

AUTOCATALYTIC SETS IN POLYMER NETWORKS WITH VARIABLE CATALYSIS DISTRIBUTIONS

WIM HORDIJK AND MIKE STEEL

ABSTRACT. All living systems – from the origin of life to modern cells – rely on a set of biochemical reactions that are simultaneously self-sustaining and autocatalytic. This notion of an autocatalytic set has been formalized graph-theoretically (as ‘RAF’), leading to mathematical results and polynomial-time algorithms that have been applied to simulated and real chemical reaction systems. In this paper, we investigate the emergence of autocatalytic sets in polymer models when the catalysis rate of each molecule type is drawn from some probability distribution. We show that although the average catalysis rate f for RAFs to arise depends on this distribution, a universal linear upper and lower bound for f (with increasing system size) still applies. However, the probability of the appearance (and size) of autocatalytic sets can vary widely, depending on the particular catalysis distribution. We use simulations to explore how tight the mathematical bounds are, and the reasons for the observed variations. We also investigate the impact of inhibition (where molecules can also inhibit reactions) on the emergence of autocatalytic sets, deriving new mathematical and algorithmic results.

1. AUTOCATALYTIC SETS AND RAF THEORY

The notion of self-sustaining autocatalytic sets was introduced and studied by Kauffman [14, 15, 16] in the context of the origin of life [4]. The concept was later formalized mathematically and further developed as RAF theory [22, 7, 2]. In this section, we briefly review the basics of RAF theory and its main results.

First, we define a *chemical reaction system* (CRS) as a tuple $\mathcal{Q} = (X, \mathcal{R}, C)$ consisting of a set X of molecule types, a set \mathcal{R} of chemical reactions, and a catalysis set C indicating which molecule types catalyze which reactions. Formally, each reaction $r \in \mathcal{R}$ can be viewed as a pair of subsets of X (the ‘reactants’ and ‘products’ of the reaction r) and C is a subset of $X \times \mathcal{R}$. We also assume throughout that X contains a distinguished subset, a *food set* $F \subseteq X$, which is a subset of molecule types that are assumed to be directly available from the environment (i.e., they do not necessarily have to be produced by any of the reactions in the system).

A CRS can be graphically represented by a *bipartite directed hypergraph*. In such a graph representation there are two types of nodes: (1) molecule types (black dots) and (2) chemical reactions (white boxes). There are also two types of (directed) edges: (1) reaction edges (solid black) and (2) catalysis edges (dashed grey). The end points (incident nodes) of these edges are always from different node sets (i.e., an edge is always directed from a molecule node to a reaction node, or vice versa). A simple example is shown in Fig. 1.

The notion of catalysis plays a central role here. A *catalyst* is a molecule type that significantly speeds up the rate at which a chemical reaction happens but without being “used up” in that reaction. Catalysis is ubiquitous in life [25]. Almost all organic reactions are catalyzed, and catalysts are essential in determining and regulating the functionality of the chemical networks that support life.

Given a CRS $\mathcal{Q} = (X, \mathcal{R}, C)$ and food set $F \subseteq X$, an autocatalytic set, or *RAF set*, is now defined as a nonempty subset $\mathcal{R}' \subseteq \mathcal{R}$ of reactions which is:

- (1) *Reflexively Autocatalytic* (RA): each reaction $r \in \mathcal{R}'$ is catalyzed by at least one molecule type that is either a product of another reaction in \mathcal{R}' or is an element F .
- (2) *Food-generated* (F): each molecule type that is a reactant of any reaction $r \in \mathcal{R}'$ can be generated from the food set F by using a series of reactions only from \mathcal{R}' itself.

A simple example of a RAF set consisting of just two reactions is presented in Fig. 1, though a RAF set can of course be of any size. A more formal mathematical definition of RAF sets was provided in [7, 5], including an efficient (polynomial-time) algorithm for finding such sets in arbitrary CRSs (or determining that there is none). This RAF algorithm is based on a simple pruning rule that iteratively removes reaction nodes from the input CRS that (by definition) cannot be part of a RAF set. If the algorithm returns a non-empty subgraph, this subgraph then represents the *maximal* RAF set (i.e., the union of all RAF sets) contained within the input CRS. Otherwise, there is no RAF set in the given CRS.




FIGURE 1. An example of a simple CRS with just two reactions, represented as a bipartite directed hypergraph. Black dots represent molecule types and white boxes represent chemical reactions. Solid (black) arrows indicate reactants going into and products coming out of a reaction. Dashed (grey) arrows indicate catalysis. With a food set consisting of the four molecule types at the bottom, this CRS forms an autocatalytic (RAF) set where the two products (at the top) mutually catalyze each other’s formation.

Note that the example RAF set in Fig. 1 requires at least one “spontaneous” (i.e., uncatalyzed) reaction event before the system forms a RAF (i.e. at least one of the two reactions needs to happen spontaneously before the first catalyst will be present when starting with only molecules from the food set). This is, of course, always possible, but at a much lower rate than when these same reactions are catalyzed. As a consequence, there may be a (stochastic) “waiting time” before this RAF set is realized in a dynamical sense. However, once a RAF is realized, it can, in principle, grow in concentration at an exponential rate, due to its (collectively) autocatalytic nature.

As argued elsewhere, this requirement for (rare) spontaneous reactions is actually a useful property for the potential evolvability of autocatalytic sets [10]. Furthermore, it is the main reason why catalysts are treated separately in RAF theory, instead of including them in a reaction as both a reactant and a product (which is often done in alternative chemical reaction network representations). Such an alternative representation requires the catalyst to always be present (as a reactant), preventing reactions from occasionally happening spontaneously.

RAF theory has been applied extensively to simple polymer-based models of chemical reaction networks (see next section), showing that autocatalytic sets are highly likely to exist in such models, and also for chemically realistic levels of catalysis. More specifically, it was shown both experimentally and theoretically that only a *linear* growth rate in the level of catalysis with increasing maximum polymer length is required to get RAF sets with high probability [7, 18], where the level of catalysis is expressed in terms of the average number of reactions catalyzed per molecule type. This is in stark contrast to the required *exponential* growth rate in the original argument of Kauffman [15, 16]. Also, as opposed to the original assumption that autocatalytic sets appear as “giant connected components” making up almost the entire reaction network [15, 16], RAF sets in the simple polymer model tend to be much smaller, consisting of only a fraction of the full reaction network [7]. In fact, the collection of RAF sets under the partial order of set inclusion is generally very large, with a unique maximal element (the maximal RAF), all the way down to the minimal RAF subsets (called *irreducible* RAFs) [12].

Furthermore, these main results have been shown to hold under a wide variety of chemically more realistic model assumptions, such as template-based catalysis (as with base-pair formation in RNA) [5], allowing only the longest polymers to be catalysts [13], allowing partitioned chemical reaction networks (as in nucleic acids vs. proteins) [20], or using power-law distributed catalysis [3].

However, autocatalytic sets are not just a theoretical construct, as they have also been created and studied in real chemical networks under controlled laboratory conditions [19, 1, 17, 24]. In fact, the formal RAF framework was used to analyze one of these real autocatalytic networks in detail [9], one consisting of 16 catalytic RNA molecules, or ribozymes [24]. Moreover, it was shown, using the RAF algorithm, that the metabolic network of *E. coli* forms a large autocatalytic set of close to 1800 reactions [21]. As far as we know, this is the first formal proof that living organisms (or at least essential parts thereof) can be realised as a RAF.

Finally, we have shown that “higher levels” of autocatalytic sets can emerge [12, 11]. For example, a boundary (such as a lipid layer) can be considered as an additional catalyst: it increases the rate at which reactions happen inside it by keeping the relevant molecules in close proximity rather than having them diffuse away, though the boundary itself is not used up in those reactions. This way, an “autocatalytic set of autocatalytic sets” emerges, which can form a simple (proto)cell-like structure [11]. This line of reasoning can, of course, be extended to the next emergent level of autocatalytic sets forming multicellular organisms, and so on, all the way up to the species level.

