

Lumpability Abstractions of Rule-based Systems*

Jerome Feret

LIENS (INRIA/ÉNS/CNRS)
Paris, France
feret@ens.fr

Thomas Henzinger

Institute and Science of Technology
Vienna, Austria
thenzinger@ist.ac.at

Heinz Koepl

School of Computer and Communication Sciences
EPFL
Lausanne, Switzerland
heinz.koepl@epfl.ch

Tatjana Petrov

School of Computer and Communication Sciences
EPFL
Lausanne, Switzerland
tatjana.petrov@epfl.ch

The induction of a signaling pathway is characterized by transient complex formation and mutual posttranslational modification of proteins. To faithfully capture this combinatorial process in a mathematical model is an important challenge in systems biology. Exploiting the limited context on which most binding and modification events are conditioned, attempts have been made to reduce the combinatorial complexity by quotienting the reachable set of molecular species, into species aggregates while preserving the deterministic semantics of the thermodynamic limit. Recently we proposed a quotienting that also preserves the stochastic semantics and that is complete in the sense that the semantics of individual species can be recovered from the aggregate semantics. In this paper we prove that this quotienting yields a sufficient condition for *weak lumpability* and that it gives rise to a *backward Markov bisimulation* between the original and aggregated transition system. We illustrate the framework on a case study of the EGF/insulin receptor crosstalk.

1 Introduction

Often a few elementary events of binding and covalent modification [31] in a biomolecular reaction system give rise to a combinatorial number of non-isomorphic reachable species or complexes [17, 18]. Instances of such systems are signaling pathway, polymerizations involved in cytoskeleton maintenance, the formation of transcription factor complexes in gene-regulation.

For such biomolecular systems, traditional chemical kinetics face fundamental limitations, that are related to the question of how biomolecular events are represented and translated into a mathematical model [24]. More specifically, chemical reactions can only operate on a collection of fully specified molecular species and each such species gives rise to one differential equation, describing the rate of change of that species' concentration. Many combinatorial systems do not permit the enumeration of all molecular species and thus render their traditional differential description prohibitive. However, even if one could enumerate them, it remains questionable whether chemical reactions are the appropriate way to represent and to reason about such systems.

As the dynamics of a biomolecular reaction mixture comes about through the repeated execution of a few elementary events one may wonder about the effective degrees of freedom of the reaction mixture's dynamics. If the velocity of all events – or their probabilities to occur per time-unit per instance –

*Jérôme Feret's contribution was partially supported by the ABSTRACTCELL ANR-Chair of Excellence. Heinz Koepl acknowledges the support from the Swiss National Science Foundation, grant no. 200020-117975/1. Tatjana Petrov acknowledges the support from SystemsX.ch, the Swiss Initiative in Systems Biology.

are different for all complexes (w.r.t. modification) and pairs of complexes (w.r.t. binding) to which the events can apply to, then the degrees of freedom would equal to the number of molecular species. However, due to the local nature of physical forces underlying molecular dynamics, the kinetics of most events appear to be ignorant with respect to the global configuration of the complexes they are operating on. More provocatively, one may say that even if there would be variations of kinetics of an event from one context to another, experimental biology does not – and most likely never will – have the means to discern between all different contexts. For instance, fluorescence resonance energy transfer (FRET), may report on a specific protein-binding event and even its velocity, however we have no means to determine whether the binding partners are already part of a protein complex – not to speak of the composition and modification state of these complexes. To this end, molecular species remain elusive and appear to be inappropriate entities of descriptions.

To align with the mentioned experimental insufficiencies and with the underlying biophysical locality, rule-based or agent-based descriptions were introduced as a framework to encode such reaction mixtures succinctly and to enable their mathematical analysis [8, 3]. Rules exploit the limited context on which most elementary events are conditioned. They just enumerate that part of a molecular species that is relevant for a rule to be applicable. Thus, in contrast to chemical reactions, rules can operate on a collection of partially specified molecular species. Recently, attempts have been made to identify the set of those partially specified species that allow for a self-consistent description of the rule-set's dynamics [11, 6]. Naturally, as partially specified species – or *fragments* – in general encompass many fully specified species, the cardinality of that set is less than of the set of molecular species. These approaches aim to obtain a self-consistent fragment dynamics based on ordinary differential equations. It represents the dynamics in the thermodynamic limit of stochastic kinetics when scaling species multiplicities to infinity while maintaining a constant concentration (multiplicity per unit volume) [22]. In many applications in cellular biology this limiting dynamics is an inappropriate model due to the low multiplicities of some molecular species – think of transcription factor - DNA binding events. We recently showed that the obtained *differential fragments* cannot be used to describe the finite volume case of stochastic kinetics [12]. Exploiting statistical independence of events we instead derived *stochastic fragments* that represent the effective degrees of freedom in the stochastic case. Conceptually, the procedure replaces the Cartesian product by a Cartesian sum for statistically independent states. In contrast to the differential case, stochastic fragments have the important property that the sample paths of molecular species can be recovered from that of partially specified species.

We believe that interdisciplinary fields, such as systems biology, can move forward quickly enough only if the well-established knowledge in each of the disciplines involved is exploited to its maximum, i.e. if the standardized, well-established theories are recognized and re-used. For that reason, in this paper we translate our abstraction method ([12]) into the language of well-established contexts of abstraction for probabilistic systems – lumpability and bisimulation. Lumpability is mostly considered from a theoretical point of view in the theory of stochastic processes [19, 13, 28, 26, 27, 4]. A Markov chain is lumpable with respect to a given aggregation (quotienting) of its states, if the lumped chain preserves the Markov property [21]. A sound aggregation for *any* initial probability distribution is referred to as *strong* lumpability, while otherwise it is termed *weak* lumpability [4, 29]. Approximate aggregation techniques for Markov chains of biochemical networks are discussed in [14]. Probabilistic bisimulation was introduced as an extension to classic bisimulation in [23]. It is extended to continuous-state and continuous-time in [9] and, for the discrete-state case, to weak bisimulation [2]. For instance, in [9] the authors use bisimulation of labelled Markov processes, the state space of which is not necessarily discrete, and they provide a logical characterization of probabilistic bisimulation. Another notion of weak bisimulation was recently introduced in [10]. Therein two labeled Markov chains are defined to

be equivalent if every finite sequence of observations has the same probability of occurring in the two chains. Herein we recognize the sound aggregations of [12] as a form of *backward Markov bisimulations* on weighted labeled transition systems (WLTS), and we show it to be equivalent to the notion of *weak lumpability* on Markov chains.

The remaining part of the paper is organized as follows. In the Section 2, we introduce weighted labeled transition systems (WLTS) and we assign it the trace density semantics of a continuous-time Markov chain (CTMC). Moreover, we define the Kappa language, and we assign a WLTS to a Kappa specification. Based on the notion of the annotated contact map we briefly summarize in Section 3 the general procedure to compute stochastic fragments, as it is offered in [12]. In Section 4, we introduce the characterizations of sound and complete abstractions on WLTS as a backward Markov bisimulation. Moreover, we define it being equivalent to the weak lumpability on Markov chains. Finally, we provide in Section 5 results for the achieved dimensionality reduction for a rule-based model of the crosstalk between the EGF/insulin signaling pathway [6]. This mechanistic model comprises 76 rules giving rise to 42956 reactions and 2768 molecular species.

2 Preliminaries

The stochastic semantics of a biochemical network is modelled by a continuous-time Markov chain (CTMC). The main object that we will use in the analysis is the weighted labelled transition system (WLTS) on a countable state space. We will assign a WLTS to a given Kappa specification, and we manipulate that object when reasoning about abstractions.

2.1 CTMC

We will observe the CTMC that is generated by the weighted labelled transition system (WLTS) on a countable state space. We define the CTMC of a WLTS, by defining the Borel σ -algebra containing all cylinder sets of traces [20] that can occur in the system, and the corresponding probability distribution among them. We also introduce the standard notation of a rate matrix, which we will use when analysing the lumpability and bisimulation properties in Sec. 4.

Definition 1. (WLTS) A weighted-labelled transition system \mathcal{W} is a tuple $(\mathcal{X}, \mathcal{L}, w, \pi_0)$, where

- \mathcal{X} is a countable state space;
- \mathcal{L} is a set of labels;
- $w : \mathcal{X} \times \mathcal{L} \times \mathcal{X} \rightarrow \mathbb{R}_0^+$ is the weighting function that maps two states and a label to a real value;
- $\pi_0 : \mathcal{X} \rightarrow [0, 1]$ is an initial probability distribution.

We assume that the label fully identifies the transition, i.e. for any $x \in \mathcal{X}$ and $l \in \mathcal{L}$, there is at most one $x' \in \mathcal{X}$, such that $w(x, l, x') > 0$. Moreover, we assume that the system is finitely branching, in the sense that (i) the set $\{x \in \mathcal{X} \mid \pi_0(x) > 0\}$ is finite, and (ii) for arbitrary $\hat{x} \in \mathcal{X}$, the set $\{(l, x') \in \mathcal{L} \times \mathcal{X} \mid w(\hat{x}, l, x') > 0\}$ is finite.

The activity of the state x_i , denoted $a : \mathcal{X} \rightarrow \mathbb{R}_0^+$ is the sum of all weights originating at x_i , i.e.

$$a(x_i) := \sum \{w(x_i, l, x_j) \mid x_j \in \mathcal{X}, l \in \mathcal{L}\}.$$

The definition of a WLTS implicitly defines a *transition relation* $\rightarrow \subseteq \mathcal{X} \times \mathcal{X}$, such that $(x_i, x_j) \in \rightarrow$, if and only if there exists a non-zero transition from state x_i to state x_j , i.e. the total weight over all labels is strictly bigger than zero, written $\sum \{w(x_i, l, x_j) \mid l \in \mathcal{L}\} > 0$. Moreover, we can differentiate the *initial set of states* $\mathcal{I} \subseteq \mathcal{X}$, such that their initial probabilities are positive, i.e. $\mathcal{I} = \{x \in \mathcal{X} \mid \pi_0(x) > 0\}$.

Definition 2. (Rate matrix of a WLTS) Given a WLTS $\mathcal{W} = (\mathcal{X}, \mathcal{L}, w, \pi_0)$, we assign it the CTMC rate matrix $R : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, given by $R(x_i, x_j) = \sum \{w(x_i, l, x_j) \mid l \in \mathcal{L}\}$.

The consequence is that we do not enforce $R(x_i, x_i) = -\sum \{R(x_i, x_j) \mid i \neq j\}$, as it is usual for the generator matrix of CTMCs. This however does not affect the transient, not the steady-state behavior of the CTMC [1]. We do so for the following reason. When abstracting the WLTS by partitioning the state space, we get another WLTS. If the two states x and x' which have a transition between each other were aggregated in the same partition class \tilde{x} , it will result as a prolongation of the residence time in the abstract state \tilde{x} , i.e. we will have a self-loop in the abstract WLTS.

If we refer to the generated stochastic Markov process, written as a continuous-time random variable $\{X_t\}_{t \in \mathbb{R}_0^+}$, over the countable state space \mathcal{X} . We write $\Pr(X_t = x_i)$, the probability that the process takes the value x_i at time point t . It thus holds that $\Pr(X_0 = x_i) = \pi_0(x_i)$, and $\Pr(X_{t+dt} = x_j \mid X_t = x_i) = R(x_i, x_j)dt$ when $i \neq j$, whereas $\Pr(X_{t+dt} = x_i \mid X_t = x_i) = R(x_i, x_i)dt + (1 - \sum \{R(x_i, x_{j'})dt \mid x_{j'} \in \mathcal{X}\})$, which gives after simplification: $\Pr(X_{t+dt} = x_i \mid X_t = x_i) = 1 - \sum \{R(x_i, x_{j'})dt \mid j' \neq i\}$.

