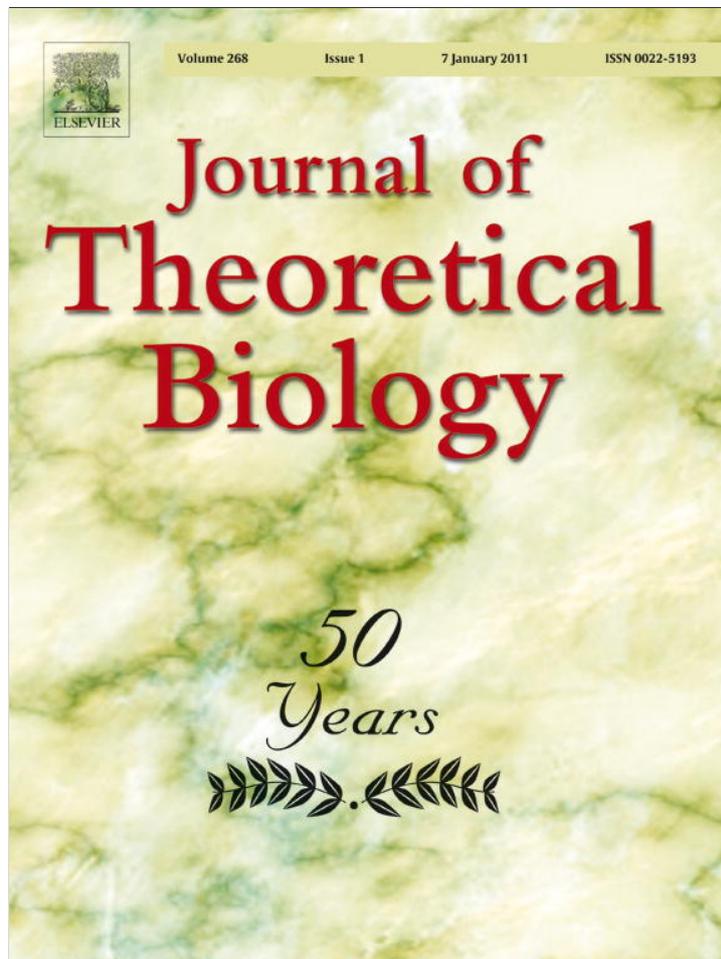


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



(This is a sample cover image for this issue. The actual cover is not yet available at this time.)

This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at SciVerse ScienceDirect

Journal of Theoretical Biology

journal homepage: www.elsevier.com/locate/jtbi

Predicting template-based catalysis rates in a simple catalytic reaction model

Wim Hordijk^{a,*}, Mike Steel^b^a University of Lausanne, Department of Ecology and Evolution, 1015 Lausanne, Switzerland^b University of Canterbury, Biomathematics Research Centre, Private Bag 4800, Christchurch, New Zealand

ARTICLE INFO

Article history:

Received 1 October 2011

Received in revised form

20 November 2011

Accepted 21 November 2011

Available online 2 December 2011

Keywords:

Catalytic reaction system

Template-based catalysis

Random autocatalytic network

Origin of life

ABSTRACT

We show that in a particular model of catalytic reaction systems, known as the binary polymer model, there is a mathematical concordance between two versions of the model: (1) random catalysis and (2) template-based catalysis. In particular, we derive an analytical calculation that allows us to accurately predict the (observed) required level of catalysis in one version of the model from that in the other version, for a given probability of having self-sustaining autocatalytic sets exist in instances of both model versions. This provides a tractable connection between two models that have been investigated in theoretical origin-of-life studies.

© 2011 Elsevier Ltd. All rights reserved.

1. Introduction

In previous work, we introduced and analyzed a model of catalytic reaction systems (CRS) (Steel, 2000; Hordijk and Steel, 2004; Mossel and Steel, 2005; Hordijk et al., 2011), based on the original model of Kauffman (1986, 1993). This model consists of molecule types (bit strings up to a given length n), reactions (ligation and cleavage), and randomly assigned catalysis, where molecule types can catalyze reactions with a certain (given) probability. Such models were developed and investigated within the context of theoretical origin-of-life studies. In this setting, a question of particular interest is the level of catalysis required to have a high probability of *autocatalytic sets* (“closed”, self-sustaining subsets of molecules and reactions) appearing in instances of this random binary polymer CRS model. We showed, both theoretically and computationally, that this level (in terms of the average number of reactions catalyzed by any one molecule) needs only grow linearly with n (the maximum molecule size in the system) to get a high probability of autocatalytic sets (Hordijk and Steel, 2004; Mossel and Steel, 2005). However, in Hordijk et al. (2011), we showed that there is still some discrepancy between the theoretically predicted and computationally observed required levels of catalysis, and we provided a partial explanation for this discrepancy. The theoretical analysis thus provides an important (linear) upper bound, but not (yet) an exact prediction.

* Corresponding author. Tel.: +41 216924268.

E-mail addresses: wim@WorldWideWanderings.net,
wim@santafe.edu (W. Hordijk), mike.steel@canterbury.ac.nz (M. Steel).

In Hordijk et al. (2011), we also analyzed an extension of the basic model, based on initial experiments reported in Kauffman (1993), where instead of completely randomly assigned catalysis, a molecule has to match the reaction template (a certain number of bits around the reaction site) to be able to be a catalyst for that particular reaction. For this more realistic version of the CRS model, we also showed (again both theoretically and computationally) that a linear growth rate (in n) in the level of catalysis suffices to get a high probability of autocatalytic sets. However, in this case there is also a discrepancy between the theoretically predicted and computationally observed values.