2. POLYMER MODELS

Polymers, such as RNA sequences and the amino acid sequences that comprise enzymes and proteins, play a central role in biochemistry and are likely to have also played an important role in early life. In previous work, we have considered a simple polymer-based chemical reaction system known as the *binary polymer model*, originally introduced by Kauffman [15, 16]. In this model, polymers are represented by bit strings that can be concatenated into longer ones or split up into shorter ones. Catalysis is assigned at random. Instances of this model then form a full CRS as follows:

- The molecule set X consists of all bit strings up to (and including) a maximum length n . In mathematical notation: $X = \{0, 1\}^{\leq n}$.
- The food set F consists of all bit strings up to (and including) length t , where $t \ll n$. In our experiments, we use $t = 2$, i.e., $F = \{0, 1, 00, 01, 10, 11\}$.
- The reaction set \mathcal{R} consists of two types of reactions:
 - (1) *Ligation*: gluing two bit strings together into a longer one, e.g., $00 + 111 \rightarrow 00111$.
 - (2) *Cleavage*: splitting a bit string into two smaller ones, e.g., $010101 \rightarrow 01 + 0101$.

All possible ligation and cleavage reactions between bit strings are included in \mathcal{R} , as long as none of their reactants or products violates the maximum molecule length n . Note that each ligation reaction has a corresponding cleavage reaction. In fact, in most of our experiments, these corresponding “forward” and “backward” reactions are counted as one bi-directional reaction.

- The catalysis set C , which consists of pairs (x, r) of a molecule types and reaction, is assigned according to some random process. In the original (uniform) model, each molecule type x has the same probability of catalyzing each reaction r . However, biochemical networks (e.g. the *E. coli* study from [21]) suggest that catalysis is far from uniform and that the distribution of catalysis events across molecule types has a longer tail that is closer to a power law than a binomial distribution. Thus a focus of this paper is to explore the different distributions of catalysis. We will present some new mathematical results that hold for a general distribution of catalysis rates, but we also focus on four particular ways of assigning catalysis:
 - (1) *Uniform*: For each molecule type $x \in X$ and reaction $r \in \mathcal{R}$, the pair (x, r) is independently included in C with probability p . This is the method used in the original model [15, 16].
 - (2) *Power law*: For each molecule type $x \in X$, first a random number $k_x \geq 1$ is drawn from a Zipf distribution with parameter a :

$$\mathbb{P}(k_x = k) = \frac{k^{-a}}{\zeta(a)},$$

where ζ is the Riemann Zeta function and k_x ranges over the positive integers. Next, molecule type x is assigned as a catalyst to $k'_x - 1$ distinct but randomly chosen reactions from the full set \mathcal{R} (which are drawn independently for each molecule type), where $k'_x = \min\{k_x, |\mathcal{R}|\}$. Note that we subtract one from the value of k'_x , so there are many molecule types that do not catalyze any reaction. This method was first investigated in [3].

- (3) *All-or-nothing*: Each molecule type $x \in X$ independently either catalyzes *all* reactions in \mathcal{R} with probability p , or x catalyzes no reaction with probability $1 - p$. This is an “extreme” model that we use for comparison below.
- (4) *Sparse*: Each molecule type $x \in X$ independently either (i) catalyzes each reaction $r \in \mathcal{R}$ independently with probability $\frac{1}{n}$ (this occurs with probability p), or (ii) x does not catalyze any reaction (this occurs with probability $1 - p$). This is a less extreme version of the all-or-nothing model.

In each assignment method, all reaction are considered to be bi-directional. In other words, if a molecule type is assigned as a catalyst to a reaction, it catalyzes both the forward and backward directions.

For a given value of the maximum molecule length n , the set of molecule types X and the set of reactions \mathcal{R} are fixed. In other words, the nodes (both the molecule nodes and the reaction nodes) and the reaction edges (represented graphically by solid black lines) are fixed in the reaction graph. However, each instance of the model (even for the same values of n and p) will have a different catalysis assignment C (i.e. a different configuration of catalysis edges, represented graphically by dashed grey lines). Fig. 2 shows an example of a (maximal) RAF set that was found by our RAF algorithm in an instance of the binary polymer model with parameters $n = 5$, $t = 2$, and $p = 0.0045$ (using the uniform catalysis assignment method).

As already mentioned above, the main results of RAF theory hold up under various alternative versions of the basic model that introduce more chemical realism. In fact, in some cases it is even possible to predict mathematically the required level of catalysis in these more elaborate versions from that in the standard model [8, 13]. So, the binary polymer model, even in its simplest version, already seems to hold much predictive and explanatory power. Moreover, the restriction to polymers that are binary is in no way serious, since, as we will see, most of the mathematical results translate directly




FIGURE 2. An example of a (maximal) RAF set as found by our algorithm in an instance of the binary polymer model with $n = 5$, $t = 2$, and $p = 0.0045$ (uniform catalysis). Note that the CRS that this RAF lies within has $|X| = 62$ molecule types and $|\mathcal{R}| = 196$ reactions.

to polymers over an arbitrary alphabet. For the simulations, however, it is convenient to stay with the binary polymer model.

In the next section, we derive new mathematical results that describe the probability of RAFs arising as the average catalysis rate varies in a class of polymer models that includes the four catalysis assignment methods described above. In Sections 4 and 5, we then use simulations of the binary polymer model to see how close these mathematical bounds come to the estimated probabilities of RAFs, and provide an explanation for the differences.

2.1. Inhibition (u -RAF). In this paper, we will also consider a version of the binary polymer model where molecules can *inhibit* reactions (i.e., prevent them from happening) in addition to catalyzing them. More specifically, similar to the catalysis set C , we define an inhibition set I of molecule–reaction pairs (x, r) , where $(x, r) \in I$ means that molecule type $x \in X$ *inhibits* the reaction $r \in \mathcal{R}$ (this could be represented by a dashed red edge, for example, in the graph representation of the CRS). These inhibition pairs are also chosen at random for each instance of the model, using the uniform method described above for assigning catalysis, and independent of the catalysis assignments.

In this setting, a CRS that allows inhibition is a tuple (X, \mathcal{R}, C, I) (with corresponding food set F). The definition of a RAF \mathcal{R}' can then be extended to require, in addition, that no reaction in \mathcal{R}' is inhibited by any molecule type that is produced by a reaction in \mathcal{R}' or present in the food set F .

We refer to such an “uninhibited” RAF as a u -RAF. We present a mathematical results on u -RAF in Section 3.1 and simulation results in Section 6.

3. MATHEMATICAL RESULTS ON RAF (AND u -RAF) SETS IN POLYMER MODELS

In this section, we allow a little more generality in considering polymers over an arbitrary alphabet of κ letters. To emphasize the dependence of the CRS on n , we will write $\mathcal{Q}_n = (X_n, \mathcal{R}_n, C)$, where X_n is the set of all polymers up to length n and \mathcal{R}_n is the associated set of all cleavage-ligation reactions. A fundamental relationship between the number of molecule types and reactions in \mathcal{Q}_n (Eqn. (4) from [18]) is:

$$(1) \quad 1 - O\left(\frac{1}{n}\right) \leq \frac{|\mathcal{R}_n|}{n|X_n|} \leq 1.$$

In particular, $|\mathcal{R}_n| \sim n|X_n|$.

We will also generalize catalysis beyond the four models described above to a setting where each molecule type x catalyzes Y_x reactions chosen uniformly at random from \mathcal{R}_n , where the random variable Y_x is independently and identically distributed (i.i.d.) according to some distribution \mathcal{D} (which may depend on n); we will call such models *distribution-based*. For example, in the uniform model, Y_x has the binomial distribution $\text{Bin}(|\mathcal{R}_n|, p)$, while in the sparse model, Y_x has the mixture distribution described by letting Y_x follow a binomial distribution $\text{Bin}(|\mathcal{R}_n|, \frac{1}{n})$ with probability p , and setting $Y_x = 0$ with probability $1 - p$.

For the uniform model, any collection of events of the form

$$(x_1, r_1) \in C, (x_2, r_2) \in C, \dots, (x_m, r_m) \in C$$

comprise stochastically independent events. However, for other distributions (including the other three mentioned above) these events can fail to be independent if the molecule types x_i are not distinct. For example, the two events $(x, r) \in C$ and $(x, r') \in C$ will be positively correlated in general (since although the events are conditionally independent given Y_x these two conditional probabilities both increase with Y_x). This subtlety means that certain mathematical arguments applicable to the uniform model need to be carefully adjusted to extend to distribution-based models (the models we study here do not fall under the umbrella of the ‘extended binary polymer model’ of [2], which required the stronger independence condition to hold).

