

Intrinsic information carriers in combinatorial dynamical systems

Russ Harmer,^{1,2} Vincent Danos,³ Jérôme Feret,⁴ Jean Krivine,² and Walter Fontana¹

¹Harvard Medical School, Boston, Massachusetts 02115, USA

²CNRS and Université Paris Diderot, F-75205 Paris Cedex 13, France

³University of Edinburgh, Edinburgh EH8 9AB, United Kingdom

⁴INRIA-École normale supérieure-CNRS, F-75230 Paris Cedex 5, France

(Received 30 June 2010; accepted 31 August 2010; published online 28 September 2010)

Many proteins are composed of structural and chemical features—“sites” for short—characterized by definite interaction capabilities, such as noncovalent binding or covalent modification of other proteins. This modularity allows for varying degrees of independence, as the behavior of a site might be controlled by the state of some but not all sites of the ambient protein. Independence quickly generates a startling combinatorial complexity that shapes most biological networks, such as mammalian signaling systems, and effectively prevents their study in terms of kinetic equations—unless the complexity is radically trimmed. Yet, if combinatorial complexity is key to the system’s behavior, eliminating it will prevent, not facilitate, understanding. A more adequate representation of a combinatorial system is provided by a graph-based framework of rewrite rules where each rule specifies only the information that an interaction mechanism depends on. Unlike reactions, which deal with molecular species, rules deal with patterns, i.e., multisets of molecular species. Although the stochastic dynamics induced by a collection of rules on a mixture of molecules can be simulated, it appears useful to capture the system’s average or deterministic behavior by means of differential equations. However, expansion of the rules into kinetic equations at the level of molecular species is not only impractical, but conceptually indefensible. If rules describe bona fide patterns of interaction, molecular species are unlikely to constitute appropriate units of dynamics. Rather, we must seek aggregate variables reflective of the causal structure laid down by the rules. We call these variables “fragments” and the process of identifying them “fragmentation.” Ideally, fragments are aspects of the system’s microscopic population that the set of rules can actually distinguish on average; in practice, it may only be feasible to identify an approximation to this. Most importantly, fragments are self-consistent descriptors of system dynamics in that their time-evolution is governed by a closed system of kinetic equations. Taken together, fragments are endogenous distinctions that matter for the dynamics of a system, which warrants viewing them as the carriers of information. Although fragments can be thought of as multisets of molecular species (an extensional view), their self-consistency suggests treating them as autonomous aspects cut off from their microscopic realization (an intensional view). Fragmentation is a seeded process that depends on the choice of observables whose dynamics one insists to describe. Different observables can cause distinct fragmentations, in effect altering the set of information carriers that govern the behavior of a system, even though nothing has changed in its microscopic constitution. In this contribution, we present a mathematical specification of fragments, but not an algorithmic implementation. We have described the latter elsewhere in rather technical terms that, although effective, were lacking an embedding into a more general conceptual framework, which we here provide.

© 2010 American Institute of Physics. [doi:10.1063/1.3491100]

Central to the rise of modern chemistry was the definition of a formal language for expressing the modular architecture of organic molecules and the rules of reaction between their constituent atoms and functional groups. In molecular biology, a similar modularity characterizes many of the proteins participating in molecular interaction networks that underlie cellular behavior. These proteins can be viewed as consisting of “sites” that abstractly represent specific capabilities of interaction, such as binding or modifying other proteins. Interactions between sites occur, to varying extents, independently of other sites, giving rise to a vast number of potential interactions that cannot be tracked by standard chemical kinetics, as

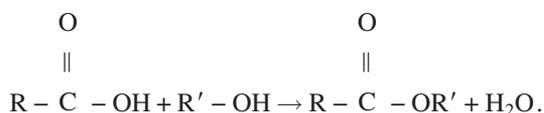
the latter requires an explicit list of all possibilities. Yet, such systems are compactly described in a rule-based format that keeps these vast possibilities implicit by only mentioning those aspects of molecules that mechanisms are known (or hypothesized) to care about. We show that such a system of rules admits a corresponding deterministic dynamical system cast in terms of coarse-grained variables. These new variables, which we call “fragments,” are entirely determined by static analysis of the rules. Fragments constitute the effective information carriers of a system in that they are those features that the system of rules can collectively distinguish on average.

I. INTRODUCTION

A. Rules: A linguistic turn

New instruments, new experimental methods, and a new language were among the tools that removed the intellectual obstruction checking the progress of chemistry in the 18th century.^{1,2} This new language became necessary, not only to remove ambiguities and anachronisms, but also to accommodate the increasing number of newly discovered elements and their combinations. In reforming chemical language to reflect more systematically the compositional nature of compounds, Lavoisier³ was inspired by a powerful idea—due to Étienne Bonnot, Abbé de Condillac—that “languages are true analytical methods.” Here we sketch the transposition of this idea to molecular biology and discuss a resultant change in perspective on the dynamics of complex interaction networks.

During the process that transformed alchemy into chemistry, the naming of compounds evolved from proper nouns, as in Fig. 1(a), to systematic naming schemes (1b) and symbolic expressions at different levels of granularity [(1c)–(1g)]. Chemical notation often emphasizes the compositional structure of molecules, not only in terms of atoms but also functional groups [(1d) and (1f)]. These are sets of atoms distinguished by characteristic reactions with other such groups, reflecting modular mechanisms that produce chemical changes local to each group. For example, a reaction of alanine, $\text{CH}_3\text{CHNH}_2\text{COOH}$, with methanol, CH_3OH , may affect only their carboxy and alcohol groups, respectively. This is an instance of a general schema according to which these groups interact in a specific manner regardless of the wider molecular context. A schema is expressed by a rule that makes explicit *only* that upon which the mechanism actually depends,



A single reaction schema compactly represents an infinity of possible reactions, depending on how the placeholders R and R' are instantiated. Although a chemical transformation is typically localized to reacting functional groups, the behavior of the latter can be controlled by other groups present in the same molecule. Thus, a bulky ligand of the

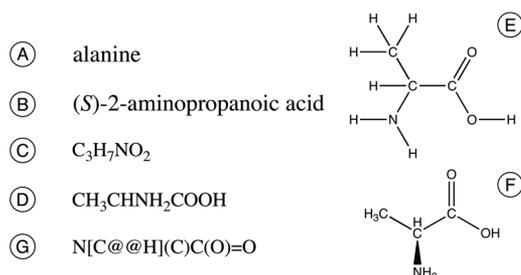


FIG. 1. Chemical names of different granularity: (a) proper noun, (b) systematic name, (c) empirical (sum) formula, (d) structural formula, (e) connectivity formula, (f) stereo formula, and (g) simplified molecular input line entry system string, which expresses the same information as (f).

nitrogen in Fig. 1(f) might affect the velocity and repertoire of reactions available to the COOH group. To account for such dependencies, a reaction schema would have to be refined into subschemata.

Many proteins contain modules with characteristic local interaction capabilities, much like functional groups in organic chemistry. These generally consist of either structurally autonomous domains that bind specific epitopes on other proteins, or short peptide motifs that undergo covalent modifications, such as phosphorylation and dephosphorylation. We shall refer to these loci of action as sites. The modular character of protein-protein interactions justifies a representation in terms of rules analogous to reaction schemata in chemistry. Because proteins are typically large molecular objects, their basic identity is not altered by these (reversible) interactions. Thus, rather than thinking of proteins as undergoing chemical transformations, biochemists think of them as undergoing state changes at their sites.

A language for expressing rules of protein-protein interaction will therefore treat proteins formally like “atoms” and protein complexes like “molecules.” As in organic chemistry, a protein-protein interaction rule only specifies which changes occur and the sites on which these changes depend; nothing else is mentioned at all. However, unlike in chemistry, such a rule is not necessarily informed by any *theory* of the mechanism being described: it is purely descriptive—empirical or hypothetical in origin. Indeed, the *formal* nature of rules gives us license to write whatever mechanisms we please, irrespective of physical plausibility or even possibility. Rules need not talk about physical events, such as electronic rearrangements; instead, they operate at a higher level of abstraction where formal “bonds” can be formed, or broken, without need for explicating what constitutes the bond physically.

In this way, rules relieve us from the need for a complete, microscopic understanding of what makes an interaction between macromolecules possible. Instead, we can directly represent the observed consequences of incompletely understood mechanisms, i.e., empirical knowledge, as rules operating at this more formal level. Clearly, having theoretical knowledge does not disbar us from using rules to represent it—we certainly could represent organic reaction schemata with rules—but side-stepping its necessity greatly expands the scope of what we can describe and even allows us to mix theoretical mechanisms with purely empirical knowledge.^{4,5} A rule is thus, at its most general, a *formal mechanistic hypothesis*, perhaps at least partially grounded in theoretical understanding—but forever subject to refinement and revision in the light of new data or improved theoretical explanation.

B. Dynamic consequences of combinatorial complexity

This formal representation of molecules uncovers an uncomfortable gap between the rules that express the *chemistry* and the equations—deterministic or stochastic—used to express the chemical *kinetics* (i.e., the changes in the abundance of molecular species resulting from their chemistry). While rules reference the structure of molecules, kinetic

equations are stuck in the alchemical naming scheme of proper nouns, as every possible molecular species must be assigned a unique variable *at the outset*. Although names of variables may be structured, this has no formal significance; we can rename variables arbitrarily, as long as we do so consistently.

Because rules describe mechanisms of a local nature, they generally only weakly constrain the overall states of the agents they apply to. As a consequence, they can elegantly describe systems—such as those often found in molecular biology^{6–8}—with very large state spaces: if a rule only tests one site out of ten, each of which has two possible binding partners, it simply does not care which of the $3^9=19\,683$ possible states the agent is in, and even a small number of such rules rapidly give rise to a system with astronomical state space. The traditional extensional view of kinetics therefore prevents us from studying the dynamic consequences of this kind of *combinatorial complexity*, simply because we cannot even write down the system, let alone integrate it numerically.

An intensional solution to this could be to forgo an explicit kinetic description and simulate the stochastic dynamics induced by a set of rules on an initial mixture. This approach, while still in its early stages, permits significant forays into the dynamics of combinatorially complex interaction networks.^{4,5,9,10} However, such direct simulation does not provide any general insight for determining the critical variables that shape the dynamics. In identifying these variables of a complex system, we are presumably seeking *in a principled way* its “carriers of information”—a set of variables that constitutes the effective units of dynamics; something that a given set of rules collectively observes. If the idea of a rule-based representation of interactions is warranted—as argued in this section—and given that the whole point of a rule is to ignore everything known (or hypothesized) to be irrelevant to the mechanism it describes, it stands to reason that most molecular species (the objects on which reactions are defined) are not meaningful units of dynamics: some species are indistinguishable simply because the system of rules is not capable of telling them apart. This raises the question whether it is possible to formally *derive* a sound system of units and their kinetic equations directly from the set of rules.

