

Systems Biology

Exact maximal reduction of stochastic reaction networks by species lumping

Luca Cardelli,¹ Isabel Cristina Perez-Verona,² Mirco Tribastone ,^{2,*}
Max Tschaikowski,³ Andrea Vandin  ⁴ and Tabea Waizmann²

¹Department of Computer Science, University of Oxford, Oxford 34127, OX1 3QD, UK, ²IMT School for Advanced Studies, Lucca 55100, Italy, ³Department of Computer Science, University of Aalborg, Aalborg 34127, 9220, Denmark and ⁴Sant'Anna School of Advanced Studies, Pisa 56127, Italy

*To whom correspondence should be addressed.

Associate Editor: Alfonso Valencia

Received on October 16, 2020; revised on January 9, 2021; editorial decision on January 27, 2021; accepted on January 28, 2021

Abstract

Motivation: Stochastic reaction networks are a widespread model to describe biological systems where the presence of noise is relevant, such as in cell regulatory processes. Unfortunately, in all but simplest models the resulting discrete state-space representation hinders analytical tractability and makes numerical simulations expensive. Reduction methods can lower complexity by computing model projections that preserve dynamics of interest to the user.

Results: We present an exact lumping method for stochastic reaction networks with mass-action kinetics. It hinges on an equivalence relation between the species, resulting in a reduced network where the dynamics of each macro-species is stochastically equivalent to the sum of the original species in each equivalence class, for any choice of the initial state of the system. Furthermore, by an appropriate encoding of kinetic parameters as additional species, the method can establish equivalences that do not depend on specific values of the parameters. The method is supported by an efficient algorithm to compute the largest species equivalence, thus the maximal lumping. The effectiveness and scalability of our lumping technique, as well as the physical interpretability of resulting reductions, is demonstrated in several models of signaling pathways and epidemic processes on complex networks.

Availability and implementation: The algorithms for species equivalence have been implemented in the software tool ERODE, freely available for download from <https://www.erode.eu>.

Contact: micro.tribastone@imtlucca.it

Supplementary information: [Supplementary data](#) are available at *Bioinformatics* online.

1 Introduction

Stochastic reaction networks are a foundational model to study biological systems where the presence of noise cannot be neglected, for instance in cell regulatory processes governed by low-abundance biochemical species (Guptasarma, 1995), which may introduce significant variability in gene expression (Elowitz, 2002). Their analysis—either by solution of the master equation or by stochastic simulation—is fundamentally hindered by a discrete representation of the state space (Van Kampen, 2007), which leads to a combinatorial growth in the number of states in the underlying Markov chain as a function of the abundances of the species.

Here, we present a method for exact reduction that preserves the stochastic dynamics of mass-action reaction networks, a fundamental kinetic model in computational systems biology (Voit et al., 2015). The method rests on a relation between species, called *species equivalence* (SE), which can be checked through criteria that depend

on the set of reactions of the network. SE gives rise to a reduced stochastic reaction network where the population of each macro-species tracks the sum of the population levels of all species belonging to an equivalence class.

As with all reduction methods, SE implies some loss of information; namely, the individual dynamical behavior of a species that is aggregated into a macro-species cannot be recovered in general. However, our algorithm for computing SE gives freedom to the modeler as to which original variables to preserve in the reduced network. Indeed, building upon a celebrated result in theoretical computer science (Paige and Tarjan, 1987), we compute SE as the coarsest partition that satisfies the equivalence criteria and that refines a given initial partition of species. Thus, a species of interest that is isolated in a singleton block is guaranteed to be preserved in the reduced network. Our partition-refinement algorithm is computationally efficient, in the sense that the algorithm runs in polynomial time as function of the number of species and reactions of the

original network. Finally, we can prove the existence of a maximal SE, i.e. the equivalence that leads to the coarsest aggregation of the reaction network.

Formally, SE can be seen a lifting to reaction networks of the notion of lumpability of Markov chains (Buchholz, 1994; Kemeny and Snell, 1976). That is, the reduced network yields a state space where each macro-state tracks the sum of the probabilities of the states in the original Markov chain. Ordinary lumpability requires the availability of the state space that underlies the master equation (Van Kampen, 2007); thus, it also requires the initial state of the Markov chain to be fixed. Instead, SE works at the structural level of the reaction network, by lumping species instead of states; thus, it involves the analysis of an exponentially smaller mathematical object in general. In addition, a practically useful consequence of reasoning at the network level is that an SE holds for any initial state. Given that a reaction network can be seen as a Petri net where each species is represented as a place (Brijder, 2019), our structural approach is close in spirit to the notion of place bisimulation (Autant and Schnoebelen, 1992). However, that induces a bisimulation over markings in the classical, non-quantitative sense (Joyal et al., 1996).

There are several methods for the reduction of the deterministic rate equations of biochemical reaction networks, e.g. Snowden et al. (2017). However, these reductions do not preserve the stochastic behavior in general. For stochastic models in systems biology, lumpability has been studied for rule-based formalisms, providing reduction methods based on rule conditions that induce a lumping of the underlying Markov chain (Feret et al., 2012, 2013).

For mass-action networks, the earlier approach to species lumping by Cardelli et al. (2017b), called syntactic Markovian bisimulation, suffers from two limitations. First, syntactic Markovian bisimulation is only a sufficient condition for lumpability. Here, we prove that SE is the coarsest possible aggregation that yields a Markov chain lumping according to an equivalence over species. We show that this yields coarser aggregations than syntactic Markovian bisimulation in benchmark models.

The second limitation is that syntactic Markovian bisimulation only supports networks where reactions involve at most two reagents. Instead, SE can be applied to arbitrary higher-order reactions. This may appear unnecessary because in models of practical relevance reactions are typically of order two at most, following the basic principle that the probability of more than two bodies probabilistically colliding at the same time can be negligible (Gillespie, 1977). However, this generalization enables the identification of ‘qualitative’ relations between species, i.e. equivalences that do not depend on the specific choice of values of the kinetic parameters. This is done by systematically turning the original network into one where each kinetic parameter appears as a further auxiliary species in a reaction, thus increasing its order by one. A number of case studies from the systems biology literature are used to show examples of parameter-independent physically intelligible model reductions.

