

Hypergraph-based models of random chemical reaction networks: Conservation laws, connectivity, and percolation

Cite as: J. Chem. Phys. 163, 194112 (2025); doi: 10.1063/5.0292914

Submitted: 24 July 2025 • Accepted: 30 October 2025 •

Published Online: 20 November 2025



Shesha Gopal Marehalli Srinivas,^{1,a)} Massimiliano Esposito,^{1,b)} and Nahuel Freitas^{1,2,c)}

AFFILIATIONS

¹Complex Systems and Statistical Mechanics, Department of Physics and Materials Science, University of Luxembourg, 30 Avenue des Hauts-Fourneaux, L-4362 Esch-sur-Alzette, Luxembourg

²Departamento de Física, FCEyN, UBA, Pabellón 1, Ciudad Universitaria, 1428 Buenos Aires, Argentina

^{a)}Author to whom correspondence should be addressed: shesha.marehalli@uni.lu

^{b)}Electronic mail: massimiliano.esposito@uni.lu

^{c)}Electronic mail: nfreitas@df.uba.ar

ABSTRACT

Random graph models have been instrumental in characterizing complex networks, but chemical reaction networks (CRNs) are better represented as hypergraphs. Traditional models of random CRNs often reduce CRNs to bipartite graphs, representing species and reactions as distinct nodes, or simpler derived graphs, which can obscure the relationship between the statistical properties of these representations and the physical characteristics of the CRN. We introduce a straightforward model for generating random CRNs that preserve their hypergraph structure and atomic composition, enabling the direct study of chemically relevant features. Notably, our approach distinguishes two notions of connectivity that are equivalent in graphs but differ fundamentally in hypergraphs. These notions exhibit percolation-like phase transitions, which we analyze in detail. The first type of connectivity has relevance to steady-state synthesis and transduction, determining the effective reactions an open CRN can perform at steady state. The second type is suitable to identify which species can be produced from a given initial set of species in a closed CRN. Our findings highlight the importance of hypergraph-based modeling for uncovering the complex behaviors of CRNs.

Published under an exclusive license by AIP Publishing. <https://doi.org/10.1063/5.0292914>

I. INTRODUCTION

Chemical reaction networks (CRNs) are defined by a set of chemical species and the reactions in which they are involved. In many instances (e.g., metabolism, biogeochemistry, atmospheric chemistry, and astrobiology), the space of possible chemical compounds is vast and complex, even when restrictions are applied to the composition, structure, and size of molecules.^{1,2} The space of possible chemical reactions is typically even larger.^{3–5} Thus, from biology to engineering, it is crucial to characterize the structural properties of large complex CRNs and their statistical features.

As for the study of complex networks with random graphs,^{6–8} or complex quantum systems with random matrix theory,⁹ establishing a meaningful notion of random CRNs (under given constraints) would be highly suitable. It would provide a useful ‘null

hypothesis’ against which actual networks can be compared. The power of this approach lies in the fact that random models filter out overly specific details and highlight generic characteristics—thus revealing classes of typicality.

Attempts have been made to apply these ideas to complex CRNs by resorting to graph-based representations of CRNs. It has been argued, for instance, that the graphs of metabolic networks and the network of organic chemistry are scale-free,^{5,10} but the biochemical relevance of such results has been questioned.^{11–13} The major issue is that CRNs are hypergraphs and not graphs.¹⁴ In graphs, edges connect two nodes, whereas in hypergraphs, edges can connect multiple nodes. Since reactions can transform multiple species into multiple products, CRNs are naturally represented as hypergraphs. Projections into bipartite graphs can be achieved by assigning nodes to every species as well as to every reaction and placing edges from

a given reaction node to every species node involved in that reaction. Since, in contrast with hypergraphs, many tools are available to study the statistical properties of graphs, most studies resort to such representations.^{5,10,13,15,16} However, it has been shown that these mappings blur the relationship between the statistical properties of these representations and the physical characteristics of the CRN and can lead to misleading conclusions.^{14,17,18} The usual graph theoretic quantities, such as degree distribution or clustering, have little chemical meaning in this context. Developing tools to study the statistical properties of hypergraphs rather than graph representations of CRNs is thus essential to characterize chemically meaningful features of complex CRNs.

Abstract mathematical models of random hypergraphs are starting to be studied.^{19–21} However, the generation of random CRNs poses specific challenges related to the chemical consistency of the network.²² Indeed, since chemical reactions conserve atomic compositions, the resulting random CRNs must display corresponding conservation laws. Random hypergraph models of CRNs have been proposed in the literature. Some of these works study the probability that deficiency (a structural property) arises in random CRNs,^{23,24} but without enforcing any conservation laws. Other studies define random CRNs that do preserve conservation laws but focus on dynamical features rather than structural ones.^{25,26}


Our contributions in this article are threefold. First, we propose a simple model of random CRNs that, by construction, enforce the conservation of atomic compositions. Starting from a representation solely in terms of the atomic composition of molecules, we assign reactions with a probability p from the set of possible reactions allowed by atomic conservation laws. This generation scheme offers practical and conceptual advantages over prior alternatives, which sample reactions in an unconstrained way and then filter for conservation.^{25,26} In contrast, every CRN in our ensemble conserves mass and atoms by construction, eliminating the filtering stage and also enabling analytical progress. Second, we identify two different notions of connectivity, the first characterizing the existence of effective reactions between a set of chosen species that are exchanged in an open CRN, the second identifying which species can be produced from a given initial set of chosen species in a closed CRN. Both notions are equivalent for graphs, but crucially different for CRNs. Third, we characterize the statistical properties of these notions of connectivity and show that they each manifest different percolation-like phase transitions wherein the CRNs transition from low connectivity to high connectivity. These findings have relevance in determining the likelihood of different types of synthetic pathways in complex CRNs.

This paper is structured as follows: Section II provides the background; we introduce CRNs and their structure in Subsection II A; we discuss issues with graph projections of CRNs in Subsection II B; we introduce our two notions of connectivity in Subsection II C; in Sec. III, we describe our method for generating chemically consistent random CRNs; in Sec. IV, we present our main results on the various statistical features of our model of random CRNs and show that they display percolation-like transitions; in Sec. V, we discuss our findings and compare our method for generating random CRNs with other methods that have been used in the literature; and in Sec. VI, we conclude with a discussion on the broader implications and future scope of our results.

II. INTRODUCTION TO CHEMICAL REACTION NETWORKS

A. Structural aspects of CRNs

We consider CRNs wherein the chemical species, labeled $\sigma \in \mathcal{Z}$, are interconverted via elementary,²⁷ reversible, and mass-balanced chemical reactions $\rho \in \mathcal{R}$ of the form



Here, $\boldsymbol{\sigma} = (\dots, \sigma, \dots)^\top$ and $\mathbf{v}_{+\rho} = (\dots, v_{\sigma,+\rho}, \dots)^\top$ (resp. $\mathbf{v}_{-\rho} = (\dots, v_{\sigma,-\rho}, \dots)^\top$) is the vector collecting the stoichiometric coefficients of the forward $+\rho$ (resp. backward $-\rho$) reaction. The structure (also called topology) of CRNs is encoded in the stoichiometric matrix ∇ whose columns, $\nabla_\rho = \mathbf{v}_{-\rho} - \mathbf{v}_{+\rho}$, capture the net change in every species due to the reaction ρ .

The linearly independent vectors that span the left null space of the stoichiometric matrix, $\boldsymbol{\ell}^\lambda \cdot \nabla = 0$, are called conservation laws: they identify parts of (or entire) molecules, called moieties that are preserved by the reactions. Since we consider CRNs with all mass-balanced chemical reactions, there is always a conservation law, called the mass conservation law and denoted $\boldsymbol{\ell}^m$, which involves all the species: $\boldsymbol{\ell}^m = (\dots, \ell_\sigma^m, \dots)^\top$ with $\ell_\sigma^m \geq 1$. We emphasize that every physically consistent CRN has at least one mass conservation law.²²

The linearly independent vectors that span the right null space of the stoichiometric matrix, $\nabla \cdot \mathbf{c} = 0$, are called (stoichiometric) cycles: they identify series of reactions that leave all the chemical species unchanged upon completion. Denoting the set of conservation laws as $\mathcal{L} = \{\boldsymbol{\ell}^\lambda\}$ and the set of cycles as $\mathcal{C} = \{\mathbf{c}\}$, the rank-nullity theorem²⁸ applied to the stoichiometric matrix ∇ and its transpose $\nabla^\top$ gives the relation,

$$\text{rk}(\nabla) = |\mathcal{Z}| - |\mathcal{L}| = |\mathcal{R}| - |\mathcal{C}|, \quad (2)$$

where $\text{rk}(\cdot)$ denotes the rank of the matrix and $|\cdot|$ denotes the cardinality of the set. We will make extensive use of this relation in Sec. IV C.

The structural considerations described so far are appropriate for *closed* CRNs, where no species are exchanged with the surroundings. For closed CRNs, conservation laws define dynamical invariants,^{29,30} and both cycles and conservation laws are essential for thermodynamic descriptions.^{31,32} For open CRNs where some chemical species are exchanged with the surroundings, we partition the set of species $\mathcal{Z}$ into distinct subsets of internal species (X) and external species (Y), where the external species are exchanged with the surroundings. We refer to the selection of a set of Y species as chemostatting the corresponding species. We emphasize that this work focuses exclusively on the structural aspects of CRNs, without considering kinetics. The open/closed distinction used here represents the physical settings we use to describe the two notions of connectivity.

The structural aspects of *open* CRNs are derived using the (sub)stoichiometric matrix with rows that correspond only to the X species and denoted by ∇^X . In open CRNs, the set of conservation laws $\mathcal{L}$ splits into the sets of unbroken (resp. broken) conservation laws, denoted $\mathcal{L}_u$ (resp. $\mathcal{L}_b$). The unbroken conservation laws

$\mathcal{L}^{\lambda_u}$ satisfy $\mathcal{L}^{\lambda_u} \cdot \nabla^X = 0$, i.e., they are conserved moieties of the internal species and are unaffected by the exchange. Broken conservation laws, denoted by $\mathcal{L}^{\lambda_b}$, on the other hand, involve the chemostatted species and are affected by the exchange with the environment.

A key structural feature of open CRNs are *emergent cycles*: effective interconversions of the external species that leave the internal species unaffected. Thus, the emergent cycles, denoted by $\mathbf{c}_e$, satisfy

$$\nabla^X \cdot \mathbf{c}_e = \mathbf{0} \quad \text{and} \quad \nabla^Y \cdot \mathbf{c}_e \neq \mathbf{0}, \quad (3)$$

where ∇^Y is the substoichiometric matrix with rows corresponding to the Y species. Denoting the number of emergent cycles as $|\epsilon|_Y$, from Eq. (3), we infer that

$$|\epsilon|_Y = \dim(\ker(\nabla^X)) - |C| \geq 0. \quad (4)$$

Given an emergent cycle $\mathbf{c}_e$, the stoichiometric coefficients of the effective reaction are obtained as the entries of the vector $\nabla^Y \cdot \mathbf{c}_e$, up to multiplication by a constant. By collecting the negative (resp. positive) elements of $\nabla^Y \cdot \mathbf{c}_e$ into $\mathbf{v}_+^e$ (resp. $\mathbf{v}_-^e$), we can write the effective reaction as



where $\sigma_y = (\sigma \mid \sigma \in Y)^T$ is the vector of external species. We emphasize that (i) emergent cycles interconvert the sets of substrates and products and (ii) the number of emergent cycles depends on the choice of the external species.

We can link the number of broken conservation laws and the number of emergent cycles by the relation,³³

$$|\mathcal{L}_b| + |\epsilon|_Y = |Y|, \quad (6)$$

where $|Y|$ denotes the number of external species. Equation (6) provides a fundamental insight: chemostatting a species either breaks a conservation law or creates an emergent cycle. The dynamical and thermodynamical importance of emergent cycles and broken conservation laws has been thoroughly explored in the literature, and we refer readers to Refs. 29, 30, 33, and 34 for more details.

B. Graph projections of CRNs

Hypergraphs are a generalization of graphs wherein (hyper)edges can connect more than two nodes. Since chemical reactions as considered in Eq. (1) generically involve multiple chemical species, they are appropriately represented by hyperedges. By mapping chemical species onto nodes and reactions onto hyperedges, CRNs can be represented as hypergraphs.^{14,35} Since we work with reversible reactions, the corresponding hyperedges are arbitrarily oriented. Reversing these orientations will not change any of the results.

Several graph representations of CRNs have been considered in the literature. We briefly describe the most important representations in the following and direct readers to Refs. 16 and 22 for a comprehensive review of the same.

1. S, R and SR graphs

Any CRN can be mapped into a bipartite graph by mapping the species and reactions onto two distinct types of nodes. In this

representation, edges connect the species nodes to a reaction node if the species participate in the corresponding reaction. These graphs are also called SR-graphs (species-reaction graphs) and can be used to extract information about the dynamics of the CRN.^{36–38} The SR graph can be further projected into the S graph (substrate graph) and R graph (reaction graph), respectively. In the S graph,³⁹ the nodes are chemical species and two species σ, σ' are connected by an edge if there is a reaction in which σ is the reactant and σ' is the product (or vice versa). In the R graph, chemical reactions are nodes and are connected by an edge if there is a chemical species involved in both reactions. Although the S and R graphs can be unambiguously derived from the SR graph, the opposite may not be true,¹⁸ implying an information loss associated with moving from the bipartite graph description to the S graph or the R graph.

2. Graph of complexes


Another way to define a graph from a CRN is to consider the vectors $\mathbf{v}_{\pm p}$ (called complexes) of Eq. (1) as nodes. Every reaction p is then an edge connecting the node $\mathbf{v}_{+p}$ to the node $\mathbf{v}_{-p}$ (or vice versa). This description has been used in chemical reaction network theory to define a topological quantity, called deficiency, which is linked to the asymptotic dynamics and thermodynamics of CRNs.^{38,40–42} The deficiency of the CRN, denoted by δ , is defined as the difference between the number of stoichiometric cycles and the number of cycles in the graph of complexes, i.e.,

$$\delta = \dim(\text{Ker } \nabla) - \dim(\text{Ker } \partial) \geq 0, \quad (7)$$

where ∂ is the (oriented) incidence matrix of the graph of complexes.

3. Example

We consider the following example, inspired by Ref. 14, to illustrate some problems with the graphical descriptions discussed above; consider the following CRN:



with the stoichiometric matrix,

$$\nabla = \begin{array}{c} \begin{matrix} A \\ B \\ C \\ D \end{matrix} \end{array} \begin{array}{ccc} \begin{matrix} \rho_1 & \rho_2 & \rho_3 \end{matrix} \\ \begin{bmatrix} -1 & -1 & 0 \\ 1 & -1 & 0 \\ 0 & 1 & -1 \\ 0 & 1 & 1 \end{bmatrix} \end{array}. \quad (9)$$

The hypergraph and graphs associated with Eq. (8) are shown in Fig. 1(a).

Suppose that we wish to answer the following question: can the species D be synthesized starting from the species A? By considering the set $Y = \{A, D\}$ as external species and carrying out the procedure outlined below Eq. (5), we find the emergent cycle $\mathbf{c}_e = (1, 1, 1)^T$ corresponding to the net reaction $2A \rightleftharpoons 2D$. This shows that starting from A, we can use the first reaction to create species B and then perform the second reaction to create the species C, which then creates species D via the third reaction. If we




FIG. 1. (a) CRN in Eq. (8) represented as a hypergraph and a bipartite (SR) graph, with reaction nodes colored. (b) Same representation of the CRN when reaction ρ_1 is removed. The blue dashed arrow depicts the path from A to D despite the fact that the CRN is unable to achieve it without ρ_1 .

now remove the first reaction, this is not possible anymore. Indeed, repeating the same procedure, the emergent cycle described above disappears. However, the SR graph still has a path starting from the node A ending at the node D, as shown in Fig. 1(b). This illustrates that fundamental graph notions, such as connectivity in the SR graph (or in the S and R graphs as well), might not be chemically meaningful.

C. Notions of connectivity

From a chemical perspective, a natural question is whether a possibly large network of elementary reactions can, based solely on its topology (i.e., independent of any kinetic or thermodynamic considerations), interconvert certain chemical species into others. More precisely, given a set of substrates and a set of desired products, we ask whether there exists a sequence of reactions that effectively consumes the substrates to produce the products. In this subsection, we show that this question can be formulated in two distinct ways.

1. Chemostatted and emergent cycles

The first scenario characterizes the capacity of an open CRN to sustain steady-state synthesis. Given a set of externally maintained (chemostatted) species, we ask whether at least one emergent cycle exists. If so, the CRN can sustain the steady-state synthesis of the chemostatted products from the chemostatted substrates associated with each emergent cycle, without generating any additional waste products.

2. Reachability from a single species

The second scenario characterizes the capacity of a closed CRN, initialized with an infinite amount of a single species, to synthesize other species as it relaxes to equilibrium. In other words, the question is: which species can be produced from a given initial species by iterating the network's reactions until no new species are produced? If we start with species σ , we refer to the set of species that is obtained by the previous procedure as the *forward reachable set* of σ , denoted

$\mathcal{F}(\sigma)$. If $\sigma' \in \mathcal{F}(\sigma)$, we say that σ' can be reached from σ . This concept has also appeared in the metabolic network literature under the name *scope*, referring to the set of metabolites synthesizable from a given compound.⁴³

Forward reachability is generically not symmetric, i.e., $\sigma' \in \mathcal{F}(\sigma)$ does not imply $\sigma \in \mathcal{F}(\sigma')$. Moreover, unlike in the emergent cycle scenario where all species are assumed to be available and every reaction can occur, here we must account for the availability of substrates. If a substrate for a reaction is absent, it cannot proceed, even if the reverse reaction is allowed. As a result, even though all reactions are reversible, the choice of the initial species introduces a directionality. Forward reachability can be viewed as a chemically grounded analog of the notion of the out-component in a directed graph. In a directed graph, the out-component of a node is the set of nodes reachable by following directed edges.⁸ We note that the forward reachable set given two or more initial species can be defined analogously, but we do not consider that notion in this work.

Importantly for CRNs that can be represented as undirected graphs (i.e., CRNs solely made of unimolecular reactions $A \rightleftharpoons B$), the two notions of connectivity introduced above are equivalent. Indeed, in this case, forward reachability is symmetric and any terminal species involved in an emergent cycle belongs to the same reachable set. This shows that the two notions of connectivity introduced above can be considered non-equivalent generalizations of connectivity in undirected graphs.