In this paper we take a different approach and ask whether it is possible to accurately predict the (observed) required level of catalysis in the template-based version of the model given the (observed) values in the original, completely random version of the model (and vice versa). At first, this may not seem obvious, as the two versions of the model have quite different constraints on which molecules can/will catalyze which reactions. However, we derive an analytical calculation that allows us to make this prediction to a very high degree of accuracy. Therefore, we can conclude that, as far as the probability of the existence of self-sustaining autocatalytic sets is concerned, the more realistic template-based version of the model is (mathematically) no different from the simpler original (completely random) version of the model.

The outline of this paper is as follows. In the next section, we briefly review the mathematical CRS model, in particular the relevant notation and definitions, and highlight the differences between the two versions of the model. In Section 3 we derive, step by step, an analytical calculation that allows us to directly relate the required level of catalysis in both model versions. In

Section 4, we then show the close match between the analytically predicted and actually observed required level of catalysis in the template-based model. Finally, Section 5 summarizes the main results and conclusions.

2. A model of catalytic reaction systems

Consider a *catalytic reaction system* (CRS) $\mathcal{Q} = (X, \mathcal{R}, C)$, where X is a set of molecule types x , R is a set of reactions r (converting reactants into products) and C is a catalysis set, i.e., a set of molecule–reaction pairs (x, r) indicating that molecule x can catalyze reaction r . We also include the notion of a *food set* $F \subset X$, which is assumed to contain molecule types that are freely available in the environment. See Hordijk and Steel (2004) and Steel (2000) for a detailed definition.

A particular CRS model that was introduced in Kauffman (1986, 1993), here referred to as the *binary polymer model*, consists of:

- A set of molecule types represented by bit strings up to a given length n , i.e., $X = \{0, 1\}^{\leq n}$.
- A food set consisting of bit strings up to a given (small) length $t < n$, i.e., $F = \{0, 1\}^{\leq t}$.
- Two types of reactions: (1) ligation, which “glues” two smaller bit strings together into one larger bit string, and (2) cleavage, which breaks a larger bit string into two smaller ones. The reaction set R consists of all such reactions that are possible within the constraint of the maximum molecule length n .
- Randomly assigned catalysis, where each possible molecule–reaction pair (x, r) is independently assigned to the set C with a given probability $p(n)$ (this probabilistic assignment is done once at the model instantiation; each such instantiation thus gives rise to a different set C).

Thus, n , t , and $p(n)$ are the model parameters.

The main question that was studied in Kauffman (1986, 1993) with this model is: under which conditions (model parameter values) is there a high probability of an autocatalytic set existing within a full CRS? An *autocatalytic set* can be (informally) described as a subset of reactions $\mathcal{R}' \subset \mathcal{R}$ (and the molecules involved in these reactions) in which (1) each reaction $r \in \mathcal{R}'$ is

catalyzed by at least one molecule that is either in the food set or can be produced from it by repeated application of reactions only from \mathcal{R}' , and (2) all reactants of the reactions $r \in \mathcal{R}'$ are either in the food set or can be produced from it by reactions from \mathcal{R}' only (a simple example is provided in Fig. 1(i)). For a formal definition, see Hordijk and Steel (2004), where we used the term RAF (reflexively autocatalytic and food-generated) for such autocatalytic (sub)sets.

Note that this CRS and RAF formalism is similar, or at least related, to alternative models in the context of the origin of life, such as Petri nets (Sharov, 1991), (M,R) systems (Rosen, 1991; Letelier et al., 2006; Jaramillo et al., 2010), the chemoton model (Gánti, 1997), other artificial chemistries and topological approaches (Benkő et al., 2009), and several other frameworks (see also Hordijk et al., 2010; Letelier et al., 2011 for a more complete overview and comparison). However, what most of these other formalisms seem to be missing, is some way of actually finding or identifying autocatalytic sets within a larger catalytic reaction system, and a thorough analysis of the probabilities of finding such subsets under different conditions (model parameters).

In Hordijk and Steel (2004), we introduced an efficient algorithm to find RAF sets in any (general) CRS, and applied it to instances of the binary polymer model. We showed both computationally (Hordijk and Steel, 2004) and theoretically (Mossel and Steel, 2005) that a linear growth rate (with n , the maximum molecule size) in the level of catalysis (in terms of the average number of reactions catalyzed per molecule) is sufficient to achieve a high probability P_n of autocatalytic sets occurring in instances of the random binary polymer CRS model.

In Hordijk et al. (2011), we then analyzed a chemically more realistic version of the binary polymer model: template-based catalysis. This differs from the original model in the way molecule–reaction pairs (x, r) are included in the catalysis set C :

- Each possible molecule–reaction pair (x, r) is *considered* for inclusion in the set C with a given probability $p'(n)$, but is only *allowed* to be included if the molecule x , somewhere along its length, matches the *reaction template* of reaction r .

