

I. TITLE PAGE

Title: A mechanism of abiogenesis based on complex reaction networks organized by seed-dependent autocatalytic systems

Author Line

Zhen Peng¹, Jeff Linderoth^{1,2}, David A. Baum^{1,3*}

Author Affiliations

¹Wisconsin Institute for Discovery, University of Wisconsin-Madison, Madison WI 53706, USA

²Department of Industrial and Systems Engineering, University of Wisconsin-Madison, Madison WI 53706, USA

³Department of Botany, University of Wisconsin-Madison, Madison WI 53706, USA

Corresponding Author

David A. Baum

Address: Department of Botany, Birge Hall 216/220, 430 Lincoln Drive, Madison, WI 53706, USA

Phone number: (+1) 608-265-5385

Email address: dbaum@wisc.edu

Preprint Servers

The early version of the manuscript was deposited to ChemRxiv (<https://doi.org/10.33774/chemrxiv-2021-1nj64>) under CC BY NC ND 4.0 License.

Classification

Major: Physical Sciences; Biological Sciences

Minor: Chemistry; Biochemistry; Evolution

Keywords

origin of life, abiogenesis, reaction networks, autocatalysis, chemical ecosystem, evolution

II. ABSTRACT

The complexity gap between the biotic and abiotic worlds (1, 2) has made explaining abiogenesis one of the hardest scientific questions. A promising strategy for addressing this problem is to identify features shared by abiotic and biotic chemical systems that permit the stepwise accretion of complexity. Therefore, we compared abiotic and biotic reaction networks (3, 4) in order to evaluate the presence of autocatalysis, the underlying basis of biological self-propagation, and to see if the organization of autocatalytic motifs permits stepwise complexification. We provide an algorithm to detect seed-dependent autocatalytic systems (SDASs), namely subnetworks that can use food chemicals to self-propagate but must be seeded by some non-food chemicals to become activated. We show that serial activation of SDASs can cause incremental complexification. Furthermore, we identify life-like features that emerge during the accretion of SDASs, including the emergence of new ecological opportunities and improvements in the efficiency of food utilization. The SDAS concept explains how driven abiotic environments, namely ones receiving an ongoing flux of food chemicals, can incrementally complexify without the need for genetic polymers. This framework also suggests experiments that have the potential to detect the spontaneous emergence of life-like features, such as self-propagation and adaptability, in driven chemical systems.

III. SIGNIFICANCE STATEMENT

The great gap between the simplest living system and the most complex non-living one makes the origin of life a challenging problem. In this paper we show that abiotic and biotic reaction networks have a remarkably similar structure. Both contain seed-dependent autocatalytic systems (SDASs), which are sets of chemicals with the potential to collectively self-propagate once activated by a set of “seed” molecules. Furthermore, the SDASs were hierarchically organized in both networks that we studied, such that the activation of a lower-tier SDAS allows a higher-tier SDAS to then be seeded. This organization enables stepwise complexification, during which abiotic entities may accumulate life-like features. Our results demonstrate the fundamental continuity of non-life and life, suggesting new classes of experimental studies.

IV. MAIN TEXT

1. INTRODUCTION

Life is the paradigmatic example of a complex system (1, 2). It consists of diverse chemical components, some of which are large, energy-expensive molecules. Furthermore, these components are not simply lumped together but are organized and coordinated in such a way that the entire system can resist environmental perturbations and grow or divide to give rise to more life. Such organization is not cheap, of course, which is why life must consume energy to maintain its internal order. It is almost magical that complex organisms can use “untargeted” energy sources such as light to convert simple materials (e.g., water, carbon dioxide, and minerals) into more living matter of the same kind. The mystery of the origin of life is how a system with sufficient complexity to conduct such improbable conversions could emerge spontaneously, when there was no prior design to follow or template to copy.

The solution to this conundrum lies, we believe, not in looking for particular molecules or reactions, but in looking at the emergent properties of chemical reaction networks. If real chemical reaction networks contain many autocatalytic motifs, systems with a pre-existing ability to sustain themselves may arise readily, and then become progressively more complicated through the accretion of more autocatalytic modules. In our recent work (5), we used analyses of “toy” reaction networks to show that interacting autocatalytic cycles can exhibit features that closely resemble those seen when species interact ecologically. This result implies that driven chemical reaction systems, or “chemical ecosystems,” can show complex dynamics, including succession. Here, we sought to extend this analysis to real chemical reaction networks to see if they have the features needed for evolution. Specifically, we focused on databases that “bracket” the origin of life, a radiolytic and geochemical reaction database (3) serving to represent chemistry without life, and a curated subset of biochemical reactions (4) serving to constrain the metabolic network of the Last Universal Common Ancestor (LUCA). Our reasoning is that features shared by both networks are very likely to also apply to systems on the path from non-life to life. Thus, if we find mechanisms or network structures that allow for the evolutionary accretion of complexity in both databases, then it is reasonable to assume that these same phenomena applied at the time when systems deserving the label “life” first emerged.