2 Results

2.1 Stochastic mass-action networks

Formally, a stochastic mass-action network is described by a set of species \mathcal{S} and a set of reactions \mathcal{R} . Each reaction is in the form $\rho \xrightarrow{\alpha} \pi$, where $\alpha > 0$ is a kinetic parameter and ρ and π are multisets of species called reagents and products, respectively. The multiplicity of species S in ρ is denoted by $\rho(S)$, which represents its stoichiometry coefficient. The set of all reagents and products across all reactions in the network are denoted by $\rho(\mathcal{R})$ and $\pi(\mathcal{R})$, respectively.

A stochastic mass-action network gives rise to a continuous-time Markov chain (CTMC) where each state σ is a multiset of species. From a state σ such that $\rho \subseteq \sigma$, a reaction $\rho \xrightarrow{\alpha} \pi$ induces a transition with mass-action propensity $\alpha \prod_{S \in \rho} \binom{\sigma(S)}{\rho(S)}$ to state $\sigma + \pi - \rho$, where the plus and minus operators indicate multiset union and difference, respectively, while $S \in \rho$ denotes that S belongs to the support of ρ , i.e. $\rho(S) > 0$. Given an initial state $\hat{\sigma}$, the state space can be derived by exhaustively applying the reactions to compute all

possible states reachable from $\hat{\sigma}$. We denote with $out(\sigma)$ the multiset of outgoing transitions from state σ :

$$out(\sigma) = \left\{ \sigma \xrightarrow{\lambda} \sigma + \pi - \rho \mid (\rho \xrightarrow{\alpha} \pi) \in \mathcal{R}, \lambda = \alpha \prod_{S \in \rho} \binom{\sigma(S)}{\rho(S)} \right\}.$$

For any two distinct states σ and θ , we denote by $q(\sigma, \theta)$ the sum of the propensities from σ to θ across all reactions, i.e.:

$$q(\sigma, \theta) = \sum_{(\sigma \xrightarrow{\lambda} \theta) \in out(\sigma)} \lambda.$$

Moreover, we set $q(\sigma, \sigma)$ to be the negative sum of all possible transitions from state σ , i.e. $q(\sigma, \sigma) = -\sum_{\theta \neq \sigma} q(\sigma, \theta)$. These values form the CTMC generator matrix, which characterizes the dynamical evolution of the Markov chain by means of the master equation $\dot{p} = p^T Q$. Each component of its solution, $p_\sigma(t)$, is the probability of being in state σ at time t starting from some initial probability distribution (Van Kampen, 2007).

Figure 1 shows a simple running example to summarize the main results of this article using the network in Figure 1a with species S_1, \dots, S_4 . The state space from the initial state $\hat{\sigma} = S_1 + 2S_4$ is in shown Figure 1b.

Ordinary lumpability is a partition of the state space such that any two states σ_1, σ_2 in each partition block H have equal aggregate rates toward states in any block H' , i.e. $\sum_{\sigma \in H'} q(\sigma_1, \sigma) = \sum_{\sigma \in H'} q(\sigma_2, \sigma)$ (Buchholz, 1994; Kemeny and Snell, 1976). 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. Distinct colored boxes in Figure 1b identify an ordinarily lumpable partition of the sample Markov chain. Ordinary lumpability preserves stochastic equivalence in the sense that the probability of each block/macro-state is equal to the sum of the probabilities in each original state belonging to that block.

2.2 Species equivalence

Verifying the conditions for ordinary lumpability requires the full enumeration of the CTMC state space, which grows combinatorially with the multiplicities of initial state and the number of reactions. Additionally, the presence of interactions such as constitutive transcription, e.g. $S_1 \alpha S_1 + S_2$, may give rise to infinite state spaces, preventing the use of ordinary lumpability altogether. SE detects ordinary lumpability at the finitary level of the reaction network by identifying an equivalence relation (i.e. a partition) of the species which induces an ordinary lumpable partition over the multisets representing CTMC states.

For this, we consider a natural lifting of a partition \mathcal{H} of species to multisets of species, which we call the *multiset lifting* of \mathcal{H} and which we denote by \mathcal{H}^\dagger . It relates multisets that have same cumulative multiplicity from each partition block. That is, two multisets/states σ_1 and σ_2 belong to the same block $M \in \mathcal{H}^\dagger$ if the condition $\sum_{S \in H} \sigma_1(S) = \sum_{S \in H} \sigma_2(S)$ is satisfied for all blocks of species $H \in \mathcal{H}$.

At the basis of SE is the notion of reaction rate $\mathbf{rr}(\rho, \pi)$ from reagents ρ to products π :

$$\mathbf{rr}(\rho, \pi) = \begin{cases} \sum_{(\rho \xrightarrow{\alpha} \pi) \in \mathcal{R}} \alpha & \text{if } \rho \neq \pi, \\ - \sum_{\pi' \in \pi(\mathcal{R}), \rho \neq \pi'} \mathbf{rr}(\rho, \pi') & \text{if } \rho = \pi. \end{cases}$$

Intuitively, it is defined as the analogue to the entries of the Markov chain generator matrix, but it is computable by only inspecting the set of reactions. SE is defined as a partition of species \mathcal{H} such that, for any two species S_i and S_j in a block of \mathcal{H} , and for any block of multisets $M \in \mathcal{H}^\dagger$ containing at least one product in $\pi(\mathcal{R})$, the condition:

