

Syntactic Markovian Bisimulation for Chemical Reaction Networks

Luca Cardelli¹, Mirco Tribastone², Max Tschaikowski², and Andrea Vandin²

¹ Microsoft Research & University of Oxford, UK

² IMT School for Advanced Studies Lucca, Italy

Abstract. In chemical reaction networks (CRNs) with stochastic semantics based on continuous-time Markov chains (CTMCs), the typically large populations of species cause combinatorially large state spaces. This makes the analysis very difficult in practice and represents the major bottleneck for the applicability of minimization techniques based, for instance, on lumpability. In this paper we present syntactic Markovian bisimulation (SMB), a notion of bisimulation developed in the Larsen-Skou style of probabilistic bisimulation, defined over the structure of a CRN rather than over its underlying CTMC. SMB identifies a lumpable partition of the CTMC state space a priori, in the sense that it is an equivalence relation over species implying that two CTMC states are lumpable when they are invariant with respect to the total population of species within the same equivalence class. We develop an efficient partition-refinement algorithm which computes the largest SMB of a CRN in polynomial time in the number of species and reactions. We also provide an algorithm for obtaining a quotient network from an SMB that induces the lumped CTMC directly, thus avoiding the generation of the state space of the original CRN altogether. In practice, we show that SMB allows significant reductions in a number of models from the literature. Finally, we study SMB with respect to the deterministic semantics of CRNs based on ordinary differential equations (ODEs), where each equation gives the time-course evolution of the concentration of a species. SMB implies forward CRN bisimulation, a recently developed behavioral notion of equivalence for the ODE semantics, in an analogous sense: it yields a smaller ODE system that keeps track of the sums of the solutions for equivalent species.

1 Introduction

Chemical reaction networks (CRNs) are a powerful model of interaction at the basis of many branches of science such as organic and inorganic chemistry, ecology, epidemiology and systems biology. In computer science, the interpretation of biological systems as computing devices has stimulated a vigorous line of research ranging from the understanding of the computational power of such models (e.g., [50, 41, 13]) to the development of formal techniques for their specification, analysis, and verification (e.g., [15, 30, 17, 48]).

Traditionally, CRNs have been equipped with the well-known quantitative semantics based on a system of ordinary differential equations (ODEs), where an ODE relates to the time-course deterministic evolution of the concentration of each species. It is

well-known, however, that such semantics may not always accurately reflect the observed behavior, for example when some species are present in low copies [22]. The alternative stochastic semantics based on continuous-time Markov chains (CTMCs) may provide more accurate estimates, but at an increased computational expense. Indeed, since each CTMC state is a population vector giving the number of copies of each species, there is a combinatorial explosion of the CTMC state space as a function of the initial population of species. In order to cope with this, it would be highly desirable to be able to perform CTMC aggregation, e.g., based on lumpability [5, 20]. However, its applicability in practice is fundamentally hampered by the fact that available methods require to explicitly enumerate the state space. This is typically infeasible for realistic CRNs sizes, or even impossible because CRNs may give rise to infinite state spaces.

Inspired from the seminal work of Larsen and Skou on probabilistic bisimulation [37], in this paper we propose a reduction technique that avoids the generation of the original state space. Instead of reasoning at the semantic level, we identify conditions on the *CRN syntax*. More precisely, we provide a new notion of equivalence over CRN species, called *syntactic Markovian bisimulation (SMB)*, based on properties that can be checked by inspecting the set of reactions only, but it induces a partition on states of the underlying CTMC: two CTMC states are related if they are invariant with respect to the total population of species in the same SMB equivalence class. To clarify this, suppose we have a CRN with species A , B , and C , and the SMB that gives the partition $\{\{A, B\}, \{C\}\}$. Then the CTMC state $(n_A = 1, n_B = 2, n_C = 1)$ belongs to the same block as state $(n_A = 2, n_B = 1, n_C = 1)$ because they have equal sums within the equivalence classes. The resulting CTMC partition is an ordinarily lumpable one [5]: in the lumped CTMC each macro-state represents the sum of the probabilities of the original states of a partition block.

Importantly, the lumped CTMC can be obtained avoiding the generation of the original state space altogether, owing to an algorithm that constructs a quotient CRN for an SMB. The possibility of such a CRN-to-CRN transformation is useful not only for model *minimization*, but also for using bisimulation as a technique for model *comparison*. This has received increased attention, largely motivated by applications to evolutionary biology [26, 7, 8, 14, 11, 9, 49, 43].

SMB turns out to be a natural extension of the ordinary lumpability condition (defined on the underlying CTMC semantics) to the CRN syntax. Ordinary lumpability relates two CTMC states whenever they have the same cumulative transition rates toward any partition block. Analogously, SMB relates two species when, roughly speaking, the cumulative kinetic parameters of the reactions where they are involved as reagents are the same for every *lifted equivalence class* of products. This lifting is defined by relating two products that are invariant up to the SMB equivalence classes, as above. An important consequence of this definition style is that it allows us to also extend the aforementioned CTMC minimization algorithms to SMB. In particular, we present an algorithm for computing the largest SMB that refines a given input partition of species in polynomial time and space.

Being syntactically driven, it is perhaps not surprising that SMB is only a sufficient condition for CTMC lumpability. As a consequence, it is important to understand to what extent it can be effectively applied in practice. On CRN models of biological systems

taken from the literature we show that SMB can achieve substantial compressions, yielding reduced CRNs with significantly fewer species and reactions in some cases. We measure the impact of SMB on the analysis of the CRN when this is done by means of stochastic simulation [28], the method of choice in realistic systems due to the large state spaces involved (e.g., [18]). We report noticeable runtime speed-ups in many cases, up to two orders of magnitude, even allowing the execution of benchmark models that would otherwise generate out of memory errors if not reduced. These numerical tests also reveal an interesting connection between SMB and the deterministic semantics of CRNs: the equivalence classes of species found in all the analyzed models coincide with those recently reported in [8] for *forward CRN bisimulation* (FB), an equivalence relation over species that aggregates related ODEs in an analogous way, exactly preserving their total concentration trajectories at all time points. We explain this fact by showing that SMB implies an FB, however the converse is not true in general. Nevertheless, in our tests FB was not able to aggregate more than SMB.

Further related work. The closest approach to ours is by Feret et al. [25] who identify *stochastic fragments* on the rule-based language κ [19]. These represent syntactic criteria that yield a sufficient condition for *weak lumpability* (see, e.g., [5]) on the CTMC. The advantage is that the rule-based model is often combinatorially smaller than its underlying CRN description; however, the approach is domain specific in that it can be applied to systems, e.g., protein-protein interaction networks, which can be conveniently expressed as rule-based systems. On the contrary, since SMB works at the level of the CRN it is more general, at the expense of a more expensive syntactic analysis in this application domain.