As an illustration, we can go back to the CRN in Eq. (8) *without* the first reaction. In this case, chemostating the species $Y = \{A, D\}$ does *not* produce an emergent cycle. However, the forward reachable set of the species D is given by $\mathcal{F}(D) = \{D, C, A, B\}$. If we start with a large amount of D, then the species C can be formed via the third reaction and together, and they can react to form A and B. In turn, the forward reachable set of the species A is $\mathcal{F}(A) = \{A\}$ because no reaction can solely run on the species A. These results illustrate that forward reachability and emergent cycle are two different notions of connectivity and that forward reachability is asymmetric. It may seem that bidirectional reachability between two species σ and σ' (i.e. if σ' can be reached from σ and vice versa) would imply the existence of an emergent cycle if one would chemostat them. However, this is not generically true, as shown by counterexamples in Appendix A.

We end this section by discussing another common notion of connectivity in graphs: that of connected components. Mathematically, the number of connected components of a graph is equal to the dimension of the left nullspace of the (oriented) incidence matrix. Therefore, a natural analog of connected components in CRNs is given by the conservation laws, which generate the left nullspace of the stoichiometric matrix (see Sec. II A).

III. GENERATION OF RANDOM CRNS

In this section, we explain our algorithm for generating random CRNs. We begin by explaining the problem to be overcome. Let $\nabla = \{\nabla_{\rho}^{\sigma}\}_{\sigma=1, \dots, |\mathcal{Z}|, \rho=1, \dots, |\mathcal{R}|}$ be the $|\mathcal{Z}| \times |\mathcal{R}|$ stoichiometric matrix of a generic system of $|\mathcal{Z}|$ chemical species and $|\mathcal{R}|$ reactions. Since we work with closed CRNs, each column of ∇ must have at least one entry of each sign. If we restrict ourselves to reactions where all species appear with a stoichiometric coefficient less than or equal to κ , then the components of ∇ are integer numbers between $-\kappa$

and κ . If the order of the reaction is restricted to be less than or equal to K , i.e., $\sum_{\sigma} v_{\sigma, \pm p} \leq K$, then the sum of positive entries in each column ∇ must be less than or equal to K , and the same must be true for $-\nabla$. The entries of ∇ could be filled with random integers in a way such that those conditions are verified. However, for consistency, the resulting CRN must also have at least a mass conservation law. This condition is not trivial to satisfy in a random assignment of elements to ∇ since it introduces a constraint between different columns or reactions (in contrast to all the previous conditions, which are internal to each column). For example, the reactions $A + B \rightleftharpoons C$ and $A + B \rightleftharpoons 2C$ cannot be part of the same system. Furthermore, if the $|\mathcal{Z}|$ molecular species are formed by different types of atoms, then one might have additional conservation laws beyond that of the total mass since chemical reactions cannot alter atomic identities.

We propose the following solution to the problem: we start by fixing the number N_A of different atomic types forming the possible molecules, and we consider that the maximum number of atoms of type $a = 1, \dots, N_A$ involved in any molecular species is Q_a . In addition, we assume that a molecular species is fully determined by the number of atoms of each type that make it (the chemical formula). In other words, we represent a chemical species by its molecular formula, completely disregarding its internal structure. For example, the molecules pentene and cyclopentane would be the same for us since both have five carbons and ten hydrogen atoms. Then, at maximum, we can have

$$|\mathcal{Z}|^{\max} = \prod_{a=1}^{N_A} (Q_a + 1) - 1 \quad (10)$$

molecular species. Each one of them is identified by the tuple $\sigma = (n_1, n_2, \dots, n_{N_A})$, where n_a is the number of atoms of each type. We then focus on the case in which the overall order of any reaction is at most two, as is typically the case for elementary reactions. Given that we are disregarding the internal structure of the chemical species, unimolecular reactions such as $A \rightleftharpoons B$ are meaningless since A and B would be the same species due to atom conservation by reactions. Thus, we consider reactions of the kind $A + B \rightleftharpoons C$ or $2A \rightleftharpoons B$. Although bimolecular reactions such as $A + B \rightleftharpoons C + D$ could also be taken into account, they can be considered as effective descriptions of the process $A + B \rightleftharpoons X \rightleftharpoons C + D$ and, therefore, for simplicity, will not be counted separately in what follows. For a molecule $(n_1, \dots, n_{N_A})$, we have the possible dissociations,

$$(n_1, n_2, \dots, n_{N_A}) \rightleftharpoons (n_1 - m_1, n_2 - m_2, \dots, n_{N_A} - m_{N_A}) + (m_1, m_2, \dots, m_{N_A}) \quad (11)$$

for any values of $m_1, m_2, \dots, m_{N_A}$, such that $\min\{m_a\} > 0$, $\min\{n_a - m_a\} > 0$, and that $0 \leq m_a \leq n_a$ for all $1 \leq a \leq A$. Noting that the substitution of $n_a - m_a$ by m_a in (11) leaves the reaction unchanged, the total number of distinct dissociations of the molecule $\sigma = (n_1, \dots, n_{N_A})$ is

$$D(\sigma) = \frac{\prod_{a=1}^{N_A} (n_a + 1)}{2} - \pi(\sigma), \quad (12)$$

where $\pi(\sigma)$ is $1/2$ if all the n_a values are even, or 1 otherwise. Since any possible reaction is a dissociation of some species, the

total number of possible reactions is [by convention, we define that $D(\sigma = 0) = 0$]

$$\begin{aligned} |\mathcal{R}|^{\max} &= \sum_{\sigma} D(\sigma) \\ &= \sum_{n_1=0}^{Q_1} \sum_{n_2=0}^{Q_2} \cdots \sum_{n_{N_A}=0}^{Q_{N_A}} D(n_1, \dots, n_{N_A}) \\ &= 2^{-(N_A+1)} \prod_{a=1}^{N_A} (Q_a + 1)(Q_a + 2) \\ &\quad + 2^{-1} \prod_{a=1}^{N_A} (\lfloor Q_a/2 \rfloor + 1) - |\mathcal{Z}|^{\max} - 1. \end{aligned} \quad (13)$$

To derive the last line in Eq. (13), observe that $\pi(\sigma) = 1$, except for those σ with even entries, for which $\pi(\sigma) = 1/2$. Hence, $\sum_{\sigma} \pi(\sigma) = 1 + |\mathcal{Z}|^{\max} - (1/2) \prod_{a=1}^{N_A} (\lfloor Q_a/2 \rfloor + 1)$, where the negative term corrects for the σ with all even entries. Note that if $Q_a \gg 1$, then we have

$$|\mathcal{R}|^{\max} \simeq \frac{|\mathcal{Z}|^{\max} (|\mathcal{Z}|^{\max} + 1)}{2^{N_A+1}} - |\mathcal{Z}|^{\max}. \quad (14)$$

All the possible $|\mathcal{R}|^{\max}$ reactions can be easily enumerated and form the total set of reactions we consider. Appendix B illustrates Eqs. (10) and (13) for the simple case of $N_A = 2$, $Q_1 = Q_2 = 2$.

To create a random stoichiometric matrix ∇ , one possible approach is to randomly select exactly $|\mathcal{R}|$ reactions from the total set of $|\mathcal{R}|^{\max}$ possible reactions. This ensures the number of active reactions is fixed and that all subsets of size $|\mathcal{R}|$ are sampled uniformly. Another possibility, which is the one we use in what follows, is to assign to each of the $|\mathcal{R}|^{\max}$ reactions an independent probability p of being active. In this case, each reaction is included independently with probability p , and the total number of active reactions $|\mathcal{R}|$ becomes a random variable that follows a binomial distribution with mean $p \cdot |\mathcal{R}|^{\max}$. Unlike the first method, this procedure does not constrain the total number of reactions to a fixed value, but instead defines it probabilistically.

The proposed model is similar to the Erdős-Rényi (ER) model of random graphs (see Appendix C for a brief description of the ER model), but extended to hypergraphs describing CRNs. Note that the scaling relation between the total number of possible reactions and the total number of chemical species in our model is similar to the ER random graph [in the sense that $|\mathcal{R}|^{\max}$ grows quadratically in $|\mathcal{Z}|^{\max}$; see Eq. (14) and compare with Appendix C].

The same process can be used to construct networks with reactions of an order higher than 2, or elementary reactions of the kind $A + B \rightleftharpoons C + D$, which were discarded before. The scaling of $|\mathcal{R}|^{\max}$ and $|\mathcal{Z}|^{\max}$ will be different and the precise counting of $|\mathcal{R}|^{\max}$ will also be more involved. Finally, we note that by appropriately defining the way reactions are selected, CRNs models analogous to well-studied random graph models (Watts-Strogatz and Barabási-Albert) can also be implemented.^{6,8,44}

IV. STATISTICAL PROPERTIES OF RANDOM CRNS

In this section, we consider ensembles of randomly generated CRNs and present our main findings of their statistical properties. We first study the statistics of the number of species and the number

of reactions they are involved in. These results will then be used to study the statistics of their number of conservation laws, cycles, and emergent cycles, as well as the statistics of the deficiency, and the size of the largest forward reachable set. As for the ER random graph (see Appendix C), we scale the probability of choosing a reaction p as

$$p = \gamma / |\mathcal{Z}|^{\max}. \quad (15)$$

This ensures the CRN we generate is *sparse*, i.e., the average number of reactions in which a given species participates will be fixed by γ , even as we scale the number of reactions and species. We also focus on the *extensive limit*, when $|\mathcal{Z}|^{\max} \rightarrow \infty$. Many of the results in this section are semi-analytic, especially for $N_A > 1$. Closed-form expressions for the case $N_A = 1$, together with additional remarks, are collected in Appendix D.

A. Selected reactions and species

Given a set of reactions $\mathcal{R}$, $\mathcal{Z} \subseteq \mathcal{Z}^{\max}$ is the set of species that are involved in at least one of the reactions in $\mathcal{R}$. From here on, we restrict our focus to $\sigma \in \mathcal{Z}$ when generating a random stoichiometric matrix ∇ . Since the set $\mathcal{R}$ is constructed by independently choosing possible reactions with probability p , the sizes $|\mathcal{R}|$ and $|\mathcal{Z}|$ of both sets will be stochastic quantities. The probability of a particular value of $|\mathcal{R}|$ is given by the binomial distribution,

$$\mathbb{P}(|\mathcal{R}|) = \binom{|\mathcal{R}|^{\max}}{|\mathcal{R}|} p^{|\mathcal{R}|} (1-p)^{|\mathcal{R}|^{\max}-|\mathcal{R}|}, \quad (16)$$

with the average number of reactions given by

$$\langle |\mathcal{R}| \rangle = p |\mathcal{R}|^{\max}. \quad (17)$$

The probability of a particular value of $|\mathcal{Z}|$ is highly non-trivial since the events $\sigma \in \mathcal{Z}$ and $\sigma' \in \mathcal{Z}$ are not independent. However the probability for a particular species σ to not be in $\mathcal{Z}$ is

$$P(\sigma \notin \mathcal{Z}) = (1-p)^{I(\sigma)}, \quad (18)$$

where $I(\sigma)$ is the number of possible reactions involving σ . This number can be written as

$$I(\sigma) = D(\sigma) + A(\sigma), \quad (19)$$

where $D(\sigma)$ is the number of possible dissociations given in Eq. (12), and

$$A(\sigma) = \prod_{a=1}^{N_A} (Q_a - n_a + 1) - 1 \quad (20)$$

is the number of possible associations of the form

$$(n_1, n_2, \dots, n_{N_A}) + (m_1, m_2, \dots, m_{N_A}) \\ \rightleftharpoons (n_1 + m_1, n_2 + m_2, \dots, n_{N_A} + m_{N_A}) \quad (21)$$

for any values of $m_1, m_2, \dots, m_{N_A}$, such that $\min\{m_a\} > 0$ and $0 \leq m_a \leq Q_a - n_a$ for all $1 \leq a \leq N_A$.

B. Number and degree distributions

Given a random set of reactions, $\mathcal{R}$, the degree $d(\sigma)$ of a given species $\sigma = (n_1, \dots, n_{N_A})$ is defined as the number of reactions in $\mathcal{R}$ in which it participates. The probability for the species σ to have degree k is given by the binomial distribution,

$$p_{\sigma}(k) = \binom{I(\sigma)}{k} p^k (1-p)^{I(\sigma)-k}. \quad (22)$$

Note that in the Erdős–Rényi model, each node in a random graph also has a binomial distribution for its degree, but this distribution is the same for all nodes.⁸ This is not the case here since Eq. (22) depends explicitly on the species σ .

We define the degree distribution as the probability that a randomly chosen species has a degree k . The analogous quantity in the context of random graphs offers useful information about the graph.⁸ We first compute a quantity similar to the degree distribution: the number distribution, which is the *number* of species with degree k . For this computation, we assume that species σ having degree k and species σ' having degree k are independent events. This approximation can be expected to hold in the limit of a large number of reactions as there can only be at most two different reactions involving the species σ and σ' [see Eq. (11)]. Thus, the probability that n_k species have degree k is given by

$$P(n_k) = \sum_{\substack{A \subset \mathcal{Z}^{\max} \\ |A|=n_k}} \prod_{\sigma \in A} p_{\sigma}(k) \prod_{\sigma' \notin A} (1-p_{\sigma'}(k)), \quad (23)$$

where the sum runs over all subsets A of $\mathcal{Z}$ with cardinality n_k . We note that the resulting probability density in Eq. (23) is a Poisson–Binomial distribution⁴⁵ with mean

$$\mu(k) = \sum_{\sigma} p_{\sigma}(k) \quad (24)$$

and variance,

$$\text{Var}(k) = \sum_{\sigma} p_{\sigma}(k)(1-p_{\sigma}(k)). \quad (25)$$

It is seen from Fig. 2(a) that Eq. (24) captures the mean very well. In the scaling regime of Eq. (15), we can expect that the probability of a randomly chosen species to have degree k , i.e., the degree distribution denoted p_k , to be equal to (in the limit of large $|\mathcal{Z}|^{\max}$)

$$p_k = \lim_{|\mathcal{Z}|^{\max} \rightarrow \infty} \frac{\mu(k)}{|\mathcal{Z}|^{\max}}. \quad (26)$$

We show in Appendix E that Eq. (26) can be expressed as

$$p_k = \frac{\gamma^k}{k!} \int_0^1 dx_1 \cdots \int_0^1 dx_{N_A} \left(\frac{1}{2} \prod_{k=1}^{N_A} x_k + \prod_{j=1}^{N_A} (1-x_j) \right)^k \\ \times e^{-\gamma \left(\frac{1}{2} \prod_{k=1}^{N_A} x_k + \prod_{j=1}^{N_A} (1-x_j) \right)}, \quad (27)$$

which is valid for any number of atomic types. Figure 2(b) plots the degree distribution obtained from simulation for two different values of γ against Eq. (27) [also see Eq. (D6) for a closed form expression for the case of $N_A = 1$]. We see that Eq. (27) captures the behavior of the degree distribution very well.




FIG. 2. (a) Average number of species with degree k as a function of k for the parameters: $N_A = 2$, $Q_1 = 27$, $Q_2 = 15$, and $p = 0.13$ averaged over 5000 reaction sets compared to the theoretical curve (black line) computed using Eq. (24). The standard error of the mean for each point is indicated by the errorbar, which is not visible in the plot. (b) Probability for a randomly chosen species to have degree k as a function of k for the parameters: $N_A = 1$ and $Q = 500$ averaged over 5000 reaction sets. Two different datasets with $\gamma = 3.1$ and $\gamma = 5.2$ are plotted against the theoretical curves (black line) computed using Eqs. (27) and (D6). For both curves, the standard error of the mean is plotted on the data and is not visible on the plot.

The average number of species participating in at least one reaction (denoted $\langle |\mathcal{Z}| \rangle$) can be computed from the average of the number distribution, $\mu(k)$, as follows:

$$\begin{aligned} \langle |\mathcal{Z}| \rangle &= |\mathcal{Z}|^{\max} - \mu(0), \\ &= \prod_{a=1}^{N_A} (Q_a + 1) - 1 - \sum_{\sigma} (1 - p)^{I(\sigma)}, \end{aligned} \quad (28)$$

where we used Eqs. (10) and (24) [also see Eq. (D7) for a closed form expression for the case of $N_A = 1$]. At very small probabilities, using the binomial approximation $(1 - x)^n \approx 1 - nx$ along with Eq. (28), we find that $\langle |\mathcal{Z}| \rangle$ scales linearly in p with

$$\langle |\mathcal{Z}| \rangle \approx |\mathcal{Z}|^{\max} - \sum_{\sigma} (1 - I(\sigma)p) = p \sum_{\sigma} I(\sigma), \quad (29)$$

whereas in the limit $p \rightarrow 1$, the expected number of species approaches the limiting value $|\mathcal{Z}|$. We compare Eq. (28) to simulations in Fig. 3.

C. Connectivity I: Emergent cycles and conservation laws

We now analyze connectivity in terms of emergent cycles, as defined in Subsection II C. Since our random CRNs may include any possible stoichiometrically allowed reaction, we focus on the probability that randomly selected substrates can interconvert to form randomly chosen products, as p and Q_a are varied. We work in the regime described by Eq. (15) and in the extensive limit.

Numerically, to estimate the probability of generating an emergent cycle, we first generate ensembles of random CRNs for given values of γ and $\{Q_1, Q_2, \dots, Q_{N_A}\}$. We then randomly select a subset of species in $\mathcal{Z}$ and check if chemostating them forms an emergent cycle. Importantly, in a CRN with N_A atoms, at least $N_A + 1$ chemostatted species are typically needed to obtain a non-negligible




FIG. 3. Average number of species vs p averaged over 10000 reaction sets for the parameters: $N_A = 2$, $Q_1 = 27$, $Q_2 = 15$, and $p = 0.13$. The theoretical curve (black line) is calculated using Eq. (28). The standard error for each point is indicated by the errorbar, which is not visible in the plot.

chance of forming an emergent cycle. Since reactions conserve atomic identities, synthesizing a species with N_A atoms generally requires N_A distinct precursors. Using fewer precursors imposes strict linear constraints on atomic balances in the products, which are unlikely to be satisfied by randomly chosen products, greatly reducing the chance of forming an emergent cycle linking them.