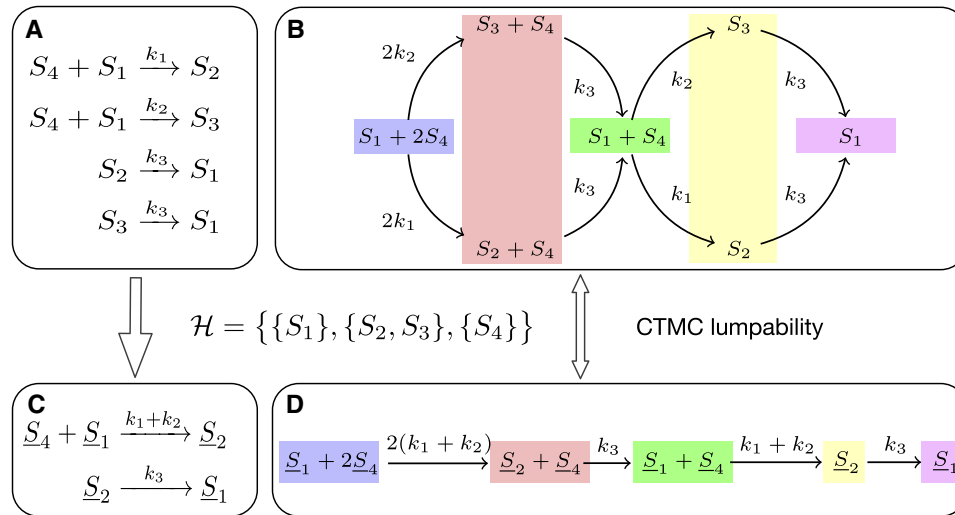


Fig. 1. Illustration of SE on a simple example. (a) Original network with four species S_1, \dots, S_4 . (b) Underlying Markov chain derived from the initial state $\hat{\sigma} = S_1 + 2S_4$. The Markov chain is shown in the customary graph representation where each node is a state and transitions between two states σ, σ' are directed arcs labeled with $q(\sigma, \sigma')$. The colored boxes represent five blocks of an ordinary lumpable partition of the Markov chain (here it suffices to check that the outgoing transitions are equal for states in blocks of size two). The partition of species denoted by \mathcal{H} can be shown to be an SE independently on the actual values of the parameters k_1, k_2 and k_3 , hence states that are equal up to the sum of the second and third coordinate form a lumpable partition block. (c) This SE gives rise to a reduced network by choosing the representatives S_1, S_2 and S_4 for each block (underlined in the figure to distinguish them from original species names). The reduced network has fewer reactions due to the fact that reactions in the original network are merged into a single one after renaming. (d) Underlying Markov chain of the reduced network derived from the matching initial state $s_1 + 2s_4$. The Markov chain of the reduced network corresponds to the lumped Markov chain of the original network (as indicated by the matching colors of the nodes)

$$\sum_{\pi \in \mathcal{M}} \mathbf{rr}(S_i + \rho, \pi) = \sum_{\pi \in \mathcal{M}} \mathbf{rr}(S_j + \rho, \pi), \quad (1)$$

holds for all ρ such that $S_i + \rho$ or $S_j + \rho$ are in the set of reagents $\rho(\mathcal{R})$.

According to this definition, species S_2 and S_3 in the sample network of Figure 1 belong to the same block of an SE. This explains why the ordinarily lumpable partition depicted in Figure 1b groups Markov chain states that have the same total multiplicities of S_2 and S_3 .

Our first result is that SE characterizes ordinary lumpability, in the sense that the multiset lifting of an SE yields an ordinarily lumpable partition of the underlying Markov chain derived from any initial state $\hat{\sigma}$; and, vice versa, if a multiset lifting of a partition of species \mathcal{H} is an ordinarily lumpable partition of the underlying Markov chain from any initial state $\hat{\sigma}$, then \mathcal{H} is an SE (proved in Supplementary Material, Section A.1). We also note that, by Rózsa and Tóth (2003), our result also applies to Markov chains with infinite state spaces, because each state has finitely many incoming and outgoing transitions due to the fact that the number of reactions is finite, and the state space is partitioned in blocks of finite size by multiset lifting.

2.3 Computation of a reduced reaction network up to SE

Analogously to the existence of a lumped Markov chain, one can build a reduced network from an SE partition. The reduction algorithm is similar to that in Cardelli et al. (2015), where it was defined for deterministic mass-action networks with a reaction-rate interpretation based on ordinary differential equations. Briefly, the reduced network is obtained by applying the following four steps: (i) choose a representative species for each block of species; (ii) discard all reactions whose reagents have species that are not representatives; (iii) replace the species in the products of the remaining reactions with their representatives; (iv) reduce the set of reactions by merging all those that have same reactants and products by summing their kinetic parameters. The correctness of this algorithm is discussed in Supplementary Material Section A.2. Following Cardelli et al. (2015), the reduced reaction network can be computed in $O(rs \log s)$ time, where s is the number of species and r is the number of reactions.

Each representative in the reduced network can be interpreted as a macro-species that tracks the sum of the populations of the distinct species in the original network that belong to the same SE partition block. Therefore, for any given initial condition $\hat{\sigma}$ of the original network, it is possible to directly generate its lumped Markov chain from the reduced network by fixing a matching initial condition up to sums of populations, as related in general by multiset lifting. The network in Figure 1c shows the reduced network up to an SE. The Markov chain obtained by ordinary lumpability of the Markov chain in Figure 1b corresponds to the Markov chain generated by the reduced network with the matching initial condition.

2.4 Computation of the maximal SE

There exist efficient algorithms that compute the coarsest ordinarily lumpable partition, i.e. the maximal aggregation, of a Markov chain with a finite state space (Derisavi et al., 2003; Valmari and Franceschinis, 2010). Here, we develop an analogous algorithm for species of a reaction network. First, we show that, indeed, there exists the largest SE (Supplementary Material, Section A.3). Then, we develop a partition refinement algorithm that takes an initial partition of species as input and computes the largest SE that refines such initial partition (Supplementary Material, Section A.4). The maximal SE is thus a special case that can be computed by initializing the algorithm with the partition with the trivial singleton block containing all species.

The algorithm maintains a reference to the current candidate SE partition and a set of splitters, i.e. blocks of products against which the candidate partition is to be checked. Both structures are initialized using the input partition. A fixed-point iteration splits a block of the current candidate SE partition whenever it falsifies the condition in Equation (1) with respect to a splitter M . If no such block is found, then the algorithm terminates and the candidate partition is proven to be the largest SE that refines the initial partition. Else, the falsifying block is split into sub-blocks that have equal values for the quantities in Equation (1). The set of splitters is recomputed as the multiset lifting to the current partition. We prove (in Supplementary Material) that the algorithm has $O(pr)$ space and $O(s^2 r^3 p(p + \log r))$ time complexity, where p is the largest number of different species appearing in the reagents or products of every reaction.

2.5 Parameter-independent SEs by network expansion

Similarly to ordinary lumpability, checking the conditions of SE by Equation (1) implicitly assumes that the values of all the kinetic parameters in the network are fixed. However, without further theory it is also possible to find equivalences that are independent from the specific values of the parameters. In order to do so, let \mathcal{P} denote the set of all kinetic parameters used in the reaction network and assume, without loss of generality, that each kinetic parameter $\alpha \in \mathcal{P}$ is a rational number n_α/d_α . Let us then consider an expanded reaction network where we take each parameter α as an additional species P_α , and every original reaction $\rho\alpha\pi$ is transformed into the reaction $P_\alpha + \rho1\pi + P_\alpha$. This is a reaction of higher order with kinetic parameter equal to one.