For process algebra with quantitative semantics based on CTMCs, several approaches have been proposed for *on-the-fly* computations of lumped chain that avoid the generation of the original state space. These are based on deriving transitions of the lumped chain from a canonical representative of an equivalence class (e.g., [31, 29, 36, 45]). Here considerable state-space compressions are owed to *symmetry reduction*, whereby identical copies of a process in parallel composition can be collapsed through a lumpable partition that contains all processes that are equal up to a permutation of the composed sub-terms. Symmetry reduction could be useful in the case that the CRN is described at the individual molecular level, as for instance in Cardelli’s Chemical Ground Form [6]. However, we remark that a CRN gives a CTMC that tracks the population sizes of each species, implicitly accounting already for symmetry due to the assumption that two molecules of the same species are identical. SMB, instead, captures structural relations, see [8] for a physical interpretation of some equivalence classes. In this sense, SMB is closer in spirit to the idea of *place bisimulation* for Petri nets, which establishes a relation over places that induces a bisimulation in the classical, non-quantitative strong sense [2].

Paper structure. Section 2 introduces the notion of CRN and defines its semantics. SMB is introduced in Section 3, while in Section 4 it is shown that SMB induces a reduced CRN whose CTMC is related via ordinary lumpability to the CTMC of the original CRN. Section 5 presents the algorithm for computing the largest SMB. Applicability and efficiency of the algorithm are demonstrated on biological models from the literature

in Section 6. A formal comparison of SMB with FB complements the experiments. Conclusions are drawn in Section 7.

2 Chemical Reaction Networks

In this paper we consider *mass-action* CRNs, where each reaction is labeled with a constant, the reaction rate. The speed of the reaction will be proportional with this rate to the product of the abundances of the reactants. In particular, we focus on basic chemistry where only *elementary reactions* are considered: *unary reactions*, involving a single reactant performing a spontaneous reaction, and *binary reactions*, where two reactants interact; we call a binary reaction a *homeoreaction* if the two reactants are of the same species. Elementary reactions pose no restrictions 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 [27]. In the rest of the paper we will refer to such *elementary mass-action CRNs* as just CRNs.

Formally, a CRN (S, R) is a set of species S and a set of chemical reactions R . Each reaction is a triple written in the form $\rho \xrightarrow{\alpha} \pi$, where ρ and π are the multi-sets of species representing the *reactants* and *products*, respectively, and $\alpha \geq 0$ is the reaction rate. We denote by $\rho(X)$ the multiplicity of species X in the multi-set ρ , and by $\mathcal{MS}(S)$ the set of finite multi-sets of species in S . To adhere to standard chemical notation, we shall also use the operator $+$ to denote multi-set union, e.g., $X + Y + Y$ (or just $X + 2Y$) denotes the multi-set of species $\{X, Y, Y\}$; similarly $\rho - X$ denotes multi-set difference $\rho \setminus \{X\}$. We also use X to denote either the species X or the singleton $\{X\}$.

Example 1. We now provide a simple CRN, (S_e, R_e) , with $S_e = \{A, B, C, D, E\}$ and

$$R_e = \{A \xrightarrow{6} D, A \xrightarrow{2} 3C, C+D \xrightarrow{5} 2C+D, B \xrightarrow{6} C, \\ B \xrightarrow{2} 3D, E+D \xrightarrow{5} 2C+D, 2D \xrightarrow{3} C\},$$

which will be used as a running example throughout the paper.

We next recall the well-known CTMC semantics of CRNs (see, e.g., [6, 28]), which allows us to associate a population-based CTMC to a given CRN and an initial population of its species. Here the state descriptor gives the number of elements for each species, hence it is formally represented as a multi-set of species. The CTMC specification is mediated by a multi-transition system (MTS), to record multiplicity of transitions. This is needed to account for two or more reactions contributing to the same CTMC transition, e.g., $A + B \xrightarrow{\alpha_1} B + C$ and $A \xrightarrow{\alpha_2} C$. The whole state space is defined by enumerating states, starting from some initial state.

Definition 1 (Multi-transition system of a CRN). Let (S, R) be a CRN. The multiset of outgoing transitions from state $\sigma \in \mathcal{MS}(S)$ is obtained as

$$\begin{aligned} out(\sigma) = & \{\!\!\{ \sigma \xrightarrow{\alpha \cdot \sigma(X)} \sigma - X + \pi \mid (X \xrightarrow{\alpha} \pi) \in R \}\!\!\} \\ & \uplus \{\!\!\{ \sigma \xrightarrow{\alpha \cdot \sigma(X) \cdot \sigma(Y)} ((\sigma - X) - Y) + \pi \mid X \neq Y \wedge (X + Y \xrightarrow{\alpha} \pi) \in R \}\!\!\} \\ & \uplus \{\!\!\{ \sigma \xrightarrow{\frac{\alpha}{2} \cdot \sigma(X) \cdot (\sigma(X) - 1)} ((\sigma - X) - X) + \pi \mid (X + X \xrightarrow{\alpha} \pi) \in R \}\!\!\} \end{aligned}$$

The set of reachable states from σ , denoted by $reach(\sigma)$, is the smallest set such that: i) $\sigma \in reach(\sigma)$; and ii) if $\sigma' \in reach(\sigma)$, then the target states of $out(\sigma')$ belong to $reach(\sigma)$. Finally, for an initial state $\sigma_0 \in \mathcal{MS}(S)$, the MTS for (S, R) and σ_0 is the union of the multi-sets of transitions outgoing from any reachable state, i.e. $MTS(\sigma_0) = \uplus_{\theta \in reach(\sigma_0)} out(\theta)$.

We note that each reaction $\rho \xrightarrow{\alpha} \pi$ can be applied to source states σ containing ρ , i.e. $\sigma = \sigma' + \rho$ for some multi-set σ' . The corresponding target state is $\sigma' + \pi$. The rate for unary reactions $X \xrightarrow{\alpha} \pi$ is $\alpha \cdot \sigma(X)$ and accounts for the fact that each instance of the reagent can perform the reaction independently. For binary reactions $X + Y \xrightarrow{\alpha} \pi$ with $X \neq Y$, instead, the transition rate is proportional to the product of the populations of the species involved, i.e. $\alpha \cdot \sigma(X) \cdot \sigma(Y)$. This corresponds to the number of possible interactions between molecules, proportionally to the reaction propensity α [35, 28]. For a homeoreaction involving X , the number of distinct interactions is given by $\binom{\sigma(X)}{2} = \frac{1}{2} \cdot \sigma(X) \cdot (\sigma(X) - 1)$.