We define the cylinder sets of traces that can be observed in the system \mathcal{W} . By observing the trace at a certain time point, we mean observing the sequence of visited states, labels that were assigned to the executed transitions, and time points of when the transition happened.

Definition 3. (A trace of a WLTS) Let us observe the WLTS $\mathcal{W} = (\mathcal{X}, \mathcal{L}, w, \pi_0)$ and its CTMC. Given a number k in \mathbb{N} , we define a trace of length k as $\tau \in (\mathcal{X} \times \mathcal{L} \times \mathbb{R}_0^+)^k \times \mathcal{X}$, written

$$\tau = x_0 \xrightarrow{l_1, t_1} x_1 \dots x_{k-1} \xrightarrow{l_k, t_1 + \dots + t_k} x_k.$$

If the trace τ is such that (i) $\pi_0(x_0) > 0$, and (ii) for all i , $0 \leq i \leq k$, we have that $w(x_i, l_i, x_{i+1}) > 0$, then we say that τ belongs to the set of traces of \mathcal{W} , and we write it $\tau \in \mathcal{T}(\mathcal{W})$.

The ‘time stamps’ on each of the transitions denote intuitively the absolute time of the transition, from the moment when the system was started ($t = 0$). We cannot assign the probability distribution to the traces in $\mathcal{T}(\mathcal{W})$, since the probability of any such trace is zero. We thus introduce the cylinder set of traces over intervals of times.

Definition 4. (Cylinder set of traces) If $\mathbb{I}\mathbb{R}$ is the set of all nonempty intervals in \mathbb{R}_0^+ , then we define the cylinder set of traces $\tau_{\mathbb{I}\mathbb{R}} \in (\mathcal{X} \times \mathcal{L} \times \mathbb{I}\mathbb{R})^k \times \mathcal{X}$, such that:

$$\tau_{\mathbb{I}\mathbb{R}} = x_0 \xrightarrow{l_1, I_1} x_1 \dots x_{k-1} \xrightarrow{l_k, I_k} x_k \tag{1}$$

denotes the set of all traces $\tau = x_0 \xrightarrow{l_1, t_1} x_1 \dots x_{k-1} \xrightarrow{l_k, t_1 + \dots + t_k} x_k$, such that $t_i \in I_i$, $1 \leq i \leq k$. If the cylinder of traces $\tau_{\mathbb{I}\mathbb{R}}$ is such that $\pi_0(x_0) > 0$, and for all $i = 0, \dots, k-1$, we have that $w(x_i, l_i, x_{i+1}) > 0$, then we say that $\tau_{\mathbb{I}\mathbb{R}}$ belongs to the cylinder set of traces of \mathcal{W} , and we write $\tau_{\mathbb{I}\mathbb{R}} \in \mathcal{T}_{\mathbb{I}\mathbb{R}}(\mathcal{W})$.

Let $\Omega(\mathcal{T}_{\mathbb{I}\mathbb{R}}(\mathcal{W}))$ be the smallest Borel σ -algebra that contains all the cylinder sets of traces in $\mathcal{T}_{\mathbb{I}\mathbb{R}}(\mathcal{W})$. We define a probability measure over $\Omega(\mathcal{T}_{\mathbb{I}\mathbb{R}}(\mathcal{W}))$ in the following way.

Definition 5. (Trace density semantics on a WLTS) Given a WLTS $(\mathcal{X}, \mathcal{L}, w, \pi_0)$, and a number k in \mathbb{N} , the probability of the cylinder set of traces $\tau_{\mathbb{I}\mathbb{R}} \in \mathcal{T}_{\mathbb{I}\mathbb{R}}(\mathcal{W})$, specified as in expression (1), is given by:

$$\pi(\tau_{\mathbb{I}\mathbb{R}}) = \pi(x_0 \xrightarrow{l_1, I_1} x_1 \dots x_{k-1} \xrightarrow{l_k, I_k} x_k) = \pi_0(x_0) \prod_{i=1}^k \frac{w(x_{i-1}, l_i, x_i)}{a(x_{i-1})} \left(e^{-a(x_{i-1}) \cdot \inf(I_i)} - e^{-a(x_{i-1}) \cdot \sup(I_i)} \right).$$

Note that $\int_{I_i} a(x_{i-1}) e^{-a(x_{i-1}) \cdot t} dt = e^{-a(x_{i-1}) \cdot \inf(I_i)} - e^{-a(x_{i-1}) \cdot \sup(I_i)}$ is the probability of exiting the state x_{i-1} in a time interval I_{i-1} , since the probability density function of the residence time of x_{i-1} is equal to $a(x_{i-1}) e^{-a(x_{i-1}) \cdot t}$.

2.2 Kappa

We present Kappa in a process-like notation. We start with an operational semantics, then define the stochastic semantics of a Kappa model.

We assume a finite set of agent names \mathcal{A} , representing different kinds of proteins; a finite set of sites \mathcal{S} , corresponding to protein domains; a finite set of internal states \mathbb{I} , and $\Sigma_{\iota}, \Sigma_{\lambda}$ two signature maps from \mathcal{A} to $\wp(\mathcal{S})$, listing the domains of a protein which can bears respectively an internal state and a binding state. We denote by Σ the signature map that associates to each agent name $A \in \mathcal{A}$ the combined interface $\Sigma_{\iota}(A) \cup \Sigma_{\lambda}(A)$.

Definition 6. (*Kappa agent*) A Kappa agent $A(\sigma)$ is defined by its type $A \in \mathcal{A}$ and its interface σ . In $A(\sigma)$, the interface σ is a sequence of sites s in $\Sigma(A)$, with internal states (as subscript) and binding states (as superscript). The internal state of the site s may be written as s_{ε} , which means that either it does not have internal states (when $s \in \Sigma(A) \setminus \Sigma_{\iota}(A)$), or it is not specified. A site that bears an internal state $m \in \mathbb{I}$ is written s_m (in such a case $s \in \Sigma_{\iota}(A)$). The binding states of a site s can be specified as s^{ε} , if it is free, otherwise it is bound (which is possible only when $s \in \Sigma_{\lambda}(A)$). There are several levels of information about the binding partner: we use a binding label $i \in \mathbb{N}$ when we know the binding partner, or a wildcard bond $-$ when we only know that the site is bound. The detailed description of the syntax of a Kappa agent is given by the following grammar:

$$\begin{array}{llll}
 a & ::= & N(\sigma) & \text{(agent)} \\
 N & ::= & A \in \mathcal{A} & \text{(agent name)} \\
 \sigma & ::= & \varepsilon \mid s, \sigma & \text{(interface)} \\
 s & ::= & n_i^{\lambda} & \text{(site)} \\
 n & ::= & x \in \mathcal{S} & \text{(site name)} \\
 \iota & ::= & \varepsilon \mid m \in \mathbb{I} & \text{(internal state)} \\
 \lambda & ::= & \varepsilon \mid - \mid i \in \mathbb{N} & \text{(binding state)}
 \end{array}$$

We generally omit the symbol ε .

Definition 7. (*Kappa expression*) Kappa expression E is a set of agents $A(\sigma)$ and fictitious agents \emptyset . Thus the syntax of a Kappa expression is defined as follows:

$$E ::= \varepsilon \mid a, E \mid \emptyset, E.$$

The structural equivalence \equiv , defined as the smallest binary equivalence relation between expressions that satisfies the rules given as follows:

$$\begin{array}{l}
 E, A(\sigma, s, s', \sigma'), E' \equiv E, A(\sigma, s', s, \sigma'), E' \\
 E, a, a', E' \equiv E, a', a, E' \\
 E \equiv E, \emptyset \\
 \frac{i, j \in \mathbb{N} \text{ and } i \text{ does not occur in } E}{E[i/j] \equiv E} \\
 \frac{i \in \mathbb{N} \text{ and } i \text{ occurs only once in } E}{E[\varepsilon/i] \equiv E}
 \end{array}$$

stipulates that neither the order of sites in interfaces nor the order of agents in expressions matters, that a fictitious agent might as well not be there, that binding labels can be injectively renamed and that dangling bonds can be removed.

Definition 8. (*Kappa pattern, Kappa mixture*) A Kappa pattern is a Kappa expression which satisfies the following five conditions: (i) no site name occurs more than once in a given interface; (ii) each site name s in the interface of the agent A occurs in $\Sigma(A)$; (iii) each site s which occurs in the interface of the agent A with a non empty internal state occurs in $\Sigma_{\iota}(A)$; (iv) each site s which occurs in the interface of the agent A with a non empty binding state occurs in $\Sigma_{\lambda}(A)$; and (v) each binding label $i \in \mathbb{N}$ occurs exactly twice if it does at all –there are no dangling bonds. A mixture is a pattern that is fully specified, i.e. each agent A documents its full interface $\Sigma(A)$, a site can only be free or tagged with a binding label $i \in \mathbb{N}$, a site in $\Sigma_{\iota}(A)$ bears an internal state in \mathbb{I} , and no fictitious agent occurs.

Definition 9. (*Kappa rule*) A Kappa rule r is defined by two Kappa patterns E_ℓ and E_r , and a rate $k \in \mathbb{R}_0^+$, and is written: $r = E_\ell \rightarrow E_r @ k$.

A rule r is well-defined, if the expression E_r is obtained from E_ℓ by finite application of the following operations: (i) creation (some fictitious agents \emptyset are replaced with some fully defined agents of the form $A(\sigma)$, moreover σ documents all the sites occurring in $\Sigma(A)$ and all site in $\Sigma_i(A)$ bears an internal state in \mathbb{I}), (ii) unbinding (some occurrences of the wild card and binding labels are removed), (iii) deletion (some agents with only free sites are replaced with fictitious agent \emptyset), (iv) modification (some non-empty internal states are replaced with some non-empty internal states), (v) binding (some free sites are bound pair-wise by using binding labels in \mathbb{N}).

From now on, we assume all rules to be well-defined. We sometimes omit the rate of a rule. Moreover, we denote by $E_\ell \leftrightarrow E_r @ k_1, k_2$ the two rules $E_\ell \leftrightarrow E_r \rightarrow @ k_1$ and $E_r \rightarrow E_\ell @ k_2$.

Definition 10. (*Kappa system*) A Kappa system $\mathcal{R} = (\pi_0^{\mathcal{R}}, \{r_1, \dots, r_n\})$ is given by finite distribution over initial mixtures $\pi_0^{\mathcal{R}} : \{M_{0_1}, \dots, M_{0_k}\} \rightarrow [0, 1]$, and a finite set of rules $\{r_1, \dots, r_n\}$.

In order to apply a rule $r := E_\ell \rightarrow E_r @ k$ to a mixture M , we use the structural equivalence \equiv to bring the participating agents to the front of E (with their sites in the same order as in E_ℓ), rename binding labels if necessary and introduce a fictitious agent for each agent that is created by r . This yields an equivalent expression E' that *matches* the left and side (lhs) E_ℓ , which is written $E \models E_\ell$ as defined as follows:

$$\begin{array}{l} E \models \varepsilon \\ a \models a_\ell \wedge E \models E_\ell \implies a, E \models a_\ell, E_\ell \\ \emptyset \models \emptyset \end{array} \quad \begin{array}{l} \sigma \models \sigma_\ell \implies N(\sigma) \models N(\sigma_\ell) \\ \sigma \models \varepsilon \\ s \models s_\ell \wedge \sigma \models \sigma_\ell \implies s, \sigma \models s_\ell, \sigma_\ell \end{array} \quad \begin{array}{l} \iota \models \iota_\ell \wedge \lambda \models \lambda_\ell \implies n_\iota^\lambda \models n_{\iota_\ell}^{\lambda_\ell} \\ \iota_\ell \in \{\varepsilon, \iota\} \implies \iota \models \iota_\ell \\ \lambda_\ell \in \{-, \lambda\} \implies \lambda \models \lambda_\ell \end{array}$$

Note that in order to find a matching, we only uses structural equivalence on E , not E_ℓ . We then replace E' by $E'[E_r]$ which is defined as follows:

$$\begin{array}{l} E[\varepsilon] = E \\ (a, E)[a_r, E_r] = a[a_r], E[E_r] \\ \emptyset[a_r] = a_r \\ a_r[\emptyset] = \emptyset \end{array} \quad \begin{array}{l} N(\sigma)[N(\sigma_r)] = N(\sigma[\sigma_r]) \\ \sigma[\varepsilon] = \sigma \\ (s, \sigma)[s_r, \sigma_r] = s[s_r], \sigma[\sigma_r] \end{array} \quad \begin{array}{l} n_\iota^\lambda[n_{\iota_r}^{\lambda_r}] = n_{\iota[\iota_r]}^{\lambda[\lambda_r]} \\ \iota_r \in \mathbb{I} \implies \iota[\iota_r] = \iota_r \\ \lambda_r \in \mathbb{N} \cup \{\varepsilon\} \implies \lambda[\lambda_r] = \lambda_r \\ \lambda[-] = \lambda \end{array}$$

This may produce dangling bonds (if r unbinds a wildcard bond or destroys an agent on one side of a bond) or fictitious agents (if r destroys agents), so we use \equiv resolve them.