The reaction template, following Kauffman (1993), is made up of the two bits on either side of the reaction site, or four bits in total (in previous studies, the *complement* of the reaction template was

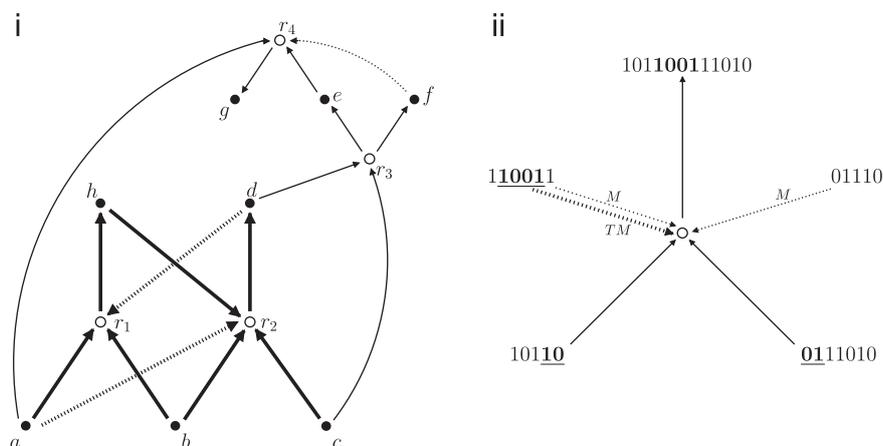


Fig. 1. (i) An example of a simple CRS in which a, b and c constitute the food set, catalysis is indicated by dashed lines, and where the reactions $\mathcal{R}' = \{r_1, r_2\}$ (shown in bold) form an RAF set (each reaction in \mathcal{R}' is catalyzed by at least one molecule that is either in the food set or can be produced from it by repeated application of reactions only from \mathcal{R}' , and all reactants of the reactions in \mathcal{R}' are either in the food set or can be produced from it by reactions from \mathcal{R}' only). This is the only RAF present in this CRS. (ii) In the original polymer model (M), either one of the potential catalysts shown (110011 and 01110) has the same probability p of catalyzing the ligation reaction $10110 + 0111010 \rightarrow 101100111010$; in the template-based model (TM), only 110011 contains a matching template (namely *1001*) for this reaction. In order for the two models to have equal probabilities of containing an RAF, the probability p' that a template-matching molecule catalyzes a reaction needs to be elevated in the TM model (indicated by the bold dashed line) by a factor m (relative to p in model M) that can be calculated analytically.

actually used; however, given that the model uses bit strings, this does not make a difference in terms of the mathematics due to the inherent symmetry). So, for example, in the ligation reaction $000 + 11111 \rightarrow 00011111$, the reaction template is 0011, i.e., the last two bits of the first reactant plus the first two bits of the second reactant (another example is provided in Fig. 1(ii)). For this template-based catalysis version of the model, we also showed that a linear growth rate (with n) in the level of catalysis suffices for autocatalytic sets to appear with high probability (although with a higher coefficient, or slope, in the linear relation than for the original model).

So, in the original (template-free) random catalysis version of the model, we have $\Pr[(x,r) \in C] = p(n)$. Let $m(n)$ be the probability that an arbitrary molecule (of length at most n) matches the four-bit reaction template of an arbitrary reaction. Then $\Pr[(x,r) \in C] = p'(n) \cdot m(n)$ in the template-based version of the model. This suggests that taking $p'(n) = p(n)/m(n)$ might lead to a similar probability P_n of finding autocatalytic sets in the template-based model as with $p(n)$ in the original model. We will show that this heuristic estimate turns out to be extremely accurate, which is not clear *a priori*, as the two models are quite different (in the original model molecules catalyze each reactions with identical probabilities, while in the template-based model, these probabilities are highly heterogeneous). So, to predict the required level of catalysis $p'(n)$ in the template-based version of the model, given $p(n)$ and a particular value of P_n , we need to know the probability $m(n)$.

3. The molecule–reaction template match probability

In this section, we derive a precise mathematical procedure for calculating the molecule–reaction template match probability $m(n)$ analytically. This procedure consists of several steps, which will be explained in detail. We then show the results of applying the derived procedure to calculate the actual template match probabilities.

3.1. The number of substring matches

As the first step, we obtain a procedure to calculate analytically the number $f_s(n)$ of bit strings of a given length n that contain a given substring s of given length. We will describe this in detail for a substring s of length four, but the procedure is applicable for any substring length. In the context of our CRS model, $f_s(n)$ is equivalent to the number of molecules of a given length n that match a given four-bit reaction template s .

We apply a mathematical technique, known as the *transfer-matrix method* (Stanley, 1986), which calculates the number $\bar{f}_s(n)$ of bit strings of length n that *do not* contain a given substring s . This is actually a very general method with a wide range of applicability, but here it will be explained explicitly in the context of counting bit strings of length n that contain a given substring of length four.