We should stress that such a coarse-graining is determined entirely from the rule set and, although it generally results in a much-reduced set of variables, it should be distinguished from a practice of model simplification that consists in aggregating molecular species from the outset, i.e., in the very description of the model. Such an approach risks missing crucial insight by simply not incorporating sufficient complexity in our description: we cannot simplify in a principled manner if we have not represented the system in full complexity in the first place. There do exist formal and numeric model reduction techniques^{7,11,12} that exploit, for example, separation of time scales and conservation constraints. However, these techniques are near powerless when it comes to combinatorially complex systems, since they require the explicit (and thus unattainable) system of kinetic equations as input. Feasibility aside, our goal fundamentally

differs from these techniques by seeking units of dynamics that are grounded in mechanism. (In fact, once such a system of units has been identified, classical reduction techniques might actually become feasible.)

In a situation where mechanisms are highly local and molecular species so numerous, an *intrinsic coarse-graining* of the state space with respect to the actual observational capabilities of the rules is therefore not only a pragmatic boon but a conceptual necessity: the *information* in a state space of astronomical size must surely be elsewhere than in the microscopic states.

C. Structure of the paper

In this paper, we investigate the requirements for such a coarse-graining to be valid. While certain aspects of this discussion are generic, we focus primarily on the coarse-graining of the *deterministic* semantics of rules, i.e., on average time-evolution. We introduce the concepts of “fragment”—a partially specified molecular species—and “fragmentation”—the process of identifying a *self-consistent* set of fragments for coarse-grained dynamics, meaning that the (average) time-evolution of a fragment depends only on other fragments. We also discuss the interpretation of (sets of) fragments as the information carriers of dynamical systems in the sense that they expose, even for highly combinatorial networks where extensional intuition and technique breaks down, what their dynamics is about.

II. RULE-BASED SYSTEMS: THE BASICS

Two sustained efforts at a rule-based framework for molecular biology have taken shape independently over the past few years, resulting in comparable languages. The biological network generator language (BNGL) originated among computational biologists,^{4,6,13,14} while Kappa (or κ -language) hatched from computer scientists.^{5,9,15,16} The convergence of these efforts onto the same level of abstraction and a similar syntax is reassuring. In this section, we review the mathematical framework of Kappa, including a fully detailed account of how one passes from rules to structureless reactions, which is the reference level for conventional representations of chemical kinetics. In Sec. II C we briefly discuss differences between Kappa and BNGL with a view on how they impact fragmentation.

In Kappa, the notion of “agent” refers to a basic unit (e.g., a protein) that can be combined into well-formed complexes according to a grammar. Such a grammar will be relatively unconstrained, since proteins do not possess “valences” determined by fundamental laws. The interaction capabilities of protein-agents are described by hypothetical or empirically grounded rules. Such a level of abstraction emphasizes *process* rather than structural detail, although the latter deeply informs the former.

A. Site graphs

The basic components of rules are graphs (Fig. 2), more precisely *site graphs*, as defined by the following data:

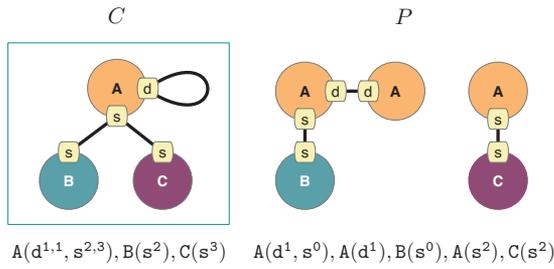


FIG. 2. (Color) Site graphs. A contact map C (left) is a site graph that exhibits all agent types, their full complement of sites, and possible binding interactions, as determined by a set of rules. A pattern P (right) is a site graph where each site engages in at most one bond. These graphs can also be expressed using an equivalent textual notation. The simple grammar defining such graphs forms a language known as Kappa (Refs. 5 and 15). Each agent (node) is represented by its name followed, in parentheses, by an unordered list of sites exposed by the agent. Agents are separated by commas to indicate that they can be shuffled without altering the underlying graph. Bonds are identified by a superscript at their end points. The label, usually a number, is arbitrary and appears exactly twice. In a pattern, like P , each site can have at most one superscript, as it can participate in at most one bond; in a contact map, like C , a site may have more than one superscript, because it may be the end point of several possible bonds. A self-loop on a site, such as d of A in C , indicates that two agents of the same type can bind one another on the same site. A contact map is meant to summarize the binding possibilities specified by a set of rules, while a pattern is meant to define a realization of these possibilities.

- (1) a finite set V of nodes (or agents);
- (2) a function $\lambda: V \rightarrow \mathcal{A}$ assigning a name to each node;
- (3) a function $\sigma: V \rightarrow \mathcal{P}(\mathcal{S})$ assigning a set of sites to each node, where $\mathcal{P}(\cdot)$ denotes the power set;
- (4) a symmetric relation π on $\sum_{v \in V} \sigma(v)$ saying which sites are connected by edges.

Site graphs fall into two major classes: contact maps and patterns (Fig. 2).

A *contact map* is characterized by every node having a different name, i.e., the function λ is injective. A *pattern* is characterized by their being at most one edge emanating from every site, i.e., the relation π is an irreflexive partial pairing. A contact map is a summary statement that specifies which agents are allowed to bind each other via which sites. In a contact map, it is therefore perfectly legitimate for one site to have edges to many others, including itself; this means that multiple possible bindings may compete for that site. A pattern is intended to represent a realizable site graph. It can have multiple nodes with the same name, representing different individuals of the same kind, but each site can sustain at most one bond at a time.

We use contact maps to classify patterns: a pattern P is said to *respect* a contact map C if there exists a function $f: V_P \rightarrow V_C$, from the nodes of P to those of C , satisfying

- (1) for all $v \in V_P$, the name of v is the same as that of $f(v)$ and the sites displayed by v are a subset of those displayed by $f(v)$;
- (2) if there is an edge from v_1 to v_2 in P , there must be an edge from $f(v_1)$ to $f(v_2)$ in C .

Such a function is a typical *homomorphism* of site graphs, analogous to the familiar concept of graph homomorphism. Note that distinct nodes of P that share the same name need

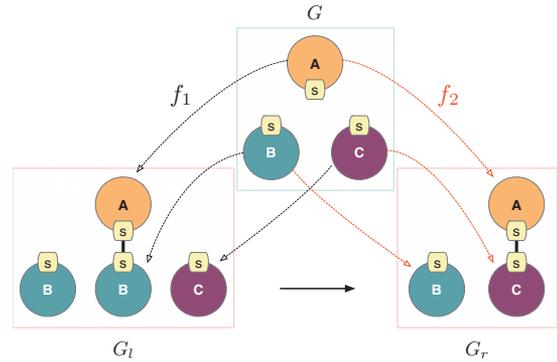


FIG. 3. (Color) Actions as spans. A rule rewrites the pattern G_l on the left into the pattern G_r on the right. Since agents can be created and destroyed, additional information relating agents in G_l to agents in G_r must be provided to unambiguously determine the action of a rule. This information is given by the (maximal) subgraph G defined to remain invariant across the rule, along with two maps, f_1 and f_2 , mapping each agent in G to a unique agent in G_l and G_r , respectively. In the case shown, the rule destroys the first agent B in G_l and shifts the bond. Without the span G , f_1 and f_2 , the action of the rule might conceivably consist of destroying all agents on the left and then recreating them on the right. We omit G whenever the intended action of the rule is obvious.

not display the same sites; but they must all display a subset of the sites displayed by the corresponding node of C . Also, many nodes of P may map to the same node of C , i.e., the function f is not generally injective; if f is injective, we say that it is a *monomorphism*. If *every* node of P displays exactly the same sites as its counterpart in C , we say that P is a *mixture*. A connected component of a mixture is known as a *complex species* or just a *complex*.

B. Rules

A rule specifies the modification of a site graph by the addition/removal of nodes and edges. This is represented by two site graphs—the “before” [left-hand side (LHS)] and the “after” [right-hand side (RHS)]—together with sufficient information to identify which nodes have *persistent identity*, i.e., are neither added nor removed by the rewriting—although they may gain or lose edges from their sites. (For the sake of simplicity, we do not consider *internal states* of sites in this paper, but the framework presented here generalizes straightforwardly to deal with that. Moreover, we assume that only complexes can be added or removed; this avoids the technical complication of *side effects* where removal of a node causes all its incident edges to be removed too.) The sum total of modifications effected by the rule is called its *action*.

This idea is rigorously defined by a span of monomorphisms, i.e., two monomorphisms $f_1: G \rightarrow G_l$ and $f_2: G \rightarrow G_r$, where the left target G_l is the rule’s LHS pattern, the right target G_r is the modified pattern and the source G is the subgraph of G_l that remains invariant under the rewrite, Fig. 3. We need injectivity to ensure that persistent nodes are unambiguously located in both G_l and G_r .

Given a rule with LHS pattern P , how do we know where to *apply* it in a mixture M ? It is not enough to take a monomorphism from P to M as a possible location as there may be additional edges in M that invalidate the pattern P .

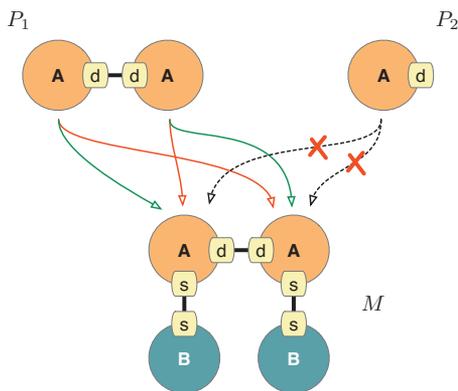


FIG. 4. (Color) Embeddings of graphs. An embedding from P into M must preserve not only agent types but also the state of each site in P . For example, P_1 has two embeddings into M , depending on whether the first agent A of P_1 is mapped to the first or second agent A in M . In contrast, P_2 does not embed into M , since site d of A is unbound in P_2 , whereas it is bound in either agent A of M .