2.2.1 Population-based stochastic semantics

In addition to the rate constants k , careful counting of the number of times each rule can be applied to a mixture is required to define the system's quantitative semantics correctly [7, 25]. Thus we define the notions of embedding between a mixture and an expression. Let $Z = a_1, \dots, a_m$ and $Z_\ell = c_1, \dots, c_n$ be two patterns with no occurrence of the fictitious agent and such that there exists a pattern $Z' = b_1, \dots, b_m$ that satisfies both $Z \equiv Z'$ and $Z' \models Z_\ell$ (and so, in particular, $n \leq m$).

The agent permutations used in the proof that $Z \equiv Z'$ allow us to derive a permutation p such that $a_{p(i)} \equiv b_i$. The restriction ϕ of p to the integers between 1 and n is called an *embedding* between Z_ℓ and Z . This is written $Z_\ell \triangleleft_\phi Z$. There may be several embeddings between Z_ℓ and Z for the same Z' ; if so, this influences the relative weight of the reaction in the stochastic semantics. We denote by $[Z, Z']$ the set

of embeddings between Z and Z' . This notion of embedding is extended to patterns (including fictitious agents) by defining $Z_\ell \triangleleft_\phi Z$ if, and only if, $(\downarrow_\emptyset Z_\ell) \triangleleft_\phi (\downarrow_\emptyset Z)$, where \downarrow_\emptyset removes all occurrences of the fictitious agent in patterns.

We assume that E_ℓ is the lhs of a rule $r := E_\ell \rightarrow E_r @ k$ and Z is a mixture such that $E_\ell \triangleleft_\phi Z$. Let $Z = a_1, \dots, a_m$ and $\downarrow_\emptyset E_\ell = c_1, \dots, c_n$. Given $Z' \equiv Z$ (we write $\downarrow_\emptyset Z' = b_1, \dots, b_m$) and a bijection p such that we have $Z' \models E_\ell$, $b_i \equiv a_{p(i)}$ for $1 \leq i \leq m$ and $\phi(j) = p(j)$ for $1 \leq j \leq n$. The result of applying r along ϕ to the mixture Z is defined (modulo \equiv) as any pattern that is \equiv -equivalent to $Z'[E_r]$. In other words the embedding ϕ between E_ℓ and Z fully defines the action of r on Z up to structural equivalence.

We are now ready to define the stochastic semantics by the mean of a WLTS. In this semantics, the state is a soup of agents, that is to say that we do not care about the order of agents in mixture. So the states of the system are the class of \equiv -mixture.

Defining species as connected mixture, the state of the system can be seen as a multi-set of species. The formal definition of a Kappa species is as follows:

Definition 11. (*Kappa species*) A pattern E is reducible whenever $E \equiv E', E''$ for some non-empty patterns E', E'' ; A Kappa species is the \equiv -equivalence class of a irreducible Kappa mixture.

As explained earlier, the action of a rule r on a mixture E is fully defined (up to \equiv) by an embedding ϕ between the lhs E_ℓ of the rule r and the mixture. So as to consider computation steps over \equiv -equivalent of mixtures, we introduce an equivalence relation $\equiv_{\mathcal{L}}$ over triples (r, E, ϕ) where ϕ is an embedding of the lhs E_ℓ of r into E . We say that $(r_1, E_1, \phi_1) \equiv_{\mathcal{L}} (r_2, E_2, \phi_2)$ if, and only if, (i) $r_1 = r_2$ and (ii) there exists an embedding $\psi \in [E_1, E_2]$ such that $\phi_2 = \psi \circ \phi_1$.

Definition 12. (*WLTS of a Kappa system*) Let $\mathcal{R} = (\pi_0^{\mathcal{R}}, \{r_1, \dots, r_n\})$ be a Kappa system. We define the WLTS $\mathcal{W}_{\mathcal{R}} = (\mathcal{X}, \mathcal{L}, w, \pi_0)$ where: (i) \mathcal{X} is the set of all \equiv -equivalent classes of mixtures; (ii) \mathcal{L} is the set of all $\equiv_{\mathcal{L}}$ -equivalence classes of triples (r, E, ϕ) such that ϕ is an embedding between the lhs E_ℓ of r and E ; (iii) $w(x, l, x') = \frac{k}{|[E_\ell, E]|}$ whenever there exist a rule $r = E_\ell \rightarrow E_r @ k$, two mixtures E and E' , and an embedding $\phi \in [E_\ell, E]$, such that $x = [E]_{\equiv}$, $l = [r, E, \phi]_{\equiv_{\mathcal{L}}}$, and E' is the result (up to \equiv) of the application of r along ϕ to the mixture E ; otherwise $w(x, l, x') = 0$; (iv) $\pi_0(x) = \sum \{\pi_0^{\mathcal{R}}(E') \mid E' \in \text{Dom}(\pi_0^{\mathcal{R}}) \cap x\}$.

The stochastic semantics of a Kappa system \mathcal{R} is then defined as the trace distribution semantics of the WLTS $\mathcal{W}_{\mathcal{R}}$.

3 Reduction procedure

In this section, we assume, without any loss of generality that Σ_i and Σ_λ are disjoint sets. This can always be achieved by taking two disjoint copies \mathcal{S}_i and \mathcal{S}_λ of \mathcal{S} and using site names in \mathcal{S}_i to bear internal states, and site names in \mathcal{S}_λ to bear binding states.

Informally, a contact map represents a summary of the agents and their potential bindings.

Definition 13. (*Contact map*) Given a Kappa system \mathcal{R} , a contact map (CM) is a graph object $(\mathcal{N}, \mathcal{E})$, where the set of nodes \mathcal{N} are agent types equipped with the corresponding interface, and the edges are specified between the sites of the nodes. Formally, we have that $\mathcal{N} = \{(A, \Sigma(A)) \mid A \in \mathcal{A}\}$ and $\mathcal{E} \subseteq \{((A, s), (A', s')) \mid A, A' \in \mathcal{A} \text{ and } s \in \Sigma(A), s' \in \Sigma(A')\}$. If the site s of an agent of type A and the site s' of an agent of type A' bear the same binding state in the rhs E_r of a rule, then there exists an edge $e = ((A, s), (A', s')) \in \mathcal{E}$ between $s \in \Sigma(A)$ and $s' \in \Sigma(A')$.

We say that a site s of the agent a is *tested* by the rule r , if it is contained in the lhs E_ℓ of the rule r .

Definition 14. (*Annotated Contact map*) Given a Kappa system \mathcal{R} , and its CM $(\mathcal{N}, \mathcal{E})$, a valid annotated contact map (ACM) $(\mathcal{N}, \mathcal{E})$ is a contact map where all agents are annotated with respect to the rule set \mathcal{R} . The annotation on the agent of type $A \in \mathcal{A}$ with respect to the rule r is given by the equivalence relation on its set of sites $\approx_A \subseteq \Sigma(A) \times \Sigma(A)$ such that:

- If a rule r tests the sites s_1 and site s_2 of agents a_1, a_2 (it is possible that $a_1 = a_2$) of type A , then $s_1 \approx_A s_2$;
- If a rule r creates an agent a of type A , then all the sites of $\Sigma(A)$ are in the same equivalence class, i.e. $\approx_A = \Sigma(A) \times \Sigma(A)$;

Note that there can be several annotations of the agent type $A \in \mathcal{A}$ which satisfy the conditions. More precisely, if the equivalence relation \approx_A meets the condition, then so does any of its refinement. This allows to define the smallest such equivalence relation \approx_A which we call the minimal annotation of agent A . An ACM is minimal whenever each agent type is annotated by its minimal annotation.

Let r be a rule and an ACM which is valid with respect to the singleton $\{r\}$. Then for any agent type $A \in \mathcal{A}$, either A does not occur in the lhs of r , or A occurs but all occurrences of A have an empty interface, or A occurs, tests some sites which all belongs to a same equivalence class C in \approx_A . In the latter case, we define $\text{test}_r^{ACM}(A) = C$, otherwise, we define $\text{test}_r^{ACM}(A) = \emptyset$.

The meaning of the ACM is to summarize the dependences between sites that can occur during the simulation of a Kappa system. If the two sites s and s' in the $\Sigma(A)$ are correlated by the relation \approx_A , i.e. $s \approx_A s'$, it suggests that they are *dependent* in the following way. We must not aggregate in the same equivalence class any two states x and x' , such that they contain the agent A in a different evaluation of the sites s and s' . On the other hand, if the two sites s and s' are not correlated by \approx_A , then we may aggregate the states by the 'marginal' criteria, i.e. the condition which involves only one of the sites. Therefore, the less states are related by $(\approx_A)_{A \in \mathcal{A}}$, the better the reduction will be. To numerically justify this, we can imagine having an agent type A whose interface has n different sites s_1, \dots, s_n , and each of them has two possible internal state modifications. Let us observe the two limiting relations \approx_A , i.e. $\approx_A = \{(s_i, s_j) \mid 1 \leq i \leq n, 1 \leq j \leq n\}$, and $\approx'_A = \{(s_i, s_i) \mid 1 \leq i \leq n\}$. The annotation \approx_A enforces at least to 2^n states to describe all modifications of the agent A , whereas the annotation \approx'_A suggests that it is enough to use only $2 \cdot n$ of them.

The ACM can be used to identify parts of Kappa species that we call fragments.

Definition 15. (*Kappa fragments*) A fragment is the \equiv -equivalent class of a non empty irreducible pattern E such that: (i) the set of sites in the interface σ of an agent $A(\sigma)$ in E is an equivalence class of \approx_A , (ii) sites can only be free or tagged with a binding label $i \in \mathbb{N}$ and sites in Σ_i are tagged with an internal state in \mathbb{I} , (iii) there is no occurrence of fictitious agent \emptyset .

We can use fragments to abstract the WLTS $\mathcal{W}_{\mathcal{R}}$, by identifying the mixtures which have the same (multi-)set of fragments. To reach that goal, we first overload the definition of \equiv in order to identify mixtures having the same fragments. We introduce the binary relation \equiv^\sharp as the smallest equivalence relation over patterns which is compatible with \equiv and such that:

$$A(\sigma), A(\sigma'), E \equiv^\sharp A(\uparrow_C \sigma', \uparrow_{\Sigma(A) \setminus C} \sigma), A(\uparrow_C \sigma, \uparrow_{\Sigma(A) \setminus C} \sigma'), E$$

for any agent type $A \in \mathcal{A}$, σ, σ' interfaces, E pattern, and C an \approx_A -equivalence class of sites. For any set of sites $X \subseteq \mathcal{S}$, the projection function \uparrow_X over interfaces keeps only the sites in X , formally \uparrow_X is defined by $\uparrow_X \varepsilon = \varepsilon$, $\uparrow_X (s_i^\lambda, \sigma') = s_i^\lambda$, $\uparrow_X \sigma'$ whenever $s \in X$, and $\uparrow_X (s_i^\lambda, \sigma') = \uparrow_X \sigma'$ otherwise.