To apply this method, we first need to construct a *state transition graph* G_s that can generate all bit strings (of any length) that *do not* contain a given substring s of length four. For this purpose, it will be instructive to start with a state transition graph G that can generate *all* possible bit strings. This graph G is shown in Fig. 2, and can be interpreted as follows. Each node in the graph represents one of eight possible states, with a state being defined (and labeled) by the last three bits that were seen while generating a bit string bit by bit. For example, the bottom-left node labeled 000 represents a state where the last three bits, of the bit string as generated so far, were all zero. The next bit that will be generated is either another zero or a one. If it is another zero, then again the last three bits seen are all zero, giving rise to a state transition represented by the arrow that loops back to node 000 itself. If the next bit is a one, then the last three bits seen are now 001, giving rise to a state transition represented by the arrow

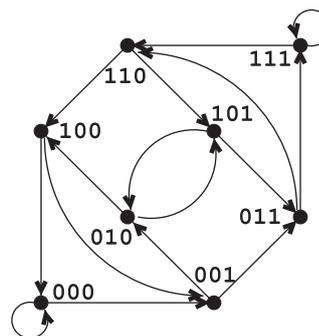


Fig. 2. The state transition graph G that generates all possible bit strings.

that points from node 000 to node 001. Obviously these are the only two possible state transitions (arrows) going out from node 000 (we cannot, for example, go from the last three bits being 000 to 111 by simply adding one bit). Similarly, there are two outgoing transitions from each of the eight nodes, corresponding to generating either a zero or a one. So, following the state transitions (arrows) as given in Fig. 2, and adding the corresponding bit (a zero or a one) to the bit string generated so far at each state transition, all possible bit strings can be produced. This state transition graph G can be represented mathematically by its *adjacency matrix* A :

$$A = \begin{matrix} & \begin{matrix} 000 & 001 & 010 & 011 & 100 & 101 & 110 & 111 \end{matrix} \\ \begin{matrix} 000 \\ 001 \\ 010 \\ 011 \\ 100 \\ 101 \\ 110 \\ 111 \end{matrix} & \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{pmatrix} \end{matrix}, \quad (1)$$

where there is a 1 in position (i,j) if there is an arrow in the corresponding graph G from state i to state j ; otherwise, there is a 0. Note that a simple state transition graph with only two nodes would actually suffice to generate all possible bit strings. However, this more elaborate eight-state graph is needed for the next step: excluding a given substring of length four.

From the general graph G (and matrix A) given above, it is now straightforward to construct the corresponding graph G_s (and matrix A_s) that generates all bit strings that *do not* contain a given substring s of length four. For example, consider $s=0010$. In this case, in the graph G of Fig. 2, making the transition from node 001 to node 010 would not be allowed (as it would produce the substring 0010). Equivalently, the entry $(001,010)$ in the adjacency matrix A should be set to zero. The resulting graph G_{0010} and adjacency matrix A_{0010} are shown in Fig. 3 and Eq. (2), respectively (the changed entry in the adjacency matrix is shown in bold).

$$A_{0010} = \begin{matrix} & \begin{matrix} 000 & 001 & 010 & 011 & 100 & 101 & 110 & 111 \end{matrix} \\ \begin{matrix} 000 \\ 001 \\ 010 \\ 011 \\ 100 \\ 101 \\ 110 \\ 111 \end{matrix} & \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \mathbf{0} & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{pmatrix} \end{matrix}. \quad (2)$$

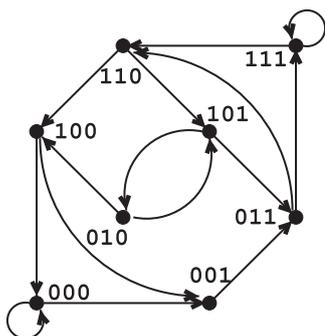


Fig. 3. The state transition graph G_{0010} that generates all bit strings that do not contain the substring 0010.

Similarly, for each of the other 15 possible substrings s of length four, there is a state transition graph G_s where one of the 16 arrows of the original graph G is left out, and a corresponding adjacency matrix A_s where one of the 16 ones in matrix A is set to zero.

Now, according to the transfer-matrix method (Stanley, 1986), the number $\bar{f}_s(n)$ of bit strings of length n that do not contain a given substring s of length four can be obtained by first computing the following generating function:

$$F_s(\lambda) = \sum_{i,j=1}^8 (I - \lambda A_s)_{ij}^{-1} \quad (3)$$

and then taking the coefficient of λ^{n-3} in $F_s(\lambda)$ as the value for $\bar{f}_s(n)$. Note that, in Eq. (3), I refers to the 8×8 identity matrix, and $(I - \lambda A_s)_{ij}^{-1}$ is the row i and column j entry of the inverse of the matrix $I - \lambda A_s$.

Given that there are 2^n bit strings of length n , the number $f_s(n)$ of bit strings of length n that contain a given substring s of length four is now simply

$$f_s(n) = 2^n - \bar{f}_s(n). \quad (4)$$

Note that this procedure can be easily extended to substrings (i.e., reaction templates) s of any length $|s| < n$. To exclude a given substring s , we simply need to remember the last $|s|-1$ bits that were seen while generating or scanning a bit string, so the required state transition graphs have $2^{|s|-1}$ nodes. However, their construction, and that of their corresponding adjacency matrices, and the subsequent calculations are similar to the example given above, where the summation in Eq. (3) is now from $ij=1$ to $2^{|s|-1}$, and the coefficient of $\lambda^{n-(|s|-1)}$ should be taken.

3.2. The template match probability

As the next step, using the results of Eq. (4), the average (or expected) number $g(n)$ of bit strings of length n that contain an arbitrary substring of length four, is now easily calculated as

$$g(n) = \frac{1}{16} \sum_{s \in \{0,1\}^4} f_s(n). \quad (5)$$

From this, the probability $h(n)$ that a bit string of length at most n contains an arbitrary substring of length four is then calculated as

$$h(n) = \frac{\sum_{i=1}^n g(i)}{2^{n+1} - 2}, \quad (6)$$

where $2^{n+1} - 2$ is the total number of bit strings of length at most n . In practice, the summation needs only run from $i=4$ to n , as any bit string of length $n < 4$ will obviously not contain a substring of length four (i.e., $g(n)=0$ for $n < 4$).

