Optimality-Preserving Reduction of Chemical Reaction Networks

Kim G. Larsen[1], Daniele Toller[1], Mirco Tribastone[2], Max Tschaikowski[1], and Andrea Vandin[3,4(✉)]

[1] Aalborg University, Aalborg, Denmark
[2] IMT School for Advanced Studies Lucca, Lucca, Italy
[3] Sant'Anna School of Advanced Studies, Pisa, Italy
andrea.vandin@santannapisa.it
[4] DTU Technical University of Denmark, Kongens Lyngby, Denmark

Abstract. Across many disciplines, chemical reaction networks (CRNs) are an established population model defined as a system of coupled nonlinear ordinary differential equations. In many applications, for example, in systems biology and epidemiology, CRN parameters such as the kinetic reaction rates can be used as control inputs to steer the system toward a given target. Unfortunately, the resulting optimal control problem is nonlinear, therefore, computationally very challenging. We address this issue by introducing an optimality-preserving reduction algorithm for CRNs. The algorithm partitions the original state variables into a reduced set of macro-variables for which one can define a reduced optimal control problem with provably identical optimal values. The reduction algorithm runs with polynomial time complexity in the size of the CRN. We use this result to reduce verification and control problems of large-scale vaccination models over real-world networks.

1 Introduction

The interplay between control theory and systems biology is instrumental to gain insights into the dynamics of natural systems across different scales (e.g., [3]). In particular, the problem of controlling a biological system is relevant in applications such as smart therapeutics and biosensors [42]. Mathematically, this can be studied as the problem of controlling a formal chemical reaction network (CRN), whereby the biological system under study is modeled as a finite set of species (i.e., agents) that interact across a finite set of reactions. This representation admits both a stochastic interpretation in terms of a continuous-time Markov chain (CTMC), where discrete changes in the population levels of each species are tracked, and a deterministic one as a system of nonlinear ordinary differential equations (ODEs), where each equation tracks the time evolution of the concentration of each species. Notably, and particularly relevant for the theoretical developments in this paper, under mild conditions the deterministic equations correspond to a limit regime of a family of CTMCs (e.g., [25,28,30,48,52]).

T. Margaria and B. Steffen (Eds.): ISoLA 2024, LNCS 15220, pp. 13–32, 2025.
https://doi.org/10.1007/978-3-031-75107-3_2

In this setting, the control inputs may be represented by the parameter values of designated reaction rates [4], such that the overall controller design may be studied as an optimal control problem. Treating certain rates as inputs can also be used for the complementary goal of studying the open-loop behavior of the system when some parameters are unknown/uncertain, by estimating reachable sets [34]; this is a pressing problem in systems biology, where rate parameters are often not directly accessible.

Controlling the biological system by studying its ODEs in place of the CTMC is appealing because the ODE system size has, in general, exponentially fewer equations. However, the control problem is computationally prohibitive in general due to the fact that it is nonlinear [14,22,37]. One approach to tackling this problem is to devise an *optimality-preserving* reduction of a control system, where the hope is to solve a reduced optimal control problem instead of the original one. While for linear systems this problem is well-understood [37], it remains challenging for nonlinear ones.

In this paper we consider CRNs with the well-known mass-action semantics (e.g., [13,54]), leading to ODE systems with polynomial right-hand sides. Here, reactions are characterized by rate parameters which can be used as inputs, taken from bounded domains. The optimal control problem consists in finding the values of those parameters such that a given cost function is minimized. We present an optimality-preserving reduction method based on a partition of the set of species, thus corresponding to a partition of the set of ODE variables. This follows a long tradition in the development of *lumping* techniques for (bio-)chemical systems (e.g., [36,43]), most of which are concerned with preserving the dynamics of the system and not of the solution of the control problem as done here.

Our reduction is exact in the sense that one can define a reduced optimal control problem whose solution can be exactly related to that of the original problem. Based on this, we develop an algorithm that finds the coarsest partition, i.e., the maximal lumping, that satisfies this property. The algorithm is based on previous recent results for lumping of uncertain Markov chains (essentially seen as linear control systems [5]). Similarly to that, the number of required computational steps is at most polynomial in the number of species and reactions of the CRN. However, the technical machinery required here is profoundly different and, importantly, identifies the polynomial ODE system of a mass-action CRN as the deterministic limit process of a family of CTMCs. Specifically, in the derivation of the main result, visualized in Fig. 1, we:

- make use of fluid limit results [7,30,51] and associate to each controlled CRN (CCRN) a family of continuous-time Markov chains (CTMCs) which, roughly speaking, converge in probability to the control system of the CCRN;
- show that the original CTMC family can be replaced by the CTMC family of a lumped CCRN while preserving optimality by relying on [10,11];
- show that the control system of the original and the lumped CCRN have common optimal values by essentially combining [7,11], as visualized in Fig. 1.

By doing so, we circumvent the problem of having to relate nonlinear control systems directly.

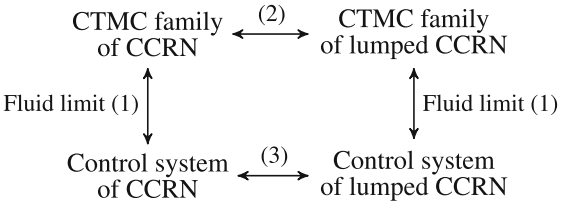

Fig. 1. Visualization of the main result. As first, we approximate the control systems of the original and the lumped CCRN by means of suitable CTMC families (1). Then, we show that the lumped CTMC family admits the same optimal value as the original one (2). Combining (1) and (2), we conclude that the lumping of the control system is optimality preserving (3).

We implement the lumping algorithm in the software tool ERODE [15] and conduct a large-scale evaluation. Specifically, we study epidemiological models over weighted networks [39], where a) each node is subject to vaccination control or b) the network weights are subject to uncertainty. We took 1558 real world networks from the Netzschleuder repository [40], where the largest network had 51919 nodes.

Related Work. While results on exact optimality-preserving lumping techniques for linear control systems have been explored (see [5,37] and references therein), nonlinear counterparts are scarce. Akin to notions of behavioral equivalence in concurrency theory [18,27,32,47,49], the bisimulation of control systems [38,44,53] is closely related but complementary to CCRN species equivalence. Specifically, for a given observation, the largest bisimulation gives rise to a lumped dynamical system which coincides with the original one up to a previously chosen observation map. Instead, CCRN species equivalence seeks to find from a family of linear observation maps the one that gives rise to the largest bisimulation. Since only observation maps expressible by equivalence relations are considered, the coarsest CCRN species equivalence can be computed in polynomial time. Apart from bisimulation/abstraction, we mention decoupling [20] that yields substantial speed-ups but may impose restrictive symmetry constraints. While this can be addressed by decoupling approaches [19], the corresponding lumping is approximative. It is worth noting that our approach is reminiscent to Koopman operator theory which expresses a nonlinear system via an infinite linear one [41]. Fluid limits are however complementary to Koopman operator theory. This is because they rely upon probabilistic arguments and their linear (Markov chain) approximations hold true for arbitrary initial conditions, rather than specific ones [41].

Paper Outline. After reviewing CRNs, Sect. 2 introduces controlled CRNs (CCRNs). Section 3 instead reviews CRN species equivalence from [12] and extends it to CCRNs. Building upon Sect. 3, Sect. 4 establishes that CCRN species equivalence allows for optimality-preserving lumping of fluid models, while Sect. 5 presents applications in nonlinear system verification and control. The paper concludes in Sect. 6.

2 Controlled CRNs

A mass-action CRN is $(\mathcal{S}, \mathcal{R}_\alpha)$ where $\mathcal{S}$ is a set of species, $\mathcal{R}_\alpha$ is a set of reactions and $\alpha = (\alpha_{i_r})_{r \in \mathcal{R}}$ is a set of kinetic parameters, with $\alpha_{i_r} \geq 0$. Each reaction $r = \pi \xrightarrow{\alpha_{i_r}} \rho$ comprises multisets π and ρ of species, denoting the reactants and products, respectively. Mass-action CRNs are traditionally given both a stochastic and a deterministic interpretation as a Markov jump process and a system of polynomial differential equations.