Now we define the relation $\sim_{\mathcal{L}^\sharp}$ which stipulates that the rule r_1 applies on E_1 along ϕ_1 the same way as the rule r_2 on E_2 along ϕ_2 . More formally, we write $(r_1, E_1, \phi_1) \sim_{\mathcal{L}^\sharp} (r_2, E_2, \phi_2)$ whenever the following properties are all satisfied:

1. $r_1 = r_2$;
2. $E_1 \equiv^\# E_2$;
3. $\phi_2 = \psi \circ \phi_1$, where ψ is the permutation which tracks how the sub-interface $\uparrow_{\text{test}_r^{ACM}(A_i)}(A_i(\sigma_i))$ is moved in the proof that $E_1 \equiv^\# E_2$, for any agent $A_i(\sigma_i)$ occurring in E_1 .

More precisely, the transposition $[\begin{smallmatrix} i \\ i+1 \end{smallmatrix}]$ is associated to an agent permutation of the agents at position i and $i+1$; the transposition $[\begin{smallmatrix} 1 \\ 2 \end{smallmatrix}]$ is associated to a step which permutes the sub-interface $\text{test}_r^{ACM}(A)$ of two agents of type A , for any agent type $A \in \mathcal{A}$; any other step is associated with the identity function (over \mathbb{N}). The function ψ is defined as the composition of all the permutations (in the reverse order) which are associated to the elementary steps in the proof that $E_1 \equiv^\# E_2$.

4. the result of the application of r_1 to E_1 along ϕ_1 is \equiv -equivalent to the result of the application of r_2 to E_2 along ϕ_2 .

Definition 16. (Abstract WLTS of a Kappa system) Let $\mathcal{R} = (\pi_0^{\mathcal{R}}, \{r_1, \dots, r_n\})$ be a Kappa system. We define the WLTS $\mathcal{W}_{\mathcal{R}} = (\tilde{\mathcal{X}}, \tilde{\mathcal{L}}, \tilde{w}, \tilde{\pi}_0)$ where:

- $\tilde{\mathcal{X}}$ is the set of all $\equiv^\#$ -equivalent class of mixture;
- $\tilde{\mathcal{L}}$ is the set of all $\equiv_{\mathcal{L}}^\#$ -equivalent class of triples (r, E, ϕ) such that ϕ is an embedding between the lhs E_ℓ of r and E ;
- $\tilde{w}(\tilde{x}, \tilde{\lambda}, \tilde{x}')$ is equal to $\frac{k}{|[E_\ell, E]|}$ whenever there exist a mixture E , a rule r , and an embeddings ϕ such that $\tilde{x} = [E]_{\equiv^\#}$ and $\tilde{\lambda} = [r, E, \phi]_{\equiv_{\mathcal{L}}^\#}$; otherwise, $\tilde{w}(\tilde{x}, \tilde{\lambda}, \tilde{x}')$ is equal to 0;
- for any $\tilde{x} \in \tilde{\mathcal{X}}$, $\tilde{\pi}_0(\tilde{x}) = \sum_{E' \in \text{Dom}(\pi_0^{\mathcal{R}}) \cap \tilde{x}} \pi_0^{\mathcal{R}}(E')$.

Let us define the relation \sim over $\tilde{\mathcal{X}}$ by $[E_1]_{\equiv^\#} \sim [E_2]_{\equiv^\#}$ if, and only if, $E_1 \equiv^\# E_2$ and the relation $\sim_{\mathcal{L}}$ over $\tilde{\mathcal{L}}$ by $[\lambda_1]_{\equiv_{\mathcal{L}}^\#} \sim_{\mathcal{L}} [\lambda_2]_{\equiv_{\mathcal{L}}^\#}$ if, and only if, $\lambda_1 \equiv_{\mathcal{L}}^\# \lambda_2$.

4 Abstraction

We introduce abstractions on WLTS by aggregating the states and labels into partition classes. We obtain a new WLTS defined over the aggregated states and labels. Each non-trivial abstraction is a loss of information. However some of them are such that it is possible to do the stochastic analysis on the aggregates rather than on concrete states. We address the problem of characterizing when this is possible, and if so, how the weights in the abstracted system are computed. We also discuss the reverse process - given the abstracted system, and a particular probability distributions over the aggregates, whether we can make conclusions about the traces in the concrete system. We do the general theoretical analysis of the abstractions on WLTS, and afterwards we show the relation with the reduction of Kappa systems, that is presented in Sec. 3.

Definition 17. (Abstraction) Consider a WLTS $\mathcal{W} = (\mathcal{X}, \mathcal{L}, w, \pi_0)$, and a pair of equivalence relations $(\sim, \sim_{\mathcal{L}}) \in \mathcal{X}^2 \times \mathcal{L}^2$, such that each \sim -equivalence class and each $\sim_{\mathcal{L}}$ -equivalence class is finite. We denote the equivalence classes by \tilde{x}, \tilde{l} , and we write $x \in \tilde{x}$, to indicate that x belongs to the equivalence class \tilde{x} , and $l \in \tilde{l}$ to indicate that the label l belongs to the equivalence class \tilde{l} .

A WLTS of the form $\tilde{\mathcal{W}} = (\tilde{\mathcal{X}}_{|\sim}, \tilde{\mathcal{L}}_{|\sim_{\mathcal{L}}}, \tilde{w}, \tilde{\pi}_0)$, where $\tilde{\pi}_0(\tilde{x}) = \sum \{\pi_0(x) \mid x \in \tilde{x}\}$ is called an abstraction of \mathcal{W} , induced by the pair of equivalence relations $(\sim, \sim_{\mathcal{L}})$. Note that more abstractions can be induced by \mathcal{W} , depending on how \tilde{w} is defined.

Moreover, for any two cylinder sets of traces $\tilde{\tau}_{\text{IR}} \in \mathcal{T}_{\text{IR}}(\tilde{\mathcal{W}})$ and $\tau_{\text{IR}} \in \mathcal{T}_{\text{IR}}(\mathcal{W})$, we say that $\tilde{\tau}_{\text{IR}} = \tilde{x}_0 \xrightarrow{\tilde{l}_1, l_1} \tilde{x}_1 \dots \tilde{x}_{k-1} \xrightarrow{\tilde{l}_k, l_k} \tilde{x}_k$ is an abstraction of $\tau_{\text{IR}} = x_0 \xrightarrow{l_1, l_1} x_1 \dots x_{k-1} \xrightarrow{l_k, l_k} x_k$, and we write it $\tilde{\tau}_{\text{IR}} \in \tilde{\tau}_{\text{IR}}$.

Definition 18. (*Sound abstraction: Aggregation*) We say that the abstraction $\tilde{\mathcal{W}}$ is a sound abstraction of \mathcal{W} , if the probability of any cylinder set of traces $\tilde{\tau}_{\mathbb{IR}} \in \mathcal{T}_{\mathbb{IR}}(\tilde{\mathcal{W}})$ is equal to the sum of the probabilities of all the cylinder sets of traces $\tau_{\mathbb{IR}} \in \mathcal{T}_{\mathbb{IR}}(\mathcal{W})$, whose abstraction is $\tilde{\tau}_{\mathbb{IR}}$:

$$\pi(\tilde{\tau}_{\mathbb{IR}}) = \sum \{ \pi(\tau_{\mathbb{IR}}) \mid \tau_{\mathbb{IR}} \in \tilde{\tau}_{\mathbb{IR}} \}.$$

We introduce a function $\gamma: \mathcal{X}_{/\sim} \rightarrow (\mathcal{X} \rightarrow [0, 1])$ which assigns to each of the partition class $\tilde{x} \in \mathcal{X}_{/\sim}$ a probability distribution over the states that belong to this partition class. The set of all such vectors γ we denote by $\Gamma_{\mathcal{X}, \sim}$ is defined as:

$$\{ \gamma \mid \gamma: \mathcal{X}_{/\sim} \rightarrow (\mathcal{X} \rightarrow [0, 1]) \wedge \forall \tilde{x} \in \mathcal{X}_{/\sim}, \sum_{x \in \tilde{x}} \gamma(\tilde{x}, x) = 1 \}.$$

We can think of the value $\gamma(\tilde{x}, x)$ as the *conditional probability* of being in the state x , knowing that we are in state \tilde{x} , i.e. $\Pr(X_t = x \mid X_t \in \tilde{x}) = \gamma(\tilde{x}, x)$. We note that, when thinking of γ as the conditional probability, it should be a time-dependent value. However, we refer to γ as to a single, constant distribution. This will be justified in Lem. 1.

Definition 19. (*Complete abstraction: Deaggregation*) We say that the abstraction $\tilde{\mathcal{W}}$ is a complete abstraction of \mathcal{W} for $\gamma \in \Gamma_{\mathcal{X}, \sim}$, if the following holds. Given the probability of an arbitrary abstract cylinder set of traces of length $k \geq 1$, that ends in the abstract state \tilde{x}_k (written $\tilde{\tau}_{\mathbb{IR}} \rightarrow \tilde{x}_k$), we can recompute the probability of ending the trace in the concrete state $x_k \in \tilde{x}_k$ in the following way:

$$\pi(\tilde{\tau}_{\mathbb{IR}} \rightarrow x_k) = \gamma(\tilde{x}_k, x_k) \cdot \pi(\tilde{\tau}_{\mathbb{IR}} \rightarrow \tilde{x}_k).$$

Sound and complete abstractions $\tilde{\mathcal{W}}$ cannot be induced by any pair of relations $(\sim, \sim_{\mathcal{L}})$, because there might not exist a weighting function $\tilde{w}: \mathcal{X}_{/\sim} \times \mathcal{L}_{/\sim_{\mathcal{L}}} \times \mathcal{X}_{/\sim} \rightarrow \mathbb{R}$, such that the conditions from Dfn. 18 and Dfn. 19 are met. Moreover, even if such \tilde{w} exists, the remaining question is whether the information on the abstract system is enough to compute them.

We now restate the main Theorem from [12], that the abstractions for Kappa systems, that we resumed in Sec. 3, are sound and complete.

Theorem 1. (*The abstraction induced by the ACM is sound and complete*) Given a Kappa system $\mathcal{R} = (\pi_0^{\mathcal{R}}, \{r_1, \dots, r_n\})$, and an ACM $(\mathcal{N}, \mathcal{E})$, an abstraction $\tilde{\mathcal{W}}_{\mathcal{R}} = (\mathcal{X}_{/\sim}, \mathcal{L}_{/\sim_{\mathcal{L}}}, \tilde{w}, \tilde{\pi}_0)$ induced by the partition classes $(\sim, \sim_{\mathcal{L}}) \subseteq \mathcal{X}^2 \times \mathcal{L}^2$, as proposed in the Def. 16 is a sound and complete abstraction of the $\mathcal{W}_{\mathcal{R}} = (\mathcal{X}, \mathcal{L}, w, \pi_0)$, provided that for any two mixtures M and M' such that $M \equiv^{\sharp} M'$, we have:

$$\pi_0([M]_{\equiv}) \cdot |[M', M']| = \pi_0([M']_{\equiv}) \cdot |[M, M]|.$$

We consider a mixture M . We denote by $x \in \mathcal{X}$ the equivalence class $[M]_{\equiv}$, and by $\tilde{x} \in \mathcal{X}_{/\sim}$ the equivalence class $[M]_{\equiv^{\sharp}} = [x]_{\sim}$. The conditional probability $\gamma(\tilde{x}, x)$ is computed as the ratio of the number of automorphisms of x (embedding between x and x) and the sum of the number of automorphisms of any \sim -equivalent state. Thus we have:

$$\gamma(\tilde{x}, x) = \frac{|[x, x]|}{\sum \{ |[x', x']| \mid x \sim x' \}}.$$

The reader can find the detailed proof in [12].