Example 2. Consider the initial population $\sigma_{0e} = 2A + C + D$ for (S_e, R_e) . Then we have $out(\sigma_{0e}) = \{\!\!\{ \sigma_{0e} \xrightarrow{6 \cdot 2} A + C + 2D, \sigma_{0e} \xrightarrow{2 \cdot 2} A + 4C + D, \sigma_{0e} \xrightarrow{5} 2A + 2C + D \}\!\!\}$. The three transitions are due, respectively, to the first, second and third reaction of R_e .

We wish to stress the difference between a CRN and its MTS. While both are collections of triples in $\mathcal{MS}(S) \times \mathbb{R} \times \mathcal{MS}(S)$, the elements of the former are *syntactic*. Instead, the nature of the latter is *semantic* because it induces the underlying CTMC. In particular, given an MTS the CTMC is obtained by collapsing all transitions between the same source and target into a single CTMC transition and summing their rates.

Definition 2 (CTMC semantics). Let (S, R) be a CRN, and σ_0 an initial population. The CTMC of (S, R) for σ_0 has states $reach(\sigma_0)$ and its transitions are given by

$$MC(\sigma_0) = \{ \sigma \xrightarrow{r} \theta \mid \sigma, \theta \in reach(\sigma_0) \wedge \sigma \neq \theta \wedge r = \sum_{\sigma \xrightarrow{r'} \theta \in MTS(\sigma_0)} r' \}.$$

For any two states $\sigma, \theta \in \mathcal{MS}(S)$ the element of the infinitesimal generator matrix of $MC(\sigma_0)$ from σ to θ is defined as:

$$q(\sigma, \theta) = \begin{cases} r & \text{if } \sigma \neq \theta \wedge \sigma \xrightarrow{r} \theta \in MC(\nu_0) \\ - \sum_{\theta' \in \mathcal{MS}(S) \text{ s.t. } \theta' \neq \sigma} q(\sigma, \theta') & \text{if } \sigma = \theta \\ 0 & \text{otherwise} \end{cases}$$

For any $\mathcal{M} \subseteq \mathcal{MS}(S)$, we define $q[\sigma, \mathcal{M}] = \sum_{\theta \in \mathcal{M}} q(\sigma, \theta)$ and $q[\mathcal{M}, \theta] = \sum_{\sigma \in \mathcal{M}} q(\sigma, \theta)$.

3 Syntactic Markovian Bisimulation

This section introduces Syntactic Markovian Bisimulation (SMB) as a sufficient condition for CTMC ordinary lumpability. We first recast this latter notion to our notation.

Definition 3. *Let (S, R) be a CRN, σ_0 an initial population, $MC(\sigma_0)$ the underlying CTMC and \mathcal{H} a partition of $\mathcal{MS}(S)$. Then, $MC(\sigma_0)$ is ordinarily lumpable with respect to \mathcal{H} iff for any σ_1, σ_2 in the same block of \mathcal{H} we have $q[\sigma_1, \mathcal{M}] = q[\sigma_2, \mathcal{M}]$ for all $\mathcal{M} \in \mathcal{H}$.*

Lumpability is given in terms of an equivalence relation among the states of a CTMC. Instead, SMB is an equivalence over the species of a CRN. Note that there is no one-to-one correspondence between species and the state space of the CTMC underlying a CRN. Indeed, the species define the state descriptor, but the cardinality of the state space is typically much larger since it depends on all the possible configurations of populations that are reachable from a given initial population. Thus we need to lift a relation over species to one over CTMC states. We do so by providing the notion of *multi-set lifting*: given a CRN (S, R) and an equivalence relation \mathcal{R} over S , the lifting of \mathcal{R} relates multi-sets with same number of \mathcal{R} -equivalent species.

Definition 4 (Multi-set Lifting). *Let (S, R) be a CRN, $\mathcal{R} \subseteq S \times S$ be an equivalence relation over S , and \mathcal{H} be the partition induced by \mathcal{R} over S . We define the multi-set lifting of \mathcal{R} on $\mathcal{MS}(S)$, denoted by $\mathcal{R}^\uparrow \subseteq \mathcal{MS}(S) \times \mathcal{MS}(S)$, as*

$$\mathcal{R}^\uparrow \triangleq \{(\sigma_1, \sigma_2) \mid \sigma_1, \sigma_2 \in \mathcal{MS}(S) \wedge \forall H \in \mathcal{H} : \sum_{X \in H} \sigma_1(X) = \sum_{X \in H} \sigma_2(X)\}$$

The multi-set lifting of \mathcal{R} can be readily seen to be an equivalence relation over $\mathcal{MS}(S)$.

Example 3. Consider the equivalence relation \mathcal{R}_m over S_e inducing $\mathcal{H}_m = \{\{A\}, \{B\}, \{C, E\}, \{D\}\}$. Examples of multi-sets related by \mathcal{R}_m^\uparrow are C and E , $2C$ and $2E$, and $C + E$ and $2E$, while $(A + C, B + C) \notin \mathcal{R}_m^\uparrow$.

The syntactic checks of SMB are performed via the notion of *reaction rate* given below. It computes, in essence, the cumulative rate that transforms a given reagent ρ into a certain product π .

Definition 5 (Reaction rate). *Let (S, R) be a CRN, and $\rho, \pi \in \mathcal{MS}(S)$. The reaction rate from ρ to π is defined as $\mathbf{rr}(\rho, \pi) = \sum_{\rho \xrightarrow{\alpha} \pi \in R} \alpha$. For any $\mathcal{M} \subseteq \mathcal{MS}(S)$, we define $\mathbf{rr}[\rho, \mathcal{M}] = \sum_{\pi \in \mathcal{M}} \mathbf{rr}(\rho, \pi)$.*

We can now define SMB.

Definition 6 (Syntactic Markovian Bisimulation). *Let (S, R) be a CRN, \mathcal{R} an equivalence relation over S , \mathcal{R}^\uparrow the multi-set lifting of \mathcal{R} and $\mathcal{H}^\uparrow = \mathcal{MS}(S)/\mathcal{R}^\uparrow$. We say that \mathcal{R} is a syntactic Markovian bisimulation (SMB) for (S, R) if and only if*

$$\mathbf{rr}[X + \rho, \mathcal{M}] = \mathbf{rr}[Y + \rho, \mathcal{M}], \text{ for all } (X, Y) \in \mathcal{R}, \rho \in \mathcal{MS}(S), \text{ and } \mathcal{M} \in \mathcal{H}^\uparrow.$$

We define the syntactic Markovian bisimilarity of (S, R) as the union of all SMBs of (S, R) .

Remark 1. Note that the multi-sets $X + \rho$ and $Y + \rho$ differ only in one species (X and Y), thus projecting comparisons involving multisets (i.e., $X + \rho$ and $Y + \rho$) onto species (i.e., X and Y). In this view, ρ plays a role similar to an action type in traditional bisimulations, since it restricts interactions with a given reagent partner (or \emptyset in case of unary reactions). Furthermore, Definition 6 entails a finite number of checks because all evaluations of \mathbf{rr} are equal to zero for multisets that are not products in the CRN, see the algorithm of Section 5.