In the context of our CRS model, this probability $h(n)$ can be taken as the probability that an arbitrary molecule up to length n matches an arbitrary four-bit reaction template.

3.3. The fraction of reactions with a four-bit reaction template

So far, we have assumed that there actually is a substring of length four to be contained (or not) in an arbitrary bit string (or, in our context, a four-bit reaction template to be matched by a molecule). However, in the template-based catalysis version of our model, this is not always the case. In particular, all reactions that have at least one molecule of length one as a reactant (or product, in case of a ligation), do not have a valid four-bit reaction template, given that we require two bits on each side of the reaction site to make up a four-bit template. However, these reactions can still be considered (with probability $p'(n)$, along with a molecule x , for inclusion in the catalysis set C (which, of course, will never actually be “approved”, given that they do not have a valid four-bit template).

So, as the final step, we need to calculate the fraction $k(n)$ of reactions that actually do have a four-bit reaction template, in order to “discount” for those reactions that might be considered, but which do not have a valid reaction template. A straightforward counting argument will provide us with a simple expression for this fraction $k(n)$.

First, recall that the total number of reactions, for a given n , is $(n-2)2^{n+1} + 4$. Next, we need to count the number of reactions that have at least one molecule of length one as a reactant (since reactions are considered bi-directional, we do not need to consider cleavage reactions separately). A molecule of length one can react (ligate) with any other molecule of length at most $n-1$ (otherwise, the maximum molecule length n would be violated). There are two molecules of length one, which can be either the first or the second reactant (i.e., four combinations) and there are $2^n - 2$ molecules of length at most $n-1$, so we have $4(2^n - 2) - 4$ reactions with at least one molecule of length one as a reactant (the ‘-4’ at the end is because otherwise we would double-count the four possible reactions where both reactants are of length one). So, the fraction $\bar{k}(n)$ of reactions without a valid four-bit reaction template is

$$\begin{aligned} \bar{k}(n) &= \frac{4(2^n - 2) - 4}{(n-2)2^{n+1} + 4} = \frac{2^{n+2} - 12}{(n-2)2^{n+1} + 4} \\ &\approx \frac{2^{n+2}}{(n-2)2^{n+1}} = \frac{2}{n-2}. \end{aligned} \quad (7)$$

Consequently, the fraction $k(n)$ of reactions with a valid four-bit reaction template is

$$k(n) = 1 - \bar{k}(n) = 1 - \frac{2}{n-2} = \frac{n-4}{n-2}. \quad (8)$$

Note that this is an approximation, and the exact number is actually equal to

$$k(n) = \frac{n-4 + \frac{1}{2^{n-3}}}{n-2 + \frac{1}{2^{n-1}}}. \quad (9)$$

However, even for relatively small values of n (10 or higher suffices for our purposes), the difference between the exact and simpler approximate values becomes negligible.

Now we can finally put everything together and calculate the probability $m(n)$ that an arbitrarily chosen molecule of length at most n will match the four-bit reaction template of an arbitrarily chosen reaction (including the possibly that a reaction does not have a valid reaction template, in which case there is no match),

Table 1
 Values for $f_s(n)$ for all eight possible four-bit templates that start with a 0, and for $4 \leq n \leq 20$, as calculated using the transfer-matrix method. By symmetry, we only need to describe templates that start with 0. Other symmetries are also apparent in the table: notice that the three columns headed by a bold template are identical, and the three columns headed by an underlined template are also identical.

n/s	0000	0001	<u>0010</u>	0011	<u>0100</u>	0101	<u>0110</u>	0111
4	1	1	1	1	1	1	1	1
5	3	4	4	4	4	4	4	4
6	8	12	12	12	12	11	12	12
7	20	32	31	32	31	28	31	32
8	48	79	75	79	75	68	75	79
9	111	186	174	186	174	158	174	186
10	251	424	393	424	393	357	393	424
11	558	944	870	944	870	792	870	944
12	1224	2065	1897	2065	1897	1731	1897	2065
13	2656	4456	4087	4456	4087	3738	4087	4456
14	5713	9512	8721	9512	8721	7996	8721	9512
15	12 199	20 128	18 463	20 128	18 463	16 972	18 463	20 128
16	25 888	42 287	38 832	42 287	38 832	35 789	38 832	42 287
17	54 648	88 310	81 222	88 310	81 222	75 052	81 222	88 310
18	114 832	183 492	169 086	183 492	169 086	156 647	169 086	183 492
19	240 335	379 624	350 571	379 624	350 571	325 616	350 571	379 624
20	501 239	782 497	724 288	782 497	724 288	674 436	724 288	782 497

by combining Eqs. (6) and (8):

$$m(n) = k(n) \cdot h(n) = \frac{n-4}{n-2} h(n). \tag{10}$$

3.4. Results

We used the software package wxMaxima (<http://wxmaxima.sourceforge.net>), which is capable of performing symbolic computation and is available for free under the GPL license, to compute the generating functions $F_s(\lambda)$ (Eq. (3)) and the corresponding values of $f_s(n)$ (Eq. (4)). For example, for $s=0010$ (using the adjacency matrix A_{0010} in Eq. (2)), we get

$$F_{0010}(\lambda) = \frac{4\lambda^3 - 2\lambda^2 - \lambda + 8}{-\lambda^4 + \lambda^3 - 2\lambda + 1}. \tag{11}$$