4.1 Lumpability

Now we define different versions of lumpability and investigate the relationship with sound and complete abstractions.

Definition 20. (*Lumped process*) Given a WLTS $\mathcal{W} = (\mathcal{X}, \mathcal{L}, w, \pi_0)$, where $\mathcal{X} = \{x_1, x_2, \dots\}$, and a partition $\sim \subseteq \mathcal{X} \times \mathcal{X}$ on its state space, we observe the continuous-time stochastic process $\{X_t\}_{t \in \mathbb{R}_0^+}$, that is generated by \mathcal{W} (Dfn. 2). We define the lumped process $\{Y_t\}$ on the state space $\mathcal{X}/\sim = \{\tilde{x}_1, \tilde{x}_2, \dots\}$ (denoted by capital indices, i.e. \tilde{x}_I, \tilde{x}_J) and with initial distribution $\tilde{\pi}_0$, so that

$$\Pr(Y_t = \tilde{x}_J \mid Y_0 = \tilde{x}_0) = \Pr(X_t \in \tilde{x}_J \mid X_0 \in \tilde{x}_0).$$

The lumped process is not necessarily a Markov process.

Definition 21. (*Lumpability*) Given a WLTS $\mathcal{W} = (\mathcal{X}, \mathcal{L}, w, \pi_0)$ that generates the process $\{X_t\}$, we say that it is lumpable with respect to the equivalence relation $\sim \subseteq \mathcal{X} \times \mathcal{X}$ if and only if its lumped process $\{Y_t\}$ has the Markov property.

The evolution of a process depends on the initial distribution, and so does the lumpability property. We thus define the set of initial distributions, for which the lumpability holds. We denote the set of all probability distributions over \mathcal{X} as $\mathcal{P}_{\mathcal{X}}$:

$$\mathcal{P}_{\mathcal{X}} = \{\pi \mid \pi : \mathcal{X} \rightarrow [0, \infty] \text{ and } \sum_{\S \in \mathcal{X}} \pi(\S) = \infty\}.$$

Moreover, we denote the set of initial distributions that produce a chain lumpable with respect to the given equivalence relation \sim by $\mathcal{P}_{\mathcal{X}, \sim}^{\mathcal{L}}$:

$$\mathcal{P}_{\mathcal{X}, \sim}^{\mathcal{L}} := \{\pi \mid \text{the lumped process initialized with } \pi \text{ is lumpable with respect to } \sim\}.$$

Whenever a distribution $\pi \in \mathcal{P}_{\mathcal{X}}$ is positive on the equivalence class \tilde{x} , i.e. $\sum\{\pi(x) \mid x \in \tilde{x}\} > 0$, we denote by $\pi|_{\tilde{x}}(x)$, the conditional distribution over the states of \tilde{x} : $\pi|_{\tilde{x}}(x) = \pi(x)/\pi(\tilde{x})$, when $x \in \tilde{x}$, and $\pi|_{\tilde{x}}(x) = 0$, otherwise.

Definition 22. (*Strong and weak lumpability*) Given a WLTS $\mathcal{W} = (\mathcal{X}, \mathcal{L}, w, \pi_0)$ that generates the process $\{X_t\}$, and an equivalence relation $\sim \subseteq \mathcal{X} \times \mathcal{X}$, we say that $\{X_t\}$ is:

- *strongly lumpable with respect to \sim* , if the lumped process $\{Y_t\}$ is Markov with respect to any initial distribution, i.e. $\mathcal{P}_{\mathcal{X}, \sim}^{\mathcal{L}} = \mathcal{P}_{\mathcal{X}}$;
- *weakly lumpable with respect to \sim* , if there exists an initial distribution that makes the lumped process $\{Y_t\}$ Markov, i.e. $\mathcal{P}_{\mathcal{X}, \sim}^{\mathcal{L}} \neq \emptyset$.

Note that the definitions of strong and weak lumpability involve the quantifiers "for all" and "exists" over the probability distributions over a set of states. Thus, checking for either of them involves in general an infinite number of checks. People have given sufficient conditions of strong and weak lumpability on discrete-time Markov chains (DTMC's) [21, 26]. The results had been extended to the continuous-time case [4, 27]. We rephrase the sufficient conditions stated therein.

In order to understand the sense of the weak lumpability characterization, we discuss the meaning of γ . We recall the semantics of a WLTS \mathcal{W} by observing the cylinder sets of traces, i.e. $\tau_{\mathbb{R}} = x_0 \xrightarrow{l_1, l_1} x_1 \dots x_{k-1} \xrightarrow{l_k, l_k} x_k \in \mathcal{T}_{\mathbb{R}}(\mathcal{W})$. The abstraction $\tilde{\mathcal{W}}$ of \mathcal{W} , induced by (\sim, \sim_L) generates an abstract cylinder set of traces, denoted $\tilde{\tau}_{\mathbb{R}} = \tilde{x}_0 \xrightarrow{l_1, l_1} \tilde{x}_1 \dots \tilde{x}_{k-1} \xrightarrow{l_k, l_k} \tilde{x}_k \in \mathcal{T}_{\mathbb{R}}(\tilde{\mathcal{W}})$.

For any cylinder set of traces $\tilde{\tau}_{\mathbb{IR}} \in \mathcal{T}_{\mathbb{IR}}(\mathcal{W})$, we denote by $\gamma_{\tilde{\tau}_{\mathbb{IR}}}$ the distribution of the conditional probabilities over the lumped state \tilde{x}_k , knowing that the abstract cylinder of traces $\tilde{\tau}_{\mathbb{IR}}$, which ends in the abstract state \tilde{x}_k , was observed, i.e.

$$\gamma_{\tilde{\tau}_{\mathbb{IR}}}(x_k) = \frac{\pi(\tilde{\tau}_{\mathbb{IR}} \rightarrow x_k)}{\pi(\tilde{\tau}_{\mathbb{IR}})}.$$

The definition of the complete abstraction suggests that, if $\gamma_{\tilde{\tau}_{\mathbb{IR}}}$ was independent of the traces on which it is conditioned, i.e. $\tilde{\tau}_{\mathbb{IR}}$, then the completeness would hold.

Theorem 2. (Lumpability on CTMCs) *Let us observe a WLTS $\mathcal{W} = (\mathcal{X}, \mathcal{L}, w, \pi_0)$ that generates the process $\{X_t\}$, and an equivalence relation $\sim \subseteq \mathcal{X} \times \mathcal{X}$. We consider the rate matrix $\mathbf{R} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$. If the lumped process is Markov, then we denote its rate matrix by $\tilde{\mathbf{R}} : \mathcal{X}_{/\sim} \times \mathcal{X}_{/\sim} \rightarrow \mathbb{R}$. Then we have the following characterizations about the lumped process $\{\tilde{X}_t\}$:*

- *If for all $x_{i_1}, x_{i_2} \in \mathcal{X}$ such that $x_{i_1} \sim x_{i_2}$, and for all $\tilde{x}_J \in \mathcal{X}_{/\sim}$, we have that*

$$\sum_{x_j \in \tilde{x}_J} R(x_{i_1}, x_j) = \sum_{x_j \in \tilde{x}_J} R(x_{i_2}, x_j), \quad (2)$$

then $\{X_t\}$ is strongly lumpable with respect to \sim ; We have:

$$\tilde{R}(\tilde{x}_I, \tilde{x}_J) = \sum \{R(x_i, x_j) \mid x_j \in \tilde{x}_J\};$$

- *If there exists a family of probability distributions over the lumped states, $\gamma \in \Gamma_{\mathcal{X}_{/\sim}}$, such that for all $x_{j_1}, x_{j_2} \in \mathcal{X}$ such that $x_{j_1} \sim x_{j_2}$ and for all $\tilde{x}_I \in \mathcal{X}_{/\sim}$, we have that*

$$a(x_{j_1}) = a(x_{j_2}) \text{ and } \frac{\sum_{x_i \in \tilde{x}_I} R(x_i, x_{j_1})}{\gamma(\tilde{x}_I, x_{j_1})} = \frac{\sum_{x_i \in \tilde{x}_I} R(x_i, x_{j_2})}{\gamma(\tilde{x}_I, x_{j_2})}, \quad (3)$$

then

1. *If the distribution γ is in accordance with π_0 , i.e. $\pi_0|_{\mathcal{X}_{/\sim}} = \gamma$, then for any finite sequence of states $(x_0, \dots, x_k) \in \mathcal{X}^{k+1}$ and any sequence of time intervals $(I_1, \dots, I_k) \in \mathbb{IR}^k$, we consider the set $\tilde{\tau}_{\mathbb{IR}}$ of the traces of the form $x'_0 \xrightarrow{l_1, t_1} x'_1 \dots x'_{k-1} \xrightarrow{l_k, t_1 + \dots + t_k} x'_k$. For all i , $0 \leq i \leq k$ and $x_i \sim x'_i$, and for all i , $1 \leq i \leq k$, $t_i \in I_i$ and $l_i \in \mathcal{L}$, we have that: if $\pi(\tilde{\tau}_{\mathbb{IR}}) > 0$ then $\gamma_{\tilde{\tau}_{\mathbb{IR}}} = \gamma$. In other words, knowing that we are in state \tilde{x}_I , the conditional probability of being in state $x \in \tilde{x}_I$ is invariant of time.*
2. *The process $\{X_t\}$ is weakly lumpable with respect to \sim . Moreover, we have:*

$$\tilde{R}(\tilde{x}_I, \tilde{x}_J) = \frac{\sum \{R(x_i, x_{j_2}) \mid x_i \in \tilde{x}_I\}}{\gamma(\tilde{x}_I, x_{j_2})};$$

One shall notice that Thm. 2 gives a weaker condition than the completeness of WLTS abstraction (eg see Dfn. 19). The main reason is that we do not 'track' transition labels, in the sense that we observe the abstraction on the cylinder sets of traces induced only by \sim , and not also by \sim_L . Yet, in the particular case when states fully define the transition labels (ie, if $w(x_1, l_1, x'_1) > 0$, $w(x_2, l_2, x'_2) > 0$, $x_1 \sim x_2$, and $x'_1 \sim x'_2$, then $l_1 \sim_{\mathcal{L}} l_2$), the given condition for weak lumpability coincides with the definition of the complete abstraction of WLTS.

The characterization of weak (resp. strong) lumpability given in Thm. 2 is sufficient, but not a necessary condition: there exist systems which are strongly or weakly lumpable, but do not satisfy the

conditions given in the theorem. Interestingly, there are systems, such that the characterization from Thm. 2 would detect as strong, but not weakly lumpable, which is counter-intuitive with the terminology. One shall also notice that the conditions of Thm. 2 imply that: in order to aggregate two states in the CTMC, they must not have different waiting times until the next transition (e.g. they should have the same activity). It is stated explicitly in the characterization of weak lumpability and it can be obtained by summation over the outgoing class in the characterization of strong lumpability.

We consider a WLTS $\mathcal{W} = (\mathcal{X}, \mathcal{L}, w, \pi_0)$, and the set of all equivalence relations \sim on \mathcal{X} , denoted $PT_{\mathcal{X}}$. We introduce the subsets of $PT_{\mathcal{X}}$, denoted PS, PW, CS, CW in the following meaning: (i) PS - the set of all equivalence relations such that $\{X_t\}$ is strongly lumpable with respect to \sim ; (ii) PW - the set of all equivalence relations such that $\{X_t\}$ is weakly lumpable with respect to \sim ; (iii) CS - the set of all equivalence relations such that $\{X_t\}$ satisfies the condition for strong lumpability given in the Thm. 2; (iv) CW - the set of all equivalence relations such that $\{X_t\}$ satisfies the condition for weak lumpability given in the Thm. 2.