Example 4. Consider again \mathcal{R}_m and \mathcal{H}_m . From $\mathbf{rr}(C + D, 2C + D) = \mathbf{rr}(E + D, 2C + D)$ can be inferred that \mathcal{R}_m is an SMB.

As usual, we are interested in the largest bisimulation. The next result ensures that syntactic Markovian bisimilarity is an SMB, thus showing that it is also the largest one. Following the approach of [32], we show this by proving that the transitive closure of a union of SMBs is an SMB.¹

Proposition 1. *Let (S, R) be a CRN, I a set of indices, and \mathcal{R}_i an SMB 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 SMB for (S, R) .*

We now provide our first major result.

Theorem 1. *Let \mathcal{R} be an SMB for the CRN (S, R) . Then, its multi-set lifting \mathcal{R}^\uparrow induces the ordinarily lumpable partition \mathcal{H}^\uparrow on $MC(\sigma_0)$ for any initial state σ_0 .*

Three remarks are in order. First, we stress that a single SMB induces infinitely many ordinarily lumpable partitions because there are no constraints on σ_0 . Second, Theorem 1 makes no assumption on the cardinality of the CTMC state space underlying the CRN. In particular, it can also be applied to infinite state spaces; indeed Example 1 is an instance of such a situation because, e.g., of the reaction $C + D \xrightarrow{5} 2C + D$, which may generate infinitely many copies of C whenever the initial state has at least one copy of species D and one of species C . The original result of ordinary lumpability applies to finite CTMCs. However, using concepts from functional analysis and the theory of linear ODEs on Banach spaces, this statement can be extended, under certain assumptions, to CTMCs with countably infinite state spaces [39]. A sufficient condition for the theory to apply is to assume that the state space of the CTMC is partitioned in blocks of finite size. Indeed, the multi-set lifting ensures that any CTMC partition stemming from SMB enjoys this property. Third, as anticipated in Section 1, SMB is only a sufficient condition for CTMC ordinary lumpability.

Example 5. Consider the CRN $(\{F, G\}, \{F \xrightarrow{\alpha_1} G, G \xrightarrow{\alpha_2} F\})$ with $\alpha_1 \neq \alpha_2$ and $\sigma_0 = F$. The underlying CTMC has the state space $\{F, G\}$ and it readily follows that $\{\{F, G\}\}$ is an ordinarily lumpable partition, while it is not an SMB.

At the same time, however, SMB can be computed efficiently and induces significant reductions to biological models from literature, as discussed in Section 6.

¹The proofs of all statements are provided in the technical report available at <http://sysma.imtlucca.it/tools/erode/>.




Fig. 1. The relation among (\mathcal{R} -reduced) CRNs and (\mathcal{R}^\uparrow -lumped) semantics, with \mathcal{R} an SMB.

4 Reduced CRN

Given a CRN (S, R) and an SMB \mathcal{R} , we next provide an algorithm that efficiently computes a \mathcal{R} -reduced CRN that induces directly the CTMC aggregated according to \mathcal{R}^\uparrow , without exploring the state space of the original CTMC. This is visualized in Figure 1.

We wish to point out that this reduction algorithm happens to coincide with the *forward reduction* of [8], which has been applied to obtain a quotient CRN up to an FB, mentioned in Section 1, defined for the ODE semantics of CRNs. For the sake of completeness we state the notion of reduced CRN according to this paper's notation. To this end, we introduce the following notions. Given a partition \mathcal{H} of S such that $\mathcal{H} = S/\mathcal{R}$, let X^H denote the canonical representative of a block $H \in \mathcal{H}$. Moreover, for any $\rho \in \mathcal{MS}(S)$, set $\rho^\mathcal{R} = \sum_{X \in \rho} X^H$ for the multiset obtained replacing each species with its canonical representative. Also, for any $\mathcal{M} \in \mathcal{H}^\uparrow$ we use $\mathcal{M}^\mathcal{R}$ for $\rho^\mathcal{R}$, with ρ any multi-set in \mathcal{M} .

Definition 7 (Reduced CRN). *Let (S, R) be a CRN, \mathcal{R} an equivalence relation on S and $\mathcal{H} = S/\mathcal{R}$. The \mathcal{R} -reduction of (S, R) is defined as $(S, R)^\mathcal{R} = (S^\mathcal{R}, R^\mathcal{R})$, where $S^\mathcal{R} = \{X^H \mid H \in \mathcal{H}\}$ and $R^\mathcal{R}$ is computed as follows: (F1) Discard all reactions $\rho \xrightarrow{\alpha} \pi$ such that $\rho \neq \rho^\mathcal{R}$, i.e. whose reagents have species that are not representatives; (F2) Replace the species in the products of the remaining reactions with their canonical representatives; (F3) Fuse all reactions that have the same reactants and products by summing their rates.*

In the case of our running example, the above definition yields the following.

Example 6. Consider the SMB \mathcal{R}_m of Example 3 and the underlying partition $\mathcal{H}_m = \{\{A\}, \{B\}, \{C, E\}, \{D\}\}$. With C being the representative of its block, the \mathcal{R}_m -reduction of (S_e, R_e) is $S_e^{\mathcal{R}_m} = \{A, B, C, D\}$, $R_e^{\mathcal{R}_m} = \{A \xrightarrow{6} D, A \xrightarrow{2} 3C, B \xrightarrow{6} C, B \xrightarrow{2} 3D, C+D \xrightarrow{5} 2C+D, 2D \xrightarrow{3} C\}$. Note that the reaction $E + D \xrightarrow{5} 2C + D$ is discarded.

Theorem 2. *Let (S, R) be a CRN, \mathcal{R} denote an SMB and $\mathcal{H} = S/\mathcal{R}$. Further, let \mathcal{H}^\uparrow denote the partition induced by \mathcal{R}^\uparrow on $\mathcal{MS}(S)$. Then, for any initial population σ_0 of (S, R) , the underlying CTMC is such that for all $\sigma \in \mathcal{MS}(S)$ it holds that $q_{(S,R)}[\sigma, \mathcal{M}] = q_{(S^\mathcal{R}, R^\mathcal{R})}(\sigma^\mathcal{R}, \mathcal{M}^\mathcal{R})$ for any $\mathcal{M} \in \mathcal{H}^\uparrow$.*


```

1 SMBisimilarity( $S, R, \mathcal{H}$ ) :=
2    $\mathcal{H}^\uparrow = \text{lift}(\mathcal{H}, \Pi(R))$ 
3    $\text{spls} = \mathcal{L}(R) \times \mathcal{H}^\uparrow$ 
4   while ( $\text{spls} \neq \emptyset$ )
5     ( $\rho, \mathcal{M}_{spl}$ ) = pop( $\text{spls}$ )
6     Split( $S, R, \rho, \mathcal{M}_{spl}, \mathcal{H}, \text{spls}$ )
1 Split( $S, R, \rho, \mathcal{M}_{spl}, \mathcal{H}, \text{spls}$ ) :=
2   forall ( $X \in S$ )
3      $X.\text{rr}0$ 
4   forall ( $\pi \in \mathcal{M}_{spl}$ )
5     forall ( $X + \rho \xrightarrow{r} \pi \in \pi.\text{inc}$ )
6        $X.\text{rr} = X.\text{rr} + r$ 
7   updatePartitionAndSplitters( $S, R, \mathcal{H}, \text{spls}$ )

```

Fig. 2. Syntactic Markovian bisimilarity

If we assume that each block of an SMB partition stores a pointer to its representative, the reduced CRN can be computed in $O(r \cdot s \cdot \log s)$ steps [8], where $s := |S|$ and $r := |R|$.