We will let f denote the (average) *catalysis rate* in a CRS:

$$f = \sum_{r \in \mathcal{R}} \mathbb{P}((x, r) \in C),$$

for any $x \in X$ (by the symmetry in the model, different choices of x lead to the same value). Notice also that we can write $f = |\mathcal{R}_n| \mathbb{P}((x, r) \in C)$ for any choice of $x \in X$ and $r \in R$. For example, in the uniform model $f = |\mathcal{R}_n|p$ while in the sparse model, $f = p|\mathcal{R}_n|/n \sim p|X_n|$.

A main goal in RAF theory has been to estimate the probability P_n that an instance of the uniform random polymer model with parameter n contains a RAF. Bounds on P_n (for the uniform model) were initiated in [15], refined in [22], and extended further by simulations in [7], which led to a more detailed mathematical analysis in [18].

A main finding from this last paper was that RAFs arise in \mathcal{Q}_n precisely when the average number of reactions each molecule type catalyzes (f) grows (at least) linearly with n . This follows directly from the following pair of inequalities, both of which hold for all values of n . Recall that F (the food source) consists of all polymers up to a fixed length t . For any distribution-based polymer model, we let \tilde{P}_n be the probability that an instance of the model has a RAF in which every molecule type in X is generated from F using just the products of ligation reactions. By definition, $P_n \geq \tilde{P}_n$. For the uniform model we have:

(i) If $f \geq \lambda n$, then

$$(2) \quad P_n \geq \tilde{P}_n \geq 1 - \frac{\kappa(\kappa e^{-\lambda})^t}{1 - \kappa e^{-\lambda}},$$

which converges quickly to 1 as λ grows, and

(ii) if $f \leq \lambda n$, then

$$(3) \quad P_n = O(\lambda) \rightarrow 0 \text{ as } \lambda \rightarrow 0.$$

This linear rate of growth in f required for RAFs to arise in the uniform polymer model contrasts sharply with the provably exponential rate required for a RAF in which every catalyst must already be available the first time the reaction proceeds (i.e., when no spontaneous reactions are allowed) [18].

What can we say about the level of catalysis required for the emergence of RAFs in other (non-uniform) distribution-based polymer models?

For the all-or-nothing model, there is a very precise expression for P_n . If we write $f = \lambda n$ then:

$$(4) \quad P_n = 1 - \exp(-\lambda) + o(1),$$

where the $o(1)$ term decays to zero at an exponential rate with increasing n . The proof of Eqn. (4) is essentially trivial, since in the all-or-nothing model we have $f = p|\mathcal{R}_n|$ and so if we set $f = \lambda n$, then $p = \lambda n/|\mathcal{R}_n|$ and a RAF exists whenever at least one molecule type catalyzes one (and therefore every) reaction. This occurs with probability

$$1 - (1 - p)^{|X_n|} = 1 - \left(1 - \frac{\lambda n}{|\mathcal{R}_n|}\right)^{|X_n|} \sim 1 - \exp(-\lambda),$$

by the relation $n|X_n| \sim |\mathcal{R}_n|$ from Eqn. (1).

What is curious here is that in the all-or-nothing model, the same (linear) dependence of P_n on n occurs, even though the proof from [18] of the inequality (2) for the uniform model breaks down in one crucial step for the all-or-nothing model. However, the resulting RAFs in the all-or-nothing model are very different from those in the uniform model. For a start, in the all-or-nothing model, RAFs arise at the same catalysis level required for all of \mathcal{R}_n to be a RAF. By contrast, in the uniform model, the catalysis rate required for \mathcal{R}_n to be a RAF grows *quadratically* with n . More precisely, in the uniform model, if we let $f = \lambda n^2$ then:

$$(5) \quad \lim_{n \rightarrow \infty} \mathbb{P}(\mathcal{R}_n \text{ is a RAF}) = \begin{cases} 1, & \text{if } \lambda > \ln \kappa; \\ 0, & \text{if } \lambda < \ln \kappa. \end{cases}$$

A proof of this sharp transition result, which improves on a result from [22], is given in the Appendix.

A further difference between the all-or-nothing and the uniform model is that, for the latter model and at the linear catalysis rate at which RAFs appear, ‘small’ RAF are highly unlikely to arise (more precisely, with probability converging to 1 as n grows, any RAF must be larger than a certain function that grows exponentially with n) [23]. By contrast, in the all-or-nothing model, as soon as there is a RAF there is one of order n : namely, any catalyst and any sequence of $O(n)$ reactions that generates it from F .

The analysis of the sparse model is more interesting, and will also see shortly how the sparse model will also have small RAFs (of size $O(n^2)$) when RAFs appear, again in contrast to the uniform model.

We now state our first main mathematical result.

Theorem 1. *For any distribution-based polymer model, a necessary and sufficient condition for RAFs to arise is that the average catalysis rate f grows (at least) linearly with n . More precisely:*

(i) *If $f \geq \lambda n$ (and $\lambda > 2 \ln \kappa$), then*

$$(6) \quad P_n \geq \tilde{P}_n \geq \left(1 - \frac{\kappa(\kappa e^{-\lambda/2})^t}{1 - \kappa e^{-\lambda/2}}\right) (1 - e^{-\lambda/2}),$$

which converges to 1 at exponential rate as λ grows.

(ii) *If $f \leq \lambda n$ then for a constant M (independent of n), we have:*

$$(7) \quad P_n \leq M\lambda \rightarrow 0 \text{ as } \lambda \rightarrow 0.$$

Proof. We first show that a key identity (derived for the uniform model in [18]) holds for all distribution-based polymer models. Let q_* be the probability that any given reaction $r \in \mathcal{R}_n$ is *not* catalyzed by any molecule type in X_n . By the assumptions of the model, this probability is the same for each choice of r and is given by:

$$(8) \quad q_* = \prod_{x \in X_n} \left(1 - \frac{f}{|\mathcal{R}_n|}\right) \sim \exp(-f/n),$$

where the asymptotic equivalence is from Eqn. (1). The proof of Eqn. (8) exploits the independence of the events $\mathcal{E}_{x,r} := “x \text{ catalyzes } r”$ when r is fixed and x varies to give:

$$q_* = \mathbb{P}\left(\bigwedge_{x \in X_n} \overline{\mathcal{E}_{x,r}}\right) = \prod_{x \in X_n} \mathbb{P}(\overline{\mathcal{E}_{x,r}}) = \prod_{x \in X_n} (1 - \mathbb{P}(\mathcal{E}_{x,r})),$$

and Eqn.(8) now follows since $p = \mathbb{P}(\mathcal{E}_{x,r})$ and $f = |\mathcal{R}_n|p$.

To establish the lower bound inequality on \tilde{P}_n (Part (i)), we first observe that for any non-negative random variable W with finite expected value $\mu = \mathbb{E}[W]$, the following inequality holds:

$$(9) \quad \mathbb{P}(W < \sqrt{\mu}) \geq 1 - \sqrt{\mu}.$$

This inequality follows from the standard inequality $\mu \geq \mathbb{P}(W \geq w)w$ applied to $w = \sqrt{\mu}$ to deduce that $\mathbb{P}(W \geq \sqrt{\mu}) \leq \sqrt{\mu}$, from which the result follows immediately. Let N be the number of reactions in \mathcal{R}_n that are not catalyzed by any molecule type in X_n , and let $W = N/|\mathcal{R}_n|$. We have $\mathbb{E}[W] = q_*$. Let $\tilde{P}_n(W)$ be the probability that the polymer model has a RAF that generates every molecule type in X_n conditional on the random variable $W (= N/|\mathcal{R}_n|)$. Given certain reactions r_1, r_2, \dots, r_s , let $\overline{\mathcal{E}_i}$ be the event that reaction r_i is not catalyzed by any molecule type in X_n . Then, by exchangeability, we have the following expression for the joint conditional probability:

$$\mathbb{P}(\overline{\mathcal{E}_1} \wedge \overline{\mathcal{E}_2} \wedge \dots \wedge \overline{\mathcal{E}_s} | N) = \frac{N}{|\mathcal{R}_n|} \times \frac{N}{|\mathcal{R}_n|} \times \dots \times \frac{N - (s-1)}{|\mathcal{R}_n|} \leq \left(\frac{N}{|\mathcal{R}_n|}\right)^s = W^s.$$

We can now apply the argument in the proof of Theorem 4.1(ii) of [18] to deduce that:

$$(10) \quad \tilde{P}_n(W) \geq 1 - \frac{\kappa(\kappa W)^t}{1 - \kappa W}.$$

By the total law of expectation, for any w , we have:

$$\tilde{P}_n = \mathbb{E}[\tilde{P}_n(W)] = \mathbb{E}[\tilde{P}_n(W) | W < w] \mathbb{P}(W < w) + \mathbb{E}[\tilde{P}_n(W) | W \geq w] \mathbb{P}(W \geq w).$$

Since the second term on the right is non-negative, we have:

$$\tilde{P}_n \geq \mathbb{E}[\tilde{P}_n(W) | W < w] \mathbb{P}(W < w).$$

If we now set $w = \sqrt{\mathbb{E}[W]} = \sqrt{q_*}$, then Inequality (9) and Eqn. (10) gives:

$$(11) \quad \tilde{P}_n \geq 1 - \frac{\kappa(\kappa\sqrt{q_*})^t}{1 - \kappa\sqrt{q_*}}(1 - \sqrt{q_*}).$$

Finally, from Eqn. (8) and $f \geq \lambda n$, we obtain:

$$(12) \quad q_* \leq e^{-\lambda},$$

Combining Inequalities (11) and (12) and noting that the expression in (11) is decreasing with q_* , we obtain the inequality required for Part (i).