Lemma 1. (*Relations on lumpability properties and conditions*) Consider an arbitrary WLTS $\mathcal{W} = (\mathcal{X}, \mathcal{L}, w, \pi_0)$ and the equivalence relation $\sim \subseteq \mathcal{X} \times \mathcal{X}$. We have the following relations: (1a) if $\sim \in PS$ then $\sim \in PW$, if $\sim \in CS$ then $\sim \in PS$; and if $\sim \in CW$ then $\sim \in PW$; The converse implication does not hold for any of the statements; (2a) If $\sim \in CW$, that does not imply $\sim \in CS$; (2b) If $\sim \in CS$, that does not imply $\sim \in CW$.

Proof. The statement in the part (1) trivially follows from the Dfn. 21 and Thm. 2. To show (2a) and (2b), we consider the WLTS \mathcal{W} specified in the Fig. 1(a), with the state space $\mathcal{X} = \{x, y_1, y_2, y_3, z_1, z_2, z_3\}$. Let \sim_1 be an equivalence relation on \mathcal{X} , such that $y_1 \sim_1 y_2$ and $z_1 \sim_1 z_2$. By lumping the states by \sim_1 , we get the system $\tilde{\mathcal{W}}_1$, as shown in Fig. 1(b). It is easy to check that $\sim_1 \in CS$. Moreover, we have that $\sim_1 \in CW$, since for

$$\gamma = \begin{pmatrix} x & y_{12} & y_3 & z_{12} & z_3 \\ 1 & (0.5, 0.5) & 1 & (0.5, 0.5) & 1 \end{pmatrix}$$

the weak lumpability condition is satisfied, so $\sim_1 \in CS \cap CW$. We further lump the states y_{12} and y_3 , by taking the transitive closure of the relation \sim_1 union (y_{12}, y_3) , denoted $\sim_2 = tc(\sim_1 \cup (y_{12}, y_3))$ (Fig. 1(c)). This lumping is such that $\sim_2 \notin CS$ because we have

$$y_1 \sim y_3, \text{ but } w(y_1, l, z_{12}) > 0, \text{ and } w(y_3, l, z_{12}) = 0.$$

On the other hand, for

$$\gamma = \begin{pmatrix} x & y_{123} & z_{12} & z_3 \\ 1 & (1/3, 1/3, 1/3) & (0.5, 0.5) & 1 \end{pmatrix}$$

we argue that $\sim_2 \in CW$. Therefore, if the initial distribution is in accordance with γ , the abstraction $\tilde{\mathcal{W}}_2$ is sound and complete.

If we rather lump z_1 and z_2 , by $\sim_3 = tc(\sim_1 \cup (z_1, z_2))$, we get the system $\tilde{\mathcal{W}}_3$ (Fig. 1(d)). This system is such that $\sim_3 \in CS \setminus CW$. More precisely, we cannot find a γ which would witness $\sim_3 \in CW$: if such a γ existed, we would have $\gamma(\{x\})(x) = 1$, and consequently $\gamma(y_{12}) = (0.5, 0.5)$, and $\gamma(y_3) = 1$. This implies that the conditional distribution $\gamma(z_{123})$ cannot be invariant of time - it will alternate between the distributions $(0, 0, 1)$ and $(0.5, 0.5, 0)$, depending on the choice made in x . Note that, since $\sim_3 \in CS$, it follows that $\sim_3 \in PS$, and this implies $\sim_3 \in PW$. \square

This discussion indicates that if we decide to check for weak lumpability instead of for strong by using the characterization from Thm. 2, it might happen that we eliminate the aggregations that are strongly lumpable. In the case of reductions of Kappa systems, we will use the weak lumpability characterization.




Figure 1: Different abstractions of system \mathcal{W}

4.2 Bisimulations

Aiming to define the algorithm that is abstracting the WLTS of a Kappa system, we start by redefining the lumpability properties in the bisimulation notions. Bisimulation is typically defined by logically characterizing the distinguishing property of the states that may be aggregated.

We define three kinds of bisimulation relations on the WLTS, which are based on the lumpability characterizations given in Thm. 2. We adopt the terminology of [5]. The *forward* bisimulations arise from the characterization for strong lumpability: the bisimilar states have the same forward behaviour in the sense that they are each targeting any other lumped state with the same total affinity (total outgoing rate). This concept is well established for dependability or performance analysis [16, 15]. What we use in the abstractions of Kappa systems is *backward* bisimulation. The bisimilar states have the same backward behaviour in the sense that they are reached by the predecessors from one lumped state with the same probabilistic quantity, which becomes the rate in the abstract system. It is however less established and only applied in very few approaches for stochastic modelling [30]. The *backward uniform* bisimulation is an instance of a backward bisimulation with an additional constraint that only the equally-probable states may be aggregated.

Definition 23. Given a WLTS $\mathcal{W} = (\mathcal{X}, \mathcal{L}, w, \pi_0)$, and (\sim, \sim_L) a pair of equivalence relations respectively over \mathcal{X} and \mathcal{L} , we define a function $\delta_F : \mathcal{X} \times \mathcal{L}_{/\sim_L} \times \mathcal{X}_{/\sim} \rightarrow \mathbb{R}_0^+$:

$$\delta_F(x_i, \tilde{l}, \tilde{x}_j) = \sum \{ |w(x_i, l, x_j)| \mid l \in \tilde{l} \text{ and } x_j \in \tilde{x}_j \}.$$

Furthermore, given a family of probability distributions over the partitions $\gamma \in \Gamma_{\mathcal{X}, \sim}$, we define the quantity $\delta_B : \mathcal{X}_{/\sim} \times \mathcal{L}_{/\sim_L} \times \mathcal{X} \rightarrow \mathbb{R}_0^+$:

$$\delta_B(\tilde{x}_i, \tilde{l}, x_j) = \frac{\sum \{ \gamma(\tilde{x}_i, x_i) \cdot |w(x_i, l, x_j)| \mid l \in \tilde{l}, x_i \in \tilde{x}_i \}}{\gamma(\tilde{x}_j, x_j)}.$$

Specifically, if we have that γ is a uniform distribution over the equivalence classes, we can express the latter expression in terms of cardinalities of the equivalence classes:

$$\delta_{BU}(\tilde{x}_i, \tilde{l}, x_j) = \frac{|\tilde{x}_j|}{|\tilde{x}_i|} \sum \{ |w(x_i, l, x_j)| \mid l \in \tilde{l}, x_i \in \tilde{x}_i \}.$$

Definition 24. (Forward and backward Markov bisimulation) Given a WLTS $\mathcal{W} = (\mathcal{X}, \mathcal{L}, w, \pi_0)$, and (\sim, \sim_L) a pair of equivalence relations respectively over \mathcal{X} and \mathcal{L} , we say that (\sim, \sim_L) is a

1. *Forward Markov Bisimulation*, if for all x_i and x_j , the following is satisfied: $x_i \sim x_j$, iff for all equivalence classes $\tilde{x} \in \mathcal{X}_{/\sim}, \tilde{l} \in \mathcal{L}_{/\sim_L}$, we have that $a(x_i) = a(x_j)$ and $\delta_F(x_i, \tilde{l}, \tilde{x}) = \delta_F(x_j, \tilde{l}, \tilde{x})$.

Remark. Note that this involves the bisimulation in the classical sense: if x_i has a successor in some class, x_j has it as well, and they are related by appropriate labels (and probabilities in this case).

2. *Backward Markov bisimulation*, if for all x_i and x_j , there exists an $\gamma \in \Gamma_{\mathcal{X}, \sim}$, such that the following is satisfied: $x_i \sim x_j$, iff for all equivalence classes $\tilde{x} \in \mathcal{X}_{/\sim}, \tilde{l} \in \mathcal{L}_{/\sim_L}$, we have that $a(x_i) = a(x_j)$ and $\delta_B(\tilde{x}, \tilde{l}, x_i) = \delta_B(\tilde{x}, \tilde{l}, x_j)$.

Theorem 3. (Forward Markov bisimulation implies sound abstraction) Let $\mathcal{W} = (\mathcal{X}, \mathcal{L}, w, \pi_0)$ be a WLTS. If (\sim, \sim_L) induces a forward Markov bisimulation, then for any aggregates \tilde{x}_i, \tilde{l} , and \tilde{x}_j , we can define

$$\tilde{w}(\tilde{x}_i, \tilde{l}, \tilde{x}_j) = \delta_F(x_i, \tilde{l}, \tilde{x}_j).$$

The so defined abstraction $\tilde{\mathcal{W}} = (\mathcal{X}_{/\sim}, \mathcal{L}_{/\sim_L}, \tilde{w}, \tilde{\pi}_0)$ is sound. We then say that \mathcal{W} refines $\tilde{\mathcal{W}}$ by a forward Markov bisimulation (\sim, \sim_L) , written $\mathcal{W} \preceq_{F, (\sim, \sim_L)} \tilde{\mathcal{W}}$.

Theorem 4. (Backward Markov bisimulation implies sound and complete abstraction) Given a WLTS $\mathcal{W} = (\mathcal{X}, \mathcal{L}, w, \pi_0)$, if (\sim, \sim_L) induces a backward Markov bisimulation with conditional probabilities over the aggregates $\gamma \in \Gamma_{\mathcal{X}, \sim}$, then for any aggregates \tilde{x}_i, \tilde{l} , and \tilde{x}_j , we can define

$$\tilde{w}(\tilde{x}_i, \tilde{l}, \tilde{x}_j) = \delta_B(\tilde{x}_i, \tilde{l}, x_j). \quad (4)$$

If $\gamma(\tilde{x}) = \pi_0|_{\tilde{x}}$, then the so defined abstraction $\tilde{\mathcal{W}} = (\mathcal{X}_{/\sim}, \mathcal{L}_{/\sim_L}, \tilde{w}, \tilde{\pi}_0)$ is sound and complete. We then say that \mathcal{W} refines $\tilde{\mathcal{W}}$ by a backward Markov bisimulation (\sim, \sim_L) with conditional distributions γ , written $\mathcal{W} \preceq_{B, (\sim, \sim_L), \gamma} \tilde{\mathcal{W}}$.

In particular, if we know that γ is uniform, the equation (4) becomes $\tilde{w}(\tilde{x}_i, \tilde{l}, \tilde{x}_j) = \delta_{BU}(x_i, \tilde{l}, \tilde{x}_j)$, written also $\mathcal{W} \preceq_{BU, (\sim, \sim_L)} \tilde{\mathcal{W}}$.

4.3 Proving bisimulations

The forward bisimulation relation for abstracting the transition systems with CTMC semantics has been established and used in applications (eg, [16, 15]). Moreover, computing the backward uniform bisimulation when γ is uniform is defined [5, 30]. It is based on an alternative characterization of the backward uniform Markov bisimulation, which eases the analysis.

Lemma 2. (Proving backward uniform Markov bisimulation) Let $\mathcal{W} = (\mathcal{X}, \mathcal{L}, w, \pi_0)$ be a WLTS and (\sim, \sim_L) be a pair of equivalence relations respectively over \mathcal{X} and \mathcal{L} . For any state $x' \in \mathcal{X}$, and any pair of classes $\tilde{x}, \tilde{l} \in \mathcal{X}_{/\sim} \times \mathcal{L}_{/\sim_L}$, let us define the set $\text{Pred}(\tilde{x}, \tilde{l}, x')$ of transitions from a state in \tilde{x} to the state x' and with a label in \tilde{l} as follows:

$$\text{Pred}(\tilde{x}, \tilde{l}, x') = \{(x, l) \in \tilde{x} \times \tilde{l} \mid w(x, l, x') > 0\}.$$

Assume that: (1) $\pi_0|_{\mathcal{X}_{/\sim}} = \tilde{\pi}_0$, and (2) for any $x'_i, x'_j \in \mathcal{X}$ such that $x'_i \sim x'_j$ and any $\tilde{x} \in \mathcal{X}_{/\sim}, \tilde{l} \in \mathcal{L}_{/\sim_L}$, there exists a bijective map ϕ between $\text{Pred}(\tilde{x}, \tilde{l}, x'_i)$ and $\text{Pred}(\tilde{x}, \tilde{l}, x'_j)$, such that for any $(x_i, l_i) \in \text{Pred}(\tilde{x}, \tilde{l}, x'_i)$, if $\phi(x_i, l_i) = (x_j, l_j)$, then we have that $w(x_i, l_i, x'_i) = w(x_j, l_j, x'_j)$.