Instead, we require a stricter notion of *embedding*—analogous to the concept of *induced subgraph*—which is a monomorphism where the free sites of every node of P remain free in M (Fig. 4). The set of embeddings from G_1 to G_2 is written $[G_1; G_2]$.

A rule application is thus determined by a choice of embedding of its LHS pattern into the mixture. The result of a rule application is to rewrite, according to the action of the rule, that part of the mixture targeted by the embedding (Fig. 5).

An embedding from a site graph G to itself is called an *automorphism* or just a *symmetry* of G ; we often write $\text{Aut}(G)$ for $[G; G]$. If G has n connected components G_i , the total number of symmetries $|\text{Aut}(G)|$ decomposes into the product of the *individual* or *intra*-symmetries of each G_i with the number

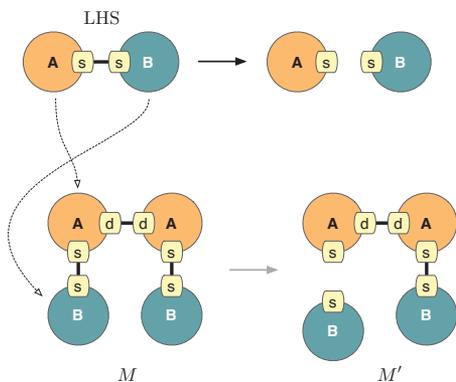


FIG. 5. (Color) Rule application. The action of a rule (top) induces a corresponding change in the mixture (bottom). The location at which the rule applies in M is determined by the choice of an embedding of the rule’s LHS in M . In the case shown, the LHS can embed in two distinct ways. The total number of embeddings in (a typically much larger) mixture M is one factor in the rate (probability) with which the rule fires. A further factor is the rate constant associated with each rule, usually declared following an @-sign. Finally, symmetries in the LHS give rise to combinatorial factors, as explained in Sec. II D.

$$\text{inter}(G) := |\text{Aut}(G)| \bigg/ \prod_{i=1}^n |\text{Aut}(G_i)|$$

of *inter*-symmetries between the G_i s.

The above formalizes the action of a rule and how it can be applied to a mixture. This purely qualitative information must be complemented by a non-negative real number, the *rate constant* of the rule. By convention, we use microscopic, i.e., “per collision,” rate constants and, as such, the choice of rate constants depends on the desired volume of reaction vessel.

C. Kappa and BNGL

The language we just defined bears substantive similarities with BNGL, a well-developed independent approach that aims at supporting the specification and execution of biological models.^{4,6,13,14} Yet, BNGL and Kappa exhibit subtle differences, reflective of their distinct origins. Kappa was conceived by computer scientists as a minimalist framework conducive to developing both models and theory.

There are two principal technical differences between Kappa and BNGL. First, *within an agent*, BNGL allows for identically named sites whereas Kappa insists that sites must be unique (Sec. II A, item 3). Second, BNGL uses two syntactic operators, + (plus) and · (dot), to combine site graphs, whereas Kappa uses the (,) (comma) only. These may seem like superficial differences but they have consequences which we now discuss briefly.

1. Identical sites

BNGL allows agents with multisets of sites for the purpose of representing symmetric proteins that contain multiple copies of the same domain. There appear to be few instances of monomeric proteins of this type. In BNGL, such a protein with two copies of a domain would be expressed as $C(c, c)$, whereas a Kappa representation would require two bound agents, $C(c, x^1), C(c, x^1)$, with no dissociation rule. Three or more identical sites can be encoded by an unbreakable ring of identical agents, giving rise to an “effective” agent with an apparent multiset of sites. Alternatively, one can use the meta-language for Kappa,¹⁷ which allows the expression of *generic* rules that apply uniformly to multiple sites (distinctly named) of the same agent once automatically translated into Kappa.

The use of multisets necessitates a more complex treatment of homomorphisms (and embeddings) since they must specify not only how nodes are mapped across the homomorphism, but also how sites are mapped. This complicates the detection of embeddings since, in general, backtracking will be necessary in order to check all possible ways of embedding one site graph into another; in contrast, no backtracking is required for Kappa site graphs. It also affects the notion of symmetry in that a single node can have nontrivial automorphisms. In Kappa, only complexes can have nontrivial automorphisms.

2. Agent composition operators

The (,) (comma) operator of Kappa gives rise to a syntax where an expression denotes a specific site graph. In a given mixture, that site graph (which need not mention all possible sites of each agent) may embed in many different ways into many different complex species (which are site graphs too). In particular, a disconnected site graph may embed into a connected site graph. This may happen because Kappa's basic syntax has no way of expressing the "relative locations" of agents. A binding rule, such as $A(x), B(y) \rightarrow A(x^1), B(y^1)$, may instantiate as a unimolecular reaction (a ring closure) or a bimolecular reaction. BNGL makes use of the + and \cdot operators to intrinsically express whether or not a disconnected site graph may embed, or not, a connected site graph: the + operator prevents this, the \cdot insists upon it, without requiring a specification of the connection. The distinction is important from a kinetic viewpoint, since bimolecular, unlike unimolecular, rate constants have an inverse volume dependence. While the BNGL operators express nonlocal constraints *directly* at the level of rules, Kappa-rules can be *annotated* to enforce restricted types of embeddings by the simulation engine. For example, a Kappa-rule can be annotated with two rate constants and the simulation engine⁹ executes mono- and bimolecular reaction instances with the proper stochastic chemical kinetics. BNGL's mixing of local and nonlocal information in the same rule expression has a subtle consequence: the syntactic expressions of BNGL do *not* denote individual site graphs: the expression $A+B$ denotes the two site graphs, A and B , together with some information on what a valid embedding of them must satisfy; the same is true of $A \cdot B$. In effect, the + and \cdot operators are syntactic *patterns* that only make sense in the context of a mixture; moreover, they may mean different things in different mixtures.

Since the present contribution is aimed at facilitating access to the broader shape of our ideas, we shall restrict ourselves to situations in which Kappa expressions have no ambiguity with respect to molecularity, that is, disconnected components of a site graph embed into disconnected components of the host site graph. This is ensured by a restriction to *acyclic contact maps*. An acyclic contact map is a site graph in which no cyclical path exists that touches the sites it contains only once. Thus, the object C in Fig. 2 is acyclic, the loop notwithstanding, as site d of A is touched twice. If agents B and C had another bond between them, e.g., $A(d^{1,1}, s^{2,3}), B(s^2, p^4), C(s^3, p^4)$, the contact map would still be acyclic because the cyclical path through agents A , B , and C touches site s of agent A twice. If, in addition, agents B and C were to bind A at distinct sites, e.g., $A(d^{1,1}, s^2, q^3), B(s^2, p^4), C(s^3, p^4)$, the contact map would be cyclic. An acyclic contact map admits no rings as realizable patterns (Sec. II A).

General contact maps and proofs are the subject of a forthcoming paper. We should note, however, that our current implementation of fragmentation,^{18,19} summarized as the annotated contact map strategy in Sec. V C, is not restricted to acyclic contact maps.

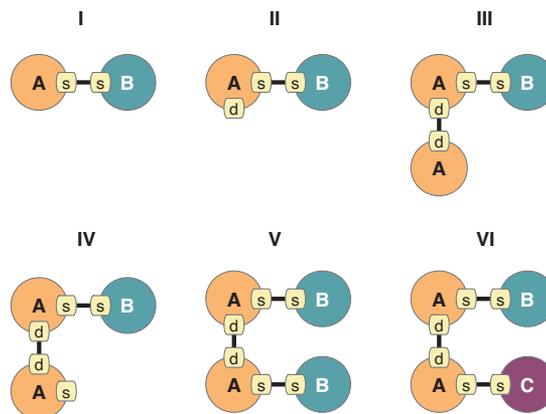


FIG. 6. (Color) Graphical Kappa expressions. A graphical presentation of the Kappa expressions mentioned in the refinement example.

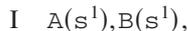
D. Reaction rules and reactions

We have seen that a rule is characterized by its LHS pattern and the rewrite action it performs. In the context of a fixed contact map, a rule may therefore apply to many (combinations of) complex species; this follows from the very fact that a pattern need not specify all sites of an agent. A rule whose LHS pattern is actually a mixture is called a *reaction rule* (or *ground rule*) as it consumes and produces only complex species.

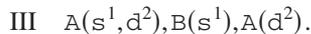
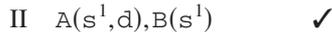
It is always possible to expand a rule to its underlying multiset of reaction rules that enumerates all possible (combinations of) complex species to which the rule can apply.^{13,14} A rule set therefore induces a multiset of reaction rules which, in general, may be infinite although, in this paper, we restrict ourselves to the case where this remains finite (acyclic contact maps). The reaction rules are then easily translated into a system of structureless reactions and then to the corresponding system of ordinary differential equations (ODEs) describing the deterministic kinetics of each species. This expansion is the most fine-grained description compatible with a given set of rules and a contact map. Although such an expansion⁴ is precisely what we would like to avoid, its explication introduces the ground-level object with whose behavior our system of new variables has to "commute" in order to qualify as sound. Some of the concepts we shall encounter in the process will be useful later on and foster a sharper appreciation of what rules are.

The expansion of a rule proceeds by identifying, for each connected component of the LHS pattern, the set of complex species that it matches. Each reaction rule is then determined by a choice of one complex species per connected component. To identify a complex species matching a connected component c , we grow c in all ways compatible with the contact map, until it only contains nodes that display all their possible sites.

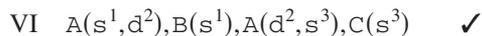
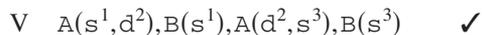
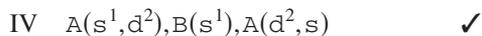
In the following example, refer to Fig. 6 to get a sense for the translation between graphical and textual expressions (labeled by roman numerals), so we can liberally use the latter thereafter. Given the contact map C from Fig. 2, if we start with the pattern



we must add to A its second site d . This site is either free or bound to a second A ,



The first of these is already a valid complex species (flagged by a checkmark). The second still lacks the site s of the second A ; so its expansion must continue, yielding



All nodes now display all sites so the expansion of $A(s^1), B(s^1)$ concludes here with a total of four complex species.