Turning to Part (ii), we now establish Inequality (7). From Eqn. (8), we have:

$$q_* = \left(1 - \frac{f}{|\mathcal{R}_n|}\right)^{|X_n|} \geq 1 - \frac{f|X_n|}{|\mathcal{R}_n|}.$$

Therefore, if $f \leq \lambda n$,

$$(13) \quad 1 - q_* \leq \lambda + o(1),$$

by Eqn. (1). The number M of reactions that take two food molecules and produce another molecule is, at most, $2|F|^2 = 2x_t^2$, where $x_t = \kappa + \kappa^2 + \dots + \kappa^t$. For there to be a RAF, at least one of these M forward reactions must be catalyzed. By Boole's inequality, the probability of this is bounded above by $M(1 - q_*)$, and so, from Eqn. (13), P_n is less than or equal to $M(\lambda + o(1))$, which converges to zero as $\lambda \rightarrow 0$. \square

Fig. 3 shows the behavior of P_n as a function of f for the binary polymer model, across our four exemplar distributions (for $n = 10$ and $n = 16$). The plots were obtained by simulations and the use of the RAF algorithm (except for the all-or-nothing model, where the exact formula was used). Two features are apparent. Firstly, the uniform model and the power law model show a much sharper transition from $P_n = 0$ to $P_n = 1$ than the other two models. Secondly, the effect of increasing n (from 10 to 16) has a larger effect on the curves for the sparse and all-or-nothing model than it does for the uniform and power law models. Theorem 1 also leads to the following interesting consequence for the structure of small RAFs in the sparse model, which is different to that for the uniform model.

Corollary 2. *Consider the sparse polymer model with $f = \lambda n$, where λ is sufficiently large such that (from Theorem 1) $\tilde{P}_n \geq 1 - \epsilon$. Then with probability at least $1 - \epsilon - o(1)$, any instance of the model will contain a RAF with less than $2\lambda n^2$ reactions (where $o(1)$ refers to a term that converges to zero exponentially with increasing n).*

Proof. In the sparse polymer model, the number of catalysts (i.e. molecule types that catalyze one or more reactions) has mean λn , and so (by the Chernoff bound) the probability that there are more than $2\lambda n$ catalysts differs from 1 by a term that converges to zero exponentially with increasing n . If $\tilde{P}_n \geq 1 - \epsilon$ then with probability at least $1 - \epsilon$, there is a RAF that constructs all molecule types in X_n from F using just the products of ligation reactions. Combining these two observations it follows that with probability at least $1 - \epsilon - o(1)$, there is a RAF whose set S of catalysts has size at most $2\lambda n$ and which constructs each molecule type $x \in X_n$ from F via some sequence $s(x) = r_1(x), r_2(x), \dots, r_m(x)$ of ligation reactions. Moreover, we must have $m < n$ (since each ligation reaction increases the size of a molecule by at least one), so the set $(r_i(x) : x \in S)$ is a RAF of size less than $2\lambda n \times n = 2\lambda n^2$. \square




FIGURE 3. Comparison across the four models of P_n (probability of a RAF) as a function of f (average catalysis rate per molecule) on the binary polymer model for $n = 10$ and $n = 16$.

3.1. The impact of inhibition. Consider now the uniform model in which, in addition to catalysis, each molecule type independently inhibits each reaction with a constant probability. Let h be the average number of reactions each molecule type inhibits, and, as before, let f be the average number of reactions each molecule type catalyzes ($f = p|\mathcal{R}_n|$). Recall that a u -RAF is a RAF \mathcal{R}' for which no reaction in \mathcal{R}' is inhibited by any molecule type in the set $S(\mathcal{R}')$ consisting of the products of reactions in \mathcal{R}' together with the molecule types in the food set F .

Clearly if $h > 0$ then the probability of a u -RAF is bounded above by a constant $\delta(h) < 1$ (even as $f \rightarrow \infty$) since there is a positive probability that all of the reactions that just involve reactants from the food set F are inhibited. Moreover $\delta(h) \rightarrow 0+$ as h grows, even if f were to grow exponentially (or faster) with h . Despite this ‘trumps-all’ feature of inhibition, one can still provide an inequality that relates the expected number of u -RAFTs in a uniform polymer model with rates $(2f, h)$ to the expected number of RAFTs in the uniform model with catalysis rate f .

Let us assume that the catalysis rate f grows linearly with n (this is a minor restriction since this rate suffices for RAFTs to arise) and allow h to also grow linearly with n but with a constant factor that decays exponentially with the corresponding constant factor for f . Under this restriction, Theorem 3 (below) shows that we expect (at least) as many u -RAFTs as there would be RAFTs if the catalysis rate had simply been halved.

Given a uniform polymer model \mathcal{Q}_n (over any finite alphabet of κ states) let $\mu(f, h)$ be the expected number of u -RAFTs as a function of f and h (the catalysis and inhibition rates, respectively) and let $\mu(f)$ be the expected number of RAFTs as a function of f .

Theorem 3. Suppose that $f \leq \lambda n \leq \frac{1}{2}|\mathcal{R}_n|$, and $h \leq \beta f$ where $\beta \leq n$ satisfies

$$(14) \quad \beta \leq \frac{1}{\lambda} \ln \left[\left(1 - \frac{\beta \lambda n}{|\mathcal{R}_n|} \right) \left(1 + e^{-\lambda} \left(1 - \frac{\lambda n}{|\mathcal{R}_n|} \right) \right) \right] \simeq \frac{\ln(1 + e^{-\lambda})}{\lambda}.$$

Then $\mu(2f, h) \geq \mu(f)$.

Proof. For a subset \mathcal{R}' of \mathcal{R} , let $\mathcal{E}_{\mathcal{R}'}$ be the event that \mathcal{R}' is a u -RAF. Then

$$(15) \quad \mu(f, h) = \sum_{\mathcal{R}' \in \mathbf{R}_{\text{FG}}} \mathbb{P}(\mathcal{E}_{\mathcal{R}'}),$$

where \mathbf{R}_{FG} be the collection of subsets \mathcal{R}' of \mathcal{R}_n that are F -generated. In the uniform model, we can write:

$$(16) \quad \mathbb{P}(\mathcal{E}_{\mathcal{R}'}) = \theta_{(f, h, s)}^{|\mathcal{R}'|},$$

where

$$\theta_{(f, h, s)} = \left[1 - \left(1 - \frac{f}{|\mathcal{R}_n|} \right)^s \right] \left(1 - \frac{h}{|\mathcal{R}_n|} \right)^s,$$

and where $s = |S(\mathcal{R}')|$ (the number of molecule types produced by reactions in \mathcal{R}' or present in the food set). If we now let $\nu = \frac{f}{|\mathcal{R}_n|}$, then

$$(17) \quad \theta_{(f, 0, s)} = [1 - (1 - \nu)^s].$$

Moreover,

$$(18) \quad \theta_{(2f, h, s)} \geq [1 - (1 - 2\nu)^s][1 - \beta\nu]^s \geq [1 - (1 - \nu)^{2s}][1 - \beta\nu]^s,$$

where the first inequality is because $h \leq \beta f$. The second inequality in (18) follows from the following generic inequality:

$$(1 - 2y)^k \leq (1 - y)^{2k},$$

for all integers $k \geq 1$ and $0 \leq y \leq \frac{1}{2}$ (this can be readily verified by induction on k) and noting that $0 \leq \nu = f/|\mathcal{R}_n| \leq \frac{1}{2}$ by the assumption on f in the statement of the theorem.