We model the outcome as a Bernoulli random variable, n_i , which is one if an emergent cycle is formed and zero otherwise, i.e.,

$$n_i = \begin{cases} 1 & \text{with prob. } \mathbb{P}_e(\gamma, \{Q_1, \dots, Q_A\}) \\ 0 & \text{with prob. } 1 - \mathbb{P}_e(\gamma, \{Q_1, \dots, Q_A\}), \end{cases} \quad (30)$$

and study how the probability $\mathbb{P}_e(\gamma, \{Q_1, \dots, Q_A\})$ changes with γ in the extensive limit. Typically, we observe that $\mathbb{P}_e(\gamma, \{Q_1, \dots, Q_A\})$ is zero below some threshold and nonzero above it, resembling a percolation transition. Below the threshold, a randomly chosen subset of species is unlikely to be connected via an emergent cycle, while above it, the probability becomes finite.

We begin with the simplest case of $N_A = 1$. For fixed values of γ and Q , we generate ensembles of random CRNs. In each CRN, we randomly select and chemostat two species and then check whether an emergent cycle is formed (note that at low values of γ , whenever $\mathcal{Z}$ involves less than two species, the answer is automatically taken to be zero). The probability $\mathbb{P}_e(\gamma, Q)$ is estimated as the mean of the observed n_i .

The result of our numerical simulations is plotted in Fig. 4. We observe the following: first, for all Q , the probability is close to zero until some γ_c , after which it rapidly increases to one. This indicates that the low γ regime is highly disconnected, as any random pair of species is unlikely to be interconvertible. On the other hand, in the high γ regime, almost any species can be synthesized from another. Second, in Fig. 4, the curves for different Q intersect each other. This suggests that in the limit of $Q \rightarrow \infty$, the probability of an emergent cycle approaches the limiting form

$$\lim_{Q \rightarrow \infty} \mathbb{P}_e(\gamma, Q) = \begin{cases} 0 & \text{if } \gamma < \gamma_c \\ 1 & \text{if } \gamma > \gamma_c. \end{cases} \quad (31)$$




FIG. 4. Estimate of the probability of having an emergent cycle when two random species involved in the CRN are chemostatted plotted as a function of the rescaled probability γ . The parameter is $N_A = 1$. The gray dashed line represents the values of γ^* obtained numerically following Eq. (46). The red dashed line represents the accurate threshold γ_c obtained from numerical interpolation of the crossing points in the data. For each value of γ , if $Q = 100, 250$, we generated 25 000 reaction sets, and if $Q = 500$, we generated 20 000 reaction sets and 10 000 reaction sets for the rest for averaging. The sample error of the mean is shown by errorbars, which are not visible on the plot.

This is similar to a discontinuous (first-order) transition at the threshold γ_c . The threshold γ_c can be estimated by extrapolating the crossing point of the curves for consecutive Q shown in Fig. 4. Carrying out this extrapolation, we find that the percolation threshold is $\gamma_c = 3.62 \pm 0.02$ (see Appendix F 1 for full details of the methodology). To further confirm the discontinuity, we analyze the variance of n_i , denoted $\tilde{\chi}(\gamma, Q)$. Since n_i is a Bernoulli variable, its variance is given by

$$\tilde{\chi} = \mathbb{P}_e(1 - \mathbb{P}_e) \xrightarrow{Q \rightarrow \infty} \Theta(\gamma - \gamma_c)\Theta(\gamma_c - \gamma), \quad (32)$$




FIG. 5. Variance of the probability of having an emergent cycle when two random species involved in the CRN are chemostatted plotted as a function of the rescaled probability γ . The parameter is $N_A = 1$. The gray dashed line represents the values of γ^* obtained numerically following Eq. (46). For each value of γ , if $Q = 100, 250$, we generated 25 000 reaction sets, and if $Q = 500$, we generated 20 000 reaction sets and 10 000 reaction sets for the rest for averaging.

where we used Eq. (31) and $\Theta(\cdot)$ denotes the Heaviside function with $\Theta(0) = 1/2$. Figure 5 confirms that the variance away from γ_c decreases with increasing Q , consistent with this expectation.

We now discuss the case of two atoms. For $N_A = 2$, we plot in Fig. 6 the probability of obtaining an emergent cycle when chemostatted three species as a function of the rescaled probability γ , for different values of $Q_1 = Q_2 = Q$. As in the case of the single atom, at low γ , the probability is nearly zero and increases after some critical probability γ_c . However, in contrast to Fig. 4, the curves in Fig. 6 do not intersect. In fact, the data in Fig. 6 indicate that for $N_A = 2$, the probability $\mathbb{P}_e(\gamma, Q_1, Q_2)$ undergoes a *continuous transition*, i.e., the probability is zero below γ_c and increases continuously afterward. Near the threshold, we expect the probability to follow a power law,

$$\lim_{Q \rightarrow \infty} \mathbb{P}_e(\gamma, Q, Q) \propto |\gamma - \gamma_c|^b. \quad (33)$$

Furthermore, in analogy with standard finite-size scaling near continuous transitions, we expect deviations from the limiting behavior to follow a universal form Ref. 46. In particular, curves for different $|\mathcal{Z}^{\max}|$ should collapse onto a master curve when suitably rescaled,

$$\mathbb{P}_e(\gamma, Q, Q)(|\mathcal{Z}^{\max}|)^{\frac{b}{a}} = f((\gamma - \gamma_c)(|\mathcal{Z}^{\max}|)^{\frac{1}{a}}). \quad (34)$$

Using Eq. (34), following the method outlined in Refs. 47 and 48, we extract the following critical exponents: $\gamma_c = 6.74 \pm 0.01$, $a = 2.219 \pm 0.131$, and $b = 0.095 \pm 0.005$. The resulting data collapse is shown in Appendix F 2, along with details of the methodology (see Figs. 21–23). To further support this picture, we analyze the derivative $\partial \mathbb{P}_e(\gamma, Q, Q)/\partial \gamma$. If the probability $\mathbb{P}_e(\gamma, Q, Q)$ follows a power law with an exponent $b < 1$, then the derivative must diverge at γ_c . Figure 7 shows the derivative (smoothed with a Gaussian filter due to numerical noise), which becomes progressively sharper with




FIG. 6. Estimate of the probability of having an emergent cycle when three random species involved in the CRN are chemostatted plotted as a function of the rescaled probability γ . The parameters are $N_A = 2$ and $Q_1 = Q_2 = Q$. The gray dashed line represents the values of γ^* obtained numerically following Eq. (47). The red line represents the accurate threshold γ_c obtained via a scaling analysis. For each value of γ , for $Q = 15$ till $Q = 45$, we generated 32 000 reaction sets for averaging. For all other Q , we generated 16 000 reaction sets for averaging. The sample error of the mean is shown by errorbars, which are not visible on the plot.




FIG. 7. Derivative of the probability in Fig. 6 of having an emergent cycle when three random species involved in the CRN are chemostatted wrt γ plotted as a function of the rescaled probability γ . The parameters are $N_A = 2$ and $Q_1 = Q_2 = Q$. The red line represents the values of γ_c obtained numerically. The gray dashed line represents the values of γ^* obtained numerically following Eq. (47). The data in Fig. 6 were smoothened via a Gaussian filter before taking a numerical derivative.

increasing Q , consistent with this expectation. Let γ_r denote the location of the peak in the derivative for a given Q . From Eqs. (33) and (34), the peak height can be expected to scale as

$$\left. \frac{\partial \mathbb{P}_e(\gamma, Q, Q)}{\partial \gamma} \right|_{\gamma=\gamma_r} \propto (|\mathcal{Z}^{\max}|)^{\frac{1-b}{a}}. \quad (35)$$

As shown in Appendix F 2, Eq. (35) fits the data well, with an exponent 0.419 ± 0.004 , close to the predicted value 0.408 ± 0.024 obtained from the scaling exponents. Furthermore, in the limit $|\mathcal{Z}^{\max}| \rightarrow \infty$, we expect $\gamma_r \rightarrow \gamma_c$. We also analyze how the peak location γ_r approaches γ_c as $|\mathcal{Z}^{\max}|$ increases in Appendix F 2. The data are consistent with the following scaling: $|\gamma_r - \gamma_c| \propto |\mathcal{Z}^{\max}|^{\frac{-1}{a}}$ with a fitted exponent $1/a = 0.499 \pm 0.011$, close to the expected value 0.451 ± 0.003 . Note that the exponents a and b cannot be directly compared to known universality classes in percolation theory, as our order parameter $\mathbb{P}_e(\gamma, \{Q_1, \dots, Q_A\})$ differs from the standard description of percolation in terms of the size of the largest connected component⁴⁶ and does not have a simple geometric interpretation.

While the transitions for $N_A = 1$ and $N_A = 2$ differ as the former is discontinuous while the latter is continuous, in both cases, the system undergoes a qualitative change from a disconnected to a connected regime. We, therefore, refer to both as percolation-like transitions.

We now turn to an analytical perspective based on the conservation laws. Indeed, as Eq. (6) showed, the number of emergent cycles is linked to the number of (broken) conservation laws. Clearly, if there are multiple conservation laws, creating an emergent cycle is difficult and vice versa. We compute the average number of conservation laws as a function of the rescaled probability γ , as shown in Fig. 8. We see that the average number of conservation laws initially increases and reaches a maximum value before rapidly falling to the value N_A . This nonmonotonic behavior is observed for both different values of $Q_1 \dots Q_{N_A}$ and different values of N_A .




FIG. 8. (a) Average number of conservation laws as a function of the probability to select a reaction γ compared to the theoretical result of Eq. (38) (solid black line). The parameters are $N_A = 1$. The gray dashed line indicates the value of γ^* obtained from Eq. (46). (b) Average number of conservation laws as a function of the probability to select a reaction γ compared to the theoretical result of Eq. (38) (solid black line). The parameters are $N_A = 2$ and $Q_1 = Q_2 = Q$. The gray dashed line indicates the value of γ^* obtained from Eq. (47). For both subplots, for each value of Q and p , we generated $\sim 10\,000$ reaction sets for averaging. The errorbars indicate sample error of the mean and are not visible in the plot.

Intuitively, for low values of p , a very low fraction of the possible reactions are selected. With each reaction added, we can consider that two or three new species are added to the network since the probability of a given species being selected twice scales as p^2 for low p [see Eq. (22)]. Thus, in the limit $p \rightarrow 0$, the generated CRN will have more species than reactions and these reactions will be independent (i.e., they will not share any species). Thus, each added reaction will typically add a conservation law. In this regime, chemostating different species thus typically breaks different conservation laws.

For high values of p , almost all possible reactions and, consequently, almost all species, will be included in the network. Thus, it is almost certain that no group of atoms will be conserved by all the reactions, except the individual atoms themselves. Therefore, the number of independent conservation laws must tend to N_A (see Appendix G 1). Consequently, chemostating $N_A + 1$ species will result in an emergent cycle with high probability.

In the following, we show how the previous arguments can be used to analytically estimate the percolation threshold.

1. Estimating the number of conservation laws

A simple bound on the number of average conservation laws is obtained from the rank of the stoichiometric matrix. Indeed, using Eq. (2) we note that the rank of the stoichiometric matrix is always upper bounded by the minimum of the number of selected species and reactions, i.e.,

$$\langle \text{rk}(\nabla) \rangle \leq \min[\langle |\mathcal{R}| \rangle, \langle |\mathcal{Z}| \rangle]. \quad (36)$$

This implies that the average number of conservation laws satisfies $\langle |\mathcal{L}| \rangle \geq \langle |\mathcal{Z}| \rangle - \min[\langle |\mathcal{R}| \rangle, \langle |\mathcal{Z}| \rangle]$ [see Eq. (2)]. This bound can be broken down further: on the one hand, at small p , while both $\langle |\mathcal{R}| \rangle$ and $\langle |\mathcal{Z}| \rangle$ scale linearly with p [see Eqs. (29) and (17)], since $\sum_{\sigma} I(\sigma) > |\mathcal{R}|^{\max}$, the rank must be upper bounded by the number

of reactions. On the other hand, at large p , $\langle |\mathcal{Z}| \rangle$ saturates (see Fig. 3) while the average number of reactions keeps increasing, implying that the rank must then be upper bounded by the number of species. Thus, we obtain a lower bound on the number of conservation laws,

$$\langle |\mathcal{L}| \rangle \geq \begin{cases} \langle |\mathcal{Z}| \rangle - \langle |\mathcal{R}| \rangle & p \leq p_1 \\ 0 & p > p_1, \end{cases} \quad (37)$$

where p_1 is a crossover probability. Actually, the previous bound can be tightened by just noting that for large p , the number of conservation laws is always larger than the number of atomic types N_A . Therefore,

$$\langle |\mathcal{L}| \rangle \geq \begin{cases} \langle |\mathcal{Z}| \rangle - \langle |\mathcal{R}| \rangle & p \leq p^* \\ N_A & p > p^*. \end{cases} \quad (38)$$

The crossover probability p^* is obtained by equating the different cases,

$$\langle |\mathcal{Z}| \rangle(p^*) - \langle |\mathcal{R}| \rangle(p^*) = N_A. \quad (39)$$

Figure 8 shows that the above-mentioned expression captures the nonmonotonicity in the number of conservation laws very well.

To analytically obtain p^* , we work in the extensive limit using Eq. (15). In first place, we derive a universal form for the behavior of the rescaled average number of conservation laws, defined as

$$\langle |\ell| \rangle = \langle |\mathcal{L}| \rangle / |\mathcal{Z}|^{\max}. \quad (40)$$

From Eq. (38), we have the following expression for the scaled number of conservation laws:

$$\langle |\ell| \rangle \geq \begin{cases} \langle |\mathcal{Z}| \rangle - \langle |\mathcal{R}| \rangle & \gamma \leq \gamma^* \\ N_A / |\mathcal{Z}|^{\max} & \gamma > \gamma^*, \end{cases} \quad (41)$$

where we have used lowercase letters to denote rescaled quantities. We now evaluate each quantity in Eq. (41). Using Eqs. (14) and (17), we can write

$$\langle |\mathcal{R}| \rangle = \frac{|\mathcal{R}|^{\max}}{(|\mathcal{Z}|^{\max})^2} \gamma \xrightarrow{|\mathcal{Z}|^{\max} \rightarrow \infty} \frac{\gamma}{2^{N_A+1}}. \quad (42)$$

Similarly, using Eq. (28), the rescaled average number of species is given by

$$\langle |\mathcal{Z}| \rangle = 1 - \frac{\mu(0)}{|\mathcal{Z}|^{\max}}. \quad (43)$$

By using Eq. (27) for the special case of $k = 0$, we have

$$\begin{aligned} \frac{\mu(0)}{|\mathcal{Z}|^{\max}} &= \int_0^1 dx_1 \dots \int_0^1 dx_{N_A} e^{-\gamma \left(\frac{1}{2} \prod_{k=1}^{N_A} x_k + \prod_{j=1}^{N_A} (1-x_j) \right)} \\ &+ \mathcal{O} \left(\frac{1}{Q_A} \right). \end{aligned} \quad (44)$$

In the limit of all the $Q_i \rightarrow \infty$, substituting Eqs. (42)–(44) in Eq. (41), we get a bound on $\langle |\ell| \rangle$ of the form

$$\langle |\ell| \rangle \geq \begin{cases} 1 - \frac{\mu(0)}{|\mathcal{Z}|^{\max}} - \frac{\gamma}{2^{N_A+1}} & \gamma \leq \gamma^* \\ 0 & \gamma > \gamma^*. \end{cases} \quad (45)$$




FIG. 9. Scaled average number of conservation laws $\langle |\ell| \rangle$ for $N_A = 1$ plotted as a function of γ and compared to the theoretical result in Eq. (46) (black line). The gray dashed line indicates the location of γ^* in Eq. (46), obtained numerically. For each data point, we generated 10 000 reaction sets for averaging. The sample error of the mean is shown by the errorbars, which are not visible in the plot.

Note that the critical value γ^* and the form of the universal curve in Eq. (45) depend only on the number of atoms. For the case of one atom, Eqs. (45) and (44) reduce to

$$\langle |\ell| \rangle \geq \begin{cases} 1 - \frac{2}{\gamma} e^{-\gamma} (e^{\gamma/2} - 1) - \frac{\gamma}{4} & \gamma \leq \gamma^* \\ 0 & \gamma > \gamma^*, \end{cases} \quad (46)$$

where the value $\gamma^* \approx 3.72$ is obtained numerically by equating the two possible cases. Clearly, γ^* is an estimate of the percolation threshold and we see that the value of $\gamma^* = 3.72$ differs from the actual percolation threshold $\gamma_c = 3.62$ by only about 3% (see the red line and the gray line in Fig. 4). We compare Eq. (46) with simulations shown in Fig. 9. We see that a very good agreement is obtained for large Q and $\gamma < \gamma^*$. For the case of two atoms, Eqs. (45) and (44) reduce to




FIG. 10. Scaled average number of conservation laws $\langle |\ell| \rangle$ for $N_A = 2$ plotted as a function of γ and compared to the theoretical result in Eq. (47) (black line). The gray dashed line indicates the location of γ^* in Eq. (47), obtained numerically. For each data point, we generated 10 000 reaction sets for averaging. The sample error of the mean is shown by the errorbars, which are not visible in the plot.




FIG. 11. Comparing the absolute difference of the simulated $\langle |\ell| \rangle(Q)$ at different values of Q and the theoretical prediction of the same in Eq. (46) at the values, (a) $\gamma = 3.5$ and (b) $\gamma = 3.8$. For both graphs, $N_A = 1$. For each data point, we generated 10 000 reaction sets for averaging. The sample error of the mean is shown by the errorbars.

$$\langle |\ell| \rangle \geq \begin{cases} 1 - \frac{\gamma}{8} - \frac{e^{-\gamma/3}}{\gamma} \left[\frac{4\text{Ei}(\gamma/3)}{3} \right] & \gamma \leq \gamma^* \\ -\frac{2\text{Ei}(-2\gamma/3)}{3} - \frac{2\text{Ei}(-\gamma/6)}{3} & \gamma > \gamma^* \\ 0 & \gamma > \gamma^* \end{cases} \quad (47)$$

where $\text{Ei}(\cdot)$ refers to the exponential integral function and the critical value $\gamma^* \approx 7.07$ is obtained numerically. Once again, the difference between the actual threshold for $N_A = 2$, which is $\gamma_c = 6.74$ (see redline in Fig. 6) and $\gamma^* = 7.07$ (gray line in Fig. 6) is about 5%. We compare Eq. (47) with simulations in Fig. 10 for $Q_1 = Q_2 = Q$. We see that as Q increases, the curves approach Eq. (47). However, since the Q values we use are quite small, the corrections to Eq. (47) of order $\mathcal{O}(1/Q)$ ensure the convergence to the regime indicated by the universal curve is quite slow.