The dissociation rule

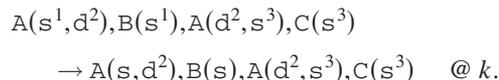
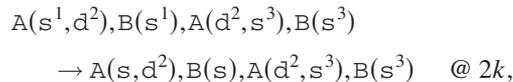
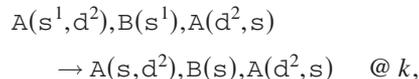
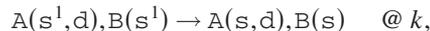


thus expands to four reaction rules. This is called a *refinement* of the original rule, because each reaction rule is a more specific instantiation of the rule. The *activity* of a rule, with LHS pattern P and rate constant k , in the mixture M is determined by mass action,

$$|[P; M]| \cdot k / |\text{Aut}(P)|.$$

The division by automorphisms in the definition of activity is justified by the fact that a rule is a mechanistic hypothesis: if we postulate a symmetric binding mechanism, this means it cannot distinguish between two complexes that match it, even if those complexes are actually *different* when taking into account the wider context that the mechanism ignores; similarly, if a postulated unbinding is symmetric, it has no way of telling if it is actually being applied to an asymmetric complex. Conversely, if we postulate an asymmetric binding mechanism, it can distinguish even between *identical* complexes that match it; and if a postulated unbinding is asymmetric, its application to a symmetric complex completely ignores that symmetry.

The refinement of a rule must be *neutral*, meaning that the overall behavior of the *family* of reaction rules must be dynamically indistinguishable from that of the original rule. A refinement is neutral if, in any mixture, the activity of the rule is the same as the sum of the activities of the cases that constitute the refinement. This requires a little care, because the neutral refinement of a rule to its reaction rules must correct the rate constants for cases where a reaction rule has lost or gained symmetry with respect to the original rule. In our example, the LHS P of rule (1) has no nontrivial automorphisms; however, one of the four reaction rules *does* have one and thus would find its contribution to the total refinement activity to be penalized by a factor of 2. To obtain a neutral refinement, we must correct for this by multiplying its rate constant by two,



In general, if the original rule has LHS pattern P and rate constant k , the induced reaction rule with LHS M has rate constant

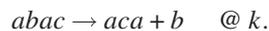
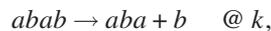
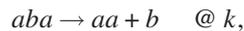
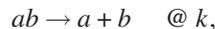
$$k_M := k \cdot |\text{Aut}(M)| / |\text{Aut}(P)|.$$

We can further convert the reaction rule to a traditional *reaction* in which complex species are replaced by unique proper nouns. These names are formally structureless; they refer to what we call plain species or simply *species*. (Throughout this paper we use combinations of slanted lower case letters to name species and upright typewriter font for expressions that formally represent the internal structure of objects.) Although we might name species cleverly to encode a reference to the object they name, that information is non-existent to the mechanism itself. As a consequence, the passage from complex to plain species eliminates all intra-symmetries and so necessitates further surgery on rate constants.

Continuing the above derivation, if M consists of n connected components C_i , its structureless version has rate constant

$$k_s := k_M / \prod_i |\text{Aut}(C_i)|.$$

In our example, the gain of symmetry of the third reaction rule is due to an intra-symmetry; so the passage to structureless reactions cancels out the need to amplify its rate constant,



It is now straightforward to write down the system of ODEs by gathering together, for each species, all the terms that consume and produce it. We refer to the fully expanded system of reaction rules (or reactions) and their corresponding complexes (or species) as the *ground system*.

III. THE PERFECT WORLD

In this section, we illustrate and discuss the idea of fragmentation with a simple example of the kind studied in multivalent ligand/receptor interactions.²⁰ We assume a central

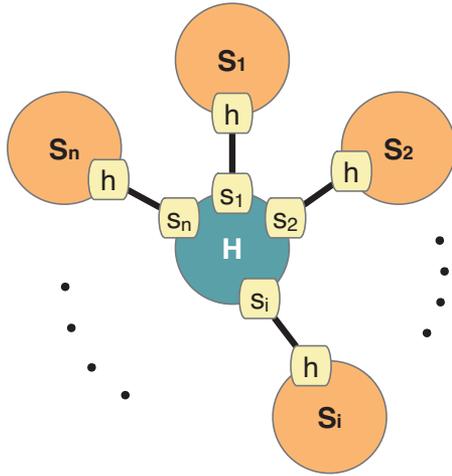
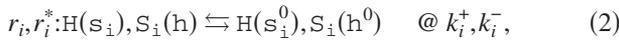


FIG. 7. (Color) The contact map of a hub-and-spoke system, Eq. (2).

“hub” agent H with n sites s_1, \dots, s_n and n distinct types of “spoke” agents S_1, \dots, S_n , each with a single site h . We have n straightforward reversible rules,

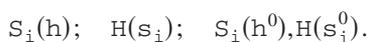


where r_i^* names the rule opposite (reverse) to r_i ; k_i^+ and k_i^- are single-site association and dissociation rate constants, respectively. These rules give rise to the contact map shown in Fig. 7. An agent H can exist in any of 2^n possible states, each of its n sites being either bound or unbound. This system of n reversible rules corresponds to a ground system of $2^n + n$ species—one for each possible state of H plus the n unbound spokes $S_i(h)$.

The salient aspect of Eq. (2) is that each of the rules observes exactly one site of H and so there is no overlap in what they depend on. We might say that the rules are *independent* of each other, e.g., firing one of the binding rules has no effect on the possibility (or not) of firing any of the others; indeed, it only affects one of the unbinding rules—its own reverse. This is a purely *static* notion of independence, based on the mechanisms expounded by the rules.

Static independence must be kept distinct from the more familiar notion of *dynamic independence*, which asserts the absence of certain correlations during simulation, e.g., a bivalent molecule might exhibit a correlation between its occupancy states, despite there being no apparent mechanistic (i.e., static) dependency, because both of its ligands independently require the molecule to adopt the same conformation in order to bind.

In this simple situation, it is easy to see that each rule is incapable of discriminating species that differ only on the sites it does not observe. This suggests that the following $3n$ patterns might be sufficient to fully capture the dynamics of the system. These patterns are of an *intensional* nature; however, they can also be viewed *extensionally* as sets of species that are indistinguishable from the vantage point of the rules,



Let us examine exactly what we mean by this, as it will help clarify the conditions that must be met for it to be true in

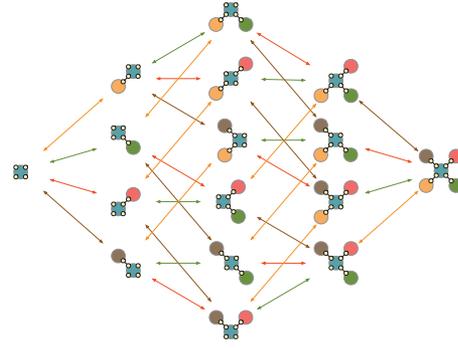
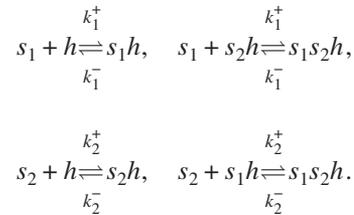


FIG. 8. (Color) The figure depicts the ground reaction network for species that is implied by the rule system (2), which has the contact map shown in Fig. 7. For the sake of a less cluttered picture we omit the free ligand-agents that participate in each reaction. Each of the n reversible rules in Eq. (2) expands into 2^{n-1} reversible reactions with the same color.

general and inform our procedure for identifying suitable sets of fragments for arbitrary rule sets where it is no longer practical to do it “by hand.”

First of all, we must specify our semantics of reference—the ground expansion of this rule set that gives rise to $n \cdot 2^{n-1}$ reversible reactions acting on the $2^n + n$ species described above, each rule instantiating to 2^{n-1} reactions (all with the same rate constant) as depicted for $n=4$ in Fig. 8. To avoid unwieldy notation, we enumerate these reactions for the case of $n=2$; the generalization to arbitrary n should be evident,



We have made the arbitrary choice to name the four species as $s_1 s_2 h$ for an H with both partners bound; $s_1 h$ (respectively $s_2 h$) for an H with just S_1 (respectively S_2) bound; and h for an H with neither partner bound. In standard fashion, these reactions give rise to the following system of ODEs; we use square brackets to denote concentrations and $'$ to denote time-derivatives,

$$[s_1]' = k_1^-([s_1 h] + [s_1 s_2 h]) - k_1^+[s_1]([h] + [s_2 h]),$$

$$[s_2]' = k_2^-([s_2 h] + [s_1 s_2 h]) - k_2^+[s_2]([h] + [s_1 h]),$$

$$[h]' = k_1^-[s_1 h] + k_2^-[s_2 h] - [h](k_1^+[s_1] + k_2^+[s_2]),$$

$$[s_1 h]' = k_1^+[s_1][h] + k_2^-[s_1 s_2 h] - [s_1 h](k_1^- + k_2^+[s_2]),$$

$$[s_2 h]' = k_1^-[s_1 s_2 h] + k_2^+[s_2][h] - [s_2 h](k_1^+[s_1] + k_2^-),$$

$$[s_1 s_2 h]' = k_1^+[s_1][s_2 h] + k_2^+[s_2][s_1 h] - [s_1 s_2 h](k_1^- + k_2^-).$$

The trajectories obtained by solving this system of equations are our semantics of reference which must be perfectly preserved by the fragmentation process.

Next, we need to specify the result of fragmentation. According to the above proposal, we have six fragments, corresponding to a linear change of variables,

$$S_1(h) := s_1,$$

$$H(s_1) := h + s_2h,$$

$$S_1(h^0), H(s_1^0) := s_1h + s_1s_2h,$$

$$S_2(h) := s_2,$$

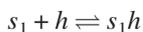
$$H(s_2) := h + s_1h$$

$$S_2(h^0), H(s_2^0) := s_2h + s_1s_2h.$$

We can think of each fragment as a multiset of species. In this case, every species belongs to at least one fragment and some species belong to several; moreover, no species appears more than once in any given fragment although this is possible in general. The fragmentation has thus defined a *covering*, not a *partition*, of the set of species. In general, a fragmentation need only be a *partial* covering of the species, i.e., not every species need belong to a fragment.

Let us note that this is a very extensional view of fragments in the sense that it defines them, as *macroscopic* species, in terms of the microscopic species. Note that, in this example, the fragments viewed as site graphs have no overlap and yet induce overlapping sets of (microscopic) species; this betrays the extensionality of defining fragments in terms of species, a point we will return to below.

We now return to the question of how this abstraction can properly account for the dynamics of the original system of ODEs. Each of the reactions of the ground expansion causes an update in the numbers of species; for example,



consumes one s_1 and one h and produces one s_1h when applied from left-to-right; consumption and production are exchanged when it is applied right-to-left.