In the stochastic interpretation [30], a state of the underlying Markov chain is a species multiset $\sigma \in \mathbb{N}_0^{\mathcal{S}}$ giving the number of molecules $\sigma(A)$ for each species $A \in \mathcal{S}$, where $\mathbb{N}_0$ are the naturals with zero. The forward equations are given by the initial value problem

$$\partial_t p_\sigma = \sum_\theta q(\sigma, \theta) p_\sigma, \tag{1}$$

where $p(0)$ is the initial probability measure, whereas the transition rate from state σ to state θ is

$$q(\sigma, \theta) = q_\alpha(\sigma, \theta) := \sum_{\substack{r = \rho \xrightarrow{\alpha_{i_r}} \pi \in \mathcal{R}_\alpha \\ \theta = \sigma + \pi - \rho}} \alpha_{r_i} \cdot \binom{\sigma}{\rho} \tag{2}$$

$q = (q(\sigma, \theta))_{\sigma,\theta}$ is the *transition rate matrix*, where $q(\sigma, \sigma) = -\sum_{\theta \neq \sigma} q(\sigma, \theta)$. The dynamics can be described as follows: when in state σ, every reaction determines a possible jump that consumes molecules according to the multiplicities of the reactants and yields new molecules according to the products; the reaction fires proportionally (via the kinetic rate α_i) to the total number of possible encounters between single molecules of the reacting species. With this in place, the CTMC described by (1) is denoted by $(X^q(t))_{t \geq 0}$.

In the deterministic interpretation [31], instead, the model is described by the system of polynomial differential equations $\partial_t v = f(v, \alpha)$, where the vector field $f : \mathbb{R}_{\geq 0}^{\mathcal{S}} \times \mathbb{R}_{\geq 0}^{|\mathcal{R}|} \to \mathbb{R}^{\mathcal{S}}$ is given, for any species $A \in \mathcal{S}$, by

$$f_A(v, \alpha) := \sum_{r = \rho \xrightarrow{\alpha_{r_i}} \pi \in \mathcal{R}_\alpha} \alpha_{r_i} (\pi(A) - \rho(A)) \prod_{B \in \mathcal{S}} \frac{v_B^{\rho(B)}}{\rho(B)!}, \tag{3}$$

with $\rho(B)!$ denoting the factorial of $\rho(B)$. Under certain assumptions, it can be shown that the stochastic model converges in probability to the deterministic

model, as the molecule counts tend to infinity. These are commonly known as fluid limit results [8,24], as discussed in Sect. 4.

We now introduce the notion of *controllable CRN* (CCRN), for which we likewise give both a stochastic and a deterministic control system. In both cases we consider two extremal CRNs $(\mathcal{S}, \mathcal{R}_{\underline{\alpha}})$ and $(\mathcal{S}, \mathcal{R}_{\overline{\alpha}})$, with $\underline{\alpha} \leq \overline{\alpha}$, which constrain the values that the decision variables (i.e., the control inputs) may attain.

- The stochastic control system is given by (1), where each $q(\sigma, \theta)$ becomes a measurable control input bounded by the corresponding values in the extremal CRNs, that is

$$q(\sigma, \theta) = q_{\underline{\alpha},\overline{\alpha}}(\sigma, \theta) : \mathbb{R}_{\geq 0} \to [q_{\underline{\alpha}}(\sigma, \theta); q_{\overline{\alpha}}(\sigma, \theta)] \tag{4}$$

  The resulting family of CTMCs is denoted by $\left(\mathbb{N}_0^{\mathcal{S}}, [q_{\underline{\alpha}}; q_{\overline{\alpha}}]\right)$ and is called the *uncertain CTMC* (UCTMC) of a CCRN.[1]
- Likewise, in the deterministic control system, each kinetic parameter in (3) becomes a control input bounded by the corresponding values in the extremal CRNs, that is a measurable $\alpha_{i_r} : \mathbb{R}_{\geq 0} \to [\underline{\alpha}_{i_r}; \overline{\alpha}_{i_r}]$. Moreover, for any bounded set of initial conditions $I \subseteq \mathbb{R}_{\geq 0}^{\mathcal{S}}$ and time $t \geq 0$, we define the set of states reachable from I at time t as

$$\mathfrak{R}(t) = \{v(t) \mid \partial_t v = f(v, \alpha), v(0) \in I, \alpha \in [\underline{\alpha}; \overline{\alpha}]\}$$

  The initial set I allows to account for uncertainty in the initial condition and encapsulates as special case the singleton set. For a given $\alpha : \mathbb{R}_{\geq 0} \to [\underline{\alpha}; \overline{\alpha}]$, we shall write v^α for the solution of $\partial_t v = f(v, \alpha)$, where $v(0) \in I$ is assumed to be given.

We shall adhere to the following notation.

Remark 1. Since q is $q_{\underline{\alpha},\overline{\alpha}}$ from (4) rather than q_α from (2) in all but few cases, q shall refer to $q_{\underline{\alpha},\overline{\alpha}}$ unless stated otherwise. Also, we shall write α_i rather than α_{i_r} to increase readability.

Remark 2. In general, ensuring that the forward Eq. (1) is regular in the sense that it admits a unique solution for every initial probability distribution $p(0)$ is nontrivial because the state space $\mathbb{N}_0^{\mathcal{S}}$ is infinite. A common way to ensure regularity is to prove that the CTMC is non-explosive by means of stochastic Lyapunov conditions [23,35].

We call a CCRN/UCTMC regular if it induces regular CTMCs only. While some of our results assume regularity, optimality-preserving lumping from Sect. 4 does not.

[1] We use the name from [10], even though the name *controlled CTMC* would be appropriate too.

3 Species Equivalence of Controlled CRNs

We use the following CCRN as running example.[2]

Example 1. Consider the CCRN $(\mathcal{S}, \mathcal{R}_{[\underline{\alpha};\overline{\alpha}]})$ with species $\mathcal{S} = \{B, A_{00}, A_{01}, A_{10}, A_{11}\}$ and reactions

$$A_{00} + B \xrightarrow{[\underline{\alpha}_1;\overline{\alpha}_1]} A_{10}, \qquad A_{10} \xrightarrow{[\underline{\alpha}_2;\overline{\alpha}_2]} A_{00} + B,$$
$$A_{00} + B \xrightarrow{[\underline{\alpha}_3;\overline{\alpha}_3]} A_{01}, \qquad A_{01} \xrightarrow{[\underline{\alpha}_4;\overline{\alpha}_4]} A_{00} + B,$$
$$A_{10} + B \xrightarrow{[\underline{\alpha}_5;\overline{\alpha}_5]} A_{11}, \qquad A_{11} \xrightarrow{[\underline{\alpha}_6;\overline{\alpha}_6]} A_{10} + B,$$
$$A_{01} + B \xrightarrow{[\underline{\alpha}_7;\overline{\alpha}_7]} A_{11}, \qquad A_{11} \xrightarrow{[\underline{\alpha}_8;\overline{\alpha}_8]} A_{01} + B.$$

The reactions model reversible binding of species B to a substrate A with two binding sites. Subscripts i, j in chemical species A_{ij} denote the availability of either binding site in the substrate A, while the value on each arrow indicates the kinetic rate parameter. For state $\sigma = A_{01} + A_{10} + B$, these yield

$$q_{\sigma,A_{11}+A_{10}}(\cdot) \in [\underline{\alpha}_7;\overline{\alpha}_7], \qquad q_{\sigma,A_{00}+A_{10}+2B}(\cdot) \in [\underline{\alpha}_4;\overline{\alpha}_4],$$
$$q_{\sigma,A_{11}+A_{01}}(\cdot) \in [\underline{\alpha}_5;\overline{\alpha}_5], \qquad q_{\sigma,A_{00}+A_{01}+2B}(\cdot) \in [\underline{\alpha}_2;\overline{\alpha}_2].$$

In the following, we make the common assumption [9] that the uncertainty intervals do not depend on the binding site, that is, $[\underline{\alpha}_i;\overline{\alpha}_i] = [\underline{\alpha}_{i+1};\overline{\alpha}_{i+1}]$ for $i \in \{1, 2, 5, 6\}$.

3.1 Lumping of CCRNs

Ordinary lumpability is a partition $\mathcal{H}$ of the state space such that any two states i, j in each partition block $H \in \mathcal{H}$ have equal aggregate rates toward states in any block $H' \in \mathcal{H}$. That is, writing $\mathfrak{q}(\cdot)$ for the transition rates of a generic CTMC that is not necessarily related to (1), it must hold that $\sum_{k\in H'} \mathfrak{q}_{i,k}(\cdot) = \sum_{k\in H'} \mathfrak{q}_{j,k}(\cdot)$, where $(\cdot)$ emphasizes the dependence on time. Given an ordinarily lumpable partition, a lumped CTMC can be constructed by associating a macro-state to each block. Transitions between macro-states are labeled with the overall rate from a state in the source block toward all states in the target block.