Combining Eqns. (17) and (18) and applying the identity $(1 - y^2)/(1 - y) = (1 + y)$ we find that the ratio of $\theta_{(2f, h, s)}$ to $\theta_{(f, 0, s)}$ satisfies

$$(19) \quad \frac{\theta_{(2f, h, s)}}{\theta_{(f, 0, s)}} \geq [1 + (1 - \nu)^s][1 - \beta\nu]^s.$$

Since $f \leq \lambda n$ (and so $\nu \leq \lambda n/|\mathcal{R}_n|$) the term on the right of this last equation is at least:

$$\left[1 + \left(1 - \frac{\lambda n}{|\mathcal{R}_n|} \right)^s \right] \left[1 - \frac{\beta \lambda n}{|\mathcal{R}_n|} \right]^s$$

and since $s \leq |X_n|$, Inequality (19) implies further that:

$$(20) \quad \frac{\theta_{(2f, h, s)}}{\theta_{(f, 0, s)}} \geq \left[1 + \left(1 - \frac{\lambda n}{|\mathcal{R}_n|} \right)^{|X_n|} \right] \left[1 - \frac{\beta \lambda n}{|\mathcal{R}_n|} \right]^{|X_n|}.$$

We now apply two inequalities: $|X_n| \geq |\mathcal{R}_n|/n$ (from Eqn. (1)) and the generic inequality $(1 - t/k)^k \geq e^{-t}(1 - t/k)$, for $t > 0$ and integer $k \geq 1$ (taking $k = |\mathcal{R}_n|/n$) to obtain:

$$\frac{\theta_{(2f, h, s)}}{\theta_{(f, 0, s)}} \geq (1 + e^{-\lambda}(1 - \lambda n/|\mathcal{R}_n|))(e^{-\beta\lambda}(1 - \beta\lambda n/|\mathcal{R}_n|)).$$

Notice that the quantity on the right of this last equation is greater or equal to 1 when β satisfies Eqn. (14). Thus, $\theta_{(2f,h,s)} \geq \theta_{(f,0,s)}$, and this holds for all possible values of $s = |S(\mathcal{R}')|$. Applying this inequality in Eqn. (16) and subsequently in Eqn. (15) leads to the inequality:

$$\mu(2f, h) \geq \mu(f, 0).$$

Since $\mu(f, 0) = \mu(f)$, this completes the proof of Theorem 3. \square

Remarks

- (i) As an application of Theorem 3, we saw in Fig. 3 that for the uniform binary polymer model with $n = 10$, the probability of a RAF is very close to 1 when $f = 2$. If we now put $\lambda = f/n = 0.2$ into Eqn. (14), this gives an upper bound value for β of $\beta \approx 3$. In other words, for the uniform binary polymer model with $n = 10$ we should expect at least as many u -RAFTs when the catalysis rate is $2f = 4$ and inhibition rate is at most $3f = 6$ as there are RAFs when $f = 2$. Indeed we will see in Section 6 that such u -RAFTs do exist.

An interesting feature of this example is that the inhibition rate is of similar magnitude (indeed larger) than the catalysis rate, and we expect to find u -RAFTs, even though inhibition is in some sense a stronger condition than catalysis (i.e. if a reaction is inhibited by one molecule type then no matter how many molecule types catalyze that reaction it cannot be in any u -RAFT). However, it is easily seen that this requires some fine balancing in the choice of h and f . Indeed, if we assume that h/f is greater or equal to some fixed value $c > 0$, then u -RAFTs can only arise when f lies within a certain window: If f is too small there will be no RAFs (and so no u -RAFTs) and when f becomes too large the growth in $h(> cf)$ will eventually ensure that all RAFs fail to be u -RAFTs.

- (ii) The approximation in the second half of Eqn. (14) becomes quickly accurate as n grows since the additional terms in the first half (involving $\lambda n/|\mathcal{R}_n|$ and $\beta \lambda n/|\mathcal{R}_n|$, where $\beta \leq n$) converge to zero at an exponential rate.

Notice also that Theorem 3 deals with the expected number of RAFs and u -RAFTs, while Theorem 1 is concerned with the probability of at least one RAF. Calculating exact non-asymptotic results for the latter quantity (for RAFs or u -RAFTs) appears to be a considerably harder exercise than calculating expected values.

4. LEVELS OF CATALYSIS WITHIN A RAF

In previous work, and above, we looked at the levels of catalysis in a CRS needed to obtain a given probability of finding RAF sets in random instances of the binary polymer model [7, 18, 4, 5, 3]. However, these analyses considered only the overall level of catalysis required in the *full* CRS, expressed as the average number f of reactions catalyzed per molecule type. One question that has remained unexplored so far, is whether this overall level of catalysis is also reflected within the actual RAF sets that occur in these model instances, or whether those RAF sets, somehow, “exploit” local variations in the overall catalysis level. Here we investigate this question in more detail.

Using the standard notation, given a CRS $\mathcal{Q} = (X, R, C)$ and a RAF $\mathcal{R}' \subseteq \mathcal{R}$, let $f_{\mathcal{R}}$ be the average number of reactions catalyzed per molecule type in the full CRS (i.e. the quantity we have so far denoted by f), and let $f_{\mathcal{R}'}$ be the average number of reactions in \mathcal{R}' that are catalyzed by molecule types involved in \mathcal{R}' (i.e. by molecule types that are produced by reactions in \mathcal{R}' or present in the food set F). We then generated 10,000 random instances of the binary polymer model for $n = 10$, $t = 2$, and various values of p (for the uniform method) or a (for the power law method) in such a way that the probability P_n of finding RAF sets ranged roughly from 0.1 to 0.9. In these model instances we

then measured and compared the levels of catalysis $f_{\mathcal{R}}$ in the full CRSs and $f_{\mathcal{R}'}$ in the RAFs that were found by our RAF algorithm within these model instances.

4.1. Uniform catalysis. Fig. 4 (left) shows the different levels of catalysis required for increasing probabilities P_n of finding RAF sets for the uniform catalysis version of the model. The black line (“CRS”) shows the average observed $f_{\mathcal{R}}$ values for those model instances that *do* have a RAF set. The red line (“RAF”) shows the corresponding average observed $f_{\mathcal{R}'}$ values (i.e., the average level of catalysis *within* those RAFs). Surprisingly, the level of catalysis within RAFs is much lower than that in the full CRS of which they are a part, and in all cases this difference is statistically highly significant. Finally, the blue line (“noRAF”) shows the average observed $f_{\mathcal{R}}$ values for those model instances that do *not* have a RAF set. Not surprisingly, this level is also lower than that of model instances *with* a RAF. If the level of catalysis is too low, RAF sets will simply not exist. However, what is striking is that the level of catalysis *within* a RAF is even lower than that in model instances that do not have a RAF set at all. Again, this difference is statistically significant.




FIGURE 4. *Left:* The average $f_{\mathcal{R}}$ (“CRS”) and $f_{\mathcal{R}'}$ (“RAF”) for the uniform binary polymer model ($n = 10$) over those model instances with a RAF set, and the average $f_{\mathcal{R}}$ (“noRAF”) over those instances without a RAF, plotted against the probability P_n of finding RAF sets. *Right:* A scatter plot of $f_{\mathcal{R}'}$ vs. $f_{\mathcal{R}}$ for $P_n \approx 0.50$ for the uniform binary polymer model with $n = 10$. The solid line represents $f_{\mathcal{R}'} = f_{\mathcal{R}}$, while the dashed line represents a linear fit.

Initially, we had expected that if there was any difference between $f_{\mathcal{R}}$ and $f_{\mathcal{R}'}$, the catalysis level $f_{\mathcal{R}'}$ within the RAF set would be *higher* than the level $f_{\mathcal{R}}$ in the full CRS, assuming that the RAF set would exploit local variations in catalysis levels where they are higher than in other parts of the network. However, exactly the opposite appears to be the case. One possible explanation could be that if catalysts are assigned randomly to reactions, as in the binary polymer model, there will be many reactions that are catalyzed by more than one molecule type. To form a RAF set, in principle only one catalyst is necessary per reaction, so a RAF set might primarily include those molecule types and reactions that happen to have a more “efficient” (local) distribution of catalysis.