According to the extensional perspective on fragments, the firing of this reaction potentially affects the number of instances of every fragment that contains at least one of these three species. As might be expected, $S_1(h)$ and $H(s_1)$ are decremented (assuming a left-to-right firing) while $S_1(h^0), H(s_1^0)$ is incremented. However, one of the other fragments is also affected, specifically $H(s_2)$. At first sight, this is a little counterintuitive since the reaction being fired is a ground instance of a rule that does not even mention s_2 . But, on closer inspection, we see that this “blindness” of the rule is precisely reflected in the fact that, while h is decremented, this is cancelled out by the increment of s_1h : in other words, although the numbers of the specific species change upon firing the reaction, the overall number of instances of the *fragment* is unchanged.

So the firing of this reaction leads to an update in the numbers of instances of each fragment. Those that gain instances are said to be *produced* by the rule; conversely, those that lose instances are *consumed*. The other fragments—such as $H(s_2)$ —remain unchanged overall and although, as we

have seen, there may (or may not) have been a redistribution of the relative abundance of their constituent species. From the intensional point of view (where a fragment is considered simply as a site graph), it is obvious in this example that $H(s_2)$ is completely unaffected by any firing of r_1 , cf. Eq. (2); guaranteeing this property in general is the principal technical difficulty of fragmentation.

It is possible, but by no means necessary, that *all* reaction instances of the rule lead to the *same* update on fragments. If this is the case, as it is in this example, we say that the rule induces an *unambiguous update*. This places a fundamental limit on how coarse-grained fragments can be. For example, if we add s_2h to $H(s_2)$, the update becomes ambiguous because only one of the two instances of r_1 updates s_2h . Moreover, it places a limit on how fine-grained fragments can be; we cannot remove s_1h from $H(s_2)$ for the same reason. This sits comfortably with the idea that fragments express all, but only, that which rules collectively observe. If all rules induce unambiguous update, we have a “perfect” fragmentation in the sense that the fragments express *exactly* (including stochasticity) what the rule set observes. A quick inspection shows us that this is true for our present example. However, we will see later more complex situations where it fails.

We conclude this example with a few remarks about the fragmented system of ODEs,

$$[S_1(h)]' = k_1^-[S_1(h^1), H(s_1^1)] - k_1^+[S_1(h)][H(s_1)],$$

$$[S_2(h)]' = k_2^-[S_2(h^1), H(s_2^1)] - k_2^+[S_2(h)][H(s_2)],$$

$$[H(s_1)]' = [S_1(h)]',$$

$$[H(s_2)]' = [S_2(h)]',$$

$$[S_1(h^1), H(s_1^1)]' = -[S_1(h)]',$$

$$[S_2(h^1), H(s_2^1)]' = -[S_2(h)]'.$$

First of all, it is a self-consistent system in the sense that the derivative of each fragment is expressed only in terms of fragments. Unlike the property of unambiguous update discussed above, self-consistency is an *absolute requirement* of any valid fragmentation; indeed, we will take this as the abstract definition of what fragmentation is. This is important because it allows us to definitively abstract away from species and think purely in terms of fragments; without this, fragmentation might be able to tell us something interesting about at least some of the information carriers of a system, but it would have no practical application to exact model reduction as one would still need to maintain information about species. More conceptually, a self-consistent set of fragments is a purely intensional object that makes no reference to, and has no dependency on, microscopic species.

Second, we see here that fragmentation does in general lead to model reduction: we replace $2^n + n$ species with $3n$ fragments which yields an actual reduction once $n > 2$. Interestingly, although the case where $n = 1$ is essentially trivial—the fragments are just the species—the case of $n = 2$ that we

have considered gives no reduction in the number of variables but does produce a nontrivial fragmentation which identifies the system's information carriers.

This leads to a final point: does fragmentation intrinsically prevent us from reconstructing the dynamics of individual species out of those fragments? Even in our current simple example, the change of variables that defines our fragments has no inverse; so we cannot recover species from fragments by inverting this fragmentation. However, the independence of the rules discussed previously suggests a nonlinear reconstruction of the trajectories of species: if we set $[H(\cdot)] := [H(s_1)] + [S_1(h^0), H(s_1^0)] = [H(s_2)] + [S_2(h^0), H(s_2^0)]$, the fraction of Hs that are fully bound is

$$[s_1 s_2 h] / [H(\cdot)],$$

while the fraction of Hs with at least S_1 bound is

$$[S_1(h^0), H(s_1^0)] / [H(\cdot)].$$

The fraction of H with at least S_2 bound is defined similarly. If the independence of the two binding rules is truly reflected in the dynamics of the system, we would expect

$$[s_1 s_2 h] \cdot [H(\cdot)] = [S_1(h^0), H(s_1^0)] \cdot [S_2(h^0), H(s_2^0)] \quad (3)$$

always to hold. In other words, the *correlation measure*

$$\chi := [s_1 s_2 h] \cdot [H(\cdot)] - [S_1(h^0), H(s_1^0)] \cdot [S_2(h^0), H(s_2^0)]$$

should be zero everywhere—and, indeed, this follows immediately from the closed formula

$$\chi' = -\chi \cdot (k_1^+ [S_1(h)] + k_2^+ [S_2(h)] + k_1^- + k_2^-)$$

for the derivative of χ , provided that $\chi=0$ in the initial conditions. Concretely, this means that we can reconstruct the species $s_1 s_2 h$ as we have a closed formula defined only in terms of fragments. The other species can be similarly recovered.

It should be noted, however, that the purpose of fragmentation is exact model reduction and/or the identification of information carriers, not to abstract away from species only then to reconstruct them. This is opportune because, as we will see in Sec. IV, it is often impossible to exactly reconstruct certain species due to correlations coming, not from the mechanisms hypothesized by the rule set, but from dynamical dependencies that cannot be deduced statically from the rules. If $\chi \neq 0$ in the initial conditions, we can no longer exactly reconstruct all species, but the fragment dynamics remains sound as it only abstracts away correlations that the rules cannot observe. By dynamic independence, we mean that, in a population of molecules undergoing reactions conforming with a given set of rules, the occupancy of one site of H tells us nothing about the likelihood that the other site is occupied. That is to say, the joint occupancy probability is induced by the marginals for each site. We will illustrate in Sec. IV that just because a rule set satisfies static independence, it does not mean it satisfies dynamic independence.

The converse is not true either: a rule set may exhibit dynamic independence while being not at all statically independent. As an extreme example, we could take the ground expansion of our rule set into reaction rules (not structureless

reactions) so that every rule explicitly mentions the binding state of every site of H. This means that the firing of a rule now affects most of the other rules in the system; the system is, in some sense, maximally statically dependent. However, the shift from the original rule set to its underlying set of reaction rules is dynamically transparent, a neutral refinement: the two rule sets define exactly the same continuous-time Markov chain. In particular, the correlation measure χ remains invariantly zero. This may seem mysterious but is explained by noting that static (in)dependence is a property of the mechanisms hypothesized by the rules, whereas dynamic independence is a property of the transition system engendered by them. So the change from the original rule set to the set of reaction rules amounts to a sea-change in the binding mechanisms we are hypothesizing. It should be noted that the system of reaction rules has many more rate constants that can be chosen to access a far greater repertoire of dynamic behaviors. Mimicking the original system is just one possibility, which is realized by inheriting the rate constants from the original rules, corrected (when necessary) by appropriate symmetry factors (Sec. II D).

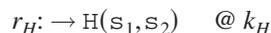
IV. THE REAL WORLD

The simple example of the previous section had two pleasing properties: (i) that rules unambiguously updated the fragments and (ii) a correspondence between static and dynamic independence. In this section, we examine these two properties in greater detail, in particular showing how even small changes to the rule set can destroy them.

A. Dynamic correlation

In Sec. III, we showed that the independence of the binding rules, as hypothesized by their mechanisms, was indeed reflected in the dynamics of the system: the correlation measure χ was everywhere zero, implying no correlation between the occupancy of the sites s_i of H. However, as we shall see, a small change to the rule set suffices to break this invariant. Despite being initially 0, χ transiently takes on positive values, indicating a positive correlation between the occupancy of H's binding sites. When this is the case, fragmentation leads to an unrecoverable loss of information and the species dynamics can no longer be reconstructed exactly. Yet, in the following modification of the example discussed in Sec. III, the system of fragments preserves both self-consistency of the fragments and the unambiguous update property and so the fragments, which remain unaffected by the modification, retain their status as perfect information carriers.

We modify the rule set r_i , given in Eq. (2), by adding the rule



for the dynamic creation of H-agents. Note that this is no way compromising the static independence of r_1 and r_2 ; however, the derivative of the correlation measure

$$\chi = [s_1 s_2 h] \cdot [H(\cdot)] - [H(s_1^1), S_1(h^1)] \cdot [H(s_2^1), S_2(h^1)]$$

changes since

$$[H(\cdot)]' = k_H$$

instead of 0. This yields

$$\chi' = k_H \cdot [s_1 s_2 h] - \chi \cdot (k_1^+[S_1(h)] + k_2^+[S_2(h)] + k_1^- + k_2^-),$$

which is not everywhere 0, even if $\chi=0$ at $t=0$. This means that there is a positive correlation of the occupancy of sites s_1 and s_2 . In effect, knowing that a given H is bound on one of its sites reveals information about how recently it was created which, in turn, affects how likely it is to be bound on its other site. This means that the nonlinear reconstruction of $[s_1 s_2 h]$, Eq. (3), is inexact; it underestimates the true value although, in this example, the discrepancy tends to diminish (the correlation tends to 0) as H swamps the system, diluting out S_1 and S_2 . Nonetheless, the dynamics of the fragment-level ODEs are still exact. What has been lost is the correspondence between static and dynamic independence—and this happens precisely because the rules are unable to observe the dynamic correlation. A finer-grained choice of fragments might recover the ability of reconstructing species exactly. For this example, it would be necessary—as is always possible—to use the underlying set of species as fragments. However, unless the rule set was refined to reaction rules as well, this choice would destroy the unambiguous update property and is neither a scalable nor insightful strategy, as discussed previously.

In summary, this example demonstrates how, even if two mechanisms are hypothesized to be independent, they might become dynamically correlated under some conditions. Moreover, the potential for this cannot be determined statically from the rules (or indeed the reactions) defining a system.