For this extended reaction network to be related to the original one, each state of its CTMC must represent a multiset of species; in particular the initial condition of each additional species P_α must be a non-negative integer, which will be fixed throughout the state space because the population of P_α does not change by construction. A suitable initialization of P_α may be for instance $\frac{d_\alpha}{\text{lcm}\{d_\alpha|\alpha \in \mathcal{P}\}}$, where lcm denotes the least common multiple of all denominators of the parameters. With this in place, the original and the expanded network will give rise to the same state space (dropping the components of the CTMC state related to P_α because they are constant, as discussed). The transition rates of the expanded network are instead all scaled up by the same factor, which can essentially be interpreted as a time rescaling of the original CTMC.

Since this rescaling is the same for all states, any ordinarily lumpable partition on the CTMC of the expanded network will be an ordinary lumpable partition on the CTMC of the original one, and vice versa. More importantly, the computation of SE on the expanded network will be made independent of the specific values chosen for the kinetic parameters. This is because the parameter values are encoded into the components of the initial CTMC state associated with the auxiliary species P_α , and SE finds equivalences that hold for all initial states of the Markov chain. Thus the computation of the largest SE in the original network may proceed by considering the initial partition consisting of two blocks, one for all the species and one for all species-parameters in the expanded network, respectively. For the example in Figure 1, the largest SE computed from the initial partition $\{\{S_1, S_2, S_3, S_4\}, \{P_{k_1}, P_{k_2}, P_{k_3}\}\}$ is $\{\{S_1\}, \{S_2, S_3\}, \{S_4\}, \{P_{k_1}, P_{k_2}\}, \{P_{k_3}\}\}$. In addition to the equivalence between the two species S_2 and S_3 , it detects that the reduced model depends only on the sum of the parameters $k_1 + k_2$, for any given value.


3 Examples

In this section, we present reductions on case studies from the literature, computed with an implementation of SE within the software tool ERODE (Cardelli et al., 2017a), available at <https://www.erode.eu>. The reported results refer to the analysis of the models with the values of the kinetic parameters as reported in the associated publications. However, the reductions are preserved also in the extended parameter-independent versions obtained as discussed earlier.

3.1 SE in multisite phosphorylation processes

Mechanistic models of signaling pathways are prone to a rapid growth in the number of species and reactions because of the combinatorial effects due to the distinct configurations in which a molecular complex can be found (Salazar and Höfer, 2009). A prototypical situation is multisite phosphorylation, a fundamental process in eukaryotic cells that is responsible for various mechanisms such as the regulation of switch-like behavior (Gunawardena, 2005; Thomson and Gunawardena, 2009). For example, let us consider a protein A with n sites that can be phosphorylated by means of kinase K according to a random mechanism, while dephosphorylation occurs as a spontaneous reaction. To describe this system one needs 2^n distinct molecular species that track the phosphorylation/dephosphorylation status of each site (Salazar and Höfer, 2009). Each species is written in the form $A(s_1, \dots, s_n)$ where $s_i = 0$

(respectively, $s_i = 1$) indicates that the i th site is dephosphorylated (respectively, phosphorylated), for all $i = 1, \dots, n$. The resulting mass-action network is given by:



for all $i = 1, \dots, n$ and for any combination of site states $s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n$. To simplify the mathematical model, it is assumed that the kinetic parameters r_1, r_2 are equal at all phosphorylation sites (Sneddon et al., 2011).

For a fixed n , the maximal SE aggregates molecular species that are equal up to the number of phosphorylated sites that they exhibit, independently of their identity. More formally, if we consider the block of species H_i that groups all configurations that have exactly i phosphorylated sites, $H_i = \{A(s_1, \dots, s_n) | s_1 + \dots + s_n = i\}$, for $i = 0, \dots, n$, then the maximal SE is given by the partition $\{\{K\}, H_0, \dots, H_n\}$.

3.2 Identification of equivalent molecular complexes in a model of synaptic plasticity

The assumption of equal kinetic parameters is not necessary to achieve aggregation with SE. We show this on a model from Pepke et al. (2010) on the interactions between calcium (Ca^{2+}), calmodulin (CaM), and the Ca^{2+} -CaM-dependent protein kinase II (CaMKII), which play a fundamental role in the mechanism of synaptic plasticity [Lisman et al., 2002 (It is available in the BioModels database; Li et al., 2010; identified as MODEL1001150000)]. The model

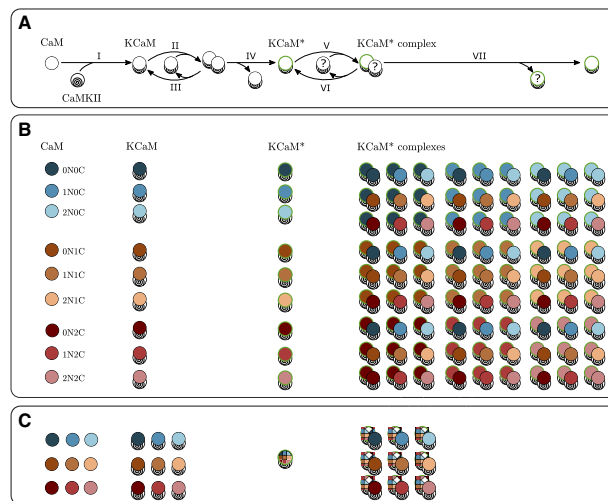


Fig. 2. (A) Kinetic scheme of the interactions between Ca^{2+} -bound-CaM and CaMKII, adapted from Pepke et al. (2010). CaM binds with the CaMKII monomer to form a KCaM complex (Reaction I). KCaM may undergo reversible dimerization (Reactions II and III), KCaM dimerization which can lead to autophosphorylation (Reaction IV). A unit of phosphorylated KCaM is labeled KCaM^* and are represented in the diagram as a green-circled KCaM. KCaM^* can interact with any unphosphorylated KCaM unit (indicated by the '?' sign) to form a KCaM^* complex reversibly (Reactions V and VI), leading to autophosphorylation (Reaction VII). (B) Molecular species that can participate in the reaction scheme. CaM units are represented with different colors to indicate the possible states of Ca^{2+} -binding, with the label $xNyC$, with $x, y = 0, 1, 2$ denoting the number Ca^{2+} -bound domains at the amino and carboxyl termini, respectively. KCaM and KCaM^* are represented similarly. The block of KCaM^* complexes contains 81 distinct molecular species obtained from all possible interactions between KCaM^* and KCaM. (C) The maximal SE yields a coarse-grained network which can be interpreted as having the same reaction scheme (A), but with fewer species. In particular CaM and KCaM complexes are not aggregated, but all the distinct KCaM^* are collapsed into the same equivalence class (indicated by the multiple-color representative). Such equivalence carries over to all KCaM^* complexes, in the sense that all dimers with the same phosphorylated form are in the same SE block. This allows the collapse of the distinct 81 KCaM^* complexes above to 9 macro-species