In this study, we focus on the concept of a seed-dependent autocatalytic system (SDAS), which is a network motif that can be activated by a rare chemical event but can, once activated, sustain itself. When a network with multiple potential SDASs is driven by a suitable flux of food, it seems possible that it could gradually acquire new components, some of which may affect the overall efficiency of food utilization and/or open-up opportunities for yet further complexification. We suggest that, when actualized in a spatially and temporally structured environment, a multi-SDAS network may be able to show adaptive evolution without producing genetic polymers.

We start by formalizing the concept of network expansion, which is the basic procedure needed to map out a network's architecture. Then, we formalize the SDAS concept and describe an algorithm for identifying them within a reaction network represented by a stoichiometric matrix. Furthermore, we show that SDASs are found in both the abiotic and biotic chemical reaction networks and are organized in a way that allows for the stepwise accretion of complexity. Finally, the larger, biotic network provides examples where later-activated processes improve the efficiency of the system as a whole and open-up new ecological opportunities. This analysis supports the contention that reaction networks can evolve adaptively, in the sense of periodically finding new quasistable states that tend to more efficiently use resources in the environment to sustain collective autocatalysis. We end by discussing some implications of our results, several remaining challenges, and how our theory can be used to guide laboratory experiments.

2. RESULTS

2.1 Databases of abiotic and biotic reactions

The abiotic reaction database we analyzed is based on a recently published reaction network (3) assembled from seven decades of publications. This database includes free radical reactions, mineral geochemical reactions, amino acid production, chloride radical and polar reactions, nitrile radical and polar reactions, RNA nucleotide assembly, nuclear decay, and physicochemical reactions. We have also added some additional well-known abiotic organic reactions to the database, including the classical formose reaction (6), but without formaldehyde dimerization because it is very slow and its reaction mechanism is yet to be determined (7, 8).

Sources of radiation such as X-rays, ultraviolet, and visible light are treated as reactants and products, which is why we will use “entity” to refer to both chemicals and electromagnetic energy sources in this specific database. In part because of these energy sources, most of the reactions in the abiotic database are irreversible. Although only covering a small portion of known abiotic chemical reactions, these reactions may still be used to test the applicability of our theoretical framework to real chemical reaction networks.

The biotic reaction database we analyzed is based on a recently published reaction network (4), which was obtained by removing reactions that only occur in eukaryotes and reactions dependent on O₂ from the KEGG reaction database (9–11). Xavier et al. (2020) (4) claimed that the resulting reaction network could be a proxy of the primordial metabolism of LUCA.

Additionally, the reaction KEGG R06974 was added because it is an important step in purine metabolism. This reaction is not oxygen-dependent and entails a reaction mechanism that is very similar to the reaction KEGG R06975 (see Materials and Methods), which is already included in the database. We also added five spontaneous reversible reactions that are missing from KEGG, such as $\text{H}_2\text{O} \leftrightarrow \text{H}^+ + \text{OH}^-$ and $\text{H}_2\text{CO}_3 \leftrightarrow \text{H}^+ + \text{HCO}_3^-$. We acknowledge that, because most of the reactions in the KEGG database are catalyzed by enzymes, it is likely that many of them could not occur at sufficiently high rates in a prebiotic world to have been relevant before biological catalysts had evolved. Nonetheless, since all these reactions are chemically feasible, we reasoned that the relationships between reactants and products in such a curated biochemical reaction network is meaningful and that features shared by it and the abiotic network should be relevant to the origin of life. Thus, we do not wish to claim that the chemical reactions in this database are those that applied during the origin of life. Rather, our goal is to use knowledge of these biotic reactions to explore how this network would be expected to change over time when driven by food flux, reasoning that any patterns seen in this chemical microcosm likely had numerous analogs in a pre-enzymatic world.