Checking the conditions for ordinary lumpability requires the full enumeration of the CTMC state space which grows combinatorially in the multiplicities of the initial state and may be even infinite in presence of species creation (e.g., $A \xrightarrow{\alpha} A + A$). Species equivalence [12] addresses this by detecting ordinary lumpability at the level of the reaction network. To this end, it identifies an equivalence relation which induces an ordinary lumpable partition over the multisets representing CTMC states. Specifically, one considers a natural lifting of a

[2] For the benefit of presentation, the appendix features a table with frequently used symbols.

partition $\mathcal{H}$ of species to multisets of species, called *multiset lifting* of $\mathcal{H}$, denoted by $\mathcal{H}^{\uparrow}$.

Definition 1 (Multiset Lifting). *Let $(\mathcal{S}, \mathcal{R}_{[\underline{\alpha};\overline{\alpha}]})$ be a CCRN, a partition $\mathcal{H}$ over $\mathcal{S}$ and let $\mathcal{E}$ be the equivalence relation of $\mathcal{H}$, i.e., $\mathcal{H} = \mathcal{S}/\mathcal{E}$. We define the multiset lifting of $\mathcal{E}$ on $\mathbb{N}_0^{\mathcal{S}}$, denoted by $\mathcal{E}^{\uparrow} \subseteq \mathbb{N}_0^{\mathcal{S}} \times \mathbb{N}_0^{\mathcal{S}}$, as*

$$\left\{(\sigma_1, \sigma_2) \in \mathbb{N}_0^{\mathcal{S}} \times \mathbb{N}_0^{\mathcal{S}} \mid \forall H \in \mathcal{H}.\ \sum_{A \in H} \sigma_1(A) = \sum_{A \in H} \sigma_2(A)\right\}$$

With this, we set $\mathcal{H}^{\uparrow} = \mathbb{N}_0^{\mathcal{S}}/\mathcal{E}^{\uparrow}$.

Intuitively, the multiset lifting relates multisets that have same cumulative multiplicity from each partition block.

Example 2. In Example 1, consider $\mathcal{H} = \{\{B\}, \{A_{00}\}, \{A_{01}, A_{10}\}, \{A_{11}\}\}$ and let $\mathcal{E}$ be such that $\mathcal{H} = \mathcal{S}/\mathcal{E}$. Then, $(A_{01}, A_{10}) \in \mathcal{E}^{\uparrow}$, $(2A_{01}, A_{01} + A_{10}) \in \mathcal{E}^{\uparrow}$, while $(A_{00}, A_{10}) \notin \mathcal{E}^{\uparrow}$ and $(2A_{01}, A_{10}) \notin \mathcal{E}^{\uparrow}$. That is, two species are equivalent w.r.t. $\mathcal{E}^{\uparrow}$ if they agree on the number of occupied binding sites. More formally, $(\sigma, \sigma') \in \mathcal{E}^{\uparrow}$ if $\sigma(C) = \sigma'(C)$ for all $C \in \{A_{00}, A_{11}\}$ and $\sigma(A_{01}) + \sigma(A_{10}) = \sigma'(A_{01}) + \sigma'(A_{10})$.

We first review the notion of CRN species equivalence from [12].

Definition 2 (CRN Species Equivalence). *Fix a CRN $(\mathcal{S}, \mathcal{R}_{\alpha})$. We call a partition $\mathcal{H}$ of $\mathcal{S}$ a CRN species equivalence if, for any two species A_i, A_j in a block of $\mathcal{H}$, any reagent $\rho \in \mathbb{N}_0^{\mathcal{S}}$, any block $H^{\uparrow} \in \mathcal{H}^{\uparrow}$, we have*

$$\sum_{\pi \in H^{\uparrow}} \mathbf{rr}_{\alpha}(A_i + \rho, \pi) = \sum_{\pi \in H^{\uparrow}} \mathbf{rr}_{\alpha}(A_j + \rho, \pi) \tag{5}$$

Here, $\mathbf{rr}_{\alpha}$ is the reaction rate from ρ to π

$$\mathbf{rr}_{\alpha}(\rho, \pi) = \begin{cases} \sum\limits_{(\rho \xrightarrow{\alpha_i} \pi) \in \mathcal{R}_{\alpha}} \alpha_i & ,\ \rho \neq \pi \\ -\sum\limits_{\pi' \neq \rho} \mathbf{rr}_{\alpha}(\rho, \pi') & ,\ \rho = \pi \end{cases}$$

For any $H^{\uparrow} \subseteq \mathbb{N}_0^{\mathcal{S}}$, we set $\mathbf{rr}_{\alpha}[\rho, H^{\uparrow}] = \sum_{\pi \in H^{\uparrow}} \mathbf{rr}_{\alpha}(\rho, \pi)$.

Any CRN species equivalence induces a lumped CRN given next.

Definition 3 (Lumped CRN). *Let $(\mathcal{S}, \mathcal{R}_{\alpha})$ be a CRN, $\mathcal{H}$ a CRN species equivalence and fix a representative $A_H \in H$ for each $H \in \mathcal{H}$. The lumped CRN is then given by $(\hat{\mathcal{S}}, \hat{\mathcal{R}}_{\hat{\alpha}})$, where the species are $\hat{\mathcal{S}} = \{A_H \mid H \in \mathcal{H}\}$, while reactions $\hat{\mathcal{R}}_{\hat{\alpha}}$ arise via*

1. *discard all reactions $\rho \xrightarrow{\alpha_i} \pi$ where ρ has a nonrepresentative species;*
2. *replace species in products of remaining reactions by their representatives;*

3. *fuse all reactions that have the same reactants and products by summing their rates.*

With the foregoing definitions in place, CCRN species equivalence is defined as the CRN species equivalence of the extremal CRNs $(\mathcal{S}, \mathcal{R}_{\underline{\alpha}})$ and $(\mathcal{S}, \mathcal{R}_{\overline{\alpha}})$.

Definition 4 (CCRN Species Equivalence). *Fix a CCRN $(\mathcal{S}, \mathcal{R}_{[\underline{\alpha};\overline{\alpha}]})$. We call a partition $\mathcal{H}$ of $\mathcal{S}$ a CCRN species equivalence whenever $\mathcal{H}$ is a CRN species equivalence of $(\mathcal{S}, \mathcal{R}_{\underline{\alpha}})$ and $(\mathcal{S}, \mathcal{R}_{\overline{\alpha}})$.*

Our example enjoys a CCRN species equivalence.

Example 3. Continuing Example 1 and 2, we note that for

$$\alpha = (\alpha_1, \ldots, \alpha_8) \in \{(\underline{\alpha}_1, \ldots, \underline{\alpha}_8), (\overline{\alpha}_1, \ldots, \overline{\alpha}_8)\},$$

it holds that $\mathbf{rr}_\alpha(A_{01}, A_{00} + B) = \mathbf{rr}_\alpha(A_{10}, A_{00} + B)$ and $\mathbf{rr}_\alpha(A_{01} + B, A_{11}) = \mathbf{rr}_\alpha(A_{10} + B, A_{11})$. Hence, $\mathcal{H}$ is a CCRN species equivalence.

The lumped CCRN is given by the lumpings of the extremals, as stated next.

Definition 5 (Lumped CCRN). *Let $(\mathcal{S}, \mathcal{R}_{[\underline{\alpha};\overline{\alpha}]})$ be a CCRN and $\mathcal{H}$ a CCRN species equivalence. The lumped CCRN $(\hat{\mathcal{S}}, \hat{\mathcal{R}}_{[\hat{\underline{\alpha}};\hat{\overline{\alpha}}]})$ arises by lumping the extremal CRNs $(\mathcal{S}, \mathcal{R}_{\underline{\alpha}})$ and $(\mathcal{S}, \mathcal{R}_{\overline{\alpha}})$ as outlined in Definition 3.*

We remark that the lumped CCRN does not depend on the choice of the representative [12]. As next, we provide the lumped CCRN of our example.

Example 4. Continuing Example 1–3, the lumped CCRN is given by $\hat{\mathcal{S}} = \mathcal{S} \setminus \{A_{10}\}$ and $\hat{\mathcal{R}}_{[\hat{\underline{\alpha}};\hat{\overline{\alpha}}]}$ such that

$$A_{00} + B \xrightarrow{[\underline{\alpha}_1+\underline{\alpha}_3;\overline{\alpha}_1+\overline{\alpha}_3]} A_{01}, \qquad A_{01} \xrightarrow{[\underline{\alpha}_4;\overline{\alpha}_4]} A_{00} + B,$$
$$A_{11} \xrightarrow{[\underline{\alpha}_6+\underline{\alpha}_8;\overline{\alpha}_6+\overline{\alpha}_8]} A_{01} + B, \qquad A_{01} + B \xrightarrow{[\underline{\alpha}_7;\overline{\alpha}_7]} A_{11}$$

The CCRN species equivalence can be computed by invoking alternately the CRN lumping algorithm from [12] on the extremal CRNs that define a CCRN, as stated next.