After applying a Taylor expansion, this yields

$$F_{0010}(\lambda) = 8 + 15\lambda + 28\lambda^2 + 52\lambda^3 + 97\lambda^4 + 181\lambda^5 + \dots \tag{12}$$

So, for example, $\overline{f_{0010}}(5) = 28$, i.e., the coefficient of $\lambda^{5-3} = \lambda^2$ in $F_{0010}(\lambda)$. From this, it can be easily calculated that $f_{0010}(5) = 2^5 - \overline{f_{0010}}(5) = 32 - 28 = 4$, i.e., four bit strings of length $n=5$ contain the substring $s=0010$. The results for all templates and $4 \leq n \leq 20$ are given in Table 1. Note that because of the inherent symmetry, we only show values for the eight templates that start with a 0. To find the value for a template s that starts with a 1, simply look up the value for its (binary) complement \bar{s} . For example, $f_{1010}(12) = f_{0101}(12) = 1731$. Notice also that other symmetries exist—in particular, $f_n(s)$ is identical to $f_n(s')$ when s' is template s in reverse order. This and further symmetries ensure that the number of distinct columns in Table 1 is only four, rather than the eight possible.

Using the values from Table 1, we can compute $g(n)$ (using Eq. (5)), then $h(n)$ (using Eq. (6)) and $k(n)$ (using Eq. (8)), and, finally, $m(n)$ (using Eq. (10)). The results of these calculations are shown in Table 2 for $4 \leq n \leq 20$. The (analytically) calculated values for $f_s(n)$ and $g(n)$ were verified by a small computer program we wrote to explicitly enumerate all possible bit string/template combinations and counting the number of matches. However, for larger values of n , this would obviously be intractable, whereas the analytical (transfer-matrix) method will still be feasible.

4. Predicting the required level of catalysis

Table 2 provides us with the molecule–reaction template match probabilities $m(n)$ which can be used to predict the required level of $p'(n)$ (the probability of considering a molecule–reaction pair for inclusion in the catalysis set C), given $p(n)$ and a probability P_n of observing autocatalytic sets.

To test these predictions, we performed computer simulations with the two versions of the binary polymer model, using our RAF algorithm to detect autocatalytic sets in instances of these CRS models. For various values of n , we identified the required levels of catalysis $p(n)$ (random catalysis model) and $p'(n)$ (template-based model) for which there was a probability $P_n=0.5$ of finding RAF sets.¹ In the template-based catalysis version of the model, we used reaction templates of four bits (two bits on each side of the reaction site) that have to be matched completely by a candidate catalyst (no partial matches allowed). We used a value of $t=4$ for the food set (i.e., molecules up to length four) to allow at least some food molecules to also be catalysts. For each combination of n and $p(n)$ (or $p'(n)$), we generated between 500 (for larger n) and 10 000 (for smaller n) instances of the binary polymer model, and applied our RAF algorithm to count the fraction (P_n) of these instances that contain an RAF set. Due to the exponential growth in the size of the reaction set R with increasing n , we went up to $n=18$ only, which already took many weeks of computing time, even on a large parallel computer cluster. Due to this computational constraint, we also restricted ourselves to reaction templates of length four. But as already mentioned in Section 3, the mathematical calculations can be easily extended to templates of any length.

With the simulation results and the analytically calculated values of $m(n)$ from Table 2, we can finally compare the predicted values $p(n)/m(n)$ (solid line) with the observed values $p'(n)$ (dots), which is shown in Fig. 4. As this figure shows, there is a very close match between the predicted and observed values (the dots fall right on the line), with increasing accuracy for increasing values of n . Fig. 5 shows the same results on a log scale, to show the results (and increasingly close fit) more clearly for the larger values of n . This confirms our proposition that the required levels of catalysis in both model versions can be accurately predicted

¹ In previous papers, we have already shown that $P_n=0.5$ provides a useful reference point.

Table 2

The analytically calculated values for $g(n)$, $h(n)$, $k(n)$, and $m(n)$, for $4 \leq n \leq 20$. The $g(n)$ values are exact, the others are rounded to three significant digits.

n	$g(n)$	$h(n)$	$k(n)^a$	$m(n)$
4	1.000	0.033	0.235	0.008
5	3.875	0.079	0.408	0.032
6	11.375	0.129	0.527	0.068
7	29.625	0.181	0.611	0.110
8	72.250	0.232	0.671	0.155
9	168.625	0.281	0.715	0.201
10	382.375	0.327	0.750	0.245
11	849.000	0.371	0.778	0.288
12	1855.125	0.412	0.800	0.329
13	4002.875	0.450	0.818	0.368
14	8551.000	0.486	0.833	0.405
15	18 118.000	0.520	0.846	0.440
16	38 129.250	0.551	0.857	0.472
17	79 787.000	0.580	0.867	0.502
18	166 151.625	0.607	0.875	0.531
19	344 567.000	0.632	0.882	0.558
20	712 003.750	0.656	0.889	0.583

^a For the calculation of $k(n)$, we used the exact formula for $n < 10$ and the approximation for $n \geq 10$. For $n = 10$, the two values differ by only 0.1%.