We preprocessed the databases such that every reaction appears only once and had clear stoichiometry. All reactions in the biotic reaction network are assumed to be reversible, so each of them was split into two unidirectional reactions (see Materials and Methods). After curation, the abiotic reaction database consists of 277 entities and 717 unidirectional reactions for a reaction density of 2.59 reactions/entity (Table S1). The resulting biotic reaction database

consists of 4216 chemical species and 8402 unidirectional reactions for a reaction density of 1.99 reactions/species (Table S2). Fig. S1 shows the histograms of the numbers of reactions that an entity or chemical species is involved in for the two databases. Both distributions fit a power law, being highly left-skewed, meaning that most entities/chemicals are involved in a small number of reactions.

2.2 Network expansion and tier-0 systems

Reaction databases are just collections of reactions allowed to occur. To generate an organized subnetwork for analysis, we define a network expansion operation, $\Xi(S_O, \mathbf{R})$, which calculates all reactions and chemical species that can be accessed given a starting set of chemical species S_O and a set of all allowed reactions \mathbf{R} . The network expansion (see Materials and Methods) starts with S_O and scans \mathbf{R} for any reactions that are not yet in the subnetwork, but whose reactants are all present in the current subnetwork. These reactions are added to the subnetwork and the expansion iterates until no more chemical reactions can be added to the subnetwork. The reaction subnetwork resulting from the expansion is described by the tuple (S_E, \mathbf{R}_E) , where S_E is the set of chemical species and \mathbf{R}_E is the set of reactions. The network expansion operation can be visualized using a stoichiometric matrix (Fig. S2A-D), where each row represents a chemical species and each column represents a unidirectional reaction, with stoichiometric coefficients as the entries. In general, reactants have negative stoichiometric coefficients while products have positive ones. However, it should be noted that some reactions may involve species presenting in both reactants and products, which could make the stoichiometric coefficients of such species non-negative even though such species are required for the reactions to occur; therefore, the columns representing such reactions should be annotated.

To apply network expansion to explore the properties of chemical systems, we need to specify a set of chemical species that are assumed to be provided by the environment, which we will call the ultimate food set. It should be noted that even though they are designated as ultimate food, they might also be produced by reactions within the subnetwork. A full expansion starting from the ultimate food set generates a reaction subnetwork including all chemicals that would be expected to be generated at a non-zero rate in an environment receiving an ongoing flux of the ultimate food species. We will call this subnetwork the tier-0 system, reflecting the fact that no

additional events are needed for its generation except for the provision of the ultimate food (Fig. S2A-D).

Although there is still no wide consensus on the environmental conditions of the prebiotic sites where life originated (12–20), to illustrate the approach, we selected sets of simple entities or chemical species as ultimate food. For all analyses with the abiotic reaction database, we chose $\{\text{H}_2, \text{CH}_4, \text{NO}, \text{FeS}_2, \text{visible light}\}$ as the ultimate food set. A full expansion starting from this set generated a tier-0 system with no additional chemicals and zero reactions. For all studies of the biotic reaction database, we chose $\{\text{H}_2\text{O}, \text{CO}_2, \text{NH}_3, \text{H}_2\text{S}, \text{H}_2\text{SO}_4, \text{H}_2\text{SO}_3, \text{HSO}_3^-, \text{H}_3\text{PO}_4, \text{H}_4\text{P}_2\text{O}_7\}$ as the ultimate food set. The full expansion starting from this set generated a tier-0 system with 30 chemical species and 44 unidirectional reactions (i.e., 22 reversible reactions). This network contains substructures that can be drawn as autocatalytic cycles (5, 21), an example of which is shown in Fig. S3. However, because all the members in the tier-0 system are either provided as the ultimate food or spontaneously synthesized from the ultimate food, these autocatalytic cycles lack significance beyond suggesting that the system might show non-linear dynamics upon initiation or perturbation.

2.3 Seeds and tier-1 seed-dependent autocatalytic systems (SDASs)

The materials and energy for building a tree already exist in the environment, but without a seed buried in the soil, materials like water, carbon dioxide, minerals, and energy from sunlight will not spontaneously form a new tree. It is the seed that provides the information for organizing the flow of simple materials and energy into a structure able to grow and self-propagate. That life must come from life has long been appreciated, at least since Louis Pasteur's famous gooseneck flask experiments. Here we will show that seeding, or at least a prototype of seeding, exists in both abiotic and biotic reaction networks, and can induce systems more complex than tier-0 systems.