It is tempting to believe that the previous results, along with Figs. 9 and 10, show that there is a transition in the number of conservation laws at $\gamma = \gamma^*$. However, this is inaccurate. We show this in the case of $N_A = 1$ shown in Fig. 11. By numerically estimating the value of $\langle |\ell| \rangle(Q)$ for different values of Q , we see that for values such as $\gamma = 3.5$, we have convergence to the predicted value, whereas for values such as $\gamma = 3.8$, we see convergence to a nonzero value, as depicted in Fig. 11. The second argument against the presence of a phase transition in the number of scaled conservation laws comes from studying the analogous quantity for ER graphs, which is given by the scaled number of connected components. As shown in Appendix G 2, the scaled number of connected components is a continuous function of γ , even at $\gamma_c = 1.0$, the critical value for the giant component transition. In fact, repeating the analysis of this section for ER graphs, one finds that the threshold predicted by this approach would be $\gamma^* = 1.49$. Thus, we conclude that, while there is an underlying percolation transition that is captured by this ansatz, it does not occur precisely at γ^* .

2. Finding the conserved moieties



While the arguments about the average rank of the stoichiometric matrix in Sec. IV C 1 allow us to predict the average number

of conservation laws, it is not clear which specific moieties are conserved in the low p regime. In this subsection, we identify the conserved moieties by introducing *root* species.

For convenience, we orient all reactions to be like an association, i.e., like Eq. (21). Note that this reorientation of reactions is always possible as we are dealing with reversible reactions. Denoting the product species of reaction ρ by $\sigma_{p(\rho)}$ and the reactant species as $\sigma_{r(\rho)}, \sigma_{r'(\rho)}$, we rewrite all reactions ρ of the form of Eq. (11) as



Consider a random CRN generated with all reactions written in the form of Eq. (48) with the corresponding stoichiometric matrix ∇ . Define the *root* species (labeled $\hat{\sigma}$) that are reactants in every reaction that they are in, i.e., any reaction involving roots would be of the form



where we denoted arbitrary non-root species with the labels σ . Operationally, given a stoichiometric matrix ∇ with reactions oriented as in Eq. (48), the root species have rows with only nonpositive entries. We direct the readers to Appendix G 3 6 for an example of a CRN and the corresponding root species.

Appendix G 3 shows that starting from the root species, we can construct every non-root species in the CRN. Thus, they play the role of effective atoms and we can expect that moieties corresponding to the roots should be conserved. However, not all the root species are independent. Indeed, if we identify the root species as the chemostatted species, it is possible that some root species can be constructed from others, i.e., that there are emergent cycles. Putting both these ideas together, we show (see Appendix G 3) that

$$|\mathcal{Z}^r| = |\mathcal{L}| + |\epsilon|_r, \quad (51)$$

where we denote the number of root species in the CRN as $|\mathcal{Z}^r|$ and $|\epsilon|_r$ is the number of emergent cycles when the root species are chosen as the chemostatted species. Equation (51) has remarkable implications. First, since it holds for any realization of a random CRN, the averages must satisfy

$$\langle |\mathcal{L}| \rangle \leq \langle |\mathcal{Z}^r| \rangle. \quad (52)$$

We thus obtain an upper bound on the average number of conservation laws [also see Eq. (D11) for a closed form expression of Eq. (52) when $N_A = 1$]. Second, from Eq. (51), if we can choose a subset of the root species that cannot be constructed from one another, then the number of these *independent* roots must be equal to the number of conservation laws. Furthermore, in Appendix G 3, we show how to identify the independent roots and provide a constructive procedure to build a basis of the conservation laws, given the set of independent roots, thus completing the identification of the independent roots with conserved moieties. This tells us the origin of the multiple conservation laws at low p . In this regime, it is more probable that many species with low atom numbers are involved in reactions that construct species with higher atom numbers, as

opposed to being involved in reactions that dissociate them. They would thus be roots and create conservation laws. However, as we increase p , either the newly added reactions are dissociations of the root species or they start creating emergent cycles between the roots, reducing the number of conservation laws.

D. Connectivity II: Forward reachability

We now discuss the notion of forward reachability, as introduced in Sec. II C. We reiterate that the notion of reachability is similar to the idea of an out-component (i.e., the set of nodes which can be reached starting from a given node) in directed networks, where every edge has a directionality associated with it. Analogous to studies of directed networks,^{49,50} we investigate reachability by studying the behavior of the largest (rescaled) forward reachable set, denoted $|f|$ as we change γ , i.e.,

$$|f|(\gamma, Q) = \frac{\max_{\sigma} |\mathcal{F}(\sigma)|}{|\mathcal{Z}|^{\max}}. \quad (53)$$

We first consider the case of $N_A = 1$. For each value of γ and Q , we first generate an ensemble of random CRNs. Then, for every realization, we compute the reachable sets of each species and record the largest reachable set (see Appendix H for a description of how to find the reachable set of a species). The resulting average reachability is plotted in Fig. 12 with the inset showing the same quantity $\langle |f| \rangle$ but over a larger range of γ . As we increase the number of species in the system, from Fig. 12, we see that a transition-like scenario emerges: there is a threshold γ below which $|f|$ is almost zero and above which $|f|$ increases to a value $\mathcal{O}(1)$.

To understand the nature of this transition, we examine the distribution of $|f|$ across realizations. In Fig. 13, we plot histograms of $|f|$ for $Q = 1200$ at four different values of γ . Initially, for low γ , the histogram is single-peaked with a mean value close to zero. As we increase γ , a second peak emerges at a higher value of $|f|$. As γ increases, the second peak begins to dominate and eventually the second peak takes over and the first peak disappears.




FIG. 12. Average size of the rescaled largest reachable set, $\langle |f| \rangle$, as a function of γ for the parameters $N_A = 1$ and Q . For each value of Q and γ , we generated 39 000 reaction sets for averaging. The errorbars indicate sample error and are not visible in the plot. The red line indicates the estimated percolation threshold, γ_c . The inset shows the variation of $\langle |f| \rangle$ over a larger range of γ . The data for the inset were averaged over 2800 reaction sets.




FIG. 13. Histogram of $|f|$ for four fixed values of γ and with the parameters $N_A = 1$ and $Q = 1000$. For each value of Q and γ , we generated 39 000 samples and the data were binned in 201 bins between zero and one and resulting normalized histogram is plotted.

This bimodal structure indicates that independent CRN realizations fluctuate between a disconnected regime (low reachability) and a connected regime (high reachability) and is a hallmark of a discontinuous (first-order) transition.⁵¹

We estimate the critical point γ_c by finding for each Q , the probability at which both peaks of the histogram are equally weighted (see Appendix H for the procedure), denoted $\gamma_{1/2}(Q)$, and then extrapolating to $Q \rightarrow \infty$. The result of this procedure is shown in Fig. 14 with the result $\gamma_c = 2.92 \pm 0.01$. The extrapolation is a linear fit to $1/Q$ ($r^2 \approx 0.94$). To further support this picture, we track the locations of the two histogram peaks, denoted by $|f|_{\pm}(Q)$, and confirm that they remain separated as Q increases. The resulting plot is shown in Appendix H.

In the limit of $Q \rightarrow \infty$, a consequence of the shift of the probability weight is a jump in the average reachability as γ is varied. However, no discontinuity can be seen since the Q values we use for simulations are quite small. To test this, we examine the derivative of the average reachability with γ , that is, $\partial \langle |f| \rangle / \partial \gamma$. A jump in the




FIG. 14. Equal weight probability $\gamma_{1/2}(Q)$ of the order parameter histogram (as in Fig. 13) plotted against $1/Q$. The parameters are $N_A = 1$. The gray dotted line represent a linear fit with intercept 2.92.




FIG. 15. Derivative of the average reachability in Fig. 12 plotted as a function of the rescaled probability γ . The parameters are $N_A = 1$. The red line indicates the value of γ_c obtained numerically. The data in Fig. 12 were smoothed via a Gaussian filter before taking a numerical derivative.

average reachability would correspond to a delta-function singularity in its derivative. Figure 15 shows the derivative (smoothed with a Gaussian filter due to numerical noise), which becomes progressively sharper with increasing Q , consistent with this expectation. The location of the peak of the derivative, denoted $\gamma_r(Q)$, also extrapolates to 2.90 as $Q \rightarrow \infty$ consistent with the previous estimate (see Appendix H). Furthermore, we note that the maximum of the derivative scales as a power law in Q with an exponent 0.35, as shown in Appendix H. We also analyze the fluctuations of the order parameter, defined as $\tilde{\chi}(\gamma, Q) = \langle |f|^2 \rangle - \langle |f| \rangle^2$. As shown in Appendix H, the fluctuations of the reachability for different Q show a peak whose height scales with Q and the location of which moves toward γ_c .

Based on our numerical results, such as a bimodal distribution, we have concluded that the reachability transition is first-order. However, we also observe some deviations from the known theory for first-order transitions. In particular, the maximum of the derivative, $\partial \langle |f| \rangle / \partial \gamma$ scales as a power law in Q with an exponent 0.35 rather than linearly.⁵¹ Similarly, the fluctuations of $|f|$ show a growing peak that shifts toward γ_c as Q increases (see Appendix H). We believe that these behaviors may indicate strong finite-size effects.

Finally, we briefly consider the case of $N_A = 2$ and investigate the reachability from a single species for the case of $Q_1 = Q_2 = Q$. We again plot the histogram of the order parameter for $Q_1 = Q_2 = 50$. The resulting histograms are shown in Fig. 16. Once again, we see the emergence of a double-peaked histogram similar to the case of $N_A = 1$ (see Fig. 13). Thus, we conclude that the reachability transition for $N_A = 2$ is discontinuous as in the case of $N_A = 1$.

E. Deficiency

We now turn to the deficiency, as defined in Sec. II B 2. First, for any CRN, constructed with reactions of the form Eq. (48), we have that $\delta = |C|$, i.e., the deficiency is identical to the number of (stoichiometric) cycles. To see this, we note that the set of complexes can be written as the union of the sets $\mathcal{Z} = \{\sigma\}$ and $\mathcal{Z}_2 = \{(\sigma, \sigma')\}$. The reactions ρ then connect complexes in the set $\mathcal{Z}_2$ with complexes in




FIG. 16. Histogram of $|f|$ for four fixed values of γ and with the parameters $N_A = 2$ and $Q_1 = Q_2 = 50$. For each value of Q and γ , we generated 10 000 samples and the data were binned in 81 bins between zero and one and resulting normalized histogram is plotted.

the set $\mathcal{Z}$ [see Eq. (48)]. A given complex in $\mathcal{Z}$ can be connected to multiple $\mathcal{Z}_2$ complexes. However, there are no connections between complexes within $\mathcal{Z}$ and $\mathcal{Z}_2$. Thus, we get a graph with a bipartite structure with the strong restriction that no two complexes of $\mathcal{Z}$ are connected to the same complex in $\mathcal{Z}_2$, removing the possibility of graph cycles. We can thus focus on the statistical properties of the number of stoichiometric cycles. Using the rank-nullity theorem Eq. (2) along with Eq. (38), we lower bound the average number of cycles as

$$\langle |C| \rangle \geq \begin{cases} 0 & p \leq p^* \\ \langle |\mathcal{R}| \rangle + N_A - \langle |\mathcal{Z}| \rangle & p > p^*. \end{cases} \quad (54)$$

Working in the extensive limit and with the probability expressed in the form of Eq. (15), we find the universal form

$$\langle |C| \rangle \geq \begin{cases} 0 & \gamma \leq \gamma^*(N_A) \\ \frac{\mu(0)}{|\mathcal{Z}|^{\max}} + \frac{\gamma}{2^{N_A+1}} - 1 & \gamma > \gamma^*(N_A) \end{cases} \quad (55)$$

analogous to Eq. (45). Figure 17 shows that Eq. (55) matches well with simulations [also see Eq. (D15) for a different bound in the case of $N_A = 1$].

From Eq. (55), we note that random CRNs tend to have zero deficiency when $\gamma < \gamma^*$. However, for $\gamma > \gamma^*$, the average deficiency increases with the probability. This observation has interesting implications for random CRN dynamics: CRNs with zero deficiency always reach stable steady states,³⁸ whereas those with higher deficiency may exhibit complex dynamics such as oscillations and multistability. Moreover, when Q_a values are large, the crossover probability $p^* = \gamma^* / |\mathcal{Z}|^{\max}$, below which the average deficiency is zero, scales inversely to the maximum possible number of species. As a result, deficiency-zero CRNs are rare, especially in the limit of large Q_a .




FIG. 17. Scaled average number of stoichiometric cycles $\langle |c| \rangle$ for $N_A = 1$ plotted as a function of γ and compared to the theoretical result in Eq. (55) (black line). The gray dashed line indicates the location of γ^* in Eq. (46), obtained numerically. For each data point, we generated 10 000 reaction sets for averaging. The sample error of the mean is shown by the errorbars, which are not visible in the plot.

V. DISCUSSION

A. Connecting to the giant component transition

We now connect the transitions in the emergent cycles and the forward reachability discussed in Secs. IV C and IV D with the percolation transition of Erdős–Rényi random graphs (see Appendix C).

We begin by noting similarities in the qualitative behavior of these transitions and those in ER random graphs. As p increases in an ER graph, small disconnected components merge into a single giant component. Since conservation laws in CRNs play a role analogous to components (see Sec. II C), we observe similar behavior in Subsection IV C 1: at small p , many species act as independent roots, yielding numerous conservation laws, while at large p , the number of conservation laws drops to N_A , akin to the formation of a giant component. A second similarity is between paths in ER graphs and emergent cycles in random CRNs. In ER graphs, two randomly chosen nodes are unlikely to be connected at small p , but are likely to lie in the same component at large p . Similarly, in CRNs, given an arbitrary subset of species, the probability of an emergent cycle is negligible at small p , but becomes finite at large p , indicating that the subset of species are connected in the sense of emergent cycles. Finally, a third similarity involves forward reachability in CRNs and out-components in directed graphs. In the latter, a giant out-component typically emerges beyond a critical probability, as observed in real-world networks such as the Internet. Since forward reachability plays an analogous role to out-components (see Sec. II C), our numerical results in Subsection IV D similarly reveal two regimes: at small p , many small reachable sets exist, while above a threshold, a giant reachable set appears.

However, there are important conceptual and quantitative differences between percolation in CRNs and in graphs. Conceptually, each node in a graph lies in exactly one component, so component size (number of nodes in the component) is unambiguous, and the size of the largest component serves as the standard order parameter used to study percolation.^{8,44,46} In contrast, species in CRNs can be part of multiple conservation laws. Furthermore, since conservation laws are a basis of the left nullspace of ∇ , any notion of “size”

for a conservation law must either use a special basis or be invariant under basis change. In addition, while paths in graphs connect pairs of nodes and carry a clear geometric interpretation, emergent cycles involve entire subsets of species and correspond to effective reactions among them. This prevents a geometric interpretation of emergent cycles analogous to paths. Quantitatively, percolation in ER graphs is a *continuous* phase transition (see Appendix C), whereas the transitions associated with reachability and emergent cycles in CRNs are both *discontinuous* for $N_A = 1$ (see Secs. IV C and IV D) and occur at distinct critical probabilities. This is especially striking in the case of reachability: for unimolecular CRNs (which correspond to undirected graphs), the reachability order parameter reduces exactly to that of ER graphs and the transition is continuous. That it becomes discontinuous in the general multimolecular setting highlights the impact of the underlying hypergraph structure. A further distinction is the role of the number of atomic types N_A , which can make striking differences to the transition and has no analog in standard graphs. For instance, as shown in Secs. IV C and IV D, the emergent cycle transition becomes continuous when $N_A = 2$, while the reachability transition remains discontinuous.

B. Relation with other CRN models

The complete set of reactions we defined in Sec. III has been used in multiple contexts in the CRN literature. The key distinction in our approach is the focus on physical consistency, through the presence of conservation laws and the study of connectivity based on hypergraph methods. We first discuss prior work that uses the same set of reactions before focusing on different algorithms for generating random CRNs.

Models of polymerization have used the same network of reactions in the case of one atom. In particular, the full set of reactions corresponds to the Smoluchowski aggregation model.^{52–54} Furthermore, Riehl *et al.*⁵⁵ have used the same one-atom CRN to study the properties of optimal balanced pathways, which are special vectors in the space of emergent cycles and have shown that this simple model can already reproduce generic properties of the metabolic network of *E. coli*. We note that our formulation is more general as we also consider the case of multiple atoms and focus on other structural quantities such as conservation laws.

In the case of $N_A > 1$, the full reaction network has also been explored under the name “cluster chemical reaction networks” (CCRN) in Ref. 56, where it was used to study the occurrence and geometric properties of autocatalytic cores. Interestingly, in their conclusion, the authors suggest future work on randomly generated CCRNs in an Erdős–Rényi framework to bound the existence and persistence of autocatalytic cycles as a function of the probability. This is a direction that we aim to address using our framework, which provides tools to analyze such questions. For $N_A = 2$, studies by Kauffman, Farmer, Bagley, Steel, and others^{57–60} have used similar models to investigate the formation of autocatalytic sets. However, two key differences must be noted here. First, many of these models focus on polymers where the precise order of monomers is important, a distinction that is not made in our model. Second, they assume that every reaction requires an enzyme, assigned probabilistically. Consequently, their arguments center on the emergence of a giant autocatalytic set as the probability of *catalysis* varies. In contrast, our approach emphasizes structural properties of the reaction

network itself, an aspect largely unexplored in these studies. Finally, we note the work of Himeoka *et al.*,⁶¹ who analyzed the response of nonequilibrium steady states in CRNs to perturbations. Their study focuses on special subnetworks that correspond to $N_A = 2$ in our terminology. However, their approach does not include a broader discussion of the network's structure, which is a major aspect of our work.

We now compare our method of generating random CRNs with other algorithms used in the literature. Anderson *et al.*^{23,24} construct random CRNs by generating random graphs in the space of complexes (see Sec. II B 2) and then study the deficiency of the associated CRN, but without taking conservation laws into account. Garcia-Chung *et al.*⁶² construct Erdős–Rényi-like CRNs by randomly choosing hyperedges between nonintersecting subsets of species. However, they do not include any stoichiometric coefficients in their analysis and thus do not explicitly conserve mass. Fischer *et al.*⁶³ construct random open CRNs by first generating S graphs randomly with different prescriptions, i.e., such as the Watts–Strogatz model or the Barabási–Albert model. They then rewire some edges in the S graph to also obtain multimolecular reactions. We note that none of the studies above have focused on constraining their reaction networks to conserve mass.

