–110, 2004.
- 19 R. De Nicola, U. Montanari, and F. Vaandrager. Back and forth bisimulations. In *CONCUR*, volume 458 of *LNCS*, pages 152–165. Springer, 1990.
- 20 J. R. Faeder, W. S. Hlavacek, I. Reischl, M. L. Blinov, H. Metzger, A. Redondo, C. Wofsy, and B. Goldstein. Investigation of early events in FcεRI-mediated signaling using a detailed mathematical model. *The Journal of Immunology*, 170(7):3769–3781, 2003.

- 21 J. Feret, T. Henzinger, H. Koepl, and T. Petrov. Lumpability abstractions of rule-based systems. *TCS*, 431:137–164, 2012.
- 22 Jerome Feret, Heinz Koepl, and Tatjana Petrov. Stochastic fragments: A framework for the exact reduction of the stochastic semantics of rule-based models. *International Journal of Software and Informatics*, 7(4):527–604, 2013.
- 23 J. Fisher and T.A. Henzinger. Executable cell biology. *Nature Biotechnology*, 25(11):1239–1249, 2007. See also correspondence in *Nature Biotechnology* 26(7):737-8;738-9, 2008.
- 24 D. Gillespie. The chemical Langevin equation. *The Journal of Chemical Physics*, 113(1):297–306, 2000.
- 25 J. Heath, M. Kwiatkowska, G. Norman, D. Parker, and O. Tymchyshyn. Probabilistic model checking of complex biological pathways. *TCS*, 391(3):239–257, 2008.
- 26 H. Hermanns and M. Rettelbach. Syntax, semantics, equivalences, and axioms for MTIPP. In *Proceedings of Process Algebra and Probabilistic Methods*, pages 71–87, Erlangen, 1994.
- 27 J. Hillston. *A Compositional Approach to Performance Modelling*. CUP, 1996.
- 28 P. Kocieniewski, J. R. Faeder, and T. Lipniacki. The interplay of double phosphorylation and scaffolding in MAPK pathways. *Journal of Theoretical Biology*, 295:116–124, 2012.
- 29 K. G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94(1):1–28, 1991.
- 30 M. S. Okino and M. L. Mavrouniotis. Simplification of mathematical models of chemical reaction systems. *Chemical Reviews*, 2(98):391–408, 1998.
- 31 R. Paige and R. Tarjan. Three partition refinement algorithms. *SIAM Journal on Computing*, 16(6):973–989, 1987.
- 32 S. W. Shin, C. Thachuk, and E. Winfree. Verifying chemical reaction network implementations: A pathway decomposition approach. In *VEMPD*, Vienna Summer of Logic, 2014.
- 33 M. W. Sneddon, J. R. Faeder, and T. Emonet. Efficient modeling, simulation and coarse-graining of biological complexity with NFsim. *Nature Methods*, 8(2):177–183, 2011.
- 34 J. Sproston and S. Donatelli. Backward Bisimulation in Markov Chain Model Checking. *IEEE Trans. Software Eng.*, 32(8):531–546, 2006.
- 35 R. Suderman and E. J. Deeds. Machines vs. ensembles: Effective MAPK signaling through heterogeneous sets of protein complexes. *PLoS Comput Biol*, 9(10):e1003278, 10 2013.
- 36 J. Toth, G. Li, H. Rabitz, and A. S. Tomlin. The effect of lumping and expanding on kinetic differential equations. *SIAM Journal on Applied Mathematics*, 57(6):1531–1556, 1997.
- 37 M. Tschaikowski and M. Tribastone. Exact fluid lumpability for Markovian process algebra. In *CONCUR*, LNCS, pages 380–394, 2012.
- 38 M. Tschaikowski and M. Tribastone. A unified framework for differential aggregations in Markovian process algebra. *JLAMP*, 84(2):238–258, 2015.