5 Computing Syntactic Markovian Bisimilarities

Syntactic Markovian bisimilarity can be encoded as a partition refinement problem [38], analogously to well-known algorithms for quantitative extensions of labeled transition systems [33, 3, 10]. Hence, we only detail the conceptually novel parts, i.e., the computation of the quantities in Definition 6 and the notion of multi-set lifting.

Notation and data structures. Our algorithm for syntactic Markovian bisimilarity, `SMBisimilarity`, is given in Figure 2, where (S, R) is the input CRN, and \mathcal{H} the initial partition to be refined up to SMB. We use $s := |S|$, $r := |R|$, and $\mathcal{L}(R)$ for the set of all *labels*, i.e., all species multi-sets ρ to be considered according to Definition 6 in the computation of `rr`. That is, $\mathcal{L}(R) = \{\{X\} \mid X \in S \wedge \exists Y \in S, \exists X + Y \xrightarrow{r} \pi \in R\} \cup \{\emptyset\}$. We use $\{\emptyset\}$ to account for unary reactions, and $\{X\} \in \mathcal{L}(R)$ for each reagent X occurring in at least one binary reaction in the CRN. We set $l := |\mathcal{L}(R)|$, which is bounded by $\min(s + 1, 2 \cdot r)$.

In this pseudo-code we assume that species and reactions are stored in data structures via pointers. Species are stored in a list, while a block of \mathcal{H} is a list of its species, each species in turn having a pointer to its block, requiring $O(s)$ space. Also R is stored in a list. Each reaction has two fields for the reagents, and a list of pairs in the form (species, multiplicity) for the products, sorted according to a total ordering on species. Thus, R requires $O(s \cdot r)$ space. Finally, $\mathcal{L}(R)$ is stored in a list too, for an overall $O(s \cdot r)$ space complexity.

Overview. `SMBisimilarity` is based on Paige and Tarjan’s classical solution to the *relational coarsest partition problem* [38] and quantitative extensions thereof (e.g., [20, 3]). A given initial partition is iteratively refined (i.e., its blocks are split), until a partition satisfying the required conditions is found. Refinements are based on the notion of *splitter*, here given by $(\rho, \mathcal{M}_{spl})$, with $\rho \in \mathcal{L}(R)$ and $\mathcal{M}_{spl} \in \mathcal{H}^\uparrow$: a block of \mathcal{H} is *split* in sub-blocks of species with same ρ -reaction rate towards a block \mathcal{M}_{spl} of *equivalent* multi-sets of species. We stress that, differently from classic bisimulations, in SMB splitters are blocks of products obtained via the multi-set lifting from a species partition \mathcal{H} , rather than blocks of \mathcal{H} itself. Note that the set $\mathcal{MS}(S)$ of all possible multi-sets of species in S is infinite. However, we can restrict to the set $\Pi(R) = \{\pi \mid (\rho \xrightarrow{\alpha} \pi) \in R\}$

only, collecting the multi-sets of species appearing as products in the reactions of the considered CRN. This is because any multi-set in $\mathcal{MS}(S) \setminus \Pi(R)$ will not contribute to the reaction rates. We store $\Pi(R)$ as a list, requiring $O(r)$ space, while a partition of $\Pi(R)$ is encoded by representing a block with a list of pointers to its products.

The SMBisimilarity procedure (Lines 1-6). The procedure starts (Line 2) by creating the partition \mathcal{H}^\uparrow of $\Pi(R)$ according to the multi-set lifting of \mathcal{H} . This requires $O(s \cdot r \cdot \log r)$ time; this is because $O(s \cdot r)$ is required to count the number of species of each block of \mathcal{H} , for each product, while $O(s \cdot r \cdot \log r)$ is required to partition the products. This is done by iteratively sorting the products according to the number of H -species they have, for each $H \in \mathcal{H}$. It requires $O(r \cdot \log r)$ per block of \mathcal{H} . Then, a set `spls` of initial candidate splitters is generated for each $\rho \in \mathcal{L}(R)$ and $\mathcal{M} \in \mathcal{H}^\uparrow$. In order to bound the size of `spls` to $O(s)$ we do not explicitly store each pair $(\rho, \mathcal{M}_{spl})$ for all ρ . Instead we store only one, initialized with a reference to the first position of $\mathcal{L}(R)$, and then update the pointer to the next position when necessary.

Lines 4-6 iterate until there are candidate splitters to be considered. One is selected and removed from `spls`, and the procedure `Split` is invoked to refine each block of \mathcal{H} according to that splitter, and to generate new candidate ones.

The Split procedure (Lines 1-7). Each species X has associated a real-valued field `X.rr`, initialized to 0 in Lines 2-3 in $O(s)$ time. Also, we assume that each product $\pi \in \Pi(R)$ is provided with a list, `inc`, which points to all the reactions that have that product. Each list `inc` has size $O(r)$, while exactly r entries appear in all `inc`. The `inc` list allows to compute the reaction rates by iterating all reactions at most only once. In fact, given an input splitter $(\rho, \mathcal{M}_{spl})$, Lines 4-6 store `rr[X + ρ, \mathcal{M}_{spl}]` in `X.rr`, for each species X , by iterating once the `inc` list of each product multi-set $\pi \in \mathcal{M}_{spl}$. In particular, we can have two cases: either $\rho = \emptyset$, when only unary reactions are considered (as $X + \emptyset = X$), or $\rho = Z$, with $Z \in S$, when only binary reactions having Z in their reagents are considered. In both cases, checking for the presence of ρ in the reagents of each reaction takes constant time. The computation has $O(r)$ time complexity, since each reaction appears in $\pi.\text{inc}$ for one π only.