describes the following processes: cooperative binding of Ca^{2+} to two pairs of domains located at the amino (N) and carboxyl (C) termini of CaM; binding of CaM to a monomeric CaMKII subunit; and autophosphorylation of a CaMKII monomer through the formation of a dimer which requires CaM to be bound to both subunits (Fig. 2A). The maximal SE finds that all phosphorylated monomers are equivalent (Fig. 2B), although their dynamics are characterized by distinct kinetic parameters to account for phosphorylation rates that depend on the number of bound Ca^{2+} (Shifman et al., 2006). Further, such equivalences carry over to all complexes where they are present as sub-units. This leads to equivalence classes consisting of nine molecular species each, with an overall reduction from 155 species and 480 reactions to 75 species and 254 reactions. Notably, important quantities to observe in this model are the amounts of free and bound CaM (Lisman et al., 2012), both recoverable from the reduced network.

3.3 Internalization of the GTPase cycle in a model of the spindle position checkpoint

In both previous examples, SE can be physically interpreted as a reduction that preserves both the structure of equivalent molecular species as well as their function. SE can also aggregate species that exhibit contrasting functionality, such as in signal transduction switches realized by GTP- and GDP-bound forms of GTPases. To show this, we consider the model in Caydasi et al. (2012) of the spindle position checkpoint (SPOC), a mechanism in the budding yeast responsible for detecting the correct alignment of the nucleus between mother and daughter cells (Lew and Burke, 2003; the BioModels identifier for this model is *BIOMD000000699*). The most upstream event of the pathway involves GTPase Tem1, which is regulated by the GTPase-activating protein (GAP) complex composed of Bfa1 and Bub2. Under correct alignment the GAP complex is inhibited by a kinase Cdc5 phosphorylating Bfa1 (Gruneberg, 2000); under misalignment, the kinase Kin4 phosphorylates Bfa1, preventing the inhibitory phosphorylation by Cdc5 (Pereira and Schiebel, 2005).

In the model, Tem1 binds to the yeast centrosomes (called spindle pole bodies, SPBs) via GAP-dependent and GAP-independent sites. The intrinsic GTPase switching cycle of Tem1 is modeled as a reversible first-order reaction that converts Tem1^{GTP} into Tem1^{GDP}

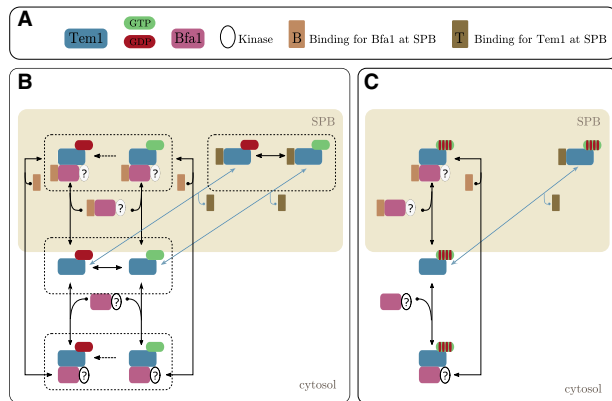


Fig. 3. SE for the SPOC dynamical model from Caydasi et al. (2012). (A) Model subunits. (B) Illustration of the pathway. Beige boxes indicate the SPB compartment. Reactions crossing the compartment boundary represent the reversible SPB association of the respective species or complexes. Blue reactions mark the intrinsic Tem1 GTPase-cycle and reversible SPB association. Tem1 which is bound directly to the SPB does not interact with Bfa1, whilst Tem1 in the cytosol interacts with cytosol-Bfa1 and SPB-bound Bfa1. These interactions occur for all instances of Bfa1 regardless of the Bfa1 phosphorylation state (indicated by the ‘?’ symbol). GTP hydrolysis by the respective Bfa1-Tem1-GTP complexes (dashed reaction arrow) is accelerated according to the GAP activity of the respective state of Bfa1. The dashed boxes represent the SE equivalence classes indicating that the two forms of the GTPase Tem1 are equal up to SE. This equivalence extends to all complexes with the same configuration of subunits, up to GTP/GDP binding state of Tem1. (C) Graphical interpretation of the network reduced by SE

and vice versa (Caydasi et al., 2012, Supplementary Material, Section S1). The maximal SE collapses complexes that are equal up to the GTP- or GDP-bound state, yielding eight equivalence classes with pairs of two molecular species (Fig. 3). The original network with 24 species and 71 reactions is reduced to 16 species and 36 reactions, from which one may recover observables of interest such as the total amount of active Bfa1 (Caydasi et al., 2012; Supplementary Material, Section S3).


3.4 SE for epidemic processes in networks

Models of epidemic processes are well established since the celebrated work by Kermack and McKendrick (1927). The availability of large datasets in a range of socio-technical systems has prompted the study of epidemic processes on complex networks that consider the heterogeneity of real-world processes, which is neglected in simpler variants that assume a well-mixed, uniform environment (Pastor-Satorras et al., 2015).

Aggregation of epidemic processes on networks has been studied in Simon et al. (2011), relating symmetries in the graph with lumping on the Markov chain. Graph symmetry is formalized in terms of nodes belonging to the same orbit, thereby satisfying the property that there exists a graph automorphism relating them. Then, the orbit partition, i.e. the partition of nodes where each block is a distinct orbit, induces a Markov chain lumping that tracks the number of nodes in each block of the orbit partition that are in any given state (Simon et al., 2011).