Theorem 1 (Computation of CCRN Species Equivalence). *Let $(\mathcal{S}, \mathcal{R}_{[\underline{\alpha};\overline{\alpha}]})$ be a CCRN. Then we have the following.*

1. *$\mathcal{H}$ is a CCRN species equivalence iff $\mathcal{H}^{\uparrow}$ is an ordinary lumpability of the CTMCs $(\mathbb{N}_0^{\mathcal{S}}, q_{\underline{\alpha}})$ and $(\mathbb{N}_0^{\mathcal{S}}, q_{\overline{\alpha}})$.*
2. *For any partition $\mathcal{G}$ of $\mathcal{S}$, Algorithm 1 computes the coarsest CCRN species equivalence of $(\mathcal{S}, \mathcal{R}_{[\underline{\alpha};\overline{\alpha}]})$ that refines $\mathcal{G}$. That is, $\mathcal{H}$ is such that*
   - *for every block $G \in \mathcal{G}$, there exist unique blocks $H_1, \ldots, H_l \in \mathcal{H}$ such that $G = H_1 \cup \ldots \cup H_l$ and;*

Algorithm 1 Algorithm for the computation of the coarsest CCRN species equivalence that refines some given partition $\mathcal{G}$.

Require: CCRN $(\mathcal{S}, \mathcal{R}_{[\underline{\alpha};\overline{\alpha}]})$ and a partition $\mathcal{G}$ of $\mathcal{S}$
  while true do
    compute using the algorithm in [12] the coarsest CCRN species equivalence of $(\mathcal{S}, \mathcal{R}_{\underline{\alpha}})$ that refines $\mathcal{G}$, **store** it in $\mathcal{H}$
    compute using the algorithm in [12] the coarsest CCRN species equivalence of $(\mathcal{S}, \mathcal{R}_{\overline{\alpha}})$ that refines $\mathcal{H}$, **store** it in $\mathcal{H}'$
    if $\mathcal{H}' = \mathcal{G}$ **then**
      return $\mathcal{G}$
    else
      $\mathcal{G} \leftarrow \mathcal{H}'$
    end if
  end while

– *$\mathcal{H}$ is a CCRN species equivalence and has a minimal number of blocks, hence a lumped CCRN of minimal size.*

The number of steps performed by Algorithm 1 is polynomial in $|\mathcal{S}|$ and $|\mathcal{R}_{\underline{\alpha}}|$.

After addressing the computation of CCRN species equivalence, we observe next that the block sums of original CCRN states are equivalent in distribution to the states of the lumped CCRN. Equivalence in distribution is denoted by $\overset{\mathrm{d}}{=}$.

Theorem 2 (Stochastic CCRN Lumping). *Let $\mathcal{H}$ be a CCRN species equivalence of $(\mathcal{S}, \mathcal{R}_{[\underline{\alpha};\overline{\alpha}]})$ and let $(\hat{\mathcal{S}}, \hat{\mathcal{R}}_{[\underline{\hat{\alpha}};\hat{\overline{\alpha}}]})$ be the respective lumped CCRN. Moreover, let X^q and $\hat{X}^{\hat{q}}$ denote family members of the respective UCTMCs. Then, if both CCRNs are regular:*

1. *For any $q(\cdot) \in [q_{\underline{\alpha}}; q_{\overline{\alpha}}]$, there exists a $\hat{q}(\cdot) \in [q_{\underline{\hat{\alpha}}}; q_{\hat{\overline{\alpha}}}]$ such that*

$$\sum_{A \in H} X^q_A(t) \overset{d}{=} \hat{X}^{\hat{q}}_{A_H}(t), \quad \forall H \in \mathcal{H} \forall\, t > 0, \tag{6}$$

   provided the statement holds for $t = 0$.
2. *Conversely, for any $\hat{q}(\cdot) \in [q_{\underline{\hat{\alpha}}}; q_{\hat{\overline{\alpha}}}]$, there is a $q(\cdot) \in [q_{\underline{\alpha}}; q_{\overline{\alpha}}]$ with (6).*

Provided the n-th order moments exist, Theorem 2 implies $\mathbb{E}\big[(\sum_{A \in H} X_A(t))^n\big] = \mathbb{E}\big[\hat{X}^n_{A_H}(t)\big]$.[3] The moments can be estimated by means of stochastic simulation [2].

Example 5. It can be shown that Example 1 is regular. With this, Theorem 2 essentially ensures that

– for any q, there exists a $\hat{q}$ such that $X^q_{A_{01}} + X^q_{A_{10}} \overset{\mathrm{d}}{=} \hat{X}^{\hat{q}}_{A_{01}}$ and $X^q_S \overset{\mathrm{d}}{=} \hat{X}^{\hat{q}}_S$ with $S \notin \{A_{01}, A_{10}\}$;

[3] Similarly to regularity, the existence of moments can be addressed by means of Lyapunov conditions [23,35].

– for any $\hat{q}$, there exists a q such that $X^q_{A_{01}} + X^q_{A_{10}} \stackrel{\mathrm{d}}{=} \hat{X}^{\hat{q}}_{A_{01}}$ and $X^q_S \stackrel{\mathrm{d}}{=} \hat{X}^{\hat{q}}_S$ with $S \notin \{A_{01}, A_{10}\}$.

That is, if one is only interested in species $S \notin \{A_{01}, A_{10}\}$ or the cumulative behavior of species $X^q_{A_{01}} + X^q_{A_{10}}$, any behavior of the original CCRN can be matched by the lumped CCRN and vice versa.

Example 5 demonstrates that CRN species equivalence allows one to lump the original CCRN to a smaller lumped CCRN at the expense of preserving sums of original species. For instance, if the modeler is interested in $X_{A_{00}}$ and X_B, partition $\mathcal{H}$ from Example 2 can be used because $\{A_{00}\}, \{B\} \in \mathcal{H}$. Instead, if a modeler is interested in $X_{A_{10}}$, it is not possible to use $\mathcal{H}$ because the lumped CCRN would only capture the cumulative behavior $X_{A_{10}} + X_{A_{01}}$. A natural question would be then if there is a CCRN species equivalence $\mathcal{H}'$ which contains $\{A_{10}\}$. This can be readily checked by applying the algorithm from Theorem 1 to $\mathcal{G}' = \{\{A_{10}\}, \mathcal{S} \setminus \{A_{10}\}\}$. This is because any CCRN species equivalence $\mathcal{H}'$ refining $\mathcal{G}'$ has to contain the block $\{A_{10}\}$. An application of the algorithm returns then the trivial CCRN species equivalence $\mathcal{H}' = \{\{S\} \mid S \in \mathcal{S}\}$. We call $\mathcal{H}'$ trivial because it does not lump the original CCRN.

4 Optimality-Preserving Lumping

We start by providing the deterministic control systems associated to our example and its CCRN lumping.

Example 6. The CCRN $(\mathcal{S}, \mathcal{R}_{[\underline{\alpha};\overline{\alpha}]})$ from Example 1 gives rise to the ODE system

$$\begin{aligned}
\partial_t v_{A_{00}} &= -(\alpha_1 + \alpha_3) v_{A_{00}} v_B + \alpha_2 v_{A_{10}} + \alpha_4 v_{A_{01}} \\
\partial_t v_{A_{10}} &= \alpha_1 v_{A_{00}} v_B - \alpha_2 v_{A_{10}} - \alpha_5 v_{A_{10}} v_B + \alpha_6 v_{A_{11}} \\
\partial_t v_{A_{01}} &= \alpha_3 v_{A_{00}} v_B - \alpha_4 v_{A_{01}} - \alpha_7 v_{A_{01}} v_B + \alpha_8 v_{A_{11}} \\
\partial_t v_{A_{11}} &= \alpha_5 v_{A_{10}} v_B + \alpha_7 v_{A_{01}} v_B - (\alpha_6 + \alpha_8) v_{A_{11}} \\
\partial_t v_B &= -(\alpha_1 + \alpha_3) v_{A_{00}} v_B + \alpha_2 v_{A_{10}} + \alpha_4 v_{A_{01}} \\
&\quad - \alpha_5 v_{A_{10}} v_B - \alpha_7 v_{A_{01}} v_B + (\alpha_6 + \alpha_8) v_{A_{11}}
\end{aligned} \tag{7}$$

Instead, the lumped CCRN $(\hat{\mathcal{S}}, \hat{\mathcal{R}}_{[\underline{\hat{\alpha}};\overline{\hat{\alpha}}]})$ from Example 4 gives rise to the ODE system

$$\begin{aligned}
\partial_t \hat{v}_{A_{00}} &= -\hat{\alpha}_1 \hat{v}_{A_{00}} \hat{v}_B + \hat{\alpha}_2 \hat{v}_{A_{01}} \\
\partial_t \hat{v}_{A_{10}} &= \hat{\alpha}_1 \hat{v}_{A_{00}} \hat{v}_B - \hat{\alpha}_2 \hat{v}_{A_{01}} - \hat{\alpha}_3 \hat{v}_{A_{01}} \hat{v}_B + \hat{\alpha}_4 \hat{v}_{A_{11}} \\
\partial_t \hat{v}_{A_{11}} &= \hat{\alpha}_3 \hat{v}_{A_{01}} \hat{v}_B - \hat{\alpha}_4 \hat{v}_{A_{11}} \\
\partial_t \hat{v}_B &= -\hat{\alpha}_1 \hat{v}_{A_{00}} \hat{v}_B + \hat{\alpha}_2 \hat{v}_{A_{01}} - \hat{\alpha}_1 \hat{v}_{A_{00}} \hat{v}_B + \hat{\alpha}_2 \hat{v}_{A_{01}},
\end{aligned} \tag{8}$$

where $\hat{\alpha} \in [\underline{\hat{\alpha}}; \overline{\hat{\alpha}}]$.