Imagine introducing a small amount of a new chemical species, a candidate seed, P , to an activated tier-0 system (Fig. S2E). P may react with some of the chemical species in the tier-0 system to generate new chemical species, which can result, through network expansion, in activation of a bigger reaction subnetwork (Fig. S2F-H). The chemical species and reactions that

are added to the tier-0 system may be called a tier-1 system (Fig. S2H), representing the fact that one seeding event is required to induce such a system. While P allows all chemicals in the tier-1 system to be formed, we will only consider P to be a true “seed” if the system it initiates is one that has the potential to sustain itself in a driven environment, namely one experiencing constant dilution and influx of the ultimate food. The only way for a tier-1 system to persist in a driven environment is to have a network topology that allows species in the tier-0 system to be converted into a stoichiometric excess of P. This means that the tier-1 system induced by P will only be viable if it contains at least one combination of reactions that are collectively autocatalytic, meaning that a net reaction equation can be written such that the reactants and products share no chemical species, only tier-0 chemicals are among the reactants, and all tier-1 chemicals are present in the products. Such a viable tier-1 system is here defined as a seed-dependent autocatalytic system (SDAS). SDASs are similar to pRAFs in the RAF theory (22), because neither can be constructed simply from the food set. However, they are not identical because SDASs require specific stoichiometric relationships among the involved reactions while RAF theory does not consider stoichiometry.

To identify SDASs, we developed a linear programming algorithm inspired by Blokhuis et al. (2020) (21). The key criterion is that for the submatrix consisting of all columns (the q th to n th columns in Fig. 1) and rows (the p th to m th rows in Fig. 1) induced by the seed, there exists a vector of non-negative elements $\mathbf{x} = (x_q, x_{q+1}, \dots, x_n)$ such that for every row in the submatrix, the dot product of this row and \mathbf{x} is positive (Fig. 1), or

$$\sum_{j=q}^n x_j s_{ij} > 0 \quad (x_j \geq 0) \quad \forall i \in [p, m], \quad [1]$$

where s_{ij} is the entry at the i th row and j th column of the stoichiometric matrix.

Whether the condition described by [1] can be satisfied can be calculated by linear programming, and we may further use integer programming to identify specific autocatalytic motifs within the SDAS, for example the autocatalytic cores that contain the fewest reactions (see Materials and Methods).

For the abiotic reaction database, 12 of the 272 non-tier-0 entities can serve as singleton seeds capable of inducing a tier-1 SDAS (Table S3). Further analysis showed that all 12 seeds induce

the same tier-1 system, which includes 91 entities and 220 unidirectional reactions that were absent from the tier-0 system (Table S4). It is worth noting that, apart from the 12 seeds, the other 79 entities in the tier-1 system cannot individually seed the system. We define entities that can each induce the same tier-1 SDAS as a “clique.” Thus, these 12 entities that can seed the same 91-entity tier-1 SDAS comprise a clique of interchangeable seeds. The smallest combination of reactions within the tier-1 SDAS that satisfies [1] entails 10 reactions. There are several alternative (largely overlapping) 10-reaction autocatalytic motifs, of which one example is provided (Fig. 2, Table S5). This autocatalytic core uses CH_4 , NO , and visible light as food and produces H_2CNH and infrared light as waste.

For the biotic reaction database, 304 of the 4186 non-tier-0 chemical species are viable singleton seeds that can induce a tier-1 system (Table S6). These 304 species belong to three cliques: 267 triggering the same 301-species (736-reaction) tier-1 SDAS (SDAS-1a), 34 triggering a 357-species (916-reaction) tier-1 SDAS (SDAS-1b) that encompasses SDAS-1a, and 3 triggering a 1414-species (4114-reaction) tier-1 SDAS (SDAS-1c) that encompasses SDAS-1b (Table S7, Table S8, Table S9). The 267-member clique includes some simple seed chemicals with as few as 2 carbon atoms, such as acetylene and glycolaldehyde, whereas every species in the 34-member clique is a pyrimidine nucleoside or a derivative thereof, meaning that it contains at least 9 carbon atoms. The 3-member clique consists of NAD^+ , NADH , and deamino- NAD^+ , meaning that every of them contains 21 carbon atoms. Since these three tier-1 systems are nested, they share the same, minimal autocatalytic core whose food is H_2O , CO_2 , and $\text{H}_4\text{P}_2\text{O}_7$ and whose waste is H_3PO_4 and H_2O_2 (Fig. 3, Table S10).