Then we have that \mathcal{W} is the backward uniform bisimulation of the abstraction $\tilde{\mathcal{W}} = (\mathcal{X}_{/\sim}, \mathcal{L}_{/\sim_L}, \tilde{w}, \tilde{\pi}_0)$, i.e. $\mathcal{W} \preceq_{BU, (\sim, \sim_L)} \tilde{\mathcal{W}}$.

$$\mathcal{W}^i \begin{array}{c} \xrightarrow{\preceq_{BU,(\sim_1, \sim_{L1})}} \mathcal{W} \\ \xrightarrow{\preceq_{BU,(\sim_2, \sim_{L2})}} \tilde{\mathcal{W}} \end{array} \Rightarrow \mathcal{W} \xrightarrow{\preceq_{B,(\sim, \sim_L), \gamma}} \tilde{\mathcal{W}}$$

Figure 2: Proving backward refinement

On the other hand, as soon as γ over the aggregates is not uniform, we cannot observe the bijection between predecessors over the states. Proving that the given abstraction is a backward bisimulation cannot be established unless we have a right 'guess' of the distributions γ . Lem. 3 states how to avoid proving backward bisimulation by instead proving two uniform backward bisimulations. More precisely, if we want to prove the backward refinement between the systems \mathcal{W} and $\tilde{\mathcal{W}}$, it is enough to observe the system \mathcal{W}^i , which is a backward uniform refinement of both of the systems \mathcal{W} and $\tilde{\mathcal{W}}$ (Fig. 4.3).

Lemma 3. (*Proving backward Markov bisimulation*) Consider a WLTS $\mathcal{W} = (\mathcal{X}, \mathcal{L}, w, \pi_0)$, and any aggregation relation (\sim, \sim_L) , such that $\tilde{\mathcal{W}} = (\mathcal{X}/\sim, \mathcal{L}/\sim_L, \tilde{w}, \tilde{\pi}_0)$. We assume that there exist a system $\mathcal{W}^i = (\mathcal{X}^i, \mathcal{L}^i, w^i, \pi_0^i)$, and the two pairs of equivalence relations (\sim_1, \sim_{L1}) , (\sim_2, \sim_{L2}) , such that $\mathcal{W}^i \preceq_{BU,(\sim_1, \sim_{L1})} \mathcal{W}$, $\mathcal{W}^i \preceq_{BU,(\sim_2, \sim_{L2})} \tilde{\mathcal{W}}$, $\sim_1 \preceq \sim_2$ (in the sense that, for any $x_1^i, x_2^i \in \mathcal{X}^i$, $x_1^i \sim_1 x_2^i \Rightarrow x_1^i \sim_2 x_2^i$), $\sim_{L1} \preceq \sim_{L2}$, and, for any x_1^i, x_2^i such that $x_1^i \sim_2 x_2^i$, the number of states which are \sim_{L1} -equivalent to x_1^i is equal to the number of states which are \sim_{L1} -equivalent to x_2^i . Under this assumption, we have that $\mathcal{W} \preceq_{B,(\sim, \sim_L), \gamma} \tilde{\mathcal{W}}$, where γ is defined as

$$\gamma(\tilde{x}, x) = \frac{|\{x^i \in \mathcal{X}^i \mid x^i \sim_1 x_0^i\}|}{|\{x^i \in \mathcal{X}^i \mid x^i \sim_2 x_0^i\}|}, \text{ for any } [x_0^i]_{\sim_1} = x.$$

This Lemma contains the key observation for the abstraction of Kappa systems, and for proving Thm.1. It thus completes the intention of the theoretical analysis in this paper. More precisely, we observe the WLTS $\mathcal{W}^{\mathcal{R}}$ of a given Kappa system \mathcal{R} , as defined in Dfn. 12 and its abstraction generated as proposed in the reduction procedure (Sec. 3, Dfn. 16). The main observation is that the system \mathcal{W} is already an abstraction. More concretely, the states of \mathcal{W} are multisets of species, and as such, they abstract the *individual* species. For example, a state that contains two agents of type $A(s_u)$ abstracts away the potential individual behavior of these two agents, for example $A_1(s_u)$ and $A_2(s_u)$. To show that the abstraction is sound and complete, we observe the system \mathcal{W}^i , which is the *individual-based* semantics of a Kappa system, where each individual agent is uniquely identified. The backward uniform refinement is established between \mathcal{W}^i and \mathcal{W} by the modeling assumptions. We are left to prove the backward uniform refinement between \mathcal{W}^i and $\tilde{\mathcal{W}}$. This is done by inspections on the ACM's (Dfn. 14).

4.4 Example

We consider the following Kappa system. We have the agent types $\mathcal{A} = \{A, B\}$, the site names $\{s, t\}$, the signatures $\Sigma_t(A) = \Sigma_t(B) = \{s\}$ and $\Sigma_\lambda(A) = \Sigma_\lambda(B) = \{t\}$, the alphabet of internal states $\mathbb{I} = \{u, p\}$. The contact map is defined by $(\mathcal{N}, \mathcal{E})$, such that $\mathcal{N} = \{(A, s), (A, t), (B, s), (B, t)\}$ and $\mathcal{E} = \{((A, t), (B, t))\}$ and the following rules:

$$\begin{aligned} r1 : A(s_u) &\leftrightarrow A(s_p) @ k_1, k_{1-} \\ r2 : B(s_u) &\leftrightarrow B(s_p) @ k_2, k_{2-} \\ r3 : A(t) , B(t) &\leftrightarrow A(t^1) , B(t^1) @ k_3, k_{3-} \end{aligned}$$

Moreover, using the minimal ACM for annotating the agents, as written in Dfn. 14, we get that \approx_A has two equivalence classes $\{s\}$ and $\{t\}$; and that \approx_B has two equivalence classes $\{s\}$ and $\{t\}$ as well.

The fragments derived from an ACM (Dfn. 15) are the following: $F_1 = A(s_u)$, $F_2 = A(s_p)$, $F_3 = A(t)$, $F_4 = A(t^1)$, $B(t^1)$, $F_5 = B(s_u)$, $F_6 = B(s_p)$, $F_7 = B(t)$.

Let us pick a (finite) initial distribution π_0 . Now we observe the WLTS $\mathcal{W} = (\mathcal{X}, \mathcal{L}, w, \pi_0)$ assigned to the Kappa system \mathcal{R}_{AB} (introduced in Dfn. 12), and the state y which is the \equiv -equivalence class of the mixture E_y defined as follows:

$$A(s_p, t^1), B(s_p, t^1), A(s_u, t^2), B(s_u, t^2), A(s_u, t^3), B(s_u, t^3)$$

The unique (up to \equiv) non \equiv^\sharp -equivalent mixtures is $E_{y'}$ which is defined as follows:

$$A(s_p, t^1), B(s_u, t^1), A(s_u, t^2), B(s_p, t^2), A(s_u, t^3), B(s_u, t^3)$$

We denote $y' = [E_{y'}]_{\equiv}$, $\tilde{y}' = [E_{y'}]_{\equiv^\sharp}$. We compute however that the distribution among state $\tilde{y} = [E_y]_{\equiv^\sharp}$.

We have: $\gamma(\tilde{y}, y) = 1/3$ and $\gamma(\tilde{y}, y') = 2/3$. Roughly speaking this comes from the fact that if we annotate fragments of type A and B in \tilde{y} with the identifiers 1, 2, 3 (there are 36 possible annotation), and if we assume that agents with the same identifiers are bound together. Then there are 12 annotations so that the phosphorylated A and B are bound together, and 24 where this is not the case. A more detailed analysis of this model is given in [12].


5 Case study

In this section, we apply our framework to a case study. We have indeed refactored in Kappa the model of a crosstalk between the *EGF* receptor and the Insulin receptor pathways which is described in [6]. In this model, two kinds of receptors, *EGF* receptor (*EGFR*) and insulin receptor (*IR*) can recruit a protein called *Sos*, which can be phosphorylated, or not. Each kind of receptor has its own pathway, and these two pathways shared some common proteins.

The contact map is given in Fig. 3. One can notice that some sites can be bound to several other sites, which denotes a competition (concurrency). Moreover, the site d of a *EGF* receptor (*EGFR*) can be bound to the site d of another *EGFR* (since there is a loop in the CM). Moreover rules are given in Table 1. We do not give the rate for rules, but we assume no hypotheses on the rates (rates are parameters, some of them might be equal, or not). Moreover, each rule is reversible.

We roughly explain how each pathway works, by focusing on the forward direction of rules. First, we describe how *EGFR* can recruit a transport molecule (*Grb*). *EGFR* can recruit a ligand (EGF) on site a (r01,r02), and two *EGFRs* can form a dimer (r03,r04,r05). We have used two rules to encode EGF/*EGFR* binding, in order to model the fact that the rate of association may depend on the fact whether *EGFR* is in a dimer, or not. The same way, the rate of dimer formation/dissociation may depend on the number of ligands that are bound to receptors. The site b of *EGFR* can be phosphorylated (r06,r07) at a rate which depends whether the receptor is in a dimer, or not. Then, *EGFR* can recruit an adapter molecule called Shc (r08). Then, *EGFR* can phosphorylate Shc (r09,r10) (the rate depend on the fact whether the receptor is still in a dimer, or not). Shc can then recruit a transport molecule (*Grb*) (r11). Yet, each receptor has a shorter way to recruit a transport molecule. The site c of *EGFR* can be phosphorylated (r12,r13), and then recruit *Grb* directly (r14,r15).