We next state our main results.

Theorem 3 (Deterministic CCRN Lumping). *Let us fix a CCRN $(\mathcal{S}, \mathcal{R}_{[\underline{\alpha};\overline{\alpha}]})$, a constant $c > 0$, assume that $\mathcal{H}$ is a CCRN species equivalence and denote the corresponding lumped CCRN by $(\hat{\mathcal{S}}, \hat{\mathcal{R}}_{[\hat{\underline{\alpha}};\hat{\overline{\alpha}}]})$. Writing $B(c)$ for the L_1 ball at the origin with radius c, let $T > 0$ be such that $\mathfrak{R}(t) \subseteq B(c)$ for any $t \in [0;T]$. Then for any initial condition $v(0) \in I$ and any $\varepsilon, \delta > 0$, the original and lumped deterministic models, $\partial_t v^\alpha = f(v^\alpha, v)$ and $\partial_t \hat{v}^{\hat{\alpha}} = \hat{f}(\hat{v}^{\hat{\alpha}}, \hat{v})$, enjoy the following.*

1. *For any $\alpha(\cdot) \in [\underline{\alpha};\overline{\alpha}]$, there is some $\hat{\alpha}(\cdot) \in [\hat{\underline{\alpha}};\hat{\overline{\alpha}}]$ such that $\partial_t v^\alpha = f(v^\alpha, v)$ and $\partial_t \hat{v}^{\hat{\alpha}} = \hat{f}(\hat{v}^{\hat{\alpha}}, \hat{v})$ satisfy*

$$\mathbb{P}\Big\{ \max_{H \in \mathcal{H}} \max_{0 \le t \le T} \Big| \sum_{A \in H} v_A^\alpha(t) - \hat{v}_{A_H}^{\hat{\alpha}}(t) \Big| > \varepsilon \Big\} < \delta$$

   provided that $\sum_{A \in \mathcal{H}} v_A(0) = \hat{v}_{A_H}(0)$ for all $H \in \mathcal{H}$.
2. *For any $\hat{\alpha}(\cdot) \in [\hat{\underline{\alpha}};\hat{\overline{\alpha}}]$, there is some $\alpha(\cdot) \in [\underline{\alpha};\overline{\alpha}]$ such that $\partial_t v^\alpha = f(v^\alpha, v)$ and $\partial_t \hat{v}^{\hat{\alpha}} = \hat{f}(\hat{v}^{\hat{\alpha}}, \hat{v})$ satisfy*

$$\mathbb{P}\Big\{ \max_{H \in \mathcal{H}} \max_{0 \le t \le T} \Big| \sum_{A \in H} v_A^\alpha(t) - \hat{v}_{A_H}^{\hat{\alpha}}(t) \Big| > \varepsilon \Big\} < \delta$$

   provided that $\sum_{A \in \mathcal{H}} v_A(0) = \hat{v}_{A_H}(0)$ for all $H \in \mathcal{H}$.

Let us next apply Theorem 3 to Example 6.

Example 7. If applied to Example 6 for given $\varepsilon, \delta > 0$, Theorem 3 implies

- for any α, there is an $\hat{\alpha}$ such that, with a probability of $1 - \delta$ or higher, it holds that $|v_{A_{01}}^\alpha + v_{A_{10}}^\alpha - \hat{v}_{A_{01}}^{\hat{\alpha}}| < \varepsilon$ and $|v_S^\alpha - \hat{v}_S^{\hat{\alpha}}| < \varepsilon$ for all $S \notin \{A_{01}, A_{10}\}$;
- for any $\hat{\alpha}$, there is an α such that, with a probability of $1 - \delta$ or higher, it holds that $|v_{A_{01}}^\alpha + v_{A_{10}}^\alpha - \hat{v}_{A_{01}}^{\hat{\alpha}}| < \varepsilon$ and $|v_S^\alpha - \hat{v}_S^{\hat{\alpha}}| < \varepsilon$ for all $S \notin \{A_{01}, A_{10}\}$.

Remark 3. In contrast to Theorem 2, Theorem 3 does not require the CCRN or its lumping to be regular. Rather, it requires that the reachable set $\mathfrak{R}$ of the CCRN does not exhibit an explosion on $[0;T]$, a property that can be often established for CRNs via conservation laws [3,12,54].

Since CCRN species equivalence essentially ensures that the trajectories of the fluid models of the original and lumped CCRN coincide, Theorem 3 yields the following.

Theorem 4 (Lumping preserves Optimality). *Additionally to the assumptions made in Theorem 3, introduce*

- *the differentiable running cost $L : \mathbb{R}^{\mathcal{S}} \times \mathbb{R}_{\ge 0} \to \mathbb{R}_{\ge 0}$ and final cost $K : \mathbb{R}^{\mathcal{S}} \to \mathbb{R}_{\ge 0}$;*
- *the functional $J_\alpha(v[0]) = \int_0^T L(t, v^\alpha(t))dt + K(v^\alpha(T))$, where $\partial_t v^\alpha = f(v^\alpha, \alpha)$, $v(0) = v[0]$ and $\alpha \in [\underline{\alpha};\overline{\alpha}]$;*
- *assume that $\partial_{v_A} L = \partial_{v_B} L$ and $\partial_{v_A} K = \partial_{v_B} K$ for all $H \in \mathcal{H}$ and $A, B \in H$.*

With this, define for any $\hat{v} \in \mathbb{R}^{\hat{\mathcal{S}}}_{\geq 0}$ *the lumped costs as* $\hat{L}(t,\hat{v}) = L(t,v)$ *and* $\hat{K}(t,\hat{v}) = K(t,v)$*, where* $v \in \mathbb{R}^{\mathcal{S}}_{\geq 0}$ *is arbitrary such that* $\sum_{A\in H} v_A = \hat{v}_{A_H}$ *for all* $H \in \mathcal{H}$*. Then, for any initial condition* $v[0] \in I$*, almost surely it holds that*

$$\inf_{\alpha} J_\alpha(v[0]) = \inf_{\hat{\alpha}} \hat{J}_{\hat{\alpha}}(\hat{v}[0]),$$

provided that $\sum_{A\in H} v_A[0] = \hat{v}_{A_H}[0]$ *for all* $H \in \mathcal{H}$*. A similar statement holds for* sup*.*

4.1 Proof of Theorem 3

We next outline the proof strategy of our main result. We start by noting that the control system (3) is known as the fluid model of a CRN. This is because it can be approximated by CTMCs that have as states, loosely speaking, fractions $\frac{1}{N}\mathbb{N}_0^{\mathcal{S}}$ rather than integers $\mathbb{N}_0^{\mathcal{S}}$ [7,30]. The discussion presented next builds on that.

Definition 6 (CTMC Approximation). *Fix a CRN* $(\mathcal{S},\mathcal{R}_\alpha)$ *and a constant* $c > 0$*. The* N*-th CTMC approximation of* $(\mathcal{S},\mathcal{R}_\alpha)$ *is* $X_N = (\frac{1}{N}\mathbb{N}_0^{\mathcal{S}}, q_\alpha^N)$*, where for two different* $\sigma,\theta \in \frac{1}{N}\mathbb{N}_0^{\mathcal{S}}$ *we have:*

$$g(\sigma) = \max\{0, \min\{1, 2 - |\sigma|/c\}\}$$
$$q_\alpha^N(\sigma,\theta) = g(\sigma)\cdot q_{\alpha^N}(N\sigma, N\theta),$$

where each $\rho \xrightarrow{\alpha_i} \pi \in \mathcal{R}_\alpha$ *induces a* $\rho \xrightarrow{\alpha_i^N} \pi \in \mathcal{R}_{\alpha^N}$ *with* $\alpha_i^N = \alpha_i/N^{|\rho|-1}$ *for* $|\rho| = \sum_{A\in\mathcal{S}} \rho(A)$*.*[4]

Generalizing the foregoing notion, we introduce a UCTMC approximation of a CCRN.