The concept of a chemical seed is most significant when seeding events are expected to be rare. For example, some reactions not included in the network might be possible in the current environment but so slow that there might be a long and uncertain duration before seeding occurs. Alternatively, seeds could disperse into the site from some another environment that allows for different chemical reactions due to a different physical environment (e.g., higher temperature), a different food set (e.g., a surficial *versus* a submarine site), or different resident catalysts (e.g., different exposed minerals). However, while we are making the implicit assumption that arrival of a seed to a driven microenvironment is stochastic, the seed needs to be produced somewhere, meaning that seeding with simpler seed molecules is *a priori* more likely. For this reason, as

discussed at greater length below, it seems probable that SDAS-1a would be initiated before SDAS-1b or SDAS-1c.

There is nothing in the concept of seeding that requires a “seed” to be a single chemical species: simultaneous seeding of multiple chemicals might be needed to initiate a SDAS, making the seeding event even less probable. Given a set of chemical species as the external food, which is the result of a full network expansion, “interdependence” between members of a set of non-food chemical species, U , can be detected if (a) U can induce a viable SDAS that sustainably synthesizes U from the external food and (b) for any non-empty $V \subset U$, V cannot induce a viable SDAS supporting V .

We can illustrate the concept of seed interdependence with a few examples. For the abiotic reaction database, neither glycolaldehyde ($C_2H_4O_2$) nor NH_3 can seed a viable tier-1 SDAS (Table S3). Nevertheless, the set $\{C_2H_4O_2, NH_3\}$ can seed a 105-entity (255-reaction) viable tier-1 SDAS (Table S11). Likewise, for the biotic reaction database, neither formaldehyde (H_2CO) nor acetate (CH_3COOH) is a viable seed (Table S6), yet together they can seed a viable tier-1 system containing 301 species and 736 reactions (Table S12). The interdependence between members of a seed set is conceptually linked to the fact that for the life as we know it, multiple chemical species need to be “seeded” together to allow the conversion of abiotic food to more cells. However, assuming equal complexity in the chemicals concerned, a SDAS requiring multiple interdependent seeds has a lower probability of emerging than one dependent on only a singleton seed.

It is also possible that a single seeding chemical can induce a viable SDAS that does not produce the seeding chemical itself. In an open environment, such chemicals, which we will call “pseudo-seeds,” can trigger a SDAS but would be expected to disappear over time. For example, in the abiotic reaction database, C_2H_3 is a pseudo-seed because it can induce a viable singleton seed, OH , but the tier-1 SDAS it triggers cannot produce C_2H_3 . Similarly, for the biotic reaction network, ATP is a pseudo-seed because it can induce a viable seed, UMP, which triggers SDAS-1b, but SDAS-1b cannot produce ATP. The existence of pseudo-seeds is worth noting because it means that, in the same way that scaffolds are not seen after an arch is built, key transitions during the origin of life might have been triggered by chemicals that are no longer produced by biochemical systems, potentially confounding historical inference.

These results show that for both the abiotic and biotic reaction networks, SDASs exist. This is significant because by receiving a rare seed, a chemical ecosystem can be pushed into a new quasistable state that is distinct from, and more complicated than, the state prior to seeding. In a sense, the ecosystem “remembers” a seeding event, implying that SDASs may be the earliest and simplest general mechanism of heritability, a prerequisite for evolution.

2.4 Higher-tier systems

An important feature of biological systems is that some events must happen in a specific temporal order. For example, in primary succession, lichens dominate the environment before grasses and trees, and herbivores can survive only when plant populations are large enough. Such necessary ordering arises when later stages require some conditions that can only be provided by the earlier stage.

Such hierarchical structuring is also manifested by reaction networks. Once a tier-1 SDAS is established, all members of the tier-1 system, together with the tier-0 members, are now available to “feed” additional higher-tier SDASs. The same procedures used to detect a tier-1 SDAS can be used to look for viable SDASs at higher trophic levels. If such additional tiers exist then there is a natural ordering: tier-1 SDASs feed on tier-0 systems, and tier-2 SDASs feeds on the tier-0 and tier-1 systems, etc.

For the abiotic reaction database, once the tier-1 SDAS has been induced by any of its potential seeds, there exists a 13-member clique that induces a 14-entity, 35-reaction tier-2 system containing an autocatalytic core (Table S13, Table S14). This tier-2 SDAS includes the formose reaction, which feeds on H_2CO generated by the tier-1 SDAS.