Both studies by Bigan *et al.*²⁵ and Nicolau *et al.*²⁶ adopt methods close to ours by enforcing conservation laws. Bigan *et al.*²⁵ construct all possible elementary reactions between species and select a subset randomly while ensuring mass conservation. Nicolau *et al.*²⁶ assign random atomic compositions to species and generate reactions by randomly choosing reactants, products, and stoichiometric coefficients. However, unlike both approaches, we do not rely on random selection or external algorithms to enforce conservation. Instead, conservation laws are directly built into the reaction-generation process, ensuring every reaction is chemically meaningful. We note that the percolation transition for a random CRN model was also studied by Nicolau *et al.* They projected the resulting CRN topology onto an S-graph for analysis (see Sec. II B 1) and showed that the critical value of the parameter $\langle |r| \rangle$ was ~ 0.15 . In our analysis, the corresponding value of the same parameter is ~ 0.93 (for $N_A = 1$ and considering emergent cycles). This difference can be attributed to the fact that adding a reaction to the CRN adds a fully connected clique of reactant nodes and product nodes to the S graph. Furthermore, it has been shown in the literature that projections of the bipartite graph may show markedly different connectivity behavior compared to the bipartite graph.¹⁷ It should be noted that both Nicolau *et al.*²⁶ and Bigan *et al.*²⁵ focus primarily on studying the dynamics of random CRNs, which is a direction we do not pursue in this work. The dynamics of randomly generated catalytic CRNs, including trimolecular or higher-order reactions, have been used to model metabolism and growth in protocells.^{64,65} These studies focus on higher-order reaction kinetics, which lie outside the scope of our framework.

VI. CONCLUSIONS AND PERSPECTIVES

In this work, we have proposed an algorithm for the generation of random CRNs that are physically consistent. We first construct a “universe” of possible reactions that respect mass balance as outlined in Sec. III and then sample reactions randomly from it. This procedure allows us to statistically study the topology of CRNs under mass

conservation constraints. We focused on two physically motivated notions of connectivity in CRNs: emergent cycles, which describe steady-state synthesis pathways in open systems, and forward reachability, which describes the set of species that can be generated from a single substrate in a closed system. Both notions exhibit percolation-like transitions as the probability of including a reaction increases.

We analyzed the probability that chemostatting an arbitrary subset of species creates an emergent cycle. Interestingly, when $N_A = 1$, the probability of an emergent cycle increased discontinuously across the transition while for the case of $N_A = 2$, the corresponding increase was continuous. We estimated the critical probability for the transitions in the emergent cycles using the number of conservation laws. The average number of conservation laws showed marked nonmonotonic behavior at low p before saturating to a constant value (the number of atoms) at high p . This directly implies that at low p , chemostatting species only breaks conservation laws, without creating effective interconversions, whereas at high p , chemostatting an arbitrary set of species is enough to create emergent cycles.

For reachability, we showed via numerical simulations that there is a critical probability below which (resp. above which) the largest reachable set involves a vanishing fraction of species (resp. an extensive fraction of species). Numerically, we further established that the largest reachable set undergoes a discontinuous transition for both $N_A = 1$ and $N_A = 2$.

Crucially, the percolation transitions for the emergent cycles and the forward reachability occur at different critical probabilities, that is, 3.62 vs 2.92 when $N_A = 1$. This stresses that both of these quantities capture different aspects of connectivity in CRNs. This is a clear departure from undirected graphs, where both notions are equivalent, which again demonstrates the problems with using graph-based representations of CRNs to study topological properties.

These transitions mark shifts in the capacity of the reaction network to take a set of substrates and create a desired product. Below the threshold, most species cannot be converted into one another. Above the threshold, the accessible space of products is greatly enhanced and products can be created, either via emergent cycles or via reachability cascades.

Our work can be extended in various directions. Our analysis of the transitions was limited to showing that percolation-like transitions exist in the emergent cycles and forward reachability and to estimating the respective critical probabilities. A classification of these transitions into universality classes is an important future direction. For the emergent cycles, a geometric perspective allowing us to define notions of (generalized) components and studying their properties is necessary. For reachability, we note that our quantitative analysis of the derivative of the order parameter and fluctuations revealed some discrepancies: the peak in the derivative grows as a power law in Q , and the susceptibility-like fluctuations also increase with Q . Although we believe that these could be strong finite-size effects, it is important to note that similar phenomena have been observed in explosive percolation models,⁶⁶ where transitions appear discontinuous at small system sizes but are actually continuous transitions that exhibit unusual finite-size properties. Thus, an important future step is to develop algorithms to calculate reachable sets in an efficient fashion, which would allow simulations

at much larger Q , as well as analytical theories to describe the transition.

A key point in the construction of our set of reactions was to focus only on the composition of the chemical species and not their structure. Consequently, unimolecular reactions such as isomerization reactions cannot be considered in this setup because they involve only structural changes. A simple solution to include unimolecular reactions is to add a new label for each species, representing possible conformations, and then add unimolecular reactions between the conformations.⁵⁶ Alternatively, more nuanced studies could start with string-based representations^{67,68} of molecules. In this context, we also mention the field of artificial chemistry, where the focus has been on studying abstract systems that are “chemistry-like.”⁶⁹ Particularly, we note the work of Benkő *et al.*^{70,71} who developed a framework to generate CRNs based on graph-based representations of molecules, with reactions formulated as procedures that rewrite molecular graphs. This approach is particularly interesting as a starting point, as the graph-rewriting procedures inherently respect conservation laws. We direct the reader to Refs. 72–74 for further details and extensions of these ideas.

The concepts and framework introduced here provide a foundation for the analysis of real complex networks, such as metabolic networks. For example, Ref. 55 showed that the typical behavior of a class of elementary flux modes (non-decomposable emergent cycles^{75,76}) in the central carbon metabolism of *E. coli* are analytically captured by a simple model equivalent to our $N_A = 1$ case. Similarly, Ref. 43 investigates the robustness of the reachable sets in metabolic networks under random deletions of reactions. Extending similar analyses to the space of emergent cycles in metabolic networks represents an important direction for future work. Finally, the power of using randomness to understand the typical dynamics of complex systems is exemplified by the application of ideas such as random matrix theory in understanding ecosystems.^{77,78} We believe that adding dynamics to our current formalism along the lines of work done in Refs. 25 and 26 and studying the response of steady states of random CRNs^{61,79} will be crucial steps for the future.

ACKNOWLEDGMENTS

S.G.M.S. thanks Nalina Vadakkayil and Benedikt Remlein for their help with the scaling analysis and Danilo Forastiere for pointing out a mistake in a previous draft. S.G.M.S. was supported by the Luxembourg National Research Fund (FNR), via the research funding schemes PRIDE (Grant No. 19/14063202/ACTIVE). M.E. and N.F. were supported by CORE project ChemComplex (Grant No. C21/MS/16356329) and by Project No. INTER/FNRS/20/15074473 funded by F.R.S.-FNRS (Belgium) and FNR. Simulations for the project were carried out using the HPC facilities of the University of Luxembourg⁸⁰ (see <https://hpc.uni.lu/>).

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions


Shesha Gopal Marehalli Srinivas: Conceptualization (equal); Data curation (equal); Formal analysis (equal); Investigation (equal); Methodology (equal); Project administration (equal); Resources (equal); Software (equal); Supervision (equal); Validation (equal); Visualization (equal); Writing – original draft (equal); Writing – review & editing (equal). **Massimiliano Esposito:** Conceptualization (equal); Data curation (equal); Formal analysis (equal); Funding acquisition (equal); Investigation (equal); Methodology (equal); Project administration (equal); Resources (equal); Software (equal); Supervision (equal); Validation (equal); Visualization (equal); Writing – original draft (equal); Writing – review & editing (equal). **Nahuel Freitas:** Conceptualization (equal); Data curation (equal); Formal analysis (equal); Funding acquisition (equal); Investigation (equal); Methodology (equal); Project administration (equal); Resources (equal); Software (equal); Supervision (equal); Validation (equal); Visualization (equal); Writing – original draft (equal); Writing – review & editing (equal).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.


APPENDIX A: ABSENCE OF LINK BETWEEN EMERGENT CYCLES AND REACHABILITY

We first note that, given a choice of chemostatted species, it is possible that there exists an emergent cycle linking them, but that no reachability relation exists between them. An example of this is provided by the Michaelis–Menten CRN,



If the substrate species S and the product species P are chemostatted, the emergent cycle $c_e = (1, 1)^T$ is created. However, the forward reachable sets of both S and P are just the species themselves (without the enzyme species E and ES , no reaction can occur).

Conversely, assuming that we choose two species σ and σ' such that $\sigma \in \mathcal{F}(\sigma')$ and vice versa, if $Y = \{\sigma, \sigma'\}$, an emergent cycle need not be created. Consider the following CRN:



Here, the forward reachable set of S is $\mathcal{F}(S) = \{S, P, A, B\}$ and is the same as the forward reachable set of P . However, if we chemostat S and P , we *do not* create an emergent cycle.

APPENDIX B: EXAMPLE

Consider the simple case of $N_A = 2$ atomic types and $Q_1 = Q_2 = 2$. The possible molecular species is given by the set,

$$\mathcal{Z}^{\max} = \{(0,1) (1,0) (0,2) (2,0) (1,1) (1,2) (2,1) (2,2)\}. \quad (\text{B1})$$

From (10), we indeed find that $|\mathcal{Z}|^{\max} = 3 \times 3 - 1 = 8$. The possible reactions are

$$\begin{aligned} (2,0) &\rightleftharpoons (1,0) + (1,0), \\ (0,2) &\rightleftharpoons (0,1) + (0,1), \\ (1,1) &\rightleftharpoons (1,0) + (0,1), \\ (1,2) &\rightleftharpoons (1,0) + (0,2), \\ (1,2) &\rightleftharpoons (1,1) + (0,1), \\ (2,1) &\rightleftharpoons (2,0) + (0,1), \\ (2,1) &\rightleftharpoons (1,1) + (1,0), \\ (2,2) &\rightleftharpoons (2,0) + (0,2), \\ (2,2) &\rightleftharpoons (1,0) + (1,2), \\ (2,2) &\rightleftharpoons (2,1) + (0,1), \\ (2,2) &\rightleftharpoons (1,1) + (1,1). \end{aligned} \quad (\text{B2})$$

Indeed, from (13) we find $|\mathcal{R}|^{\max} = (3 \times 4)^2/8 + 2^2/2 - 8 - 1 = 11$.

APPENDIX C: ERDŐS-RÉNYI RANDOM GRAPHS

The Erdős–Rényi (ER) random graph is the simplest and most well-studied random graph model. We briefly describe it here and direct readers to Refs. 8, 44, and 81 for comprehensive introductions. In the ER model, a graph is generated with a given number of nodes, denoted by $|\mathcal{N}|$, where each possible pair of nodes is connected by an edge with a probability p . The total number of potential edges in such a graph is given by $|\mathcal{E}|^{\max} = |\mathcal{N}|(|\mathcal{N}| - 1)/2 \sim |\mathcal{N}|^2/2$. One of the most well-studied properties of the ER model is its component structure as a function of the connection probability p ,^{6,8,82} in the limit of large $|\mathcal{N}|$. As p varies, the graph transitions between different connectivity regimes. When $p < 1/|\mathcal{N}|$, the graph is fragmented into many small components, each with a maximum size of order $\mathcal{O}(\ln|\mathcal{N}|)$. However, as p exceeds $1/|\mathcal{N}|$, a “giant component” of size $\mathcal{O}(|\mathcal{N}|)$ suddenly emerges, containing an extensive fraction of the nodes. This emergence of a giant component is a second-order phase transition, belonging to the universality class of mean field percolation.^{46,83} The control parameter for this transition is given by $\gamma = (|\mathcal{N}| - 1)p$, while the order parameter is the fraction of nodes in the largest component. As we vary γ , the order parameter remains zero below the critical value $\gamma_c = 1$, beyond which it is nonzero.

APPENDIX D: ONE ATOM

In this section, we collect all the results for the case of $N_A = 1$, with the maximum possible number of atoms being Q . There are $|\mathcal{Z}|^{\max} = Q$ species labeled by the index σ , i.e., the set of species is of the form $\{(1), (2), \dots, \sigma, \dots, (Q)\}$. Using Eqs. (12) and (20), the total number of reactions that a species σ can participate in is given by the sum of the number of dissociations $D(\sigma) = \lfloor n_\sigma/2 \rfloor$ (with $\lfloor \cdot \rfloor$ being the floor function) and n_σ being the number of atoms in σ . Similarly, the number of associations $A(\sigma) = Q - n_\sigma$ and

$$I(\sigma) = Q - \lceil n_\sigma/2 \rceil, \quad (\text{D1})$$

where $\lceil \cdot \rceil$ is the ceiling function. The total number of possible reactions thus follows from Eq. (13),

$$|\mathcal{R}|^{\max} = \lfloor (Q/2)^2 \rfloor. \quad (\text{D2})$$

1. Number distribution and average species

The probability for a species σ to have degree k is given by the binomial distribution,

$$p_\sigma(k) = \binom{Q - \lceil n_\sigma/2 \rceil}{k} p^k (1-p)^{Q - \lceil n_\sigma/2 \rceil - k}, \quad (\text{D3})$$

where we also used Eq. (D1). Using the independence approximation stated in Sec. IV B, the number distribution given by Eq. (23), takes the form

$$P(n_k) = \sum_{\substack{A \subset \{(1), (2), \dots, (Q)\} \\ |A| = n_k}} \prod_{\sigma \in A} p_\sigma(k) \prod_{\sigma' \notin A} (1 - p_{\sigma'}(k)), \quad (\text{D4})$$

with the mean value,

$$\begin{aligned} \mu(k) &= \sum_{\sigma} p_\sigma(k), \\ &= \frac{1}{k!} \left(\frac{p}{1-p} \right)^k \sum_{n_\sigma=1}^Q \frac{I(\sigma)!}{(I(\sigma) - k)!} (1-p)^{I(\sigma)}. \end{aligned} \quad (\text{D5})$$

In the regime of Eq. (15) and in the limit of $Q \rightarrow \infty$, we obtain a closed form expression for the degree distribution given in Eq. (27),

$$p_k = \frac{2}{\gamma k!} [g(k+1, \gamma) - g(k+1, \gamma/2)], \quad (\text{D6})$$

with $g(k, \gamma)$ denoting the lower incomplete gamma function.⁸⁴

We now turn to the expected number of species. Using Eq. (28) with Eq. (D5), the average number of species $\langle |\mathcal{Z}| \rangle$ participating in a reaction takes the form

$$\begin{aligned} \langle |\mathcal{Z}| \rangle &= Q - \sum_{\sigma} (1-p)^{Q - \lceil n_\sigma/2 \rceil}, \\ &\approx Q - (1-p)^{Q-1/2} \left[\frac{1 - (1-p)^{-Q/2}}{1 - (1-p)^{-1/2}} \right], \end{aligned} \quad (\text{D7})$$

where we have replaced $\lceil n_\sigma/2 \rceil$ with $n_\sigma/2$ in the second line. In the limit of $p \rightarrow 0$, Eq. (D7) goes to the value $Q(Q-1/2)p$, which also vanishes. This is consistent as at very low probabilities, no reactions are chosen, implying the average number of chosen species goes to zero. Similarly, in the limit of $p \rightarrow 1$, Eq. (D7) goes to the value Q , implying all species are chosen on average.

2. Upper bound on conservation laws

As argued in Sec. IV C 2, the number of conservation laws is upper bounded by the number of root species [see Eqs. (52) and (51)]. When $N_A = 1$, we can use this result to find a closed form upper bound on the average number of conservation laws. We

denote the probability that a species σ is a root as $\mathbb{P}[\sigma \in \mathcal{Z}^r]$. Assuming that $[\sigma \in \mathcal{Z}^r]$ and $[\sigma' \in \mathcal{Z}^r]$ are independent events, the average number of roots is the sum,

$$\langle |\mathcal{Z}^r| \rangle = \sum_{\sigma} \mathbb{P}[\sigma \in \mathcal{Z}^r]. \quad (\text{D8})$$

We now condition the probability that σ is a root on the number of reactions $|\mathcal{R}|$, i.e.,

$$\mathbb{P}[\sigma \in \mathcal{Z}^r] = \sum_{|\mathcal{R}|=1}^{|\mathcal{R}|^{\max}} \mathbb{P}[\sigma \in \mathcal{Z}^r \mid |\mathcal{R}|] \binom{|\mathcal{R}|^{\max}}{|\mathcal{R}|} p^{|\mathcal{R}|} (1-p)^{|\mathcal{R}|^{\max}-|\mathcal{R}|}. \quad (\text{D9})$$

The conditional probability that a species σ is a root, given a fixed number of reactions $|\mathcal{R}|$, is the fraction of possible CRNs with $|\mathcal{R}|$ reactions such that there is at least one association reaction involving σ but no dissociation reaction of σ i.e.,

$$\mathbb{P}[\sigma \in \mathcal{Z}^r \mid |\mathcal{R}|] = \frac{\binom{|\mathcal{R}|^{\max} - \lfloor n_{\sigma}/2 \rfloor}{|\mathcal{R}|} - \binom{|\mathcal{R}|^{\max} - Q + \lfloor n_{\sigma}/2 \rfloor}{|\mathcal{R}|}}{\binom{|\mathcal{R}|^{\max}}{|\mathcal{R}|}}. \quad (\text{D10})$$

Substituting Eq. (D10) in Eqs. (D8) and (D9), we get the result,

$$\langle |\mathcal{L}| \rangle \leq \langle |\mathcal{Z}^r| \rangle = \sum_{\sigma} \sum_{|\mathcal{R}|=1}^{|\mathcal{R}|^{\max}} p^{|\mathcal{R}|} (1-p)^{|\mathcal{R}|^{\max}-|\mathcal{R}|} \times \left\{ \binom{|\mathcal{R}|^{\max} - \lfloor n_{\sigma}/2 \rfloor}{|\mathcal{R}|} - \binom{|\mathcal{R}|^{\max} - Q + \lfloor n_{\sigma}/2 \rfloor}{|\mathcal{R}|} \right\}. \quad (\text{D11})$$

Equation (D11) captures the nonmonotonic behavior of the average number of conservation laws, as shown in Fig. 18, and is a good approximation at low p but is not a tight bound at higher p .

3. Cycles

In this subsection, we first build a basis set for cycles when $p = 1$, i.e., all reactions are considered inductively. Using this






FIG. 18. Average number of conservation laws as a function of the probability to select a reaction (dotted) compared to the theoretical result in Eq. (D11). The parameters are $N_A = 1$ and $Q = 36$. For each value of p , we generated 5000 reaction sets for averaging.

basis set, we construct a lower bound for the average number of cycles.