Here, we show that SE can be seen as a complementary, exact aggregation method for epidemic processes on complex networks. As an example, we study the well-known susceptible-infected-susceptible (SIS) model, where each node in the network in the susceptible state can be infected with a rate proportional to the number of infected neighbors, and recover from the infection according to an independent Poissonian process. Let $A = (a_{ij})$, with $A \in \mathbb{R}^{N \times N}$, define the adjacency matrix of a graph with N nodes representing the network topology, with $a_{ij} > 0$ denoting the presence of a possibly weighted edge between node i and j .

The SIS epidemic process can be described by the network:



where the first reaction models infections by neighbors and the second reaction is the spontaneous recovery, with parameters λ and γ . In a similar fashion, different variants of the process, such as SIR, SIRS and SEIR (Pastor-Satorras et al., 2015), can be described. Any physically meaningful initial state $\hat{\sigma}$ for this network must be such that each node i is initially in infected ($\hat{\sigma}_{S_i} = 0, \hat{\sigma}_{I_i} = 1$) or susceptible ($\hat{\sigma}_{S_i} = 1, \hat{\sigma}_{I_i} = 0$). This setting makes stochastic models of epidemics spreading on complex networks difficult to study exactly because the state of each individual node is tracked explicitly (Wang




Fig. 4. Example of SE reduction of SIS dynamics on a coarse-grained network. (A) Star network over which an SIS process evolves according to Equation (2), starting from an initial condition where the infection starts at node 0. (B) Reduced network (species representatives are underlined in the figure for clarity) according to the largest SE refinement of the initial partition with blocks $\mathcal{S} = \{\underline{S}_0, S_1, S_2, S_3, S_4\}$ and $\mathcal{I} = \{\underline{I}_0, I_1, I_2, I_3, I_4\}$. This SE has blocks $\{S_0\}, \{I_0\}, \{S_1, S_2, S_3, S_4\}$ and $\{I_1, I_2, I_3, I_4\}$. (C) The SE partition induces a partition on the graph with blocks $\{0\}$ and $\{1, 2, 3, 4\}$. The reduced network corresponds to the description of the SIS dynamics on the quotient graph. The lumpability relation holds for an initial condition of the reduced network that is consistent with the initial condition of the original network up to SE

Table 1. Aggregation of SIS dynamics on benchmark networks

Network	Reference	Original size		Reduced size		Orbits
		N	E	N	E	
tntp-ChicagoRegional	Eash et al. (1983)	1467	2596	635	932	166
ego-facebook	McAuley and Leskovec (2012)	2888	5962	35	104	35
as20000102	Leskovec et al. (2007)	6474	27 790	3885	19 437	3690
arenas-pgp	Boguñá et al. (2004)	10 680	48 632	8673	44 074	7944
web-webbase-2001	Boldi et al. (2004)	16 062	51 186	5253	24 232	3574
as-caida20071105	Leskovec et al. (2007)	26 475	106 762	13 393	69 184	13 252
ia-email-EU	Leskovec et al. (2007)	32 430	108 794	6262	53 228	6259
topology	Zhang et al. (2005)	34 761	215 440	19 246	168 782	19 128
douban	Zafarani and Liu (2009)	154 908	654 324	59 524	462 128	59 493

Note: Models tntp-ChicagoRegional, ego-facebook, as20000102, arenas-pgp, as-caida20071105, topology and douban, are taken from the Koblenz Network Collection (Kunegis, 2013); web-webbase-2001 and ia-email-EU are taken from the Network Data Repository (Rossi and Ahmed, 2015). N = number of vertices; E = number of edges in the network.

et al., 2017), leading to a state space size with 2^N distinct configurations (Simon et al., 2011). SE provides an ordinary lumpability of the underlying Markov chain, without ever generating it, on the network of Equation (2), which has exponentially smaller size because it has $2N$ species and $E+N$ reactions, where E is the number of nonzero entries in the adjacency matrix of the graph.

For the SIS model, the maximal SE is the trivial partition where all the species are in a single block. This is an invariant property stating that the total population of individuals in the system is constant (Simon et al., 2011). Thus, we consider non-degenerate reductions using initial partitions with two blocks, $\{S_1, \dots, S_n\}$ and $\{I_1, \dots, I_n\}$, that separate species associated with nodes in the susceptible state from those in the infected state, respectively. As an illustrative example, let us consider the simple star graph (Fig. 4). An inspection of the obtained SE equivalence classes reveals that each refinement of the initial block $\{S_1, S_2, S_3, S_4\}$ matches a refinement of $\{I_1, I_2, I_3, I_4\}$ for the same subset of nodes of the graph. Such an SE naturally induces a partitioning of the graph, and the reduction can be understood as an SIS dynamics on the quotient graph where each macro-node subsumes a partition block of nodes induced by SE.

We performed a systematic analysis of SIS processes evolving on several real-world benchmark networks (Table 1), which confirms the observation made on the simple star graph. Since in all cases the reduced model is interpretable as an epidemic process, it is still amenable to a wide range of analysis techniques developed for such models (Pastor-Satorras et al., 2015; Wang et al., 2017). These include mean-field and pair approximation (Cator and Van Mieghem, 2012; Mata and Ferreira, 2013; Van Mieghem, 2011), whose computational cost for the generation and solution of the resulting nonlinear differential equations may benefit from the availability of a stochastically equivalent reduced model.

Coarser aggregations of the Markov chain state space could be obtained in principle. For example, the line graph in Figure 4C admits the orbit partition that collapses nodes 0 and 1, thereby inducing a lumping following (Simon et al., 2011). However, this is not detected by SE. Importantly, this does not contradict our characterization result. The reason is that the lumpability relation induced by SE must hold for all population vectors that are equal up to SE. However, the lumpable partition derived with the approach in Simon et al. (2011) violates this property because it does not aggregate states $S_0 + I_0 + S_1 + I_1$ and $S_0 + S_0 + I_1 + I_1$, which preserve the sums of infected and susceptible individuals. Indeed, in the real-world networks in Table 1 we found that SE always induces a partition on the nodes of the graph which is finer than the orbit partition [whose size is listed in the last column, as reported in Ball and Geyer-Schulz (2018)], albeit not considerably so in some cases. On the other hand, SE can be applied to models that do not satisfy the conditions in Simon et al. (2011). Indeed, the star network of Figure 4 can be lumped also in the case of node-specific parameters (Supplementary Material, Section SA.5), while the results in Simon

et al. (2011) require equal transmission and recovery rates at every node.