The actual splitting (Line 7). Using the computed rates, Line 7 then performs the actual splitting. We do not detail this part, as it is inspired by the usual approach, e.g., [20, 3], consisting of the following three steps: (i) Each block is split using an associated balanced binary search tree (BST) in which each species X of the block is inserted providing `rr[X + ρ, \mathcal{M}_{spl}]` as key, and a new block is added to \mathcal{H} for each leaf of the BST; this requires $O(s \cdot \log s)$ time, as there are at most s insertions in the BSTs, each having size at most s ; (ii) If at least one block has been split, all candidate splitters must be discarded; this takes $O(r)$ time, as `spls` contains at most an entry per product $\pi \in \Pi(R)$,² while deletion from `spls` takes constant time assuming that it is implemented as a linked list; (iii) If at least a block has been split, all splitters have to be recomputed. This is because another multi-set lifting must be considered from the new partition. It takes $O(s \cdot r \cdot \log r)$ to do so. Thus, overall `Split` has time complexity $O(s \cdot (\log s + r \cdot \log r))$. Also, note that the BSTs do not worsen the space complexity, as only one for a block is built at a time.

² Recall that, given a block \mathcal{M}_{spl} , only one entry is stored to represent all pairs $(\rho, \mathcal{M}_{spl})$.

<i>Id</i>	<i>Int.</i>	<i>Original model</i>			<i>SMB reduction</i>				
		$ R $	$ S $	<i>CTMC (s)</i>	<i>SMB (s)</i>	<i>Red. (s)</i>	$ R $	$ S $	<i>CTMC (s)</i>
M1	50	3538944	262146	—	2.68E+5	2.04E+1	990	222	1.77E+1
M2	50	786432	65538	—	6.67E+3	4.09E+0	720	167	1.15E+1
M3	50	172032	16386	—	1.95E+2	3.4E−1	504	122	7.96E+0
S1	50	36864	4098	9.12E+2	9.38E+0	1.01E−1	336	86	5.28E+0
S2	50	7680	1026	1.03E+2	7.33E−1	3.00E−2	210	58	3.69E+0
M5	3600	194054	14531	3.54E+5	5.88E+2	1.20E+0	142165	10855	3.29E+5
M6	3840	187468	10734	1.79E+4	1.96E+2	5.68E−1	57508	3744	1.47E+3
M7	3840	32776	2506	1.34E+3	8.80E+0	2.68E−1	16481	1281	4.88E+2
S3	500000	284	143	4.13E+2	3.30E−2	1.40E−2	142	72	1.39E+2

Table 1. SMB reductions and corresponding speed-ups in CTMC simulation.

Complexity. We observe that `Split` is invoked $O(l \cdot s \cdot r)$ times. This is because, initially, $l \cdot r$ candidate splitters have to be considered. At every step where some blocks of \mathcal{H} get split (which happens at most s times), all splitters are removed, and at most $l \cdot r$ new candidate splitters are added to `spls`. In conclusion, syntactic Markovian bisimilarity takes $O(l \cdot s^2 \cdot r \cdot (\log s + r \cdot \log r))$ time and $O(s \cdot r)$ space.

Remark 2. We wish to stress that step (iii) of the actual splitting phase is not necessary in classic partition refinement algorithms [20, 3, 10]. This is because only blocks of the current partition are used as splitters in those algorithms. Hence, splitters are computed and maintained at no additional cost. It is exactly due to this reason that the time complexity of our algorithm exceeds those from [20, 3, 10].

6 Evaluation

In this section we experimentally evaluate SMB studying its effectiveness to reduce a number of biochemical models from the literature. All experiments were performed on a 2.6 GHz Intel Core i5 with 4 GB of RAM, and are replicable using a prototypal tool available at <http://sysma.imtlucca.it/tools/erode/samba/>. The tool takes in input CRNs specified in the `.net` format of BioNetGen [4], version 2.2.5-stable. The CRN reduced by SMB is then converted back to the BioNetGen format to perform stochastic simulations using Gillespie’s stochastic simulation algorithm [28].

SMB was tested in terms of model reduction capabilities and corresponding CTMC analysis speed-ups on the collection of models listed in Table 1. Models with labels starting with “M” are the largest models also considered in [8]. Additionally, in this paper we analyze models S1–S3; M1–M3 and S1–S2 belong to the same family of synthetic benchmarks that are generated by varying the number of phosphorylation sites in a complex described in BioNetGen’s rule-based format [40]. S3 arises by studying ultrasensitivity in multisite proteins [21]. In all cases we applied our reduction technique starting from the trivial partition with one block only (i.e., $\{S\}$ for every CRN (S, R)).

Column “*Int.*” shows the units of time used for the simulation, taken from the respective papers. This information was missing for M1–M3 and S1–S2, for which we used an estimate of their steady states. Initial populations for the simulations were taken as well from the respective papers. Under “*Original model*” are listed the number of reactions ($|R|$) and species ($|S|$) of the CRN, and the overall time to perform 10 simulations. The same information is given in the columns under “*SMB reduction*” for the corresponding CRNs reduced up to SMB, providing in addition the time necessary to compute the SMB partition (*SMB* (s)) and to perform the reduction (*Red.* (s)).

The results indicate that SMB can find equivalences in a significant number of models concerning different biological mechanisms. In the three largest models, M1–M3, SMB was able to provide a compact aggregated CRN which could be straightforwardly analyzed, while the simulations of the original models did not terminate due to out-of-memory errors in our experimental set-up. This is consistent with [40], where the same issue was reported for model M1. Models with more sensible reductions in the number of reactions gave better speed-ups. For example, for S1–S2 and M6–M7, the reduced CRN could be analyzed in about one tenth of the time necessary for the original one. We attribute this to the fact that at every simulation step, Gillespie’s algorithm scans all reactions (in the worst case) to decide which one to fire next. Also, we note that typically many simulations, often in the order of hundreds or thousands, are necessary to satisfy a given confidence interval (or precision); hence, even small speed-ups per single run may turn into consistent gains in the overall simulation runtimes. Finally, as can be expected from their respective computational complexities, the runtimes to reduce a CRN by SMB according to Definition 7 are considerably smaller compared to the runtimes for computing the largest SMB.

Comparison with κ -based reduction techniques. SMB and stochastic fragmentation can be experimentally compared in rule-based biochemical models with finite underlying CRNs, like those in Table 1, where both techniques can be applied. In [8] we have shown that FB and *differential fragmentation*, a variant of fragmentation defined for the ODE semantics of κ , are not comparable. The same holds for SMB and stochastic fragmentation. For example, SMB reduces M12 from [8] to 56 species, while fragmentation does not. Conversely, the κ model of cross-talk between a model of the early events of the EGF pathway and the insulin receptor of [16] can be reduced by stochastic fragmentation, but not by SMB.