We illustrate the inductive approach via an example: when $Q < 4$, the CRN has no cycles. The first cycle emerges when $Q = 4$ and is given by



If we increase the number of species Q to 5, any new cycle that emerges must involve the species (5). There is once again one new cycle,



We see that both the cycles have exactly 4 reactions. We now extend the approach to a general Q . First, a general set of 4 reactions, which form a cycle are of the form,



Finally, let the number of cycles for $Q-1$ species be $N_c(Q-1)$. If we now add another species, the number of cycles increases to $N_c(Q)$. The extra cycles must all involve reactions of species (Q), which are $\lfloor Q/2 \rfloor$ in number. The first reaction is fixed to be $(1) + (Q-1) \rightleftharpoons (Q)$. Choosing any of the other $\lfloor Q/2 \rfloor - 1$ reactions gives different independent cycles. Thus, we have

$$N_c(Q) = N_c(Q-1) + \lfloor Q/2 \rfloor - 1, \quad (\text{D13})$$

with solution,

$$N_c(Q) \approx \frac{Q^2 - 3Q}{4}. \quad (\text{D14})$$

The same result was also derived differently in Ref. 56. We compare Eq. (D13) at different Q to the exact number of cycles that can be obtained numerically in Fig. 19(a). We see that at large Q , we get a very good agreement. While the construction above gives an independent basis for the kernel of ∇ , the different cycles in the basis share reactions. Thus, the average number of cycles at any p is lower bounded by the probability of choosing cycles in the basis independently, which is given by

$$\langle |\mathcal{C}| \rangle \geq N_c(Q) p^4, \quad (\text{D15})$$




FIG. 19. (a) Total possible number of cycles plotted for $N_A = 1$ as a function of the maximum size of a molecule Q (dotted) compared to the theoretical result in Eq. (D14), (b) average number of cycles scaled by the maximum possible number of cycles for $N_A = 1$ plotted as a function of the probability to select a reaction (dotted) compared to the theoretical result in Eq. (D15). For each value of p , we generated 1000 reaction sets for averaging.

where we used Eq. (D12) along with the fact that the probability to choose each reaction is p . Comparing Eq. (D15) with simulations in Fig. 19(b), we see that the lower bound is only tight for small Q and small p . We note that, in the extensive limit with the rescaling of Eq. (15), the lower bound in Eq. (D15) goes to zero, as anticipated in Eq. (55).

APPENDIX E: LIMITING FORM OF DEGREE DISTRIBUTION

We want to compute the limiting form of Eq. (26) in the limit of $|\mathcal{Z}|^{\max} \rightarrow \infty$ and with the rescaled variable $p = \gamma / |\mathcal{Z}|^{\max}$. Using Eq. (24),

$$\begin{aligned} \mu(k) &= \sum_{\sigma} \binom{I(\sigma)}{k} p^k (1-p)^{I(\sigma)-k}, \\ &= \gamma^k \sum_{\sigma} \binom{I(\sigma)}{k} \left(\frac{1}{|\mathcal{Z}|^{\max}} \right)^k \left(1 - \frac{\gamma}{|\mathcal{Z}|^{\max}} \right)^{I(\sigma)-k}, \\ &\simeq \frac{\gamma^k}{k!} \sum_{\sigma} \left(\frac{I(\sigma)}{|\mathcal{Z}|^{\max}} \right)^k e^{-\gamma I(\sigma) / |\mathcal{Z}|^{\max}}, \end{aligned} \quad (\text{E1})$$

where $I(\sigma)$ is the number of reactions a species participates in [see Eq. (19)] and where we used the Poisson approximation of the Binomial. Using Eqs. (10), (12), and (20), we can write

$$\begin{aligned} \frac{I(\sigma)}{|\mathcal{Z}|^{\max}} &= \frac{D(\sigma)}{|\mathcal{Z}|^{\max}} + \frac{A(\sigma)}{|\mathcal{Z}|^{\max}}, \\ &= \frac{1}{2} \prod_{a=1}^{N_A} \frac{n_a + 1}{Q_a + 1} + \prod_{a=1}^{N_A} \left(1 - \frac{n_a}{Q_a + 1} \right) + \frac{\delta(\sigma)}{|\mathcal{Z}|^{\max}}, \end{aligned} \quad (\text{E2})$$

where in the quantity $\delta(\sigma)$ we have combined $\mathcal{O}(1)$ terms from Eqs. (20) and (12). Defining the fractions $x_a = n_a / (Q_a + 1) \approx n_a / Q_a$ and substituting Eq. (E2) in (E1), we find

$$\begin{aligned} \frac{\mu(k)}{|\mathcal{Z}|^{\max}} &\simeq \frac{\gamma^k}{k!} \sum_{x_1=0}^1 \cdots \sum_{x_{N_A}=0}^1 \left(\frac{1}{2} \prod_{i=1}^{N_A} x_i + \prod_{j=1}^{N_A} (1-x_j) \right)^k \\ &\times \exp \left[-\gamma \left(\frac{1}{2} \prod_{i=1}^{N_A} x_i + \prod_{j=1}^{N_A} (1-x_j) \right) \right] \\ &\simeq \frac{\gamma^k}{k!} \int_0^1 dx_1 \cdots \int_0^1 dx_{N_A} \left(\frac{1}{2} \prod_{i=1}^{N_A} x_i + \prod_{j=1}^{N_A} (1-x_j) \right)^k \\ &\times \exp \left[-\gamma \left(\frac{1}{2} \prod_{i=1}^{N_A} x_i + \prod_{j=1}^{N_A} (1-x_j) \right) \right], \end{aligned} \quad (\text{E3})$$

where we converted the discrete sum in the first equation into an integral.

APPENDIX F: SIMULATIONS FOR EMERGENT CYCLES

The simulations in Sec. IV C comprised of two steps. First, we generated random CRNs for prescribed values of γ and $Q_1, Q_2, \dots, Q_{N_A}$. For each random CRN, we then randomly selected a subset of species in $\mathcal{Z}$ and asked whether chemostating them produced an emergent cycle by evaluating the difference $\dim(\ker(\nabla^X)) - \dim(\ker(\nabla))$. To perform this rank calculation efficiently for large ∇ , we used the package SuiteSparseQR⁸⁵ through the Python wrapper sparseqr. In the following, we describe the specifics and the analysis of the resulting data for the cases $N_A = 1$ and $N_A = 2$.

1. One atom

We outline here the procedure used to estimate the threshold γ_c for the single-atom case. For each value of Q , samples of n_i (as defined in Sec. IV C) were simulated for $\gamma \in [3.0, 4.5]$ with a step size of 0.01. The Q values used were (100, 250, 500, 750, 1000, 1250, 1500, 1750, 2000), with the number of samples per γ set to (25 000, 25 000, 20 000) for the three smallest Q values, and 10 000 samples for all larger Q . The corresponding probabilities $\mathbb{P}_e(\gamma, Q)$ were computed as the sample mean of n_i and are plotted in Fig. 4, along with the standard error of the mean for selected Q values.

In principle, γ_c can be estimated by extrapolating the crossing points of the $\mathbb{P}_e(\gamma, Q)$ curves for successive values of Q (e.g., between $Q = 100$ and 250, 250 and 500, and so on), followed by extrapolation to $Q \rightarrow \infty$. However, in our case, the crossing points were all tightly clustered around $\gamma \approx 3.62$, and applying standard extrapolation techniques did not yield reliable results. We speculate that Q dependent corrections are already small for $Q \geq 100$.

To obtain a more accurate estimate, we instead used the peaks of the variance $\tilde{\chi}(\gamma, Q)$. For each Q , we focused on the range $\gamma \in (3.5, 3.74)$ and interpolated the data using a cubic spline on a refined grid with step size 0.0001. The location of the peak was then identified for each Q and these peak positions were extrapolated using the BST algorithm to get an estimate of γ_c .⁸⁶

To estimate uncertainty, we repeated the entire procedure with 1000 bootstrap resamples of the underlying data. This was followed by computing $\tilde{\chi}(\gamma, Q)$ for the resampled data and the procedure outlined above was followed. The resulting distribution had a mean of 3.6242 and a standard deviation of 0.0198. Since the original data




FIG. 20. Data from Fig. 6 smoothed by a Gaussian filter and the difference between the smoothed and original data plotted as a function of γ . The parameters are $N_A = 2$ and $Q_1 = Q_2 = Q$ and $\sigma = 4$ for the filtering.

were defined on a grid with spacing 0.01, we report the threshold as $\gamma_c = 3.62 \pm 0.02$.

2. Two atom

We now describe the procedure used to extract the critical exponents for the two-atom case ($N_A = 2$) with $Q_1 = Q_2 = Q$. For each Q , samples of n_i (as defined in Sec. IV C) were simulated for $\gamma \in [5.5, 7.5]$ using a step size of 0.01. The Q values used were (10, 15, 20, 24, ..., 70), with 32 000 samples per γ for all Q except the three largest, for which 16 000 samples were generated. The corresponding probabilities $\mathbb{P}_e(\gamma, Q, Q)$ were computed as the sample mean of n_i and are plotted in Fig. 6, along with the standard error of the mean for selected values of Q . To compute the numerical derivative, we first smoothed the data using the `gaussianfilter1d` routine from `scipy.ndimage` with a smoothing parameter $\sigma = 4.0$. The derivative was then calculated using a forward Euler scheme. We show the magnitude of the difference between $\mathbb{P}_e(\gamma, Q, Q)$ for $Q = 25, 35, 45$ and the resulting data after smoothing in Fig. 20.

For the scaling analysis, we followed the approach of Refs. 47 and 48 that defined a collapse quality function S , where a good collapse corresponds to $S \approx 1$. To determine the critical exponents, we minimized S using the Nelder–Mead algorithm from `scipy.optimize`, starting from the initial guess $\gamma_c = 6.76$, $a = 2.24$, and $b = 0.09$, chosen by eye.

A crucial point here is to estimate the limits of the rescaled variable, denoted $x = (\gamma - \gamma_c)|Z^{\max}|^{1/a}$ in which the collapse is valid. If we denote the upper limit (resp. lower limit) of this interval as $x_{\max}$ (resp. $x_{\min}$), the collapse is perfectly valid in the limit of $x_{\max}, x_{\min} \rightarrow 0$. However, choosing a narrow interval may lead to overfitting and unphysical values of the exponents (e.g., $a < 0$). Typically, in such cases, both $x_{\min}$ and $x_{\max}$ lie on the same side of zero. We scanned over values of $x_{\max}$ from 1.5 to 15 in steps of 0.5. For each candidate window, we minimized S with respect to the five parameters $(\gamma_c, a, b, x_{\min}, x_{\max})$. Fits yielding unphysical exponents were discarded, and the best-fit parameters were selected from the remaining cases.




FIG. 21. Data from Fig. 6 rescaled as in Eq. (34) and plotted as a function of $(\gamma - \gamma_c)(|Z|^{\max})^{1/a}$. The parameters are $N_A = 2$ and $Q_1 = Q_2 = Q$. The dotted black lines indicate the regime of the best fit.

This procedure yielded the estimates $\gamma_c = 6.7415$, $a = 2.2191$, and $b = 0.0949$, with corresponding limits $x_{\min}$, $x_{\max}$, and a minimum collapse score of $S_{\min} = 1.454$. To estimate uncertainties, we varied the exponents until $S = S_{\min} + 1$, resulting in: $\delta\gamma_c = 0.0057$, $\delta a = 0.1315$, and $\delta b = 0.0052$. Given the original grid spacing of 0.01, we report the final estimates as $\gamma_c = 6.74 \pm 0.01$, $a = 2.219 \pm 0.131$, and $b = 0.095 \pm 0.005$. This rounding changes the collapse score only slightly to $S = 1.56$, which still indicates a good fit.

The data in Fig. 6 are rescaled according to Eq. (34) using the above-mentioned exponents and plotted in Fig. 21. The black dashed lines denote the best-fit region, which includes $x = 0$. Figure 22 shows the peak height of the smoothed numerical derivative $\partial\mathbb{P}_e(\gamma_r, Q, Q)/\partial\gamma$ plotted against $|Z|^{\max}$. A linear fit on a log-log scale (excluding the first three data points) yields $r^2 > 0.99$, and the extracted exponent is consistent with the scaling exponents reported above. The error in the exponent was calculated by bootstrap resampling and repeating the smoothing and differentiation steps. Finally, the location of the peak for each $|Z|^{\max}$ approaches γ_c , as shown in Fig. 23. Once again, a linear




FIG. 22. Peak of the derivative of the probability of having an emergent cycle wrt γ plotted as a function of $|Z|^{\max}$ in the log-log scale. The parameters are $N_A = 2$ and $Q_1 = Q_2 = Q$. The red line represents a linear fit.




FIG. 23. Magnitude of the difference between the location of the peak of the derivative of the probability of having an emergent cycle wrt γ from γ_c , plotted as a function of $|Z|^{\max}$ in the log-log scale. The parameters are $N_A = 2$ and $Q_1 = Q_2 = Q$. The red line represents a linear fit.

fit on a log-log scale (excluding the first three data points) yields $r^2 > 0.99$ and the extracted exponent is consistent with the scaling analysis.

We note two caveats here: first, while $r^2 > 0.99$ indicates a good linear fit, it does not in itself confirm a power-law relationship; second, the derivative peaks and their locations were extracted after applying smoothing. However, given that the resulting exponents are consistent with the scaling analysis, we believe any bias introduced by this step is negligible.

APPENDIX G: PROOFS

1. Conservation laws at $p = 1$

Given an arbitrary conservation law $\ell^\lambda = (\dots, \ell_{\sigma_a}^\lambda, \dots)^\top$, its components satisfy a set of equalities of the form,

$$\ell_{r(\sigma)}^\lambda + \ell_{r'(\sigma)}^\lambda = \ell_{p(\sigma)}^\lambda, \quad (\text{G1})$$

where we used the definition $\ell^\lambda \cdot \nabla = 0$ along with Eq. (48). A solution to this set of equations is given by the set of vectors $\{\ell^a\}$ with entries,

$$\ell_\sigma^a = n_a, \quad (\text{G2})$$

where $\sigma = (\dots, n_a, \dots)$ and the index a goes from 1 to N_A . Here, each vector ℓ^a counts the number of atoms of type a in σ . By construction, each vector ℓ^a is a conservation law: see Eqs. (G1), (G2), and (11). As each of the vectors, ℓ^a , are independent, the number of conservation laws of ∇ when $p \rightarrow 1$ are lower bounded by N_A .

Now, consider a generic conservation law ℓ^1 and a set of species σ_a , which are just singlet atoms, i.e., $\sigma_a = (0, 0, 0, \dots, 1, 0, 0, \dots)^\top$. Since $p \rightarrow 1$, we can expect that any realization of the random CRN will include all reactions of the form Eq. (11). Starting with this set of species, we can construct a chain of reactions, taking us to any

arbitrary species $\sigma = (n_1, n_2, \dots, n_{N_A})$. Iterating Eq. (G1) along this chain of reactions, the entries of ℓ^1 satisfy

$$\ell_\sigma^1 = \sum_{a=1}^{N_A} n_a \ell_{\sigma_a}^1 = \sum_{a=1}^{N_A} \ell_{\sigma_a}^1 \ell_{\sigma_a}^1, \quad (\text{G3})$$

which implies that arbitrary conservation laws can be expressed as linear combinations of the set $\{\ell^a\}$, showing that at probabilities close to one, the number of conservation laws should be exactly N_A .

2. Applying our ansatz to the ER transition

We use the approach of Subsection IV C 1 to show here that our approach can only approximate the percolation threshold in the ER case. We are interested in the number of connected components scaled by the number of vertices (denoted c) as a function of the rescaled probability γ . We first show the exact result⁸² and then our ansatz for comparison.

a. Exact result

The order parameter for the ER graph, denoted here with $\mathcal{S}$, is the (average) fraction of nodes in the largest component of the random graph and it is known to be the (stable) root of the equation,^{8,82}

$$1 - \mathcal{S}(\gamma) - e^{-\gamma \mathcal{S}(\gamma)} = 0. \quad (\text{G4})$$

From Ref. 82, the scaled average number of components satisfies the equation,

$$\langle c \rangle = (1 - \mathcal{S}(\gamma)) \left(1 - \frac{(1 - \mathcal{S}(\gamma))\gamma}{2} \right). \quad (\text{G5})$$

However, the quantity $\langle c \rangle$ also includes nodes that are not connected to any other component in the graph. While this is a common practice in complex networks, we exclude all such nodes in our analysis. The probability of a node having degree zero at all is $e^{-\gamma}$. Thus, we get a modified number of components denoted $\langle \bar{c} \rangle$, which takes the form

$$\langle \bar{c} \rangle = (1 - \mathcal{S}(\gamma)) \left(1 - \frac{(1 - \mathcal{S}(\gamma))\gamma}{2} \right) - e^{-\gamma}. \quad (\text{G6})$$

Substituting Eq. (G4) in Eq. (G6), close to the critical probability $\gamma_c = 1$, we have the expression,

$$\langle \bar{c} \rangle = \begin{cases} 1 - \frac{\gamma}{2} - e^{-\gamma} & \gamma \leq 1, \\ (3 - 2\gamma) \left(1 - \gamma \left(\frac{3}{2} - \gamma \right) \right) - e^{-\gamma} & \gamma > 1. \end{cases} \quad (\text{G7})$$

Figure 24 shows a simulation of $\langle \bar{c} \rangle$ as a function of γ . Crucially, both the simulation shown in Fig. 24 and the theory in Eq. (G7) show no discontinuities or kinks at $\gamma = 1$. In fact, the same is true for the first derivative of $\langle \bar{c} \rangle$ and there is only a kink in the second derivative $\langle \bar{c} \rangle$.




FIG. 24. Average number of components (scaled) in the ER random graph for $|\mathcal{N}| = 1000$ and $|\mathcal{N}| = 10000$ plotted as a function of the scaled probability to select a reaction compared to the exact theory (blue solid line) from Eq. (G6) and the rank ansatz in Eq. (G9) (dashed red line). For each value of p , we generated 1000 graphs for averaging. The gray line represents $\gamma_c = 1$.

b. Using the ansatz

We can write an ansatz with the rank of the incidence matrix ∂ similar to Eq. (36) and the analysis in Sec. IV C 1. For an ER graph, this ansatz would take the form

$$\langle \text{rk}(\partial) \rangle \leq \begin{cases} \langle |\mathcal{E}| \rangle & p \leq p^* \\ \langle |\mathcal{N}| \rangle - 1 & p > p^*. \end{cases} \quad (\text{G8})$$

Intuitively, the term N_A in Eq. (38) gets replaced by 1 as in the fully connected graph, there is only one conservation law which is the vector $(1, 1, 1, \dots)^T$. Transforming into the rescaled variables $p = \gamma / (|\mathcal{N}| - 1)$ and using Eq. (G8) along with the rank nullity theorem, we find that $\langle \bar{c} \rangle$ must satisfy

$$\langle \bar{c} \rangle \geq \begin{cases} 1 - e^{-\gamma} - \frac{\gamma}{2} & \gamma \leq \gamma^* \\ 0 & \gamma > \gamma^*, \end{cases} \quad (\text{G9})$$

where $\gamma^* \approx 1.493$. Note that while Eq. (G9) matches Eq. (G7) exactly below the percolation transition, it is a lower bound above it. Particularly, the predicted transition point γ^* has no physical meaning in the ER context. We compare both Eqs. (G7) and (G9) along with simulations in Fig. 24 to show the difference.