Fig. 4 (right) also shows a scatter plot of the $f_{\mathcal{R}'}$ values (“RAF set”) vs. the corresponding $f_{\mathcal{R}}$ values (“Full network”) for the $\sim 5,000$ model instances that contained a RAF set around $P_n = 0.50$. The solid line indicates where $f_{\mathcal{R}'} = f_{\mathcal{R}}$. Apart from a handful of exceptions, the dots are all well below the line (i.e., $f_{\mathcal{R}} > f_{\mathcal{R}'}$ in almost all instances). The dashed line shows the linear fit from a standard regression analysis.

4.2. Power law catalysis. The same analysis was performed for the power law catalysis version of the binary polymer model. Fig. 5 (left) shows the results for the average catalysis levels f . Here the




FIGURE 5. *Left:* The average $f_{\mathcal{R}}$ (“CRS”) and $f_{\mathcal{R}'}$ (“RAF”) for the power-law binary polymer model ($n = 10$) over those model instances with a RAF set, and the average $f_{\mathcal{R}}$ (“noRAF”) over those instances without a RAF, plotted against the probability P_n of finding RAF sets. *Right:* A scatter plot of $f_{\mathcal{R}'}$ vs. $f_{\mathcal{R}}$ for $P_n \approx 0.50$ for the power-law binary polymer model with $n = 10$. The solid line represents $f_{\mathcal{R}'} = f_{\mathcal{R}}$.

results are partly in line with our initial expectation, in that RAF sets do indeed seem to “exploit” differences in the catalysis level where it is locally higher. With power law catalysis, there is always a chance that there are a few molecule types that catalyze many reactions (so-called “hubs”, or highly connected nodes). Averaged over the entire network (with many molecule types that do not catalyze any reactions at all), this may not have a very large impact. However, a RAF set that is primarily constructed around these hubs will indeed have a higher level of catalysis than the full network of which they are part. This seems to happen up to a level of about $P_n = 0.40$. After that, the differences between the full CRS and the RAF are not statistically significant anymore (at the 5% level).

What is even more striking here is the difference in the average level of catalysis $f_{\mathcal{R}}$ between model instances that *do* (black line) and those that do *not* (blue line) contain a RAF set. Again, this is mostly due to assigning catalysis according to a power law, which results in a much larger variation in $f_{\mathcal{R}}$ values among model instances compared to the uniform catalysis method.

Also in Fig. 5 (right), the larger variance can be clearly seen in the scatter plot of the $f_{\mathcal{R}'}$ values (“RAF set”) vs. the corresponding $f_{\mathcal{R}}$ values (“Full network”) for the $\sim 5,000$ model instances that contained a RAF set around $P_n = 0.50$. Note the much larger range of f values compared to the uniform catalysis case above. The solid line again indicates where $f_{\mathcal{R}'} = f_{\mathcal{R}}$. Here, the majority of the dots are close to or slightly above this line, with only a large portion of dots below the line (i.e., $f_{\mathcal{R}} > f_{\mathcal{R}'}$) for smaller values of f .

5. GROWTH RATES IN LEVEL OF CATALYSIS

Two other issues that were only partially addressed in previous work are related to the growth rate in the level of catalysis f needed for RAFs to arise with increasing maximum molecule length n in

the binary polymer model. First, although the mathematical results show that linear growth with n is both necessary and sufficient, the constants involved in the upper and lower bounds are quite far apart. Second, in the power law catalysis version of the model it seemed that there is no increase in the level of catalysis required at all (with increasing n), raising the question of whether this phenomenon persists for large values of n . Both these issues are resolved here.

5.1. Theoretical vs. experimental rates. There is a clear gap between the actual growth rate in level of catalysis needed for autocatalytic sets to form, as experimentally observed [7] and the lower bounds predicted by theory (e.g. Eqns. (2) and (6)). For the uniform catalysis version of the binary polymer model, both the theoretical and experimental rates were shown to follow a linear relationship in n (the maximum molecule length) but the actual slopes differ. This is due to several simplifying assumptions in the theoretical proof. We showed previously that at least one of those assumptions can partially explain this difference [5]. Here, we show that a second simplifying assumption in the proof can explain most of the remaining difference.

The slope of the linear relationship as calculated by the lower bound on P_n given by Theorem 4.1 (ii) in [18], is given by $\ln\left(\frac{-1+\sqrt{17}}{16}\right) \approx 1.63$ (setting $P_n = 0.50$ and $\kappa = t = 2$). However, the slope derived from the original experiments in [7] is just under 0.02, a difference of two orders of magnitude. Notice, however, that the bound in Theorem 4.1(i) in [18] is for a RAF that contains all the molecule types in X , a much stronger requirement than is necessary in practice. Indeed, in [5], we showed that this assumption can explain a large part of this difference by performing experiments that require all molecule types in X to be in the RAF. This provided a slope of 0.70 [5].

Also, the bound in Theorem 4.1(i) requires only *forward* (or ligation) reactions to be considered (indicated by \mathcal{R}_+ in [18]). Recall that in the binary polymer model each ligation reaction has a corresponding cleavage (or *backward*) reaction, and if a molecule type is assigned as a catalyst to a forward reaction, it is also automatically assigned as a catalyst to the corresponding backward reaction. In other words, a forward reaction and the corresponding backward reaction always come in pairs, and are catalyzed by the same molecule type(s). Backward (cleavage) reactions are useful though, as they can potentially provide shorter molecule types (bit strings) that can then be used to create necessary longer molecule types through one ligation reaction, whereas otherwise these longer molecule types may have to be built up through several consecutive ligation reactions, starting all the way from the food set.

Indeed, it turns out that when only forward reactions are allowed in the binary polymer model, a higher level of catalysis f is required to get the same probability of finding RAF sets compared to the model with both forward and backward reactions. Redoing the experiments using only forward reactions *and* requiring all molecule types in X to be in the RAF, results in a slope of 1.51.

Case	Slope
standard	0.02
all- X	0.70
F & all- X	1.51
theory	1.63

TABLE 1. The slopes of the linear relationship for the growth rate in required level of catalysis, with increasing maximum molecule length n , as derived from the four different cases described in the text.

Table 1 summarizes these results, comparing the slopes for the four cases described: (i) the original experiments with the standard (uniform catalysis) binary polymer model (“standard”), (ii) the

experiments requiring all molecule types in X to be in the RAF (“all- X ”), (iii) the experiments using only forward reactions *and* requiring all molecule types in X to be in the RAF (“F & all- X ”), and (iv) the theoretical proof (“theory”). Clearly, the additional assumption of only forward reactions can explain most of the remaining difference between the theoretical and the experimental slopes in the linear relationship for the growth in level of catalysis with increasing n .

5.2. Power law catalysis. In contrast to the uniform catalysis version of the binary polymer model, the power law catalysis version seemed to require *no* increase in the level of catalysis, with increasing n , to get a probability $P_n \approx 0.50$ of finding RAF sets. In Fig. 9 in [3], P_n vs. f is plotted for various




FIGURE 6. The probability P_n of finding RAF sets against the average level of catalysis f for various values of the maximum molecule length n in the binary polymer model with power law distributed catalysis.

values of n (up to $n = 12$) for the power law case. However, rather than slowly moving to the right (i.e. towards higher f values, as is the case for the uniform model [7]), the curves for increasing n for the power law case cross each other around $f = 1.5$ [3]. This result seems somewhat counter-intuitive, especially as Theorem 1(ii) above implies that the curves must eventually start moving to the right as n grows. It turns out that this phenomenon is indeed a result of the relatively small values of n (maximum molecule length) that were used in the initial experiments. We have redone the experiments with the power law catalysis version of the model, and also including larger values of n (up to $n = 20$). In this case, the curves stop crossing and start to move slowly to the right for $n \geq 16$, as shown in Fig. 6.

6. INHIBITION

In this final section, we return to the issue of inhibition, discussed earlier. In biological systems, inhibitors play an important role in regulation (genetic, metabolic, or otherwise). As described in Section 2 above, the notion of inhibition can be included in the RAF framework as well. However, in

[18] it was shown that the problem of finding a u -RAF (or even determining whether or not one exists in a CRS allowing inhibition) is NP-hard. So, unlike the case with “regular” RAFs, we cannot expect to have an efficient algorithm for detecting u -RAFTs in general CRSs that allow inhibition.