3.5 Relationship with syntactic Markovian bisimulation

Applied to the models presented in this section, the earlier variant of SE, syntactic Markovian bisimulation (Cardelli et al., 2017b), yields the same reductions when applied to networks where the kinetic parameters are fixed. In Supplementary Material, Section A.6, we present further models from the literature where SE yields maximal aggregations that are coarser than syntactic Markovian bisimulation, with up to about one order of magnitude fewer species.

3.6 Speeding up stochastic simulations

In Supplementary Material, Section A.7, we use the same set of models to also provide evidence of the computational savings when analyzing by stochastic simulation the reduced network in place of the original one. We report runtime speed-ups of up to three orders of magnitude using state-of-the-art algorithms as implemented in the StochKit simulation framework (Sanft et al., 2011).

4 Conclusion

Stochasticity is a key tool to understand a variety of phenomena regarding the dynamics of reaction networks, but the capability of exactly analyzing complex models escapes us due to the lack of analytical solutions and the high computational cost of numerical simulations in general. SE enables aggregation in the sense of Markov chain lumping by identifying structural properties on the set of reactions, without the need of costly state-space enumeration. Owing to the polynomial space and time complexity of the reduction algorithm, it can be seen as a universal pre-processing step that exactly preserves the stochastic dynamics of species of interest to the modeler. Since it gives rise to a network where the reactions preserve the structure (up to a renaming of the species into equivalence classes), the reduction maintains a physical interpretation in terms of coarse-grained interactions between populations of macro-species. The possibility of computing reductions that are not dependent from specific values of the kinetic parameters may reveal structural aggregations in the network, in addition to making the reduced model reusable across different parameter settings, e.g. when performing sensitivity analyses.

Being exact, our method can be combined with other techniques for the analysis of stochastic reaction networks. For instance, when feasible, one can generate the underlying Markov chain to be further analyzed or reduced (Henzinger et al., 2009; Munsky and Khammash, 2006; Valmari and Franceschini, 2010); the reduced network can be subjected to complementary coarse-graining techniques concerned with time-scale separation (e.g. Bo and Celani, 2017; Cappelletti and Wiuf, 2016; Gómez-Uribe et al., 2008; Kang

and Kurtz, 2013; Smith et al. 2015; Sinitzyn et al., 2009). More generally, since the reduced network preserves the stochastic dynamics in the sense specified above, it can be used as the basis for other methods such as linear noise-or moment-closure approximation (Schnoerr et al., 2017), where the complexity of the resulting system of equations depends on the network size.

Funding

This work has been partially supported by Italian Ministry for Research under the PRIN project ‘SEDUCE’ [no. 2017TWRCNB], by the Independent Research Fund Denmark under the DFF RP1 Project REDUCTO [no. 9040-00224B] and the Danish Poul Due Jensen Foundation [grant 883901].

Conflict of Interest: none declared.