Definition 7 (UCTMC Approximation). *Fix a CCRN* $(\mathcal{S},\mathcal{R}_{[\underline{\alpha};\overline{\alpha}]})$ *and a constant* $c > 0$*. The* N*-th UCTMC approximation of* $(\mathcal{S},\mathcal{R}_{[\underline{\alpha};\overline{\alpha}]})$ *is* $X_N = (\frac{1}{N}\mathbb{N}_0^{\mathcal{S}}, q^N_{[\underline{\alpha};\overline{\alpha}]})$*, where* $q^N_{[\underline{\alpha};\overline{\alpha}]} = [q^N_{\underline{\alpha}}; q^N_{\overline{\alpha}}]$*, with* $q^N_{\underline{\alpha}}$ *and* $q^N_{\overline{\alpha}}$ *as in Definition 6.*

We next prove that the UCTMCs X_N converge to the fluid CCRN model of $(\mathcal{S},\mathcal{R}_{[\underline{\alpha};\overline{\alpha}]})$. To this end, we first show that for any $\alpha(\cdot) \in [\underline{\alpha};\overline{\alpha}]$ there exists a $q(\cdot) \in q^N_{[\underline{\alpha};\overline{\alpha}]}$ such that the ODE solution v^α is sufficiently close to the CTMC simulation X_N^q, provided that N is large enough and X_N^q denotes the CTMC induced by q. This follows from standard fluid limit results [24, §11.1-§11.2].

[4] The CTMC approximation could be given without a cutoff function g. It will be mainly needed for the UCTMC counterpart from Definition 7.

Proposition 1. *Fix a CCRN $(\mathcal{S}, \mathcal{R}_{[\underline{\alpha};\overline{\alpha}]})$, a time $T > 0$ and assume that $\mathfrak{R}(t) \subseteq B(c)$ for all $t \in [0;T]$. Assume further that the UCTMC approximations $(X_N(0))_{N\geq 1}$ satisfy $X_N(0) = \lfloor Nv(0) \rfloor / N$ for some $v(0) \in I$ in the set of initial conditions I of the CCRN. Then, for any $\varepsilon, \delta > 0$, there exists an $N \geq 1$ such that for any $\alpha(\cdot) \in [\underline{\alpha};\overline{\alpha}]$, there is a $q(\cdot) \in q^N_{[\underline{\alpha};\overline{\alpha}]}$ such that*

$$\mathbb{P}\Big\{ \sup_{0\leq t\leq T} |X^q_N(t) - v^\alpha(t)| > \varepsilon \Big\} < \delta.$$

Our second approximation result ensures, conversely, that for any $q(\cdot) \in q^N_{[\underline{\alpha};\overline{\alpha}]}$ there exists a $\alpha(\cdot) \in [\underline{\alpha};\overline{\alpha}]$ such that the ODE solution v^α is sufficiently close to the CTMC simulation X^q_N, provided that N is large enough.

Proposition 2. *Under the same assumptions as Proposition 1 and for any $\varepsilon, \delta > 0$, there exists an $N \geq 1$ such that for any $q(\cdot) \in q^N_{[\underline{\alpha};\overline{\alpha}]}$, there exists $\alpha(\cdot) \in [\underline{\alpha};\overline{\alpha}]$ such that*

$$\mathbb{P}\Big\{ \sup_{0\leq t\leq T} |X^q_N(t) - v^\alpha(t)| > \varepsilon \Big\} < \delta.$$

Before proving the main result, we establish our last auxiliary result which ensures that the transient probabilities of the N-th UCTMC approximation of the original and the lumped CCRN coincide on the blocks of $H^\uparrow$, if N is large enough and $\mathcal{H}$ is a CCRN species equivalence. The proof relies on [10].

Proposition 3. *Let $\mathcal{H}$ be a CCRN species equivalence of $(\mathcal{S}, \mathcal{R}_{[\underline{\alpha};\overline{\alpha}]})$ and let X_N and $\hat{X}_N$ denote, respectively, the N-UCTMC approximation of the original and the lumped CCRN, see Definition 5 and 7. Then, we have the following.*

- *For any $q(\cdot) \in q^N_{[\underline{\hat{\alpha}};\overline{\hat{\alpha}}]}$ there is a $\hat{q}(\cdot) \in q^N_{[\underline{\hat{\alpha}};\overline{\hat{\alpha}}]}$ such that*

$$\forall H^\uparrow \in \mathcal{H}^\uparrow. \sum_{\sigma\in H^\uparrow} p_{\frac{1}{N}\sigma}(t) = \hat{p}_{\frac{1}{N}\sigma_{H^\uparrow}}(t) \tag{9}$$

  holds for all $t > 0$, provided it holds for $t = 0$. Here, p and $\hat{p}$ is the transient probability of X^q_N and $\hat{X}^{\hat{q}}_N$, respectively, while $\sigma_{H^\uparrow} \in H^\uparrow \cap \mathbb{N}^{\hat{\mathcal{S}}}_0$ is the unique representative of $H^\uparrow$.
- *Conversely, for any $\hat{q}(\cdot) \in q^N_{[\underline{\hat{\alpha}};\overline{\hat{\alpha}}]}$ there is a $q(\cdot) \in q^N_{[\underline{\hat{\alpha}};\overline{\hat{\alpha}}]}$ such that (9) holds for all $t > 0$, if it holds for $t = 0$.*

With Proposition 1–3, Theorem 3 can be proven following the idea shown in Fig. 2.

5 Evaluation

We evaluate our framework by studying epidemiological network dynamics but applications in performance modelling [48,50,56], engineering [13] and biology [16,25] are also possible. Here, we consider two settings, one from optimal control [33,34,55] and one from reachability analysis [1,17,22]. Algorithm 1 has been implemented in ERODE [15] which supports [12]. The experiments were run on a 3.22 GHz machine assigning 6 GB of RAM to ERODE. In all cases, two iterations of Algorithm 1 were sufficient.

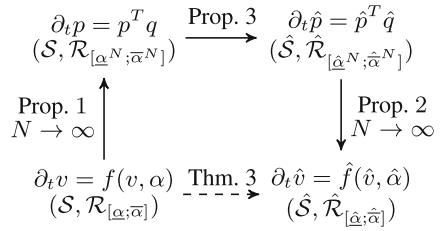

Fig. 2. Proof strategy of Theorem 3, part 1). The result is proven by approximating the deterministic control systems of the original and the lumped CCRN by means of UCTMCs (Propositions 1 and 2). These, in turn, are shown in Prop. 3 to coincide on the blocks (of the multiset lifting) of an ordinary lumpable partition. Part 2) is proven in a similar fashion by reversing the directions.

SIR Over Networks. Disease spread over networks is often modeled by variants of the susceptible-infected-recovered (SIR) model [39] over graphs [12]. Here, we study an SIR variant with vaccination [29], the respective reactions are

$$S_i \xrightarrow{[\underline{\alpha};\overline{\alpha}]} R_i + V_i, \qquad S_i + I_j \xrightarrow{w_{ij}\beta} I_i + I_j, \qquad I_i \xrightarrow{\gamma} R_i, \qquad R_i \xrightarrow{\eta} S_i,$$

where $i, j \in \{1, \ldots, n\}$. The first reaction models the vaccination, the second captures the infection across different locations, the third recovery, while the forth corresponds to the loose of immunity. Subscripts denote locations and $W = (w_{i,j})$ is the adjacency matrix of the graph representing the network topology, with $w_{ij} > 0$ denoting the presence of a possibly weighted edge between node i and j. In other words, in our experiment we interpret the presence of an edge from node j to i as the possibility for individual j (I_j) to infect individual i (S_i). Parameters β, γ, η were chosen as in [12], while vaccination bounds were set to $\underline{\alpha} = 0$ and $\overline{\alpha} = 1$ for lack of better alternative. The auxiliary species V_i keep track of the vaccinated.