B. Degrees of observation

In between a rule set and its expansion to reaction rules (Sec. II D), there is in general an entire spectrum of intermediate “granularities” of (in)dependence where some but not all rules depend on other rules. As one sweeps across this spectrum, the fragmentation process produces quite different results: a system of reaction rules can only have (complex) species as its fragments, whereas rule sets with less static dependency will generally produce smaller and fewer fragments.

Let us illustrate this with a simple variant of our running example. We consider the case of $n=3$, leaving r_2 and r_3 untouched, but replacing rule r_1 by

$$r_{1a}, r_{1a}^*: H(s_1, s_2), S_1(h) \rightleftharpoons H(s_1^0, s_2), S_1(h^0),$$

$$r_{1b}, r_{1b}^*: H(s_1, s_2^1), S_1(h), S_2(h^1) \rightleftharpoons H(s_1^0, s_2^1), S_1(h^0), S_2(h^1),$$

with the forward and backward rate constants for r_{1a} given by k_{1a}^+, k_{1a}^- and for r_{1b} by k_{1b}^+, k_{1b}^- . This change makes the binding state of s_2 visible to rule r_1 with several important consequences.

First of all, we have modified our hypothesis of the (un)binding mechanism of S_1 to H . Previously, it was assumed not to depend on whether or not S_2 (or S_3) were already bound; with this change, only independence of S_3 remains. This means that the unchanged rule r_2 , for S_2 (un)binding H ,

now impacts upon r_{1a} and r_{1b} : binding an S_2 deactivates r_{1a} and activates the previously inactive r_{1b} . Clearly, if the rate constants of r_{1a} and r_{1b} are chosen so as to furnish a neutral refinement of r_1 (Sec. II D), this will have no effect on the dynamics of the system; any other choice of rate constants yields a system with a more complex dependency between S_1 and S_2 .

Second, the fact that r_{1a} and r_{1b} now observe site s_2 means that the previous triple of fragments

$$H(s_1); S_1(h); H(s_1^1), S_1(h^1)$$

is no longer valid. This is because we have no way of knowing what proportion of, say, $H(s_1)$ is a target of r_{1a} as opposed to r_{1b} . Specifically, if we were to write an ODE for $H(s_1)$, it would be of the following form:

$$[H(s_1)]' = \dots - k_{1a}^+ \cdot [H(s_1, s_2)] \cdot [S_1(h)] \\ - k_{1b}^+ \cdot [H(s_1, s_2^1), S_2(h^1)] \cdot [S_1(h)].$$

But if we knew only $[H(s_1)]$, we could not solve this equation unless $k_{1a}^+ = k_{1b}^+$, i.e., the case where r_{1a} and r_{1b} constitute a neutral refinement of r_1 ; this is a more subtle example of dynamic independence despite static dependence. With any other choice of rate constants, we need to additionally keep track of the relative rates of r_{1a} and r_{1b} over time; this forces us to refine our fragments to

$$H(s_1, s_2); H(s_1, s_2^1), S_2(h^1); H(s_1^1, s_2), S_1(h^1)$$

$$H(s_1^1, s_2^2), S_1(h^1), S_2(h^2); S_1(h)$$

that enumerate the joint binding state of s_1 and s_2 . It should be noted that this is somehow the dual of the problem of ambiguous update: that is the situation where two reaction instances of the same rule update a fragment differently; whereas here, we have two different rules updating different instances of a single (candidate) fragment. The former does not invalidate the fragmentation, but the latter clearly invalidates the candidate fragment, $H(s_1)$, and necessitates a finer-grained fragmentation. Note, however, that the new fragmentation is still coarser-grained than the set of species: the triple of fragments

$$H(s_3); S_3(h); H(s_3^1), S_3(h^1)$$

is completely unaffected by the modification of r_1 . If r_1 were further modified to observe s_3 as well as s_2 , this would result in the fragments becoming the species.

Finally, we note a tension between $r_{1a/b}$ and r_2 in that the refined fragments demanded by $r_{1a/b}$ lead to ambiguous update for r_2 : some instances of r_2 update $H(s_1, s_2)$ while others update $H(s_1^1, s_2), S_1(h^1)$. This has arisen due to the asymmetric conditions for S_1 and S_2 binding H : S_1 depends on S_2 but not vice versa. Clearly, refining r_2 into the analogous r_{2a} and r_{2b} would restore unambiguous update with respect to the set of refined fragments. However, it is important to note that the triple

$$H(s_2); S_2(h); H(s_2^1), S_2(h^1)$$

remains a perfectly valid set of fragments for the rule set with $r_{1a/b}$ and r_2 and, moreover, all the rules have unambigu-

ous update with respect to them. This tells us something interesting: these three fragments are information carriers that self-consistently describe, not the full system but, a subsystem thereof. This means that they define a *module* in the sense that their behavior is isolated from the surrounding context. While in this case, modularity arises rather obviously from the asymmetric dependency—and could have been deduced by hand—fragmentation provides a general way of identifying this kind of situation.

More generally, by identifying information carriers in a (sub)system, fragmentation provides a measure of how complicated the system hypothesized by the rules actually is. A rule set inducing very few information carriers obviously describes a simple system; perhaps more importantly, it also implies that the system is rather inflexible. For example, the rule set of Sec. III cannot accommodate any regulation of the binding of H to its S_i ligands; the refinement of r_1 to $r_{1a/b}$ increases the number and size of fragments, a sign that the system has become more sophisticated. Indeed, this slight change in the mechanism of S_1 's binding to H allows the system to regulate their association as a function of S_2 's presence—either positively or negatively, depending on the choice of rate constants for r_{1a} and r_{1b} . On the other hand, a system with very many information carriers may well be hypothesizing an unrealistic degree of self-observation that violates the local character of interactions between macromolecules. It seems that, for a system to be both realistic and flexible, it needs to find an appropriate middle ground between too many and too few information carriers.

V. FRAGMENTATION

A. The nature of fragments

Let us review the preceding examples and summarize what we have learned from them. We have seen that two contrasting perspectives can be taken of fragments: the *extensional*, which defines and discusses fragments in terms of microscopic species; and the *intensional*, which adopts a more abstract point of view that, incidentally, has certain pragmatic benefits in a world where there are typically an astronomical, if not infinite, number of possible species. There is an irrevocable tension between these two points of view, as is well-illustrated by the decoupling, in general, of static and dynamic independence that we saw in Sec. IV A: aside from the problem of combinatorial complexity, the distinction between extensional and intensional remains fairly anodyne until such time as one can no longer recover the former from the latter. When this breaks down, we might feel we are “losing something” by adopting the intensional perspective. We advocate here the contrary: since that which is lost cannot be observed by the system, nothing important can actually depend on it: far from losing anything, we are gaining clarity by ignoring, in a principled manner, unnecessary complexity engendered by the extensional viewpoint.

A related point concerns the property of unambiguous update. Initially, we formulated this in extensional terms but it is clearly a fundamentally intensional notion. Indeed, a

fragment enjoying unambiguous update with respect to the ambient rule set is rather indifferent to its underlying set of microscopic species: the actual dynamic make-up of the fragment can be in continual flux due to the action of apparently unrelated rules, but at the serene macroscopic level of the fragment, none of this frenetic activity will ever be seen. Only when a rule is fired that explicitly depends on the information represented by the fragment does this effect any real change to the system; moreover, the resulting change is entirely effected at the macroscopic level and requires no microscopic knowledge.

However, we have also seen that this property is not necessary for fragmentation; this is because it is a special case of self-consistency and, as such, guarantees more than is strictly necessary. Specifically, unambiguous update allows for a coarse-graining of the *stochastic* semantics of the rule set; it thus proposes information carriers for a world where mechanisms themselves are stochastic. The (much) weaker property of self-consistency only suffices to give a self-consistent set of information carriers that “average out” mechanistic stochasticity, in essence making the assumption that—despite microscopic stochasticity—the system is inherently trying to implement something deterministic. It is conceivable that both kinds of information carrier are important in signaling networks.

B. Self-consistency

As we have seen, the only truly non-negotiable aspect of fragmentation is the requirement that it produces a self-consistent set of patterns whose average concentrations can be tracked by a system of ODEs. Moreover, fragmentation must be *sound*: first converting species concentrations into fragment concentrations at time t_0 and subsequently letting the (macroscopic) fragment dynamics evolve to a time $t > t_0$ must get us exactly to the same place as first allowing the (microscopic) species dynamics to evolve to time t and then converting the species concentrations into fragment concentrations. In this section we sketch the construction of a self-consistent set of patterns. The proof that this construction is sound will appear in a forthcoming paper; but see Sec. V C.

Conceptually, it is convenient to consider two stages: (i) the writing of ODEs for *any* patterns—typically representing observables we are interested in—in terms of species, and (ii) an iterative process of pattern construction that starts with the desired observables and saturates the set of patterns with respect to the rules. The resulting set will be self-contained, in the sense that no reference to any pattern outside that set is required for tracking the average time-evolution of the initially posited observables.

The first stage consists in identifying, for any pattern F , all the ways in which the application of a rule r can consume or produce it. Let us discuss the slightly simpler case of consumption. The example depicted in Fig. 9 shows two patterns, F and P . Let P be the LHS of rule r ; so an embedding of P into M identifies a reaction instance of r . An embedding of F into M identifies an occurrence of F in the mixture; this may, or may not, overlap—in the mixture—with our reaction instance. If it does, firing *this* instance of r will consume *this*

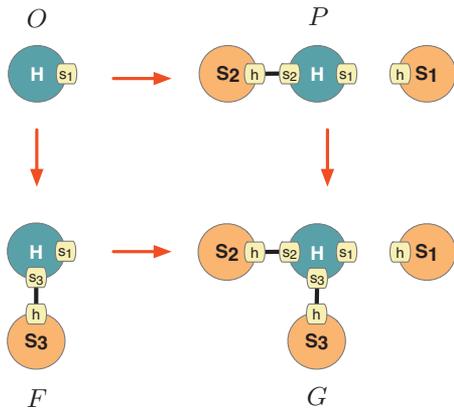


FIG. 9. (Color) Gluing. A gluing is a diagram asserting that two patterns, P and F , can be joined into a pattern G on the basis of a region O they have in common. See text for details.

instance of F . To know how the instantaneous activity of rule r affects the instantaneous consumption rate of F , we therefore need to count all the pairs (p, f) of embeddings (of P and F , respectively, into M) that overlap in this way. We will do this using the concept of a *gluing* which, for a given F and P , identifies the pattern G whose embeddings correspond exactly to the (p, f) pairs of interest.