Comparison with Forward Bisimulation. We now relate SMB with FB, introduced in [8] for the ODE semantics of CRNs. For this, it is convenient to recall such semantics. The ODE system $\dot{V} = F(V)$ underlying a CRN (S, R) (where the dot notation indicates derivative with respect to time) is given by $F : \mathbb{R}_{\geq 0}^S \rightarrow \mathbb{R}^S$, where each component F_X , with $X \in S$ is defined by the expression

$$F_X(V) := \sum_{\rho \xrightarrow{\alpha} \pi \in R} (\pi(X) - \rho(X)) \cdot \alpha \cdot \prod_{Y \in S} V_Y^{\rho(Y)}. \quad (1)$$

This provides the well-known *mass-action* kinetics, where the reaction rate is proportional to the concentrations of the reactants involved.

Example 7. The ODE system associated to our running example is as follows.

$$\begin{aligned}\dot{V}_A &= -8 V_A & \dot{V}_C &= 6 V_A + 5 V_C V_D + 6 V_B + 10 V_D V_E + 3 V_D^2 & \dot{V}_E &= -5 V_D V_E \\ \dot{V}_B &= -8 V_B & \dot{V}_D &= 6 V_A + 6 V_B - 6 V_D^2\end{aligned}$$

As for SMB, FB is an equivalence over the species of a CRN computed by looking at the reactions only. Also, FB induces a notion of lumpability similar in spirit to that of Definition 3, as it allows to rewrite the ODEs underlying a CRN in terms of macrovariables that govern the evolution of the cumulative concentrations of the species of each block.

Example 8. Consider the partition $\{\{A, B\}, \{C, E\}, \{D\}\}$ for our running example. This can be shown to be an FB. Indeed, the ODEs of (S, R) can be rewritten, under the variable renaming $V_{AB} = V_A + V_B, V_{CE} = V_C + V_E$, as

$$\dot{V}_{AB} = -8 V_{AB} \quad \dot{V}_{CE} = 6 V_{AB} + 5 V_{CE} V_D + 3 V_D^2 \quad \dot{V}_D = 6 V_{AB} - 6 V_D^2$$

Although SMB and FB work on different semantics, the fact that they are both equivalences over species that induce analogous aggregations at the semantic level calls for the question of establishing a formal relation between the two equivalences.

Theorem 3. *Let (S, R) be a CRN, and \mathcal{R} an equivalence relation over S . Then, if \mathcal{R} is an SMB for (S, R) , it also is an FB for (S, R') , with*

$$\begin{aligned}R' &= \{X \xrightarrow{\alpha} \pi \mid (X \xrightarrow{\alpha} \pi) \in R\} \cup \\ &\quad \{X + Y \xrightarrow{\alpha} \pi \mid (X + Y \xrightarrow{\alpha} \pi) \in R \wedge X \neq Y\} \cup \\ &\quad \{X + X \xrightarrow{\alpha/2} \pi \mid (X + X \xrightarrow{\alpha} \pi) \in R\}\end{aligned}$$

When R has singleton products only, then \mathcal{R} is an SMB for (S, R) iff it is an FB for (S, R') .

An important remark to be made regarding this result is that it requires to halve the rates of homeoreactions. This is due to an inherent, well-known inconsistency existing between the CTMC and ODE semantics of CRNs. While, as discussed in Section 2, homeoreactions are treated specially in the CTMC semantics in order to capture the combinatorial nature of the discrete molecular interactions, the ODE semantics does not make such difference, e.g., [24], see also (1). We refer to [6] for a more in depth discussion on this. It is interesting to note that a *different* ODE semantics would be possible, grounded on a limit result by Kurtz which establishes the ODE solution as the asymptotic behavior of a sequence of infinitely large CTMCs induced by the same CRN with increasing volumes of a solution having given initial concentrations of species. This interpretation would lead to a $1/2$ coefficient in the rates of homeoreactions also in the ODE case [35]. We leave it for future work to understand if, by appropriately adapting FB to this different ODE semantics, Theorem 3 can be stated so as to relate SMB and FB on the same CRN also in presence of homeoreactions.

We also remark that the converse of the theorem does not hold in general for CRNs with non-singleton products: in our running example $\{\{A, B\}, \{C, E\}, \{D\}\}$ has been

discussed to be an FB. However, given that $\mathbf{rr}(A + \emptyset, D) = 6$ and $\mathbf{rr}(B + \emptyset, D) = 0$, it is not an SMB.

FB has been applied to the models of Table 1 in [8], where a biological interpretation of the obtained aggregations was also provided. Interestingly, the largest SMBs of Table 1 correspond to the largest FBs of [8] on these models. Since none of these models have homeoreactions, we conclude that for these CRNs FB has the same discriminating power as SMB. Since S3 has unary products only, in that case this is guaranteed by Theorem 3.

7 Conclusion

Syntactic Markovian bisimulation (SMB) is an equivalence relation operating at the syntactic level of a chemical reaction network that induces a reduced one in the sense of the theory of Markov chain lumpability. A numerical evaluation has demonstrated its usefulness in practice by showing significant reductions in a number of models available in the literature, even if SMB is only a sufficient condition for aggregation. A partition-refinement algorithm computes the largest SMB that refines a given input partition. The freedom in choosing such input may be exploited to single out certain *observable species*. Thus, SMB may give a reduced model that exactly preserves the dynamics of interest. Since the CRN syntax is often combinatorially smaller than the underlying CTMC, we envisage SMB to be used as a pre-processing stage for CTMC analyses or for further reduction techniques on the semantics, either in exact or approximate form (e.g., [1]). Indeed, it would be interesting to conduct further experiments in order to understand how tight the lumping is with respect to the coarsest one obtained by applying Markov chain minimization algorithms on the fully enumerated state space; for this, we plan to develop fully integrated support for SMB into our software tool for model reduction techniques, *ERODE* [12] (<http://sysma.imtlucca.it/tools/erode/>). We have shown that SMB is stricter than forward bisimulation (FB), a recently introduced bisimulation for the ODE semantics of CRNs. In related line of research, we have developed behavioural equivalences for the ODE semantics of process algebra [44, 46, 47, 34]. In the future, we plan to investigate variants which imply lumpability for the CTMC semantics underlying stochastic process algebra.

Acknowledgments. This work was partially supported by the EU project QUANTICOL, 600708. L. Cardelli is partially funded by a Royal Society Research Professorship.

References

1. A. Abate, L. Brim, M. Ceska, and M.Z. Kwiatkowska. Adaptive aggregation of markov chains: Quantitative analysis of chemical reaction networks. In *CAV*, pages 195–213, 2015.
2. C. Autant and Ph. Schnoebelen. Place bisimulations in Petri nets. In *Application and Theory of Petri Nets*, 1992.
3. 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.
4. 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.