Real-World Networks. To initialize weight matrices W, we use networks taken from the Netzschleuder repository [40]. We consider all weighted networks from the repository with at most 52000 nodes. Part of the networks are directed, while the others are undirected. We implicitly transform the latter ones by replacing every undirected edge with two corresponding directed ones with same weight. On these models, we applied our lumping algorithm starting from an initial partition with blocks that separate the types of variables across all nodes:

$$\mathcal{H} = \{\{S_i \mid i \leq n\}, \{I_i \mid i \leq n\}, \{R_i \mid i \leq n\}, \{V_i \mid i \leq n\}\},$$

with n being the number of nodes in the considered network. For any reported lumping $\mathcal{H}$, any costs L, K satisfying the assumptions of Theorem 4 are applicable. This includes costs that try to minimize the cumulative infection in some

specific block $H \in \mathcal{H}$, i.e.

$$J = \omega_1 \int_0^T \sum_{i \in H} v_{I_i}(s)ds + \omega_2 \sum_{i \in H} v_{V_i}(T),$$

where ω_1 and ω_2 are non-negative weights. Intuitively, the costs aim at minimize the spread of infection while using a minimal amount of vaccination.

Results. Overall, we considered 1558 real networks from the repository. The largest considered network contains 51919 nodes, corresponding to a CCRN with 155757 variables and 330149 reactions on which our reduction algorithm took about 50 min on a standard laptop machine. The results are summarized in Fig. 3. Here, we define the reduction ratio of a model as the number of reduced variables over that of original ones (the auxiliary species V_i have been dropped for providing a cleaner picture). Overall, 877 models could be reduced (i.e., have a reduction ratio smaller than 1), while 681 were not reduced (reduction ratio = 1). Figure 3(a) focuses on the 877 models that admitted reduction, sorted by reduction ratio. We can see that about 250 models could be reduced to less than half the original number of variables. This is visualized better in Fig. 3(b). Here we count how many models have a reduction ratio within ten intervals from [0.0;0.1] (the bar from 0 to 10), to [0.9;1.0] (the bar from 90 to 100). The rightmost bar refers to the 681 models that did not admit any reduction. Overall, more than 56% of the models admitted reduction. Among these, about 28% admitted substantial reductions obtaining a reduction ratio smaller than 0.4.

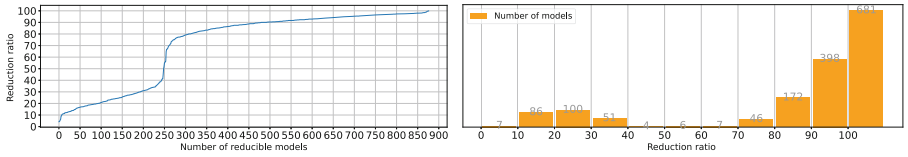

(a) The 877 reducible models sorted by reduction ratio.

(b) All 1558 models grouped by reduction ratio. The gray numbers count the models in the corresponding range.

Fig. 3. CCRN lumping of SIR-vaccination model over weighted networks from [40]. Reduction ratios are given as number of reduced variables over original ones.

Impact of Uncertain Weights on Reduction Power. In this experiment, rather than focusing on the control problem of vaccination, we study the impact of weights' uncertainty on the reduction power of our technique. To this end, we perform a new analysis of the SIR vaccination model over networks from [40] by fixing the vaccination rate (to 1), while assuming that there is uncertainty in the weights of the 1558 networks considered (we use an arbitrary interval of 0.05 centered at weights's values, to ensure that intervals remain positive).

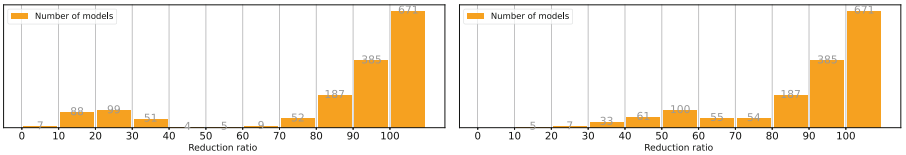

(a) Models without any uncertainty. (b) Models with uncertainty on weights.

Fig. 4. CCRN lumping of SIR-vaccination model over weighted networks from [40]. Reduction ratios are given as number of reduced variables over original ones. All 1558 models grouped by reduction ratio. The gray numbers count the models in the corresponding range.

The results are summarized in Fig. 4. Similarly to Fig. 3(b), we group models by reduction ratio. In particular, Fig. 4(a) considers models without uncertainty of the weights, while Fig. 4(b) with uncertainty on the weights. We can see that the absolute number of reducible models is not affected (in both cases, 671 models could not be reduced at all). Likewise, mild reductions with reduction ratios from 0.7 to 1.0 are not affected either. Considering the cases with lower reduction ratio, we can clearly see a similar pattern in the two figures, shifted to the right in the case of uncertain weights: reduction ratios from 0.1 to 0.4 appear to get shifted from 0.3 to 0.7.

6 Conclusion

We introduced a model reduction technique for controlled chemical reaction networks (CCRNs) whose kinetic reaction parameters are subject to control or distrubance. The smallest (lumped) CCRN can be computed in polynomial time and is shown to preserve the optimal costs of the original CCRN. The applicability has been demonstrated by reducing the verification and control problems of vaccination models over real world networks with more than 50000 nodes.

The proposed framework is holistic in that it can be used as a precomputation step before any optimization approach. In case the reduced model is sufficiently small, global optimization techniques such as the Hamilton-Jacobi-Bellman equations [20,34] or reachability analysis tools such as [1,6,22] can be invoked. If the reduced model is still too large for global optimization techniques, local optimization approaches such as the functional gradient descent, also known as Pontryagin's maximum principle [33], can be invoked. While the principle has gained recently momentum in AI by training so-called neural ordinary differential equations [21,26], its computational complexity is at least quadratic in the size of the model, thus justifying the need for optimality-preserving model reduction techniques. Likewise, heuristic approaches involving sampling and simulation, as commonly used in systems biology [45,46], can profit from optimality-preserving reductions as well.

Acknowledgments. This work was partially supported by Poul Due Jensen Foundation grant no. 883901; the S40S Villum Investigator Grant nr. 37819 from Villum Fonden; the Fsc regional Tuscan project AISLEA2 J54D23000780005; the project SERICS (PE00000014); the Tuscany Health Ecosystem (THE) project with CUP B83C22003920001, under the MUR National Recovery and Resilience Plan funded by the European Union - NextGenerationEU; the project SMaRT COnSTRUCT (CUP J53C24001460006), in the context of FAIR (PE0000013, CUP B53C22003630006) under the National Recovery and Resilience Plan (Mission 4, Component 2, Line of Investment 1.3) funded by the European Union - NextGenerationEU.

References

1. Althoff, M.: An introduction to CORA 2015. In: Proceedings of the Workshop on Applied Verification for Continuous and Hybrid Systems (2015)
2. Anderson, D.F.: A modified next reaction method for simulating chemical systems with time dependent propensities and delays. J. Chem. Phys. **127**(21), 214107 (2007)
3. Angeli, D., Ferrell, J.E., Sontag, E.D.: Detection of multistability, bifurcations, and hysteresis in a large class of biological positive-feedback systems. Proc. Natl. Acad. Sci. **101**(7), 1822–1827 (2004)
4. Angulo, M.T., Moog, C.H., Liu, Y.Y.: A theoretical framework for controlling complex microbial communities. Nat. Commun. **10**(1), 1045 (2019)
5. Antoulas, A.C.: Approximation of large-scale dynamical systems. SIAM (2005)
6. Bogomolov, S., Frehse, G., Grosu, R., Ladan, H., Podelski, A., Wehrle, M.: A box-based distance between regions for guiding the reachability analysis of SpaceEx. In: Computer Aided Verification, CAV, pp. 479–494 (2012)
7. Bortolussi, L., Gast, N.: Mean field approximation of uncertain stochastic models. In: DSN, pp. 287–298 (2016)
8. Bortolussi, L., Hillston, J., Latella, D., Massink, M.: Continuous approximation of collective system behaviour: a tutorial. Perform. Eval. **70**(5), 317–349 (2013)
9. Cardelli, L., Tribastone, M., Tschaikowski, M., Vandin, A.: Maximal aggregation of polynomial dynamical systems. Proc. Natl. Acad. Sci. **114**(38), 10029–10034 (2017)
10. Cardelli, L., Grosu, R., Larsen, K.G., Tribastone, M., Tschaikowski, M., Vandin, A.: Lumpability for uncertain continuous-time Markov chains. In: Abate, A., Marin, A. (eds.) QEST 2021. LNCS, vol. 12846, pp. 391–409. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-85172-9_21
11. Cardelli, L., Grosu, R., Larsen, K.G., Tribastone, M., Tschaikowski, M., Vandin, A.: Algorithmic minimization of uncertain continuous-time Markov chains. IEEE Trans. Autom. Control. **68**(11), 6557–6572 (2023)
12. Cardelli, L., Perez-Verona, I.C., Tribastone, M., Tschaikowski, M., Vandin, A., Waizmann, T.: Exact maximal reduction of stochastic reaction networks by species lumping. Bioinformatics **37**(15), 2175–2182 (2021)
13. Cardelli, L., Tribastone, M., Tschaikowski, M.: From electric circuits to chemical networks. Nat. Comput. (2019). https://doi.org/10.1007/s11047-019-09761-7
14. Cardelli, L., Tribastone, M., Tschaikowski, M., Vandin, A.: Comparing chemical reaction networks: a categorical and algorithmic perspective. In: Symposium on Logic in Computer Science, LICS, pp. 485–494 (2016)