In [6], we introduced a fixed-parameter tractable algorithm for finding u -RAFTs (including a maximal u -RAF), with the total number m of inhibitors in the system as a parameter. If m is small enough, the problem can still be solved efficiently, as we showed with an application to the binary polymer model and using $m = 10$ [6].

Here, we introduce an efficient iteration algorithm for finding u -RAFTs, independent of the value of m . This algorithm is polynomial in the size $|\mathcal{R}|$ of the CRS, and provides a “sufficient condition” for a u -RAF to exist. In other words, if the iteration algorithm returns a non-empty result, the CRS contains a u -RAF, and therefore also a maximal u -RAF (although this could be larger than the u -RAF found by the algorithm). However, if the iteration algorithm returns an empty set, the CRS may still contain one or more u -RAFTs (i.e., more work needs to be done to find out).

The iteration algorithm works as follows. Given a “regular” CRS plus an inhibition set I (i.e., a set of molecule type–reaction pairs (x, r) such that molecule type x inhibits reaction r), perform the following steps:

- (1) Let \mathcal{R}_1 be the “regular” RAF for the given CRS (i.e., not taking inhibition into account and using the standard RAF algorithm).
- (2) Let \mathcal{R}_2 be the subset of \mathcal{R}_1 consisting of all reactions in \mathcal{R}_1 that are *not* inhibited by any of the molecule types produced by reactions in \mathcal{R}_1 or present in the food set F .
- (3) Let \mathcal{R}_3 be the “regular” RAF contained within \mathcal{R}_2 (if any).

Clearly, if \mathcal{R}_3 is non-empty, then it is a u -RAF for the given CRS (although it may not be a maximal u -RAF).

As a test of this algorithm, we searched for u -RAFTs in the binary polymer model ($n = 10$) for the setting discussed in Remark (i) (following the proof of Theorem 3). This involves a catalysis rate of 4

(1)	(2)
6393	2547
6415	2586
6447	2729
6410	2716
6351	2581
6359	2667
6395	2718
6397	2512
6418	2631
6438	2625

TABLE 2. The sizes of (1) the maximal RAF not taking inhibitors into account, and (2) the u -RAF from the iterative algorithm, across 10 random instances of the uniform binary polymer model containing a RAF (with $n = 10$, catalysis rate $f = 4$ and inhibition rate $h = 6$).

and an inhibition rate of 6. Theorem 3 coupled with Fig. 3 suggested that u -RAFTs would likely exist, and applying the iteration algorithm over 10 random instances succeeded in detecting u -RAFTs in all

cases, as shown in Table 2. The fixed-parameter tractable u -RAF algorithm would be quite infeasible here, as $m = |X_{10}| = 2046$.

6.1. Efficiency of iteration algorithm. Since the “standard” RAF algorithm has a running time that is polynomial in the size $|\mathcal{R}|$ of the given CRS, the iteration algorithm above for finding u -RAFs is clearly also polynomial-time. We now compare the results of the iteration algorithm with that of the earlier exact algorithm on model instances in which the number m of inhibitor molecule types is bounded (this exact algorithm, by contrast, has a complexity that grows exponentially with m). We then apply the iteration algorithm to model instances that are beyond the feasible reach of the exact method.

Note that in this section we continue to work in the uniform model, but now there are m molecule types that are inhibitors, and for each of these the probability the molecule type inhibits any particular reaction is q . This is therefore quite different to the set-up in Section 3.1 (and discussed above) where each molecule type has a constant probability of inhibiting each reaction.

Using the same parameter values as in the previous study (i.e., $n = 10$, $t = 2$, $p = 0.0000792$ [giving a probability P_n of 0.50 to get “regular” RAFs]), $m = 10$, and inhibition probabilities $q = 10 \times p$ and $q = 100 \times p$ [6], the results are shown in Table 3 (including the additional parameter value $q = 50 \times p$).

For each of the three cases (different values for the inhibition probability q), 10 instances of the model that contain a “regular” RAF were taken and both the exact and the iteration algorithm for finding u -RAFs were applied. In each table, the column labeled (1) shows the size of the (regular) maximal RAFs for each of the 10 instances. Column (2) shows the size of the maximal u -RAF that was found by the exact algorithm. Column (3) shows the size of the u -RAF found by the iteration algorithm.

$q = 10 \times p$			$q = 50 \times p$			$q = 100 \times p$		
(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
1544	1510	1506	1339	1182	0	1424	1290	1
1459	1451	1451	1395	0	0	1645	1608	0
1238	1206	0	1196	1181	0	1526	1473	0
1732	1728	1728	1507	1469	1198	1447	0	0
1447	1423	1423	1481	1461	1347	1441	1277	1
1466	1457	1457	1480	1234	1150	1606	1583	0
1478	1443	1306	1501	1469	0	1419	1397	0
1516	1508	1508	1596	1509	1442	1457	1348	0
1302	1293	1272	1319	1290	0	1147	1117	0
1490	1444	1392	1642	1585	1350	1184	0	0

TABLE 3. For $m = 10$, the sizes of (1) the maximal RAF not taking inhibitors into account, (2) the maximal u -RAF, and (3) the u -RAF from the iteration algorithm, for $q = 10 \times p$ (left) and $q = 50 \times p$ (middle) and $q = 100 \times p$ (right), across ten random instances containing a RAF.

For the case where $q = 10 \times p$ (Table 3, left), the iteration algorithm clearly performs almost as well as the exact algorithm. In all but one case, the iteration algorithm also finds a u -RAF, and its size is close to or often even equal to that of the maximal u -RAF found by the exact algorithm.

However, for $q = 50 \times p$ (Table 3, middle), the situation is already quite different, with the iteration algorithm only finding a u -RAF in about half of the cases, and with a size significantly smaller than

that of the maximal u -RAF found by the exact algorithm. Note that in four of the five cases where the iteration algorithm did *not* find a u -RAF, there actually was one.

Finally, for $q = 100 \times p$ (Table 3, right) the situation is much worse still. In most cases the iteration algorithm does not find a u -RAF at all (or otherwise only a “trivial” one of size 1), while in 8 out of the 10 cases there actually is one.

Clearly, the performance of the iteration algorithm depends on the ratio of the rates of catalysis to inhibition. Recall that the catalysis rate f is defined as the average number of reactions catalyzed per molecule type, which, in the binary polymer model, can be expressed as [7]:

$$f = \frac{p|X||\mathcal{R}|}{|X|} = p|\mathcal{R}|.$$

Similarly, the inhibition rate is the average number of reactions inhibited per molecule type:

$$h = \frac{qm|\mathcal{R}|}{|X|} = \frac{spm|\mathcal{R}|}{|X|},$$

where s is the factor by which the catalysis probability p is multiplied to get the inhibition probability q (i.e., $s = 10, 50, 100$ in our experiments). Therefore, the ratio f/h of the level of catalysis to the level of inhibition is:

$$\frac{f}{h} = \frac{|X|}{sm}.$$

For $n = 10$, the number of molecule types in the CRS is $|X| = 2046$. For $s = 10$ and $m = 10$, this ratio f/h is about one order of magnitude in size, and the iteration algorithm still works quite well. But for $s = 100$ and $m = 10$, this ratio is close to one (i.e., the two levels are similar), and the iteration algorithm may fail to detect u -RAFTs when they exist, or find ones that are much smaller than the maximal ones.

However, the main advantage of the iteration algorithm is that it can be used for any value of m (number of inhibitors). For example, for $m = 100$, the exact algorithm is not feasible anymore, but the iteration algorithm can still be used as before (its running time does not depend on m). Using the same parameter values, but with $m = 100$ and $q = p$ (giving a ratio f/h about one order of magnitude in size), the results are as shown in Table 4. In all cases, the iteration algorithm found a u -RAF, although it is not known whether its size is close to the maximal u -RAF.

$m = 100, q = p$		
(1)	(2)	(3)
1417	?	1402
1278	?	1161
1225	?	1211
1445	?	1330
1347	?	1335
1451	?	1389
1344	?	1285
1531	?	1492
1508	?	1451
1312	?	1175

TABLE 4. The sizes of (1) the maximal RAF without taking inhibitors into account, (2) the maximal u -RAF, and (3) the u -RAF from the iteration algorithm, for $m = 100$, $q = p$ (right), for 10 random instances containing a RAF.