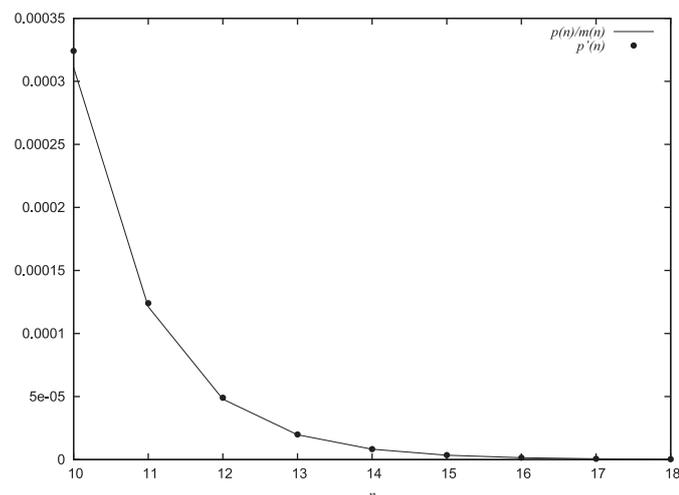


Fig. 4. Comparison of the predicted ($p(n)/m(n)$; solid line) and observed ($p'(n)$; dots) values of the required level of catalysis in the template-based version of the binary polymer CRS model, for various values of n .

from each other using the molecule–reaction template match probabilities $m(n)$.

Next, we performed the same comparison for a fixed value of n but varied the probability P_n of finding RAF sets. For two given maximum molecule lengths ($n = 12$ and $n = 15$), we identified the values of $p(n)$ and $p'(n)$ to get probabilities of finding RAF sets of $P_n \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$. Fig. 6 shows the predicted values $p(n)/m(n)$ (solid lines) and the observed values $p'(n)$ (dots) for the template-based version of the model. Again, there is a close agreement (for $n = 12$ the discrepancy is about 2%; for $n = 15$ it is merely 0.3%). So, the relationship also seems to hold for different values of P_n as well.

With these results, we can conclude to have established a direct relationship between two structurally distinct versions of the binary polymer model. More precisely, we have shown how these two models lead to identical predictions concerning the probability P_n of finding autocatalytic sets, when the probabilities of catalysis are related by the $m(n)$ scaling factor, which can be calculated analytically.

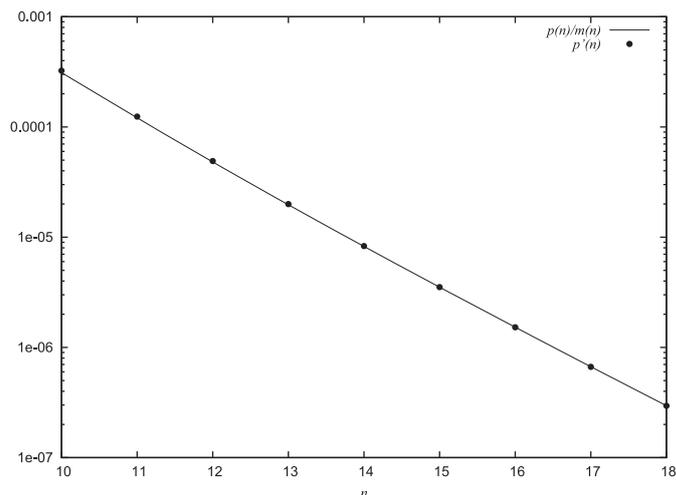


Fig. 5. Comparison of the predicted ($p(n)/m(n)$; solid line) and observed ($p'(n)$; dots) values of the required level of catalysis in the template-based version of the binary polymer CRS model, for various n , using a log scale.

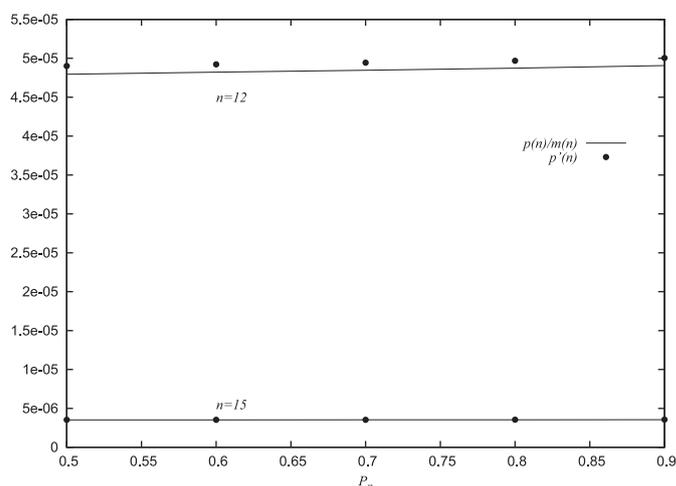


Fig. 6. Comparison of the predicted ($p(n)/m(n)$; solid lines) and observed ($p'(n)$; dots) values of the required level of catalysis in the template-based version of the binary polymer CRS model, for various values of P_n .

5. Conclusions and discussion

We have compared two different versions of a well-known binary polymer model of catalytic reaction systems as introduced in Kauffman (1986, 1993) and investigated in detail in our own previous work (Hordijk et al., 2011; Hordijk and Steel, 2004; Mossel and Steel, 2005; Steel, 2000). In the first version of the model, the probability p that an arbitrary molecule will catalyze an arbitrary reaction is equal and independent for each possible molecule–reaction pair. In the second version, a molecule will catalyze a reaction with similar probability p' only if it matches the four-bit template around the reaction site.