Then, we describe how an insulin receptor (*IR*) can recruit *Grb*. *IR* can recruit insulin molecules (Ins) on two sites a (r16,r17) and b (r18,r19) (the rate may depend on the fact whether an insulin molecule

Figure 3: Contact map for the *EGFR*/Insulin crosstalk.

r01: $EGF(a), EGFR(a,d) \leftrightarrow EGF(a^1), EGFR(a^1,d)$	r20: $IR(a,b,c_u) \leftrightarrow IR(a,b,c_p)$
r02: $EGF(a), EGFR(a,d^-) \leftrightarrow EGF(a^1), EGFR(a^1,d^-)$	r21: $IR(a^-,b,c_u) \leftrightarrow IR(a^-,b,c_p)$
r03: $EGFR(a,d), EGFR(a^-,d) \leftrightarrow EGFR(a,d^1), EGFR(a^-,d^1)$	r22: $IR(a,b^-,c_u) \leftrightarrow IR(a,b^-,c_p)$
r04: $EGFR(a,d), EGFR(a,d) \leftrightarrow EGFR(a,d^1), EGFR(a,d^1)$	r23: $IR(a^-,b^-,c_u) \leftrightarrow IR(a^-,b^-,c_p)$
r05: $EGFR(a^-,d), EGFR(a^-,d) \leftrightarrow EGFR(a^-,d^1), EGFR(a^-,d^1)$	r24: $IR(c_p), Shc(a) \leftrightarrow IR(c_p^1), Shc(a^1)$
r06: $EGFR(b_u,d) \leftrightarrow EGFR(b_p,d)$	r25: $IR(a^-,b^-,c^1), Shc(a^1,b_u) \leftrightarrow IR(a^-,b^-,c^1), Shc(a^1,b_p)$
r07: $EGFR(b_u,d^-) \leftrightarrow EGFR(b_p,d^-)$	r26: $IR(a,b,d_u) \leftrightarrow IR(a,b,d_p)$
r08: $EGFR(b_p), Shc(a) \leftrightarrow EGFR(b_p^1), Shc(a^1)$	r27: $IR(a^-,b,d_u) \leftrightarrow IR(a^-,b,d_p)$
r09: $EGFR(b_p^1,d), Shc(a^1,b_u) \leftrightarrow EGFR(b_p^1,d), Shc(a^1,b_p)$	r28: $IR(a,b^-,d_u) \leftrightarrow IR(a,b^-,d_p)$
r10: $EGFR(b_p^1,d^-), Shc(a^1,b_u) \leftrightarrow EGFR(b_p^1,d^-), Shc(a^1,b_p)$	r29: $IR(a^-,b^-,d_u) \leftrightarrow IR(a^-,b^-,d_p)$
r11: $Grb(a), Shc(b_p) \leftrightarrow Grb(a^1), Shc(b_p^1)$	r30: $IR(d_p), IRS(a) \leftrightarrow IR(d_p^1), IRS(a^1)$
r12: $EGFR(c_u,d) \leftrightarrow EGFR(c_p,d)$	r31: $IR(a^-,b^-,d^1), IRS(a^1,b_u) \leftrightarrow IR(a^-,b^-,d^1), IRS(a^1,b_p)$
r13: $EGFR(c_u,d^-) \leftrightarrow EGFR(c_p,d^-)$	r32: $Grb(a), IRS(b_p) \leftrightarrow Grb(a^1), IRS(b_p^1)$
r14: $EGFR(c_p,d), Grb(a) \leftrightarrow EGFR(c_p^1,d), Grb(a^1)$	r33: $Grb(b), Sos(d_u) \leftrightarrow Grb(b^1), Sos(d_u^1)$
r15: $EGFR(c_p,d^-), Grb(a) \leftrightarrow EGFR(c_p^1,d^-), Grb(a^1)$	r34: $Grb(b), Sos(d_p) \leftrightarrow Grb(b^1), Sos(d_p^1)$
r16: $IR(a,b), Ins(a) \leftrightarrow IR(a^1,b), Ins(a^1)$	r35: $Grb(b^1), Sos(d_u^1) \leftrightarrow Grb(b^1), Sos(d_p^1)$
r17: $IR(a,b^-), Ins(a) \leftrightarrow IR(a^1,b^-), Ins(a^1)$	r36: $Sos(d_u) \leftrightarrow Sos(d_p)$
r18: $IR(a,b), Ins(a) \leftrightarrow IR(a,b^1), Ins(a^1)$	r37: $Shc(b_u) \leftrightarrow Shc(b_p)$
r19: $IR(a^-,b), Ins(a) \leftrightarrow IR(a^-,b^1), Ins(a^1)$	r38: $IRS(b_u) \leftrightarrow IRS(b_p)$

Table 1: Rule set for the *EGFR*/Insulin crosstalk.

has already been recruited). The site c of the *IR* can be phosphorylated (r20,r21,r22,r23) at a rate which depends on the number of recruited insulin molecules (in practice the rates of rules r21 and r22 are the same). Then, *IR* can recruit an adapter *Shc* (r24). Whenever *IR* is also bound to two insulin molecules, *Shc* can be phosphorylated (r25). *Shc* can then recruit *Grb* (r11). Yet, *IR* has another way of recruiting a *Grb*. The site d of *IR* (r26,r27,r28,r29), and then recruit another adapter called *IRS* (r30) which can be activated when the insulin receptor is bound to two insulin molecules (r31). Then, *IRS* can recruit *Grb* (r32).

Independently, *Grb* can recruit a protein *Sos* (r33,r34). And *Sos* can be activated (r35,r36) at the rate which may depend on the fact whether it is bound to a *Grb*, or not. Other rules describe the recruitment of *Sos* by *Grb*. And spontaneous (de)phosphorylation of *Shc* (r37) and *IRS* (r38).

In this model, 2,768 different complexes may occur. This number is mainly due to the fact that each dimer made of two proteins *EGFR* has 4 sites (the sites b and c for each *EGFR*) to recruit a *Grb*, which induces a small combinatorial blow up. Scanning the set of rules, one can notice that no rule tests both the site a and b of some proteins *Grb*. Thus, the partition $\{\{a\}, \{b\}\}$ can be used safely for the sites of *Grb*, in the annotated contact map. As a consequence, the number of fragments is only 609. Unfortunately,

this is the only reduction that we can do (ie, the partition for the sites of any other kind of proteins, is the coarsest one).

Last, one can notice that, given some additional hypotheses on the rate of some rule, that the sites a and b of IR have a symmetric role in the system. We could consider this symmetry to reduce the set of considered fragments, by identifying two symmetric fragments, such as $IR(a^1, b)$, $Ins(a^1)$ and $IR(a, b^1)$, $Ins(a^1)$.

6 Conclusions

Reducing the complexity of combinatorial reaction mixtures is an important milestone towards simulation and analysis of large-scale realistic models of cellular signal transduction. In this paper we study a scalable reduction method, that is applicable to any rule-based specification. The reduction is sound and moreover complete, i.e. the sample traces of individual molecular species can be reconstructed from the traces of aggregated species in the reduced model. We put this method into the general context of abstractions of probabilistic transition system and show that it yields a sufficient condition for weak lumpability and that it is equivalent to backward Markov bisimulation. The reduction factor depends on the number of independent molecular events and is strictly smaller than that of the less-demanding reduction based on the differential semantics.

A compelling problem for future work is thus to analyze differential fragments in the context of stochastic semantics and to obtain error bounds for this reduction as a function of the kinetic parameters of the system.

References

- [1] Christel Baier, Boudewijn Haverkort, Holger Hermanns & Joost-Pieter Katoen (2003): *Model-checking algorithms for continuous-time Markov chains*. *IEEE Transactions on Software Engineering* 29(7), p. 2003.
- [2] Christel Baier & Holger Hermanns (1999). *Weak Bisimulation for Fully Probabilistic Processes*.
- [3] Michael L. Blinov, James R. Faeder & William S. Hlavacek (2004): *BioNetGen: software for rule-based modeling of signal transduction based on the interactions of molecular domains*. *Bioinformatics* 20, pp. 3289–3292.
- [4] Peter Buchholz (1994): *Exact and Ordinary Lumpability in Finite Markov Chains*. *Journal of Applied Probability* 31, no1, pp. 59–75.
- [5] Peter Buchholz (2008): *Bisimulation relations for weighted automata*. *Theoretical Computer Science* Volume 393, Issue 1-3, pp. 109–123.
- [6] Holger Conzelmann, Dirk Fey & Ernst D. Gilles (2008): *Exact model reduction of combinatorial reaction*. *BMC Syst Biol* 2(78), pp. 342–351.
- [7] Vincent Danos, Jérôme Feret, Walter Fontana, Russel Harmer & Jean Krivine (2008): *Rule-based modelling, symmetries, refinements*. In: *FMSB'08*, LNBI.
- [8] Vincent Danos & Cosimo Laneve (2003): *Core Formal Molecular Biology*. *Theoretical Computer Science* 325, pp. 69–110.
- [9] Josée Desharnais, Abbas Edalat & Prakash Panangaden (2002): *Bisimulation for labelled Markov processes*. *Inf. Comput.* 179(2), pp. 163–193.
- [10] Laurent Doyen, Thomas A. Henzinger & Jean-Francois Raskin (2008): *Equivalence of Labeled Markov Chains*. *Int. J. Found. Comput. Sci.* 19(3), pp. 549–563. Available at <http://dx.doi.org/10.1142/S0129054108005814>.

- [11] Jérôme Feret, Vincent Danos, Jean Krivine, Russ Harmer & Walter Fontana (2009): *Internal coarse-graining of molecular systems*. *Proc. Natl. Acad. Sci. USA* 106(16), pp. 6453–6458. Available at <http://dx.doi.org/10.1073/pnas.0809908106>.
- [12] Jerome Feret, Heinz Koepl & Tatjana Petrov: *Stochastic fragments: A framework for the exact reduction of the stochastic semantics of rule-based models*. *International Journal of Software and Informatics* To appear.
- [13] Donald Gross & Douglas R. Miller (1984): *The Randomization Technique as a Modeling Tool and Solution Procedure for Transient Markov Processes*. *Operations Research* 32, no2, pp. 343–361.
- [14] Thomas A. Henzinger, Maria Mateescu & Verena Wolf (2009): *Sliding Window Abstraction for Infinite Markov Chains*. In: *CAV*, pp. 337–352.
- [15] Holger Hermanns (2002): *Interactive Markov Chains And the Quest for Quantified Quality*. Ph.D. thesis, University of Marburg.
- [16] Jane Hillston (1996): *A compositional approach to performance modelling*. Cambridge University Press, New York, NY, USA.
- [17] William S. Hlavacek, James R. Faeder, Michael L. Blinov, Alan S. Perelson & Byron Goldstein (2005): *The complexity of complexes in signal transduction*. *Biotechnol. Bio-eng.* 84, pp. 783–794.
- [18] William S. Hlavacek, James R. Faeder, Michael L. Blinov, Richard G. Posner, Michael Hucka & Walter Fontana (2006): *Rules for Modeling Signal-Transduction Systems*. *Science's STKE* 2006(344).
- [19] D. Kannan Jianjun Paul Tian (2006): *Lumpability and Commutativity of Markov Processes*. *Stochastic analysis and Applications* 24, no3, pp. 685–702.
- [20] J. G. Kemeny, J. L. Snell & A. W. Knapp (1976): *Denumerable Markov Chains*. Springer-Verlag, New York, NY, USA.
- [21] John Kemeny & James L. Snell (1960): *Finite Markov Chains*. Van Nostrand.
- [22] Thomas G. Kurtz (1971): *Limit Theorems for Sequences of Jump Markov Processes Approximating Ordinary Differential Processes*. *Journal of Applied Probability* 8(2), pp. 344–356.
- [23] Kim G. Larsen & Arne Skou (1989): *Bisimulation through probabilistic testing (preliminary report)*. In: *POPL '89*, ACM, New York, NY, USA, pp. 344–352.
- [24] Donald A. McQuarrie (1967): *Stochastic approach to chemical kinetics*. *Journal of Applied Probability* 4(3), pp. 413–478.
- [25] Elaine Murphy, Vincent Danos, Jerome Feret, Russell Harmer & Jean Krivine (2009): *Rule Based Modelling and Model Refinement*. In: *Elements of Computational Systems Biology*, Wiley.
- [26] Gerardo Rubino & Bruno Sericola (1991): *A Finite Characterization of Weak Lumpable Markov Processes. Part I: The Discrete Time Case*. *Stochastic processes and their applications* 38, pp. 195–204.
- [27] Gerardo Rubino & Bruno Sericola (1993): *A Finite Characterization of Weak Lumpable Markov Processes. Part II: The Continuous Time Case*. *Stochastic processes and their applications* 45, pp. 115–125.
- [28] Roger B. Sidje, Kevin Burrage & Shev MacNamara (2007): *Inexact Uniformization Method for Computing Transient Distributions of Markov Chains*. *SIAM Journal on Scientific Computing* 29, issue 6, pp. 2562–2580.
- [29] Ana Sokolova & Erik P. de Vink (2003): *On relational properties of lumpability*. In: *Proceedings of the 4th PROGRESS*.
- [30] Jeremy Sproston & Susanna Donatelli (2004): *Backward Stochastic Bisimulation in CSL Model Checking*. In: *QEST '04*, Washington, DC, USA, pp. 220–229.
- [31] Christopher T. Walsh (2006): *Posttranslation Modification of Proteins: Expanding Nature's Inventory*. Roberts and Co. Publisher.