3. Results related to roots

Consider a CRN with the stoichiometric matrix ∇ and all reactions of the form of Eq. (11) and each reaction oriented as in Eq. (48). We first identify the root species by finding those species such that all the entries in the corresponding row of the stoichiometric matrix ∇ are non-positive. We then split the set of species $\mathcal{Z}$ as $\mathcal{Z} = \mathcal{Z}^r \cup \mathcal{Z}^l$, where the set of roots (resp. non-roots) is denoted $\mathcal{Z}^r$ (resp. $\mathcal{Z}^l$). Correspondingly, the stoichiometric matrix ∇ splits into

$$\nabla = \begin{pmatrix} \nabla^l \\ \nabla^r \end{pmatrix}, \quad (\text{G10})$$

where ∇^r (resp. ∇^l) is the stoichiometric matrix with rows corresponding only to the root species (resp. non-root species). The columns of ∇^l can take the following four forms (indicating only the nonzero elements):

$$\nabla_\rho^l = \begin{pmatrix} 1 \\ \cdot \\ \cdot \\ \cdot \\ \cdot \end{pmatrix}, \begin{pmatrix} \cdot \\ -1 \\ 1 \\ \cdot \\ \cdot \end{pmatrix}, \begin{pmatrix} -2 \\ \cdot \\ \cdot \\ \cdot \\ 1 \end{pmatrix}, \begin{pmatrix} \cdot \\ \cdot \\ -1 \\ -1 \\ 1 \end{pmatrix}, \quad (\text{G11})$$

where the first two columns are reactions involving root species [see Eqs. (49) and (50)] while the other two columns involve only non-root species.

a. Result I

We now argue that the left nullspace of ∇^l is trivial.

Consider any vector ℓ such that $\ell \cdot \nabla^l = 0$ where ∇^l has been defined in Eq. (G10). Using the forms of the columns of ∇^l as shown in Eq. (G11), we show that the entries of ℓ are identically zero. Consider a given species σ , which is not a root. Since it is not a root, it must be the product of at least one reaction. Consider the corresponding column of ∇^l . If it is of the first form in Eq. (G11), it is clear that $\ell_\sigma = 0$. Otherwise, assume it is of the second form in Eq. (G11). Then, we get a linear relation between the entries $\ell_\sigma = \ell_{\sigma'}$, where the species σ' is the non-root reactant of the same reaction. Starting with (σ') , we look for a reaction ρ where (σ') is a product (if no such reaction were to exist, σ' would be a root) and repeat the process. Eventually, this process will terminate in a reaction of the form of column one (in each step, the atom numbers of the species reduce—eventually the process has to stop), implying $\ell_\sigma = 0$. If the reaction is instead of the third or the fourth form in Eq. (G11), then we get an equation of the form $\ell_\sigma = \ell_{\sigma_{r(\rho)}} + \ell_{\sigma_{r'(\rho)}}$. We repeat the process for both the species $\sigma_{r(\rho)}$ and $\sigma_{r'(\rho)}$ and arrive at the result.

b. Result II

We show that the number of roots is always greater than or equal to the number of conservation laws.

Since the left nullspace of ∇^l is trivial, applying the rank-nullity theorem to its transpose, we find $\text{rk}(\nabla^l) = |\mathcal{Z}| - |\mathcal{Z}^r|$. Substituting the above in Eq. (2), we find

$$|\mathcal{Z}| + \dim(\ker(\nabla^l)) = |\mathcal{R}| + |\mathcal{Z}^r|. \quad (\text{G12})$$

Substituting for $|\mathcal{Z}|$ from Eq. (2) into Eq. (G12), we get

$$|\mathcal{L}| + \dim(\ker(\nabla^l)) - |\mathcal{C}| = |\mathcal{Z}^r|. \quad (\text{G13})$$

Finally, using Eq. (3) with Eq. (G13), we obtain Eq. (51), which directly implies the result.

c. Result III

Here, we define the construction of independent roots.

Consider the matrix $\mathbb{L}$ whose rows are the conservation laws and denote its submatrix where the columns are restricted to the root species as $\mathbb{L}^r$. The rank of $\mathbb{L}^r$ must be equal to $|\mathcal{L}|$.

If not, there is a basis of the conservation laws, where at least one row of $\mathbb{L}^r$ has all entries zero. Call this conservation law ℓ' . Since $\ell' \cdot \nabla = 0$, we have the relation,

$$\ell'_i \cdot \nabla^i + \ell'_r \cdot \nabla^r = 0, \quad (\text{G14})$$

where we wrote the conservation law as $\ell' = (\ell'_i, \ell'_r)^\top$. Since $\ell'_r = 0$, ℓ'_i must satisfy $\ell'_i \cdot \nabla^i = 0$, which is a contradiction. Since the matrix $\mathbb{L}^r$ has full rank, it must have at least one invertible square submatrix, denoted $\mathbb{L}^{\mathcal{I}}$. We call the subset of root species associated with this square invertible submatrix independent roots and denote the set of independent roots by $\mathcal{I} \subseteq \mathcal{Z}^r$. The number of independent roots must, therefore, be equal to the number of conservation laws. Assume that we now choose the set of independent roots $\mathcal{I}$ (resp. $\mathcal{Z}/\mathcal{I}$) as the external species (resp. internal species) with the stoichiometric matrices $\nabla^{\mathcal{I}}$ (resp. $\nabla^{\mathcal{X}}$). Then, as in [Appendix G 3 1](#), the left null space of $\nabla^{\mathcal{X}}$ must be trivial. If not, given a vector ℓ_0 such that $\ell_0 \cdot \nabla^{\mathcal{X}} = 0$, we can then define a new conservation law of the form,

$$\ell' \cdot \nabla = \begin{matrix} X & \mathcal{I} \\ [\ell_0 & \mathbf{0}] \end{matrix} \cdot \begin{pmatrix} \nabla^{\mathcal{X}} \\ \nabla^{\mathcal{I}} \end{pmatrix} = 0. \quad (\text{G15})$$

Finally, using this conservation law ℓ' , we see that a linear combination of the rows of $\mathbb{L}^{\mathcal{I}}$ should be zero—which contradicts the assumption that $\mathbb{L}^{\mathcal{I}}$ is invertible. Since the left nullspace of $\nabla^{\mathcal{X}}$ is trivial, following the steps outlined in [Appendix G 3 2](#), it is seen that $|e|_{\mathcal{I}} = 0$, i.e., there are no emergent cycles. We note that there might be multiple possible choices of the set of independent roots, and the results do not depend on the choice.

d. Result IV

We now show how to construct any species from the independent roots via a net reaction.

Assume we wish to construct an arbitrary species σ_0 from the set of independent roots. Consider as the set of external species $Y' = \mathcal{I} \cup \{\sigma_0\}$ and the remaining species as the internal species X' with stoichiometric matrices $\nabla^{Y'}$ and $\nabla^{X'}$. However, $\nabla^{X'} \subset \nabla^{\mathcal{X}}$, where X is the set of internal species if only the independent roots are considered the external species. In [Appendix G 3 3](#), we showed that the left nullspace of $\nabla^{\mathcal{X}}$ is trivial. Thus, the left nullspace of $\nabla^{X'}$ is also trivial (else, vectors in the left nullspace of $\nabla^{X'}$ can be extended into vectors in the left nullspace of $\nabla^{\mathcal{X}}$). Thus, the rank of $\nabla^{X'}$ must be equal to $|\mathcal{Z}| - |\mathcal{I}| - 1$. Applying the rank-nullity theorem to $\nabla^{X'}$, we find

$$|\mathcal{Z}| + \dim(\ker(\nabla^{X'})) = |\mathcal{R}| + |\mathcal{I}| + 1. \quad (\text{G16})$$

Substituting for $|\mathcal{Z}|$ from Eq. (2) into Eq. (G16), we get

$$|\mathcal{L}| + \dim(\ker(\nabla^{\mathcal{I}})) - |\mathcal{C}| = |\mathcal{I}| + 1 \implies |e|_{Y'} = 1, \quad (\text{G17})$$

where we used Eq. (3) (also see [Appendix G 3 3](#)). Thus, we have a unique emergent cycle $\mathbf{c}_{\sigma_0}$ associated with this set of external species. Finally, using Eq. (5), the net reaction constructing σ_0 is obtained from the vector $\nabla^{Y'} \cdot \mathbf{c}_{\sigma_0}$ upon suitable rescaling in the form

$$\sigma_0 \rightleftharpoons \sum_{i \in \mathcal{I}} v_{\sigma_0, i} \hat{\sigma}_i. \quad (\text{G18})$$

Unlike the coefficients in Eq. (5), all the coefficients in Eq. (G18) are integers but need not be positive. We note that the coefficients $v_{\sigma_0, i}$ give us the representation of the species σ_0 when we regard the independent roots as atoms.

e. Result V

We now derive a basis set of conservation laws starting from the set of independent roots.

For every species $\sigma \in \mathcal{Z}$, following the prescription of [Appendix G 3 4](#), we can write a representation as shown in Eq. (G18). Associated with a given independent root $\hat{\sigma}_i$, we define the vector ℓ^i whose entry for the species σ is the coefficient $v_{\sigma, i}$. Given ℓ^i , the entries corresponding to the independent roots $\hat{\sigma}_j$ are given by the Kronecker symbol δ_{ij} . Thus, we have $|\mathcal{I}| = |\mathcal{L}|$ independent vectors. If a vector ℓ_j is a conservation law, given a reaction of form Eq. (48), the coefficients of Eq. (G18) need to satisfy $v_{\sigma_{r(\rho)}, i} + v_{\sigma_{r'(\rho)}, i} = v_{\sigma_{p(\rho)}, i}$ for every reaction ρ . To prove this, consider a particular reaction ρ and consider as the external species $Y = \mathcal{I} \cup \{\sigma_{r(\rho)}, \sigma_{r'(\rho)}, \sigma_{p(\rho)}\}$ and the remaining species as internal species. This splits the stoichiometric matrix as

$$\nabla = \begin{matrix} X & \begin{bmatrix} \nabla^{\mathcal{X}} \\ \nabla^{\sigma_{r(\rho)}} \\ \nabla^{\sigma_{r'(\rho)}} \\ \nabla^{\sigma_{p(\rho)}} \\ \nabla^{\mathcal{I}} \end{bmatrix} \\ \sigma_{r(\rho)} & \\ \sigma_{r'(\rho)} & \\ \sigma_{p(\rho)} & \\ \mathcal{I} & \end{matrix} \quad (\text{G19})$$

Arguments along the lines in [Appendixes G 3 2](#) and [G 3 4](#) show us that the number of emergent cycles in this case must be exactly three. Consider the three appropriately scaled vectors $\mathbf{c}_{\sigma_{r(\rho)}}$, $\mathbf{c}_{\sigma_{r'(\rho)}}$, $\mathbf{c}_{\sigma_{p(\rho)}}$ defined as in [Appendix G 3 4](#). Each of these vectors are emergent cycles and construct the substrates and products of the reaction ρ from the independent roots. Finally, consider the vector $\mathbf{c}$, which is zero for all reactions and is one for the reaction ρ . This is also an emergent cycle as

$$\nabla \cdot \mathbf{c} = \begin{matrix} X & \begin{bmatrix} \mathbf{0} \\ -1 \\ -1 \\ 1 \\ \mathbf{0} \end{bmatrix} \\ \sigma_{r(\rho)} & \\ \sigma_{r'(\rho)} & \\ \sigma_{p(\rho)} & \\ \mathcal{I} & \end{matrix}. \quad (\text{G20})$$

Since $\mathbf{c}$ is an emergent cycle, we can decompose it as $\mathbf{c} = k_1 \mathbf{c}_{\sigma_r(\rho)} + k_2 \mathbf{c}_{\sigma_{r'}(\rho)} + k_3 \mathbf{c}_{\sigma_p(\rho)}$. Then, we must have

$$\nabla^Y \cdot \mathbf{c} = \begin{matrix} \sigma_r(\rho) \\ \sigma_{r'}(\rho) \\ \sigma_p(\rho) \\ \mathcal{I} \end{matrix} \begin{bmatrix} -1 \\ -1 \\ 1 \\ 0 \end{bmatrix}, \quad (\text{G21})$$

$$= k_1 \nabla^Y \cdot \mathbf{c}_{\sigma_r(\rho)} + k_2 \nabla^Y \cdot \mathbf{c}_{\sigma_{r'}(\rho)} + k_3 \nabla^Y \cdot \mathbf{c}_{\sigma_p(\rho)}.$$

However, from Eq. (G18), the form of the vectors $\nabla^Y \cdot \mathbf{c}_{\sigma_r(\rho)}$ is given by


$$\nabla^Y \cdot \mathbf{c}_{\sigma_r(\rho)} = \begin{matrix} \sigma_r(\rho) \\ \sigma_{r'}(\rho) \\ \sigma_p(\rho) \\ \mathcal{I} \end{matrix} \begin{bmatrix} 1 \\ 0 \\ 0 \\ -\nu_{\sigma_r(\rho),1} \\ \vdots \\ -\nu_{\sigma_r(\rho),i} \end{bmatrix}. \quad (\text{G22})$$

Substituting Eq. (G22) in Eq. (G21), we find that $k_1 = -1, k_2 = -1, k_3 = 1$, which directly also implies $\nabla^Y \cdot \mathbf{c}_{\sigma_r(\rho)} + \nabla^Y \cdot \mathbf{c}_{\sigma_{r'}(\rho)} = \nabla^Y \cdot \mathbf{c}_{\sigma_p(\rho)}$ as needed.

f. Example

Here, we show an example to illustrate the results of the subsection.

Example a: consider the following CRN, generated randomly with $N_A = 1$ and $Q = 10$:



with eight species and six reactions. The stoichiometric matrix of the CRN above is given by

$$\nabla = \begin{matrix} & \rho_1 & \rho_2 & \rho_3 & \rho_4 & \rho_5 & \rho_6 \\ \begin{matrix} (1) \\ (4) \\ (5) \\ (3) \\ (7) \\ (8) \\ (9) \\ (10) \end{matrix} & \begin{bmatrix} -1 & 0 & -1 & -1 & -1 & 0 \\ -1 & -1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & -2 \\ 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & 1 & -1 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \end{bmatrix} \end{matrix}. \quad (\text{G24})$$

Note that the rank of the stoichiometric matrix is six implying that there are two conservation laws and zero cycles. Furthermore, it is seen from the stoichiometric matrix that there are three root species (1), (3), and (4), respectively. Taking the root species as the Y species, from Eq. (G24), we get the substoichiometric matrix ∇^I (see Appendix G 3 1) to be

$$\nabla^I = \begin{matrix} & \rho_1 & \rho_2 & \rho_3 & \rho_4 & \rho_5 & \rho_6 \\ \begin{matrix} (5) \\ (7) \\ (8) \\ (9) \\ (10) \end{matrix} & \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & -2 \\ 0 & 1 & -1 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \end{bmatrix} \end{matrix}. \quad (\text{G25})$$

First, we note that the left nullspace is trivial, as expected from Appendix G 3 1. Second, we find the emergent cycle $\mathbf{c}_r = (-2, 1, 1, 1, 1, -1)^T$. This immediately verifies Eq. (51) as there are three roots and one emergent cycle, i.e., two independent roots. The net reaction connecting the roots is given by

$$\nabla^Y \cdot \mathbf{c}_r = \begin{matrix} (1) \\ (4) \\ (3) \end{matrix} \begin{bmatrix} -1 \\ 1 \\ -1 \end{bmatrix}, \quad (\text{G26})$$

corresponding to the net reaction $(4) \rightleftharpoons (1) + (3)$. Considering the independent roots as $\mathcal{I} = \{(1), (3)\}$, we now start constructing the representation of different species as outlined in Appendix G 3 4. Assume we want to find the net reaction constructing the species (7). Taking the Y species as $\{(1), (3), (7)\}$, we find the substoichiometric matrix ∇^X from Eq. (G24) to be of the form

$$\nabla^X = \begin{matrix} & \rho_1 & \rho_2 & \rho_3 & \rho_4 & \rho_5 & \rho_6 \\ \begin{matrix} (4) \\ (5) \\ (8) \\ (9) \\ (10) \end{matrix} & \begin{bmatrix} -1 & -1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & -2 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \end{bmatrix} \end{matrix}, \quad (\text{G27})$$

with the emergent cycle $\mathbf{c}' = (2, -2, -1, -1, -1, 1)^T$. The corresponding net reaction is seen to be

$$\nabla^Y \cdot \mathbf{c}' = \begin{matrix} (1) \\ (3) \\ (7) \end{matrix} \begin{bmatrix} 1 \\ 2 \\ -1 \end{bmatrix}, \quad (\text{G28})$$

corresponding to the net reaction $(7) \rightleftharpoons (1) + 2(3)$. We emphasize again that if we had represented the species $\{(1), (3), (4), (5), (7) \dots\}$ by the labels $\{A, B, C, D, E \dots\}$, the above-mentioned net reaction corresponds to stating that the species E can be written as AB_2 , which is nontrivial. Finally, following Appendix G 3 5, we derive the basis of conservation laws,

$$\ell^1 = \begin{matrix} (1) & (4) & (5) & (3) & (7) & (8) & (9) & (10) \\ [& 1 & 1 & 2 & 0 & 1 & 2 & 3 & 4 &], \end{matrix} \quad (\text{G29})$$

$$\ell^2 = \begin{matrix} (1) & (4) & (5) & (3) & (7) & (8) & (9) & (10) \\ [& 0 & 1 & 1 & 1 & 2 & 2 & 2 & 2 &]. \end{matrix} \quad (\text{G30})$$

corresponding to the conservation of the moieties (1) and (3).