At first glance, these two versions seem to impose rather different constraints, resulting in rather different (actual) molecule–reaction catalysis sets C . However, we have shown that it is possible to calculate analytically a factor m which directly (and accurately) relates the catalysis probabilities p and p' in terms of the probabilities P_n of finding autocatalytic sets in instances of both model versions. In other words, given the required level of catalysis in one model version to find autocatalytic sets with a given probability P_n , it is possible to accurately predict the

required level of catalysis in the other model version to get the same probability P_n of finding autocatalytic sets.

Because of this mathematical concordance between the two model versions, we can conclude that results obtained with the original (purely random and less realistic) model can still be considered relevant in the more realistic context of template-based catalysis. We have shown that the mathematical relationship holds for various n (with fixed P_n) as well as for various P_n (with fixed n), with increasing accuracy for increasing n , so it appears to be a very robust relationship. Furthermore, the analytical method for calculating the scaling factors $m(n)$ can be easily generalized to, for example, non-binary molecules or different template lengths, and therefore does not depend on specific values of the model parameters.

This result clearly (and substantially) counters one potential and important objection that the binary polymer model in its original form (with purely randomly assigned catalysis) is too simplistic. We already showed earlier that it is straightforward to add more realistic extensions to the original model, such as template-based catalysis, which still results in similar behavior (in particular, a linear growth rate in the level of catalysis is sufficient for a high probability of finding RAF sets) (Hordijk et al., 2011). In the current paper, we established a different, more direct and analytical relationship between the required levels of catalysis in these two model versions. Furthermore, it is interesting to note that recently, an *in vitro* “implementation” of the binary polymer model was realized (Taran et al., 2010), where the food set consists of dinucleotides such as CA and TG, which can be interpreted as 0's and 1's in the binary model (von Kiedrowski, 2011). The authors conclude that “different building blocks can be incorporated into the strand, promoting the formation of combinatorial libraries of oligonucleotides long enough to be folded into specific catalytically active structures and to potentially form initial autocatalytic sets” (Taran et al., 2010). This is a further indication that the binary polymer model (and its results) has direct (bio)chemical relevance, including in the context of origin of life studies.

Acknowledgments

The computations were performed at the Vital-IT (<http://www.vital-it.ch>) Center for high-performance computing of the Swiss Institute of Bioinformatics. M.S. thanks the Allan Wilson Centre for Molecular Ecology and Evolution and the Royal Society of NZ for funding.

References

- Benkő, G., Centler, F., Dittrich, P., Flamm, C., Stadler, B., Stadler, P.F., 2009. A topological approach to chemical organizations. *Alife* 15, 71–88.
- Gánti, T., 1997. Biogenesis itself. *J. Theor. Biol.* 187, 583–593.
- Hordijk, W., Hein, J., Steel, M., 2010. Autocatalytic sets and the origin of life. *Entropy* 12 (7), 1733–1742.
- Hordijk, W., Kauffman, S.A., Steel, M., 2011. Required levels of catalysis for emergence of autocatalytic sets in models of chemical reaction systems. *Int. J. Mol. Sci.* 12 (5), 3085–3101.
- Hordijk, W., Steel, M., 2004. Detecting autocatalytic, self-sustaining sets in chemical reaction systems. *J. Theor. Biol.* 227 (4), 451–461. < <http://wxmaxima.sourceforge.net> >.
- Jaramillo, S., Honorato-Zimmer, R., Pereira, U., Contreras, D., Reynaert, B., Hernández, V., Soto-Andrade, J., Cárdenas, M., Cornish-Bowden, A., Letelier, J., 2010. (M,R) systems and RAF sets: common ideas, tools and projections. In: *Proceedings of the Alife XII Conference*, pp. 94–100.
- Kauffman, S.A., 1986. Autocatalytic sets of proteins. *J. Theor. Biol.* 119, 1–24.
- Kauffman, S.A., 1993. *The Origins of Order*. Oxford University Press.
- Letelier, J.-C., Cárdenas, M.L., Cornish-Bowden, A., 2011. From L'Homme Machine to metabolic closure: steps towards understanding life. *J. Theor. Biol.* 286, 100–113.
- Letelier, J.-C., Soto-Andrade, J., Guíñez-Abarzúa, F., Cornish-Bowden, A., Cárdenas, M.L., 2006. Organizational invariance and metabolic closure: analysis in terms of (M,R) systems. *J. Theor. Biol.* 238, 949–961.
- Mossel, E., Steel, M., 2005. Random biochemical networks: the probability of self-sustaining autocatalysis. *J. Theor. Biol.* 233 (3), 327–336.
- Rosen, R., 1991. *Life Itself*. Columbia University Press.
- Sharov, A., 1991. Self-reproducing systems: structure, niche relations and evolution. *BioSystems* 25, 237–249.
- Stanley, R.P., 1986. *Enumerative Combinatorics*, vol. 1. Wadsworth & Brooks/Cole.
- Steel, M., 2000. The emergence of a self-catalysing structure in abstract origin-of-life models. *Appl. Math. Lett.* 3, 91–95.
- Taran, O., Thoennessen, O., Achilles, K., von Kiedrowski, G., 2010. Synthesis of information-carrying polymers of mixed sequences from double stranded short deoxynucleotides. *J. Syst. Chem.* 1 (9).
- von Kiedrowski, G., 2011. Personal communication.