To construct the gluing of two patterns F and P , we need to specify two pieces of information. First, we need an overlap region O [Fig. 9] on which the patterns agree and on the basis of which we can combine them. This is specified by a span $F \leftarrow O \rightarrow P$ of embeddings that fix exactly how O fits into F and P . Note that, for given F and P , there could be many such spans. Second, we need a larger region G into which F and P both embed. This is specified dually by a cospan $F \rightarrow G \leftarrow P$. A potential gluing is then the choice of a span and a cospan giving rise to a commuting square of embeddings—as in Fig. 9. In order to correctly count our (p, f) pairs, we cannot take just any such span and cospan; the cospan must be appropriately minimal, otherwise G might be overly restrictive—so its extension would not be as large as possible—and we would end up undercounting the consumed instances of F . We have a little more flexibility in choosing the span although care must still be taken to avoid overcounting; the easiest way to achieve this is to always use the (unique) maximal span.

Let us now formalize consumption and production with respect to a rule r , Fig. 10. Consumption of F (pattern F_1 in Fig. 10) happens if the gluing G of r 's LHS with F is modified by the action of r ; this clearly destroys an instance of F and, for a given mixture M , can happen in $[[G; M]]$ ways. We call G a *left-gluing* of F and r . Conversely, production of F (pattern F_2 , unrelated to F_1 , in Fig. 10) happens if the gluing of r 's RHS with F is modified by the action of r . To know in how many ways this can happen, we need to count not $[[G; M]]$, but $[[G'; M]]$, where G' is what G looked like *before* the rule was applied. We call G' a *right-gluing* of F and r . In the simple example of Sec. IV B, this would lead us to write for $F := S_1(h)$,

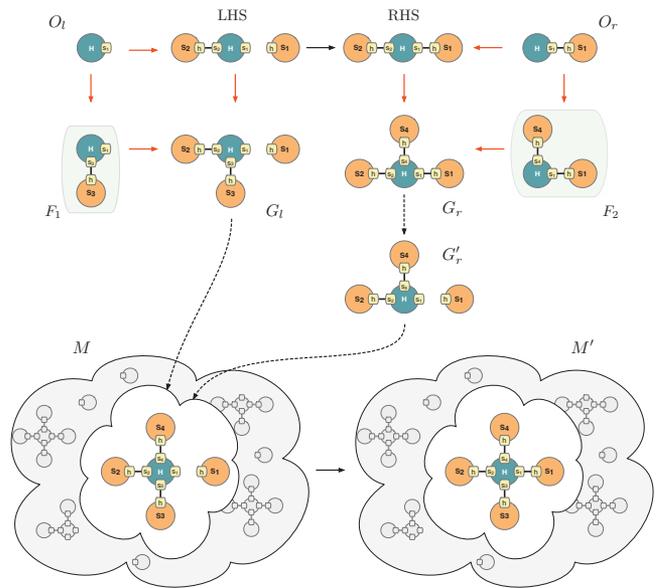


FIG. 10. (Color) Production and consumption of patterns by rules acting on a mixture. The action of a rule $LHS \rightarrow RHS$ (black arrow at top) transforms the mixture M into M' (stylized at the bottom). The resultant instantaneous rate of consumption of pattern F_1 is determined by the number of embeddings in the mixture M of the gluing G_1 (“left-gluing”), as given by the joining of F_1 with the LHS of the rule (square diagram of red arrows on the top left)—provided the shared region O_1 is modified by the rule. Similar for the production of a pattern F_2 (top right square diagram; unrelated to F_1), except that we must undo the action of the rule in the gluing G_r (“right-gluing”), resulting in pattern G'_r whose embeddings in M determine the rate of production of F_2 . The mixture M (M') highlights a particular molecular configuration (among the greyed-out remainder of the mixture) that exemplifies an embedding instance for LHS, G_1 and G'_r . Site and agent names correspond to our simple running example, as depicted in Fig. 7.

$$\begin{aligned}
 [S_1(h)]'(M) = & -k_{1a}^+ \cdot [H(s_1, s_2), S_1(h); M] \\
 & -k_{1b}^+ \cdot [H(s_1, s_2^1), S_2(h^1), S_1(h); M] \\
 & +k_{1a}^- \cdot [H(s_1^1, s_2), S_1(h^1); M] \\
 & +k_{1b}^- \cdot [H(s_1^1, s_2^2), S_1(h^1), S_2(h^2); M].
 \end{aligned}$$

Note, however, that in this differential equation the site graphs must be understood as names of *variables*; $H(s_1), S_1(h)$ could equally well be renamed x . Moreover, $[H(s_1), S_1(h); -]$ must be understood as a function from site graphs to real numbers; applying this function to M yields precisely $[H(s_1), S_1(h); M]$, the instantaneous average concentration value taken on by the variable named $H(s_1), S_1(h)$.

Let \mathcal{F} be a set of patterns; we write $[\mathcal{F}; -]$ for the vector space spanned by the set of functions $[F; -]$ for all $F \in \mathcal{F}$. We say that \mathcal{F} is self-consistent with respect to the rule set \mathcal{R} if, for all $F \in \mathcal{F}$ and all $r \in \mathcal{R}$, the functions $[G; -]$ and $[G'; -]$ tracking the left-gluing and right-gluing, respectively, of F and r are in $[\mathcal{F}; -]$.

In order to obtain a self-consistent set of patterns, it is necessary to *saturate* a *seed* set \mathcal{F} by left- and right-gluing the elements F of \mathcal{F} in all possible ways to the LHSs and RHSs of the rules. In our example, if we seed with just $F_1 := S_1(h)$, we obtain

$$F_2 := H(s_1^1, s_2), S_1(h^1)$$

by right-gluing it with r_{1a}^* and

$$F_3 := H(s_1^1, s_2^2), S_1(h^1), S_2(h^2)$$

by right-gluing it with r_{1b}^* ; we then get

$$F_4 := H(s_1, s_2)$$

by right-gluing F_2 with r_{1a} and finally

$$F_5 := H(s_1, s_2^1), S_2(h^1)$$

by right-gluing F_3 with r_{1b} . No left- or right-gluing of these five candidate fragments with r_2 and r_3 generates new candidates; so the saturation process terminates here.

By construction, this process always produces a self-consistent set of patterns; in this particular case, they allow us to track the average concentration of $S_1(h)$ over time, exactly as described in Sec. IV B. We call the elements of this set the fragments and note that any connected component of a rule LHS that intersects a fragment must—if that intersection is modified by the rule—be contained within the fragment.

The result of the saturation process obviously depends on its seed. It may also produce redundant fragments, expressible as a convex combination of others and so eliminable. For example, seeding our example with $S_1(h)$ and $S_2(h)$ produces the four fragments above plus

$$S_2(h); S_2(h^1), H(s_2^1); H(s_2).$$

The latter two can be eliminated by noting that

$$H(s_2) := H(s_1, s_2) + H(s_1^1, s_2), S_1(h^1),$$

$$S_2(h^1), H(s_2^1) := S_2(h^1), H(s_1, s_2^1) \\ + S_2(h^1), H(s_1^2, s_2^1), S_1(h^2).$$

As mentioned in Sec. IV B, there is an asymmetric dependency between S_1 's association with H and S_2 's: the refinement of r_1 to $r_{1a/b}$ is of no interest to S_2 . If we were to seed saturation with only $S_2(h)$, we would obtain—and be content with—the simple triple of fragments. This optimization could be performed automatically. However, we stress that the fragmentation procedure presented here is a mathematical procedure, a specification. In particular, although saturation always terminates for systems with acyclic contact map, a naive implementation need not terminate for more general systems. In Sec. V C, we describe an efficient, implemented procedure that closely approximates this specification and, moreover, does not depend on acyclicity.

C. Current art and the future

Our current implementation, as described in Refs. 18 and 19, does not calculate fragments exactly according to the specification in Sec. V B; instead, it efficiently approximates this ideal by applying a *dependency analysis* to the rule set (Table I) which it then uses to produce an annotated contact map (aCM).

The annotations are of two kinds: sites are grouped into *covering classes* with the requirement that every site belongs to at least one class; and edges are either *solid* or *dotted*. These annotations are derived by a static analysis of the rule set reminiscent of dependency analyses for detecting unsafe

TABLE I. Fragmentation of test models using an annotated contact map (Refs. 18 and 19). The models are rule-based representations of molecular signaling pathways described in the literature. The columns list the number of rules in the model, the number of distinct molecular (ground) species that can be generated by the rules given the initial mixture, the number of fragments resulting from our automated procedure, the time (in seconds) for generating the full system of ODEs for the ground species (s-ODE), and the time for generating the full set of ODEs for the fragments (f-ODE). (Half the indicated time is spent on writing a LaTeX output and generating an input file for numerical integration with OCTAVE or MATLAB.)

Model	Rules	Species	Fragments	s-ODE (s)	f-ODE (s)
EGF ^a	39	356	38	2.85	0.13
INS1 ^b	76	2899	208	27	0.72
INS2 ^c	74	2899	88	27	0.28
SFB ^d	69	$\approx 2 \times 10^{19}$	$\approx 2 \times 10^5$	Unfeasible	871

^aModel of early events in the epidermal growth factor (EGF) pathway (Ref. 21).

^bModel of cross-talk between insulin and EGF receptors (Ref. 22).

^cSame as INS1, but removing certain dependencies in unbinding rules to study the effect on fragmentation (Ref. 18).

^dPilot study of a larger slice of the EGF pathway (Refs. 5, 21, 23, and 24).

information flows, e.g., higher-security variables that depend on lower-security variables, or the use of Bayesian networks in statistical modeling. A fragment is then read off by picking a starting node of the contact map, choosing a covering class to specify which sites are to be displayed and picking, for each of these sites, a binding state. When a bound state is chosen, the procedure continues recursively if the chosen bindee is connected via a solid edge of the aCM; it terminates if the edge is dotted. The set of fragments is generated by an exhaustive enumeration over all possible choices.