5. P. Buchholz. Exact and Ordinary Lumpability in Finite Markov Chains. *Journal of Applied Probability*, 31(1):59–75, 1994.
6. L. Cardelli. On process rate semantics. *Theor. Comput. Sci.*, 391(3):190–215, 2008.
7. L. Cardelli. Morphisms of reaction networks that couple structure to function. *BMC Systems Biology*, 8(1):84, 2014.
8. L. Cardelli, M. Tribastone, M. Tschaikowski, and A. Vandin. Forward and backward bisimulations for chemical reaction networks. In *CONCUR*, pages 226–239, 2015.
9. L. Cardelli, M. Tribastone, M. Tschaikowski, and A. Vandin. Comparing chemical reaction networks: A categorical and algorithmic perspective. In *LICS*, pages 485–494, 2016.
10. L. Cardelli, M. Tribastone, M. Tschaikowski, and A. Vandin. Efficient syntax-driven lumping of differential equations. In *TACAS*, pages 93–111, 2016.
11. L. Cardelli, M. Tribastone, M. Tschaikowski, and A. Vandin. Symbolic computation of differential equivalences. In *POPL*, pages 137–150, 2016.
12. L. Cardelli, M. Tribastone, M. Tschaikowski, and A. Vandin. ERODE: a tool for the evaluation and reduction of ordinary differential equations. In *TACAS*, pages 310–328, 2017.
13. L. Cardelli and G. Zavattaro. Turing universality of the Biochemical Ground Form. *Mathematical Structures in Computer Science*, 20(1):45–73, 2010.
14. Luca Cardelli, Attila Csikász-Nagy, Neil Dalchau, Mirco Tribastone, and Max Tschaikowski. Noise reduction in complex biological switches. *Scientific Reports*, 6:20214, 2016.
15. F. Ciocchetta and J. Hillston. Bio-PEPA: A framework for the modelling and analysis of biological systems. *Theor. Comput. Sci.*, 410(33-34):3065–3084, 2009.
16. H. Conzelmann, D. Fey, and E. D. Gilles. Exact model reduction of combinatorial reaction networks. *BMC Systems Biology*, 2:78, 2008.
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*, 2010.
18. V. Danos, J. Feret, W. Fontana, and J. Krivine. Scalable simulation of cellular signaling networks. In *APLAS*, 2007.
19. V. Danos and C. Laneve. Formal molecular biology. *TCS*, 325(1):69–110, 2004.
20. S. Derisavi, H. Hermanns, and W.H. Sanders. Optimal state-space lumping in Markov chains. *Inf. Process. Lett.*, 87(6):309–315, 2003.
21. O. Dushek, P.A. van der Merwe, and V. Shahrezaei. Ultrasensitivity in multisite phosphorylation of membrane-anchored proteins. *Biophysical Journal*, 100(5):1189–1197, 2011.
22. M. B. Elowitz, A. J. Levine, E. D. Siggia, and P. S. Swain. Stochastic gene expression in a single cell. *Science*, 297(5584):1183–1186, 2002.
23. 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.
24. M. Feinberg. Lectures on chemical reaction networks. Technical report, University of Wisconsin, 1979.
25. J. Feret, T. Henzinger, H. Koepl, and T. Petrov. Lumpability abstractions of rule-based systems. *TCS*, 431:137–164, 2012.
26. S. Gay, S. Soliman, and F. Fages. A graphical method for reducing and relating models in systems biology. *Bioinformatics*, 26(18), 2010.
27. D. Gillespie. The chemical Langevin equation. *The Journal of Chemical Physics*, 113(1):297–306, 2000.
28. D.T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *Journal of Physical Chemistry*, 81(25):2340–2361, December 1977.
29. S. Gilmore, J. Hillston, and M. Ribaud. An efficient algorithm for aggregating PEPA models. *IEEE Transactions on Software Engineering*, 27(5):449–464, May 2001.
30. J. Heath, M. Kwiatkowska, G. Norman, D. Parker, and O. Tymchyshyn. Probabilistic model checking of complex biological pathways. *TCS*, 391(3):239–257, 2008.

31. H. Hermanns and M. Ribardo. Exploiting symmetries in stochastic process algebras. In *European Simulation Multiconference*, pages 763–770, Manchester, UK, June 1998.
32. J. Hillston. *A Compositional Approach to Performance Modelling*. CUP, 1996.
33. D. T. Huynh and L. Tian. On some equivalence relations for probabilistic processes. *Fundam. Inform.*, 17(3):211–234, 1992.
34. Giulio Iacobelli, Mirco Tribastone, and Andrea Vandin. Differential bisimulation for a Markovian process algebra. In *MFCS*, pages 293–306, 2015.
35. T. G. Kurtz. The Relationship between Stochastic and Deterministic Models for Chemical Reactions. *Journal of Chemical Physics*, 57(7), 1972.
36. M. Kwiatkowska, G. Norman, and D. Parker. Symmetry reduction for probabilistic model checking. In *CAV*, pages 234–248, Berlin, Heidelberg, 2006. Springer-Verlag.
37. K. G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94(1):1–28, 1991.
38. R. Paige and R. Tarjan. Three partition refinement algorithms. *SIAM Journal on Computing*, 16(6):973–989, 1987.
39. Z. Rózsa and J. Tóth. Exact linear lumping in abstract spaces. *Electronic Journal of Qualitative Theory of Differential Equations [electronic only]*, 2003.
40. 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.
41. D. Soloveichik, M. Cook, E. Winfree, and J. Bruck. Computation with finite stochastic chemical reaction networks. *Natural Computing*, 7(4):615–633, 2008.
42. 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.
43. S. Tognazzi, M. Tribastone, M. Tschaikowski, and A. Vandin. EGAC: A Genetic Algorithm to Compare Chemical Reaction Networks. In *The Genetic and Evolutionary Computation Conference (GECCO)*, 2017. To appear.
44. M. Tschaikowski and M. Tribastone. Exact fluid lumpability for Markovian process algebra. In *23rd International Conference on Concurrency Theory (CONCUR)*, LNCS, pages 380–394, 2012.
45. M. Tschaikowski and M. Tribastone. Generalised Communication for Interacting Agents. In *QEST*, pages 178–188, September 2012.
46. M. Tschaikowski and M. Tribastone. Tackling continuous state-space explosion in a markovian process algebra. *Theor. Comput. Sci.*, 517:1–33, 2014.
47. M. Tschaikowski and M. Tribastone. A unified framework for differential aggregations in Markovian process algebra. *JLAMP*, 84(2):238–258, 2015.
48. M. Tschaikowski and M. Tribastone. Approximate Reduction of Heterogenous Nonlinear Models With Differential Hulls. *IEEE Trans. Automat. Contr.*, 61(4):1099–1104, 2016.
49. A. Vandin and M. Tribastone. Quantitative abstractions for collective adaptive systems. In *SFM 2016, Bertinoro Summer School*, pages 202–232, 2016.
50. G. Zavattaro and L. Cardelli. Termination problems in chemical kinetics. In *CONCUR*, pages 477–491, 2008.