In conclusion, the performance of the iteration algorithm depends largely on the ratio f/h of the level of catalysis relative to the level of inhibition. If the ratio is large enough (one order of magnitude or larger), the iteration algorithm performs quite well. If this ratio is too small (close to equality), the iteration algorithm may not detect u -RAFs when they exist. However, the advantage of the iteration algorithm is that it can be used for any value of m (i.e., for any number of inhibitors) to determine whether there exist uninhibited RAFs in CRSs that also allow inhibition.

7. FUTURE WORK

Our results suggest a number of plausible conjectures and questions for further study. Firstly, the sharpness of the transition from $P_n = 0$ to $P_n = 1$ for the uniform model (and perhaps also the power law model), evident in Fig. 3, suggests that a 0-1 law (analogous to Eqn.(5)) may hold. Namely, for some constant γ under the uniform model, we have:

$$(21) \quad \lim_{n \rightarrow \infty} \mathbb{P}(\mathcal{Q}_n \text{ has a RAF}) = \begin{cases} 1, & \text{if } f \geq \lambda n \text{ for } \lambda > \gamma; \\ 0, & \text{if } f \leq \lambda n \text{ for } \lambda < \gamma. \end{cases}$$

Settling this conjecture seems a reasonable objective. Clearly, such a sharp transition does not hold for all distribution-based models, since it Eqn.(4) shows that it fails for the all-or-nothing model and it might possibly also fail for the sparse model. Simulations from [7] suggest that γ might be close to 0.02.

The sizes of the smallest RAFs at catalysis levels where RAFs first arise is also an interesting question. Here, the uniform model behaves quite differently from the sparse and all-or-nothing models: small (sub-exponential size) RAFs in the uniform model are highly unlikely when they first emerge, while in the other two models they will always be present. The situation for the power law model is less clear. Simulations and the application of the algorithm to find irreducible RAFs suggest that there are smaller RAFs in this model than in the uniform [3], but a mathematical analysis of the existence (or not) of small RAFs in the power law setting could be of interest.

Finally, although we showed that the sparse model is likely to have RAFs of, at most, quadratic size in n (when RAFs first appear), it is not clear if this can be improved to order n . In particular, is it likely that there are RAFs in which a single molecule type catalyzes some sequence of reactions that suffice to construct that molecule? A simple mathematical argument shows that this becomes increasingly unlikely as n grows if we restrict ourselves to just ligation reactions (results not shown). However, the situation is less clear without this restriction.

ACKNOWLEDGMENTS

WH thanks the Konrad Lorenz Institute Klosterneuburg for financial support in the form of a fellowship. We also thank the NLHPC, Santiago, Chile (www.nlhpc.cl) for allowing us to use their high performance computing infrastructure.

REFERENCES

- [1] G. Ashkenasy, R. Jegasia, M. Yadav, and M. R. Ghadiri. Design of a directed molecular network. *PNAS*, 101(30):10872–10877, 2004.
- [2] W. Hordijk. Autocatalytic sets: From the origin of life to the economy. *BioScience*, 63(11):877–881, 2013.
- [3] W. Hordijk, L. Hasenclever, J. Gao, D. Mincheva, and J. Hein. An investigation into irreducible autocatalytic sets and power law distributed catalysis. *Natural Computing*, 13(3):287–296, 2014.

- [4] W. Hordijk, J. Hein, and M. Steel. Autocatalytic sets and the origin of life. *Entropy*, 12(7):1733–1742, 2010.
- [5] W. Hordijk, S. A. Kauffman, and M. Steel. Required levels of catalysis for emergence of autocatalytic sets in models of chemical reaction systems. *International Journal of Molecular Sciences*, 12(5):3085–3101, 2011.
- [6] W. Hordijk, J. I. Smith, and M. Steel. Algorithms for detecting and analysing autocatalytic sets. *Algorithms for Molecular Biology*, 10:15, 2015.
- [7] W. Hordijk and M. Steel. Detecting autocatalytic, self-sustaining sets in chemical reaction systems. *Journal of Theoretical Biology*, 227(4):451–461, 2004.
- [8] W. Hordijk and M. Steel. Predicting template-based catalysis rates in a simple catalytic reaction model. *Journal of Theoretical Biology*, 295:132–138, 2012.
- [9] W. Hordijk and M. Steel. A formal model of autocatalytic sets emerging in an RNA replicator system. *Journal of Systems Chemistry*, 4:3, 2013.
- [10] W. Hordijk and M. Steel. Conditions for evolvability of autocatalytic sets: A formal example and analysis. *Origins of Life and Evolution of Biospheres*, 44(2):111–124, 2014.
- [11] W. Hordijk and M. Steel. Autocatalytic sets and boundaries. *Journal of Systems Chemistry*, 6:1, 2015.
- [12] W. Hordijk, M. Steel, and S. Kauffman. The structure of autocatalytic sets: Evolvability, enablement, and emergence. *Acta Biotheoretica*, 60(4):379–392, 2012.
- [13] W. Hordijk, P. R. Wills, and M. Steel. Autocatalytic sets and biological specificity. *Bulletin of Mathematical Biology*, 76(1):201–224, 2014.
- [14] S. A. Kauffman. Cellular homeostasis, epigenesis and replication in randomly aggregated macromolecular systems. *Journal of Cybernetics*, 1(1):71–96, 1971.
- [15] S. A. Kauffman. Autocatalytic sets of proteins. *Journal of Theoretical Biology*, 119:1–24, 1986.
- [16] S. A. Kauffman. *The Origins of Order*. Oxford University Press, 1993.
- [17] T. A. Lincoln and G. E. Joyce. Self-sustained replication of an RNA enzyme. *Science*, 323:1229–1232, 2009.
- [18] E. Mossel and M. Steel. Random biochemical networks: The probability of self-sustaining autocatalysis. *Journal of Theoretical Biology*, 233(3):327–336, 2005.
- [19] D. Sievers and G. von Kiedrowski. Self-replication of complementary nucleotide-based oligomers. *Nature*, 369:221–224, 1994.
- [20] J. Smith, M. Steel, and W. Hordijk. Autocatalytic sets in a partitioned biochemical network. *Journal of Systems Chemistry*, 5:2, 2014.
- [21] F. L. Sousa, W. Hordijk, M. Steel, and W. F. Martin. Autocatalytic sets in *E. coli* metabolism. *Journal of Systems Chemistry*, 6:4, 2015.
- [22] M. Steel. The emergence of a self-catalysing structure in abstract origin-of-life models. *Applied Mathematics Letters*, 3:91–95, 2000.
- [23] M. Steel, W. Hordijk, and J. Smith. Minimal autocatalytic networks. *Journal of Theoretical Biology*, 332:96–107, 2013.
- [24] N. Vaidya, M. L. Manapat, I. A. Chen, R. Xulvi-Brunet, E. J. Hayden, and N. Lehman. Spontaneous network formation among cooperative RNA replicators. *Nature*, 491:72–77, 2012.
- [25] R. A. van Santen and M. Neurock. *Molecular Heterogeneous Catalysis*. Wiley-VCH, 2006.

APPENDIX: PROOF OF EQN. (5)

It suffices to establish the following stronger result: if $\lambda = \ln \kappa + c > 0$, and $f = \lambda n^2$, then for $\beta := \frac{\kappa}{\kappa-1}$, we have:

$$(22) \quad \mathbb{P}(\mathcal{R}_n \text{ is a RAF in } \mathcal{Q}_n) \sim \exp(-(n\beta - \beta^2)e^{-cn}),$$

which converges to 0 or 1 depending on whether or not $c < 0$ or $c > 0$ respectively. To establish Eqn. (22), note that for $f = (\ln \kappa + c)n^2$, Eqn. (8) gives: $q_* \sim \kappa^{-n}e^{-cn}$. Now, the entire reaction set \mathcal{R}_n is a RAF precisely if every reaction from \mathcal{R}_n is catalyzed by at least one molecule type, and so

$$\mathbb{P}(\mathcal{R}_n \text{ is a RAF in } \mathcal{Q}_n) \sim \left[1 - \frac{1}{\kappa^n}e^{-cn}\right]^{|\mathcal{R}_n|} \sim \exp\left(-\frac{|\mathcal{R}_n|}{\kappa^n}e^{-cn}\right).$$

Eqn. (22) follows by noting that $|\mathcal{R}_n|/\kappa^n \sim n\beta - \beta^2$ from Eqn. (3) of [18]. □