15. Cardelli, L., Tribastone, M., Tschaikowski, M., Vandin, A.: ERODE: a tool for the evaluation and reduction of ordinary differential equations. In: Legay, A., Margaria, T. (eds.) TACAS 2017. LNCS, vol. 10206, pp. 310–328. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54580-5_19
16. Cardelli, L., Tribastone, M., Tschaikowski, M., Vandin, A.: Syntactic Markovian bisimulation for chemical reaction networks. In: Aceto, L., Bacci, G., Bacci, G., Ingólfsdóttir, A., Legay, A., Mardare, R. (eds.) Models, Algorithms, Logics and Tools. LNCS, vol. 10460, pp. 466–483. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63121-9_23
17. Cardelli, L., Tribastone, M., Tschaikowski, M., Vandin, A.: Guaranteed error bounds on approximate model abstractions through reachability analysis. In: McIver, A., Horvath, A. (eds.) QEST 2018. LNCS, vol. 11024, pp. 104–121. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-99154-2_7
18. Cardelli, L., Tribastone, M., Tschaikowski, M., Vandin, A.: Symbolic computation of differential equivalences. Theoret. Comput. Sci. **777**, 132–154 (2019)
19. Chen, M., Herbert, S.L., Tomlin, C.J.: Fast reachable set approximations via state decoupling disturbances. In: CDC, pp. 191–196 (2016)
20. Chen, M., Tomlin, C.J.: Exact and efficient Hamilton-Jacobi reachability for decoupled systems. In: CDC, pp. 1297–1303 (2015)
21. Chen, T.Q., Rubanova, Y., Bettencourt, J., Duvenaud, D.K.: Neural ordinary differential equations. In: Conference on Neural Information Processing Systems, NIPS, pp. 6572–6583 (2018)
22. Chen, X., Ábrahám, E., Sankaranarayanan, S.: Taylor model flowpipe construction for non-linear hybrid systems. In: Real-Time Systems Symposium, RTSS, pp. 183–192 (2012)
23. Eberle, A.: Lecture notes in Markov processes, January 2008
24. Ethier, S.N., Kurtz, T.G.: Markov Processes - Characterization and Convergence. Wiley, New York (1986)
25. Gast, N., Bortolussi, L., Tribastone, M.: Size expansions of mean field approximation: transient and steady-state analysis. ACM SIGMETRICS Perform. Eval. Rev. **46**(3), 25–26 (2019)
26. Hasani, R., et al.: Closed-form continuous-depth models (2021)
27. Iacobelli, G., Tribastone, M., Vandin, A.: Differential bisimulation for a Markovian process algebra. In: Italiano, G.F., Pighizzini, G., Sannella, D.T. (eds.) MFCS 2015. LNCS, vol. 9234, pp. 293–306. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48057-1_23
28. Jaeger, M., Bacci, G., Bacci, G., Larsen, K.G., Jensen, P.G.: Approximating euclidean by imprecise Markov decision processes. In: Margaria, T., Steffen, B. (eds.) ISoLA 2020. LNCS, vol. 12476, pp. 275–289. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-61362-4_15
29. Kopfova, J., Nabelkova, P., Rachinskii, D., Rouf, S.C.: Dynamics of SIR model with vaccination and heterogeneous behavioral response of individuals modeled by the Preisach operator. J. Math. Biol. **83**(11), 1 (2021)
30. Kurtz, T.G.: The relationship between stochastic and deterministic models for chemical reactions. J. Chem. Phys. **57**, 2976–2978 (1972)
31. Kurtz, T.G.: Solutions of ordinary differential equations as limits of pure jump Markov processes. J. Appl. Probab. **7**(1), 49–58 (1970)
32. Larsen, K.G., Skou, A.: Bisimulation through probabilistic testing. Inf. Comput. **94**(1), 1–28 (1991)
33. Liberzon, D.: Calculus of Variations and Optimal Control Theory: A Concise Introduction. Princeton University Press, Princeton (2011)

34. Lygeros, J.: On reachability and minimum cost optimal control. Automatica **40**(6), 917–927 (2004)
35. Meyn, S.P., Tweedie, R.L.: Stability of Markovian processes III: foster-lyapunov criteria for continuous-time processes. Adv. Appl. Probab. **25**(3), 518–548 (1993)
36. Okino, M.S., Mavrovouniotis, M.L.: Simplification of mathematical models of chemical reaction systems. Chem. Rev. **2**(98), 391–408 (1998)
37. Pappas, G.J., Lafferriere, G., Sastry, S.: Hierarchically consistent control systems. IEEE Trans. Autom. Control **45**(6), 1144–1160 (2000)
38. Pappas, G.J., Simic, S.: Consistent abstractions of affine control systems. IEEE TAC **47**(5), 745–756 (2002)
39. Pastor-Satorras, R., Castellano, C., Van Mieghem, P., Vespignani, A.: Epidemic processes in complex networks. Rev. Mod. Phys. **87**(3), 925 (2015)
40. Peixoto, T.P.: The netzschleuder network catalogue and repository (2020)
41. Rowley, C.W., Mezić, I., Bagheri, S., Schlatter, P., Henningson, D.S.: Spectral analysis of nonlinear flows. J. Fluid Mech. **641**, 115–127 (2009)
42. Seeman, N.C., Sleiman, H.F.: DNA nanotechnology. Nat. Rev. Mater. **3**, 1–23 (2017)
43. Snowden, T.J., van der Graaf, P.H., Tindall, M.J.: Methods of model reduction for large-scale biological systems: a survey of current methods and trends. Bull. Math. Biol. **79**(7), 1449–1486 (2017)
44. Tabuada, P., Pappas, G.J.: Abstractions of Hamiltonian control systems. Automatica **39**(12), 2025–2033 (2003)
45. Tognazzi, S., Tribastone, M., Tschaikowski, M., Vandin, A.: EGAC: a genetic algorithm to compare chemical reaction networks. In: Genetic and Evolutionary Computation Conference, GECCO, pp. 833–840 (2017)
46. Toni, T., Stumpf, M.P.H.: Simulation-based model selection for dynamical systems in systems and population biology. Bioinformatics **26**(1), 104–110 (2009)
47. Tribastone, M.: Behavioral relations in a process algebra for variants. In: Proceedings of the 18th International Software Product Line Conference-Volume 1, pp. 82–91 (2014)
48. Tribastone, M., Mayer, P., Wirsing, M.: Performance prediction of service-oriented systems with layered queueing networks. In International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISOLA), pp. 51–65 (2010)
49. Tschaikowski, M., Tribastone, M.: Exact fluid lumpability in Markovian process algebra. Theoret. Comput. Sci. **538**, 140–166 (2014)
50. Tschaikowski, M., Tribastone, M.: Tackling continuous state-space explosion in a Markovian process algebra. Theoret. Comput. Sci. **517**, 1–33 (2014)
51. Tschaikowski, M., Tribastone, M.: A unified framework for differential aggregations in Markovian process algebra. J. Logical Algebraic Methods Program. **84**(2), 238–258 (2015)
52. Tschaikowski, M., Tribastone, M.: Spatial fluid limits for stochastic mobile networks. J. Perform. Eval. **109**, 52–76 (2017)
53. van der Schaft, A.: Equivalence of dynamical systems by bisimulation. IEEE TAC **49**, 2160–2172 (2004)
54. Voit, E.O., Martens, H.A., Omholt, S.W.: 150 years of the mass action law. PLOS Comput. Biol. **11**(1), 1–7 (2015)

55. Whitby, M., Cardelli, L., Kwiatkowska, M., Laurenti, L., Tribastone, M., Tschaikowski, M.: PID control of biochemical reaction networks. IEEE TAC **67**(2), 1023–1030 (2022)
56. Wirsing, M., et al.: SENSORIA patterns: augmenting service engineering with formal analysis, transformation and dynamicity. In: Margaria, T., Steffen, B. (eds.) ISoLA 2008. CCIS, vol. 17, pp. 170–190. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-88479-8_13