This construction of fragments via annotation of the contact map generally leads to a highly efficient model reduction. However, it does have one significant drawback that is being addressed in current work: fragments are sometimes unnecessarily fine-grained. We consider the following rules defined on the agents of Fig. 2:

$$A(d), A(d) \rightarrow A(d^0), A(d^0),$$

$$A(d^1, s), A(d^1), B(s) \rightarrow A(d^1, s^0), A(d^1), B(s^0).$$

If we seed the generic fragmentation procedure with $B(s)$, we obtain

$$A(d^1, s), A(d^1)$$

by left-gluing it to the second rule; then

$$A(d, s); A(d)$$

by right-gluing that to the first rule. Note the need for a left-gluing to get the process started; this is because our seed is only consumed in this oversimplified example, no rule *produces* it. This defines a self-consistent set of fragments that cannot be generated by any annotation of the contact map: the second rule tests A 's site d and modifies s so they must be in the same covering class—which forces fragments to unnecessarily enumerate the binding state of site s on

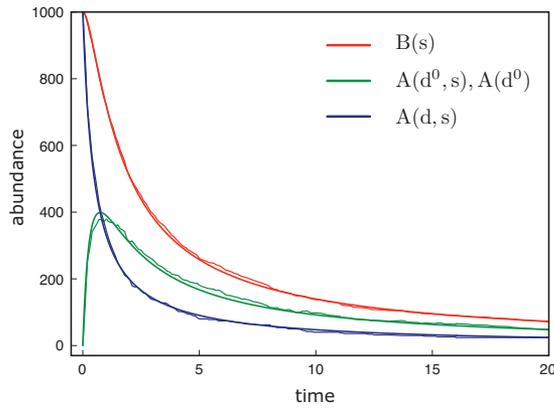


FIG. 11. (Color) Soundness. The microscopic dynamics of the system defined by the two rules discussed in Sec. V C is simulated with a continuous time Markov chain algorithm (Ref. 9). The time traces of the five possible species are then aggregated according to the fragmentation: $B(s)$; $A(d^1, s), A(d^1)$; $A(d, s)$; $A(d)$. These aggregated stochastic traces are shown in the plot as wiggly curves. [Given the initial condition, the trace for fragmentation $A(d)$ is identical to that of $A(d, s)$, as indicated in the text.] The smooth curves are the output of the ODE system of Sec. V D, as generated by fragmentation. Fluctuations aside, the curves match, illustrating soundness as defined in the text. Parameters: $k_1=k_2=0.001$; at $t=0$, $B(s)$ and $A(d, s)$ have 1000 particles each. The model can be executed online (Ref. 16).

both sides of the A dimer,

$$A(d^1, s^2), A(d^1, s), B(s^2); A(d^1, s^2), A(d^1, s^3), B(s^2), B(s^3).$$

In essence, aCM-based fragmentation forces us to make a purely local choice when picking a covering class; so the two occurrences of A are treated homogeneously whereas the rules treat them heterogeneously. A refinement of the aCM-based approach is currently being developed where nonlocal context can be exploited in order to define a heterogeneous growth procedure; this will better handle cases like the above producing smaller and, in general, fewer fragments.

D. Soundness

Let us finally illustrate numerically the *soundness*¹⁹ of our approach, meaning that first following the microscopic dynamics and then aggregating molecular species as prescribed by fragmentation will yield exactly the same outcome as first coarse-graining and then following the coarse-grained dynamics (as prescribed by fragmentation). This is shown in Fig. 11 for the simple example of Sec. V C, whose fragmentation leads to a system of four ODEs,

$$[B(s)]' = k_2[B(s)][A(d^1, s), A(d^1)],$$

$$[A(d^1, s), A(d^1)]' = 2k_1[A(d, s)][A(d)] - k_2[B(s)][A(d^1, s), A(d^1)],$$

$$[A(d, s)]' = -2k_1[A(d, s)][A(d)],$$

$$[A(d)]' = -2k_1[A(d)]^2,$$

where $[A(d, s)]$ and $[A(d)]$ are closely related: if they ever become equal, they will forever after remain so. Note also the combinatorial factor of 2 in the second equation; it arises from there being two distinct ways of gluing $A(d^1, s), A(d^1)$

on the RHS of the first rule. The other factors of 2 come about for similar reasons.

VI. EPILOGUE

The concept of “collective variables” is fundamental to many areas of theoretical physics, such as superfluidity, ferromagnetism, and hydrodynamics to mention a few. These variables are often associated with new collective properties that a many-body system acquires as a result of a phase transition. Some of these properties can be stable to the point of being “universal,” i.e., independent of the details of the material in which they occur. When collective variables fully determine each other’s dynamics, a description has been achieved that is independent of the underlying microscopic definition of the system. This autonomy justifies phrases such as “new level of description” or “emergence.” Understanding such emergence of organized behavior means clarifying the process by which new kinds of collective variables spring from low-level dynamics.²⁵

Fragments share with collective variables the property of self-consistency—being a set of mutually sufficient higher-level descriptors of system dynamics. Yet, they differ from collective variables in that, starting from a set of rules representing local mechanisms of interaction, we distill a self-consistent set of fragments proceeding *purely by static examination of the rule set*; no observation of the dynamics is involved. In contrast, collective variables are typically justified by virtue of dynamics. Nonetheless, as discussed in Sec. V D, this does not compromise soundness of fragmentation with respect to dynamics.

Importantly, fragments differ from collective variables by an intriguing “instability.” Fragmentation is a seeded process that depends on a starter set of fragments, which might be desired observables. It proceeds iteratively by left- and right-gluing already-identified fragments with rules, as described in Sec. V B. At the fixed point, we can express the dynamics of each fragment, in particular the initially declared observables, in terms only of other fragments. Fragmentation tells us the granularity that suffices to exactly describe the dynamics of the chosen observables, regardless of how the microscopic system evolves. Any further fine-graining would not add actionable information from the system’s vantage point. It is in this sense that fragments are information carriers and the dynamical system of fragments defines what we mean when we say that a system “processes information.”

If we change the observables, fragmentation will produce different fragments, even though the underlying microscopic system has not changed at all. In our toy example of Sec. V B, choosing $S_1(h)$ as an observable returns four fragments whereas choosing $S_2(h)$ returns just three, “blinking out” a whole subsystem that never touches $S_2(h)$. But who is doing the observing? It is the system itself, such as when a signal is intercepted by a receptor. (If we insist on an external observer, the system must be amended by the rules that describe the observation mechanism.) Depending on which signal is observed, different fragmentations are induced. To a molecular biologist, the microscopic system has not changed

constitution; same players, same interactions. However, the concentration profiles of molecular species are changed in response to the signal, but the meaning of these changes might remain inscrutable unless we realize that the way the system processes information has changed.

- ¹H. Butterfield, *The Origins of Modern Science* (The Free Press, New York, 1957).
- ²J. V. Golinski, in *Science in the Eighteenth Century*, The Cambridge History of Science Vol. 4, edited by R. Porter (Cambridge University Press, Cambridge, England, 2003), pp. 375–396.
- ³A. Lavoisier, *Elements of Chemistry, in a New Systematic Order, Containing All the Modern Discoveries* (William Creech, Edinburgh, 1790).
- ⁴M. L. Blinov, J. R. Faeder, and W. S. Hlavacek, *Bioinformatics* **20**, 3289 (2004).
- ⁵V. Danos, J. Feret, W. Fontana, R. Harmer, and J. Krivine, *Proceedings of the 18th International Conference on Concurrency Theory*, Lecture Notes in Computer Science Vol. 4703 (Springer, Berlin, 2007), pp. 17–41.
- ⁶W. S. Hlavacek, J. R. Faeder, M. L. Blinov, A. S. Perelson, and B. Goldstein, *Biotechnol. Bioeng.* **84**, 783 (2003).
- ⁷J. R. Faeder, M. L. Blinov, B. Goldstein, and W. S. Hlavacek, *Syst. Biol.* **2**, 5 (2005).
- ⁸W. S. Hlavacek, J. R. Faeder, M. L. Blinov, R. G. Posner, M. Hucka, and W. Fontana, *Sci. STKE* **2000**, re6 (2006).
- ⁹V. Danos, J. Feret, W. Fontana, and J. Krivine, *Proceedings of APLAS 2007*, Lecture Notes in Computer Science Vol. 4807 (Springer, Berlin, 2007), pp. 139–157.
- ¹⁰J. Yang, M. I. Monine, J. R. Faeder, and W. S. Hlavacek, *Phys. Rev. E* **78**, 031910 (2008).
- ¹¹R. Krüger and R. Heinrich, *Genome Informatics* **15**, 138 (2004).
- ¹²A. Ciliberto, F. Capuani, and J. J. Tyson, *PLoS Comput. Biol.* **3**, e45 (2007).
- ¹³J. R. Faeder, M. L. Blinov, B. Goldstein, and W. S. Hlavacek, *Complexity* **10**, 22 (2005).
- ¹⁴M. L. Blinov, J. Yang, J. R. Faeder, and W. S. Hlavacek, *Lect. Notes Comput. Sci.* **4230**, 89 (2006).
- ¹⁵V. Danos and C. Laneve, *Theor. Comput. Sci.* **325**, 69 (2004).
- ¹⁶A web application implementing the Kappa modeling platform can be found at www.rulebase.org. Open-source access to KAPPA software is located at www.kappalanguage.org.
- ¹⁷V. Danos, J. Feret, W. Fontana, R. Harmer, and J. Krivine, *Transactions on Computational Systems Biology* **11**, 116 (2009).
- ¹⁸J. Feret, V. Danos, J. Krivine, R. Harmer, and W. Fontana, *Proc. Natl. Acad. Sci. U.S.A.* **106**, 6453 (2009).
- ¹⁹V. Danos, J. Feret, W. Fontana, R. Harmer, and J. Krivine, “Abstracting the differential semantics of rule-based models: Exact and automated model reduction,” *Symposium on Logic in Computer Science* (IEEE Computer Society, Los Alamitos, CA, 2010), pp. 362–381.
- ²⁰A. S. Perelson, in *Cell Surface Dynamics: Concepts and Models*, edited by A. S. Perelson, C. DeLisi, and F. W. Wiegel (Marcel Dekker, New York, 1984), pp. 223–276.
- ²¹M. L. Blinov, J. R. Faeder, B. Goldstein, and W. S. Hlavacek, *BioSystems* **83**, 136 (2006).
- ²²H. Conzelmann, D. Fey, and E. D. Gilles, *BMC Syst. Biol.* **2**, 78 (2008).
- ²³F. A. Brightman and D. A. Fell, *FEBS Lett.* **482**, 169 (2000).
- ²⁴B. Schoeberl, C. Eichler-Jonsson, E. D. Gilles, and G. Müller, *Nat. Biotechnol.* **20**, 370 (2002).
- ²⁵R. B. Laughlin, D. Pines, J. Schmalian, B. P. Stojković, and P. Wolynes, *Proc. Natl. Acad. Sci. U.S.A.* **97**, 32 (2000).