References

- Autant,C. and Schnobelen,P. (1992). Place bisimulations in Petri nets. In: *Application and Theory of Petri Nets*, Lecture Notes in Computer Science, vol. 616, Springer: Berlin Heidelberg, pp. 45–61
- Ball,F. and Geyer-Schulz,A. (2018) How symmetric are real-world graphs? A large-scale study. *Symmetry*, 10(1), 29.
- Bo,S. and Celani,A. (2017) Multiple-scale stochastic processes: decimation, averaging and beyond. *Phys. Rep.*, 670, 1–59.
- Boguñá,M. et al. (2004) Models of social networks based on social distance attachment. *Phys. Rev. E*, 70, 056122.
- Boldi,P. et al. (2004) UbiCrawler: a scalable fully distributed web crawler. *Softw. Pract. Exp.*, 34, 711–726.
- Brijder,R. (2019) Computing with chemical reaction networks: a tutorial. *Nat. Comput.*, 18, 119–137.
- Buchholz,P. (1994) Exact and ordinary lumpability in finite Markov Chains. *J. Appl. Prob.*, 31, 59–75.
- Cappelletti,D. and Wiuf,C. (2016) Elimination of intermediate species in multiscale stochastic reaction networks. *Ann. Appl. Prob.*, 26, 2915–2958.
- Cardelli,L. et al. (2015). Forward and backward bisimulations for chemical reaction networks. In: *26th International Conference on Concurrency Theory, CONCUR*, Lecture Notes in Computer Science (LNCS), vol. 10206, Springer: Berlin Heidelberg, pp. 226–239.
- Cardelli,L. et al. (2017a). ERODE: a tool for the evaluation and reduction of ordinary differential equations. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pp. 310–328.
- Cardelli,L. et al. (2017b). Syntactic Markovian bisimulation for chemical reaction networks. In: *Models, Algorithms, Logics and Tools*, pp. 466–483.
- Cator,E. and Van Mieghem,P. (2012) Second-order mean-field susceptible-infected-susceptible epidemic threshold. *Phys. Rev. E*, 85, 056111.
- Caydasi,A.K. et al. (2012) A dynamical model of the spindle position checkpoint. *Mol. Syst. Biol.*, 8, 582.
- Derisavi,S. et al. (2003) Optimal state-space lumping in Markov chains. *Inf. Process. Lett.*, 87, 309–315.
- Eash,R.W. et al. (1983) Equilibrium traffic assignment on an aggregated highway network for sketch planning. *Transp. Res. Rec.*, 994, 30–37.
- Elowitz,M.B. (2002) Stochastic gene expression in a single cell. *Science*, 297, 1183–1186.
- Feret,J. et al. (2012) Lumpability abstractions of rule-based systems. *Theor. Comput. Sci.*, 431, 137–164.
- Feret,J. et al. (2013) Stochastic fragments: a framework for the exact reduction of the stochastic semantics of rule-based models. *Int. J. Softw. Inform.*, 7, 527–604.
- Gillespie,D. (1977) Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.*, 81, 2340–2361.
- Gómez-Uribe,C.A. et al. (2008) Enhanced identification and exploitation of time scales for model reduction in stochastic chemical kinetics. *J. Chem. Phys.*, 129, 244112.
- Gruneberg,U. (2000) Nud1p links astral microtubule organization and the control of exit from mitosis. *EMBO J.*, 19, 6475–6488.
- Gunawardena,J. (2005) Multisite protein phosphorylation makes a good threshold but can be a poor switch. *Proc. Natl. Acad. Sci. USA*, 102, 14617–14622.
- Guptasarma,P. (1995) Does replication-induced transcription regulate synthesis of the myriad low copy number proteins of *Escherichia coli*? *BioEssays*, 17, 987–997.
- Henzinger,T.A. et al. (2009). Sliding window abstraction for infinite Markov chains. In: *International Conference on Computer Aided Verification*, pp. 337–352. Springer.
- Joyal,A. et al. (1996) Bisimulation from open maps. *Inform. Comput.*, 127, 164–185.
- Kang,H.-W. and Kurtz,T.G. (2013) Separation of time-scales and model reduction for stochastic reaction networks. *Ann. Appl. Prob.*, 23, 529–583.
- Kemeny,J. and Snell,J. (1976). *Finite Markov Chains*. Springer New York, Heidelberg, Berlin.
- Kermack,W.O. and McKendrick,A.G. (1927) A contribution to the mathematical theory of epidemics. *Proc. R. Soc. Lond. Ser. A*, 115, 700–721.
- Kunegis,J. (2013). KONECT – The Koblenz Network Collection. In: *Proceedings of the International Conference on World Wide Web Companion*, pp. 1343–1350.
- Leskovec,J. et al. (2007) Graph evolution: densification and shrinking diameters. *ACM Trans. Knowl. Discov. Data*, 1, 2–40.
- Lew,D.J. and Burke,D.J. (2003) The spindle assembly and spindle position checkpoints. *Annu. Rev. Genet.*, 37, 251–282.
- Li,C. et al. (2010) BioModels database: an enhanced, curated and annotated resource for published quantitative kinetic models. *BMC Syst. Biol.*, 4, 92.
- Lisman,J. et al. (2002) The molecular basis of CaMKII function in synaptic and behavioural memory. *Nat. Rev. Neurosci.*, 3, 175–190.
- Lisman,J. et al. (2012) Mechanisms of CaMKII action in long-term potentiation. *Nat. Rev. Neurosci.*, 13, 169–182.
- Mata,A.S. and Ferreira,S.C. (2013) Pair quenched mean-field theory for the susceptible-infected-susceptible model on complex networks. *Urophys. Lett.*, 103, 48003.
- McAuley,J. and Leskovec,J. (2012). Learning to discover social circles in ego networks. In: *Advances in Neural Information Processing Systems*, pp. 548–556.
- Munsky,B. and Khammash,M. (2006) The finite state projection algorithm for the solution of the chemical master equation. *J. Chem. Phys.*, 124, 044104.
- Paige,R. and Tarjan,R. (1987) Three partition refinement algorithms. *SIAM J. Comput.*, 16, 973–989.
- Pastor-Satorras,R. et al. (2015) Epidemic processes in complex networks. *Rev. Modern Phys.*, 87, 925–979.
- Pepke,S. et al. (2010) A dynamic model of interactions of Ca²⁺, calmodulin, and catalytic subunits of Ca²⁺/calmodulin-dependent protein kinase II. *PLoS Comput. Biol.*, 6, e1000675.
- Pereira,G. and Schiebel,E. (2005) Kin4 kinase delays mitotic exit in response to spindle alignment defects. *Mol. Cell*, 19, 209–221.
- Rossi,R.A. and Ahmed,N.K. (2015). The network data repository with interactive graph analytics and visualization. In: *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*.
- Rózsa,Z. and Tóth,J. (2003) Exact linear lumping in abstract spaces. *Electron. J. Qual. Theory Differ. Equ.*, 21, 1–20.
- Salazar,C. and Höfer,T. (2009) Multisite protein phosphorylation – from molecular mechanisms to kinetic models. *FEBS J.*, 276, 3177–3198.
- Sanft,K.R. et al. (2011) Stochkit2: software for discrete stochastic simulation of biochemical systems with events. *Bioinformatics*, 27, 2457–2458.
- Schnoerr,D. et al. (2017) Approximation and inference methods for stochastic biochemical kinetics—a tutorial review. *J. Phys. A Math. Theor.*, 50, 093001.
- Shifman,J.M. et al. (2006) Ca²⁺/calmodulin-dependent protein kinase II (CaMKII) is activated by calmodulin with two bound calciums. *Proc. Natl. Acad. Sci. USA*, 103, 13968–13973.
- Simon,P.L. et al. (2011) Exact epidemic models on graphs using graph-automorphism driven lumping. *J. Math. Biol.*, 62, 479–508.
- Sinitzyn,N.A. et al. (2009) Adiabatic coarse-graining and simulations of stochastic biochemical networks. *Proc. Natl. Acad. Sci. USA*, 106, 10546–10551.
- Smith,S. et al. (2015) Model reduction for stochastic chemical systems with abundant species. *J. Chem. Phys.*, 143, 214105.
- Sneddon,M.W. et al. (2011) Efficient modeling, simulation and coarse-graining of biological complexity with NFsim. *Nat. Methods*, 8, 177–183.
- Snowden,T.J. et al. (2017) Methods of model reduction for large-scale biological systems: a survey of current methods and trends. *Bull. Math. Biol.*, 79, 1449–1486.
- Thomson,M. and Gunawardena,J. (2009) Unlimited multistability in multisite phosphorylation systems. *Nature*, 460, 274–277.
- Valmari,A. and Franceschini,G. (2010). Simple $O(m \log n)$ time Markov Chain lumping. In: *International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pp. 38–52.

- Van Kampen, N.G. (2007). *Stochastic Processes in Physics and Chemistry*, 3rd edn. Elsevier: North Holland.
- Van Mieghem, P. (2011) The n -intertwined SIS epidemic network model. *Computing*, **93**, 147–169.
- Voit, E.O. et al. (2015) 150 years of the mass action law. *PLoS Comput. Biol.*, **11**, e1004012–7.
- Wang, W. et al. (2017) Unification of theoretical approaches for epidemic spreading on complex networks. *Rep. Prog. Phys.*, **80**, 036603.
- Zafarani, R. and Liu, H. (2009). Social computing data repository at ASU, <http://socialcomputing.asu.edu/pages/home>.
- Zhang, B. et al. (2005) Collecting the Internet AS-level topology. *SIGCOMM Comput. Commun. Rev.*, **35**, 53–61.