APPENDIX H: SIMULATIONS FOR REACHABILITY

We first outline the method to determine the forward reachable set of a species given a stoichiometric matrix ∇ . Let S be a set of species, initially containing only the species σ whose reachable set we wish to compute. We then loop over all reactions: if all substrates of a reaction ρ are in S , we add all the products of ρ to S and then do the same for the reverse reaction. This process is iterated until S no longer changes. The final set S is defined as the forward reachable set $\mathcal{F}(\sigma)$. This procedure is repeated for each species σ , and the largest reachable set is recorded. The resulting order parameter is the rescaled size of the largest reachable set, $|f|$. Note that this procedure depends only on the stoichiometric matrix and is independent of the number of atoms.

1. One atom

We generated ensembles for $Q = 100, 200, \dots, 1400$, but restricted analysis to $Q \geq 500$. For each value of Q , we generated CRNs for $\gamma \in [2.0, 4.0]$ in steps of 0.01. Each γ - Q pair had 39 000 samples for $Q \leq 1200$ and 25 000 samples for $Q = 1300, 1400$. The sample mean of $|f|$ is plotted with standard errors in Fig. 12, for selected values of Q .

To estimate γ_c , we analyzed histograms of $|f|$ for $\gamma \in [2.5, 3.0]$. For each value of γ , the histogram was made using 201 bins between 0 and 1 (via histogram method in NumPy). We located the histogram peaks (maximum counts) and then found the histogram valley (minimum counts) between the peaks. The probability weight assigned to the left peak was the area from the leftmost bin, i.e., zero until the histogram valley. We plot the left probability weight, denoted $w_-(\gamma, Q)$ as a function of γ in Fig. 25 for different Q . We then interpolated between the two closest γ values where w_- crosses 0.5 to obtain $\gamma_{1/2}(Q)$, the equal-weight probability.

Uncertainty in $\gamma_{1/2}(Q)$ arises from both histogram binning and the discretization of γ . To estimate binning uncertainty, we repeated the procedure for 181, 191, 201, 211, 221 bins and computed the standard deviation of results. To estimate γ_c , we linearly fit $\gamma_{1/2}(Q)$ (averaged over bin counts) vs $1/Q$. The fit yields $\gamma_c = 2.921 \pm 0.008$




FIG. 25. Probability weight $w_-(\gamma, Q)$ of the left peak of the order parameter histogram plotted against γ for different Q . The parameters are $N_A = 1$ and Q . The horizontal dashed line is 0.5, which represents the equal weight point.




FIG. 26. Location of the left (resp. right) peak of the order parameter histogram, denoted $|f|_-(Q)$ (yellow stars) [resp. $|f|_+(Q)$ (blue dots)] plotted against Q . Note that the two y axes start at different values. The peaks are evaluated at $\gamma \approx \gamma_{1/2}(Q)$. The parameters are $N_A = 1$ and Q . The difference is indicated by the green triangles.

with $r^2 \approx 0.94$. Taking into account the γ discretization, we report $\gamma_c = 2.92 \pm 0.01$.

For each Q , we selected the γ value closest to $\gamma_{1/2}(Q)$ and extracted the locations of the two histogram peaks, denoted $|f|_{\pm}(Q)$. The peak locations and their separation are shown in Fig. 26. We observe that $|f|_+(Q)$ increases and appears to saturate, while $|f|_-(Q)$ decreases, confirming separation between the two regimes.

To confirm the discontinuous nature of the transition, we computed the derivative $\partial(|f|)/\partial\gamma$ using the data in Fig. 12. We first smoothed the data using the `gaussianfilter1d` routine from `scipy.ndimage` with a smoothing parameter $\sigma = 4.0$. The percent difference due to smoothing is shown in Fig. 27 and does not exceed four percent. The peak height of the derivative, as a function of Q , is shown in Fig. 28(a) and fits well to a power law with exponent 0.35. Figure 28(b) shows the extrapolation of the location of




FIG. 27. Data from Fig. 12 smoothed by a Gaussian Filter and the ratio of the difference between the smoothed and original data and the original data (percentage change) plotted as a function of γ . The parameters are $N_A = 1$ and $\sigma = 4$ for the filtering.




FIG. 28. (a) Peak of the derivative of the average reachability wrt γ (data from Fig. 12) plotted as a function of $1/Q$ in the log-log scale. The red line indicated a linear fit in the log-log scale. (b) Location of the peak of the derivative of the average reachability, denoted $\gamma_r^*(Q)$ wrt γ (data from Fig. 12) plotted as a function of $1/Q$. The red line represents a linear fit. For both graphs, the parameters are $N_A = 1$.




FIG. 29. Fluctuations of the average reachability shown in Fig. 12 labeled $\tilde{\chi}(\gamma, Q)$ plotted as a function of γ . The parameters are $N_A = 1$. The red line represents the estimated location of the percolation threshold.




FIG. 30. (a) Peak of the fluctuations of the reachability (data from Fig. 29) plotted as a function of $1/Q$ in the log-log scale. The red line represents a linear fit in the log-log scale. (b) Location of the peak of $\tilde{\chi}(\gamma, Q)$, denoted $\gamma_s^*(Q)$ (data from Fig. 29) plotted as a function of $1/Q$. The red line represents a linear fit. For both graphs, the parameters are $N_A = 1$.

the derivative peak (denoted $\gamma_r^*(Q)$) to $Q \rightarrow \infty$, using a linear fit in $1/Q$. The intercept is 2.92 ± 0.02 , consistent with the estimate from histogram analysis ($r^2 \approx 0.93$). We define the fluctuations of the order parameter as $\tilde{\chi}(\gamma, Q) = \langle |f|^2 \rangle - \langle |f| \rangle^2$. The behavior of $\tilde{\chi}(\gamma, Q)$ is shown in Fig. 29. As Q increases, the peak in $\tilde{\chi}$ becomes sharper and shifts toward γ_c . The peak height scales with Q as a power law with exponent 0.52, as shown in Fig. 30(a). Figure 30(b) shows that the location of the fluctuation peak [denoted $\gamma_s^*(Q)$] also extrapolates to $\gamma_c = 2.94 \pm 0.02$, consistent with the other estimates. The extrapolation is a linear fit in $1/Q$.

REFERENCES

- C. M. Dobson, *Nature* **432**, 824 (2004).
- J.-L. Reymond, *Acc. Chem. Res.* **48**, 722 (2015).
- S. Stocker, G. Csányi, K. Reuter, and J. T. Margraf, *Nat. Commun.* **11**, 5505 (2020).
- P. M. Murray, S. N. G. Tyler, and J. D. Moseley, *Org. Process Res. Dev.* **17**, 40 (2013).
- B. A. Grzybowski, K. J. M. Bishop, B. Kowalczyk, and C. E. Wilmer, *Nat. Chem.* **1**, 31 (2009).
- R. Albert and A.-L. Barabási, *Rev. Mod. Phys.* **74**, 47 (2002).
- M. E. J. Newman, *SIAM Rev.* **45**, 167 (2003).
- M. Newman, *Networks*, 2nd ed. (Oxford University Press, Oxford, New York, 2018).
- T. Guhr, A. MüllerGroeling, and H. A. Weidenmüller, *Phys. Rep.* **299**, 189 (1998).
- H. Jeong, B. Tombor, R. Albert, Z. N. Oltvai, and A.-L. Barabási, *Nature* **407**, 651 (2000).
- R. Tanaka, *Phys. Rev. Lett.* **94**, 168101 (2005).
- B. Ø. Palsson, *Systems Biology: Properties of Reconstructed Networks*, reprinted ed. (Cambridge University Press, Cambridge, 2008).
- H. B. Smith, H. Kim, and S. I. Walker, *Sci. Rep.* **11**, 6542 (2021).
- S. Klamt, U.-U. Haus, and F. Theis, *PLoS Comput. Biol.* **5**, e1000385 (2009).
- P. Holme, *J. R. Soc. Interface* **6**, 1027 (2009).
- C. I. Sandefur, M. Mincheva, and S. Schnell, *Mol. BioSyst.* **9**, 2189 (2013).
- R. Montañez, M. A. Medina, R. V. Solé, and C. Rodríguez-Caso, *BioEssays* **32**, 246 (2010).
- R. Fagerberg, C. Flamm, D. Merkle, P. Peters, and P. F. Stadler, *Math. Comput. Sci.* **7**, 275 (2013).
- M. Barthélemy, *Phys. Rev. E* **106**, 064310 (2022).
- G. Ghoshal, V. Zlatić, G. Caldarelli, and M. E. J. Newman, *Phys. Rev. E* **79**, 066118 (2009).
- G. Bianconi and S. N. Dorogovtsev, *Phys. Rev. E* **109**, 014306 (2024).
- S. Müller, C. Flamm, and P. F. Stadler, *J. Cheminf.* **14**, 1 (2022).
- D. F. Anderson and T. D. Nguyen, *J. Math. Chem.* **59**, 2063 (2021).
- D. F. Anderson and T. D. Nguyen, *J. Appl. Probab.* **59**, 384 (2022).
- E. Bigan, J.-M. Steyaert, and S. Douady, "Properties of random complex chemical reaction networks and their relevance to biological toy models," [arXiv:1303.7439](https://arxiv.org/abs/1303.7439) [q-bio.MN] (2013).
- Z. G. Nicolaou, S. B. Nicholson, A. E. Motter, and J. R. Green, *J. Chem. Phys.* **158**, 225101 (2023).
- G. Svehla, *Pure Appl. Chem.* **65**, 2291 (1993).
- C. D. Meyer, *Matrix Analysis and Applied Linear Algebra* (Society for Industrial and Applied Mathematics, Philadelphia, 2000).
- R. Rao and M. Esposito, *Phys. Rev. X* **6**, 041064 (2016).
- F. Avanzini, G. Falasco, and M. Esposito, *New J. Phys.* **22**, 093040 (2020).
- R. Rao and M. Esposito, *J. Chem. Phys.* **149**, 245101 (2018).
- M. Shimada, P. Behrad, and E. De Giuli, *Phys. Rev. E* **109**, 044105 (2024).
- M. Polettini and M. Esposito, *J. Chem. Phys.* **141**, 024117 (2014).
- F. Avanzini, N. Freitas, and M. Esposito, *Phys. Rev. X* **13**, 021041 (2023).

- ³⁵S. Dal Cengio, V. Lecomte, and M. Polettini, *Phys. Rev. X* **13**, 021040 (2023).
- ³⁶G. Shinar and M. Feinberg, *Math. Biosci.* **241**, 1 (2013).
- ³⁷G. Craciun and M. Feinberg, *SIAM J. Appl. Math.* **66**, 1321 (2006).
- ³⁸M. Feinberg, *Foundations of Chemical Reaction Network Theory* (Springer, 2019), pp. 79–80.
- ³⁹H.-M. Kaltenbach, in *Electronic Notes in Theoretical Computer Science Proceedings of the SASB 2018, the Ninth International Workshop on Static Analysis and Systems Biology, Freiburg, Germany–August 28th, 2018* (Elsevier, 2020), p. 350.
- ⁴⁰D. F. Anderson and T. G. Kurtz, *Stochastic Analysis of Biochemical Systems* (Springer International Publishing, Cham, 2015).
- ⁴¹M. Polettini, A. Wachtel, and M. Esposito, *J. Chem. Phys.* **143**, 184103 (2015).
- ⁴²S. G. Marehalli Srinivas, M. Polettini, M. Esposito, and F. Avanzini, *J. Chem. Phys.* **158**, 204108 (2023).
- ⁴³T. Handorf, O. Ebenhöf, and R. Heinrich, *J. Mol. Evol.* **61**, 498 (2005).
- ⁴⁴S. N. Dorogovtsev and J. F. F. Mendes, *The Nature of Complex Networks*, 1st ed. (Oxford University Press, Oxford, 2022).
- ⁴⁵W. Tang and F. Tang, *Stat. Sci.* **38**, 108 (2023).
- ⁴⁶D. Stauffer and A. Aharony, *Introduction to Percolation Theory: Second Edition*, 2nd ed. (Taylor & Francis, London, 2017).
- ⁴⁷J. Houdayer and A. Hartmann, *Phys. Rev. B* **70**, 014418 (2004).
- ⁴⁸O. Melchert, “autoScale.py—A program for automatic finite-size scaling analyses: A user’s guide,” [arXiv:0910.5403](https://arxiv.org/abs/0910.5403) [physics] (2009).
- ⁴⁹M. E. J. Newman, S. H. Strogatz, and D. J. Watts, *Phys. Rev. E* **64**, 026118 (2001).
- ⁵⁰S. N. Dorogovtsev, J. F. F. Mendes, and A. N. Samukhin, *Phys. Rev. E* **64**, 025101 (2001).
- ⁵¹K. Binder, K. Vollmayr, H.-P. Deutsch, J. D. Reger, M. Scheucher, and D. P. Landau, *Int. J. Mod. Phys. C* **03**, 1025 (1992).
- ⁵²M. V. Smoluchowski, *Ann. Phys.* **353**, 1103 (1916).
- ⁵³A. Sharko, D. Livitz, S. De Piccoli, K. J. M. Bishop, and T. M. Hermans, *Chem. Rev.* **122**, 11759 (2022).
- ⁵⁴J. A. D. Wattis, *Physica D* **222**, 1 (2006).
- ⁵⁵W. J. Riehl, P. L. Krapivsky, S. Redner, and D. Segrè, *PLoS Comput. Biol.* **6**, e1000725 (2010).
- ⁵⁶P. Gagrani, V. Blanco, E. Smith, and D. Baum, *J. Math. Chem.* **62**, 1012 (2024).
- ⁵⁷S. A. Kauffman, *J. Theor. Biol.* **119**, 1 (1986).
- ⁵⁸R. J. Bagley and J. D. Farmer, in *Artificial Life II*, edited by C. Langton, C. Taylor, J. Farmer, and S. Rasmussen (Addison-Wesley, Reading, Mass., 1994), Chap. 2, pp. 93–141.
- ⁵⁹M. Steel, *Appl. Math. Lett.* **13**, 91 (2000).
- ⁶⁰A. Filisetti, M. Villani, C. Damiani, A. Graudenzi, A. Roli, W. Hordijk, and R. Serra, “Advances in Artificial Life and Evolutionary Computation,” in *Communications in Computer and Information Science*, edited by C. Pizzuti and G. Spezzano (Springer International Publishing, Cham, 2014), pp. 113–126.
- ⁶¹Y. Himeoka, J. B. Kirkegaard, N. Mitarai, and S. Krishna, *Sci. Rep.* **14**, 22187 (2024).
- ⁶²A. Garcia-Chung, M. Bermúdez-Montaña, P. F. Stadler, J. Jost, and G. Restrepo, *J. Math. Chem.* **62**, 1357 (2024).
- ⁶³J. Fischer, A. Kleidon, and P. Dittrich, *PLoS One* **10**, e0117312 (2015).
- ⁶⁴C. Furusawa and K. Kaneko, *Phys. Rev. Lett.* **90**, 088102 (2003).
- ⁶⁵Y. Kondo and K. Kaneko, *Phys. Rev. E* **84**, 011927 (2011).
- ⁶⁶N. Araújo, P. Grassberger, B. Kahng, K. Schrenk, and R. Ziff, *Eur. Phys. J. Special Top.* **223**, 2307 (2014).
- ⁶⁷W. A. Warr, *WIREs Comput. Mol. Sci.* **1**, 557 (2011).
- ⁶⁸D. S. Wigh, J. M. Goodman, and A. A. Lapkin, *WIREs Comput. Mol. Sci.* **12**, e1603 (2022).
- ⁶⁹P. Dittrich, J. Ziegler, and W. Banzhaf, *Artif. Life* **7**, 225 (2001).
- ⁷⁰G. Benkő, C. Flamm, and P. F. Stadler, *J. Chem. Inf. Comput. Sci.* **43**, 1085 (2003).
- ⁷¹G. Benkő, C. Flamm, and P. F. Stadler, *Advances in Artificial Life* (Springer, Berlin, Heidelberg, 2003), pp. 10–19.
- ⁷²J. L. Andersen, C. Flamm, D. Merkle, and P. F. Stadler, *J. Syst. Chem.* **4**, 4 (2013).
- ⁷³J. L. Andersen, C. Flamm, D. Merkle, and P. F. Stadler, *Graph Transformation, Lecture Notes in Computer Science* (Springer International Publishing, 2016), pp. 73–88.
- ⁷⁴A. Arya, J. Ray, S. Sharma, R. Cruz Simbron, A. Lozano, H. B. Smith, J. L. Andersen, H. Chen, M. Meringer, and H. J. Cleaves, *Chem. Sci.* **13**, 4838 (2022).
- ⁷⁵J. Zanghellini, D. E. Ruckerbauer, M. Hanscho, and C. Jungreuthmayer, *Biotechnol. J.* **8**, 1009 (2013).
- ⁷⁶S. Schuster, C. Hilgetag, J. H. Woods, and D. A. Fell, *J. Math. Biol.* **45**, 153 (2002).
- ⁷⁷R. M. May, *Nature* **238**, 413 (1972).
- ⁷⁸W. Cui, R. Marsland, and P. Mehta, *Phys. Rev. E* **104**, 034416 (2021).
- ⁷⁹T. Aslyamov and M. Esposito, *Phys. Rev. Lett.* **133**, 107103 (2024).
- ⁸⁰S. Varrette, H. Cartiaux, S. Peter, E. Kieffer, T. Valette, and A. Olloh, in *Proceedings of the 6th ACM High Performance Computing and Cluster Technologies Conf. (HPCCT 2022)* (Association for Computing Machinery (ACM), Fuzhou, China, 2022).
- ⁸¹B. Bollobás, *Random Graphs, Cambridge Studies in Advanced Mathematics No. 73*, 2nd ed. (Cambridge University Press, Cambridge; New York, 2001).
- ⁸²A. Engel, R. Monasson, and A. K. Hartmann, *J. Stat. Phys.* **117**, 387 (2004).
- ⁸³K. Christensen, R. Donangelo, B. Koiller, and K. Sneppen, *Phys. Rev. Lett.* **81**, 2380 (1998).
- ⁸⁴*Handbook of Mathematical Functions: With Formulas, Graphs, and Mathematical Tables*, *Dover Books on Mathematics*, edited by M. Abramowitz and I. A. Stegun (Dover Publishers, New York, 2012).
- ⁸⁵T. A. Davis, “Algorithm 915, SuiteSparseQR: Multifrontal multithreaded rank-revealing sparse QR factorization,” *ACM Trans. Math. Software* **38**(1), 1–22 (2011).
- ⁸⁶C. G. West, A. Garcia-Saez, and T.-C. Wei, *Phys. Rev. B* **92**, 115103 (2015).