For the biotic reaction database, there are three nested tier-1 SDASs, each induced by a different clique. Although SDAS-1a can be induced by simple singleton seeds like acetylene or pyruvate, the simplest seeds that can induce SDAS-1b or SDAS-1c are much more complex. Therefore, we will assume that the seeding of SDAS-1a is more likely to occur first. According to this rationale, we examined whether higher-tier SDASs can be seeded following activation of SDAS-1a. We discounted singleton seeds sufficient to induce SDAS-1b or SDAS-1c, under the assumption that these molecules are too complicated to arise until after metabolic complexity had accreted.

We found a 6-member clique that seeds a 56-species, 180-reaction viable tier-2 SDAS (SDAS-2a) (Table S15, Table S16). Not surprisingly, 5 out of the 6 clique members contain pyrimidine moieties, consistent with the fact that the additional reactions and chemicals in this tier-2 system are nested within the pyrimidine-nucleoside-containing SDAS-1b (Fig. 4B,D). The smallest chemical in the clique is cytosine that contains just 4 carbon atoms, which is much simpler than the seeds of SDAS-1b, the smallest of which is the 9-carbon-atom molecule deoxycytidine.

We did not find any singleton seed able to induce a viable SDAS that recovers the complexity of SDAS-1c. However, allowing for interdependent seeds, we observed that adenine, a 5-carbon-atom molecule, together with picolinic acid, a 6-carbon-atom molecule, can seed a viable tier-2 SDAS (SDAS-2b), which has 1113 chemicals and 3378 reactions (Table S17).

SDAS-2b is nested within SDAS-1c (Fig. 4C,E) and is able to synthesize purines and their derivatives, including the hydrogen donor and cofactor NADH. Within SDAS-2b, purine synthesis is achieved via an autocatalytic cycle feeding on tier-0 and tier-1 chemical species to produce ATP (Fig. 5, Fig. S4). SDAS-2b shows that even if the direct seeding of NADH in an NADH-less world were unlikely, sequential seeding by much simpler molecules (pyruvate, adenine and picolinic acid) could nonetheless achieve an autocatalytic system synthesizing complex molecules including NADH.

2.5 Modification of system-wide productivity by the activation of a new SDAS

In a biological ecosystem, ecological interactions can be complicated and not easily organized in a simple trophic hierarchy. One consequence of this is that the introduction of a species to a community can alter the overall flux of energy and matter into and through the ecosystem (e.g., (23)). Here we seek to ask whether higher-tier SDASs have the potential to affect the efficiency or productivity of a chemical ecosystem as a whole. To explore this possibility, we examined the possible effects of the activation of SDAS-2b on overall ecosystem function.

The increased diversity of chemicals, which is maintained once higher-tier SDASs are activated, increases the chances of finding a chemical that catalyzes reactions at lower tiers. Catalysts are usually molecules that enable a series of intermediate chemical reactions whose collective effect is to lower kinetic barriers to produce the same products as an uncatalyzed reaction while

regenerating the catalyst. To identify potential catalysts we can, thus, look for SDAS-2b chemicals that provide a new path for completing a pre-existing reaction. For example, the SDAS-2b chemical FAD could catalyze assimilation of phosphate moieties via the reaction $\text{CH}_3\text{COOH} + \text{H}_4\text{P}_2\text{O}_7 \rightarrow \text{CH}_3\text{COOPO}_3\text{H}_2 + \text{H}_3\text{PO}_4$ (Fig. 6). Assuming that the composite reaction rate of the FAD-mediated reaction mechanism is higher than the direct reaction, which seems plausible, FAD would catalytically increase the rate at which the entire ecosystem could acquire phosphate from the environment.

It is also possible that a higher-tier chemical allows for a chemical ecosystem to use previously unexploited food and/or use food in new, more efficient ways. For example, without SDAS-2b chemicals like NAD^+ , the ultimate food H_2S serves as a source of thiol group but does not function as a hydrogen donor. However, due to the reaction, $\text{NAD}^+ + \text{H}_2\text{S} \leftrightarrow \text{NADH} + \text{H}^+ + \text{S}$, the presence of NAD^+ allows H_2S to serve as a terminal hydrogen donor. As a result, H_2S should provide a stronger driver of carbon fixation than the thermodynamically less favorable carbon fixation pathway within the tier-1 SDAS (Fig. 3). This way of exploiting H_2S was not possible before the “discovery” of a molecule like NAD^+ , making this another example of feedback on overall ecosystem function. Furthermore, the presence of NADH also allows a new carbon-fixing autocatalytic cycle (Fig. 7) that is shorter, and thus potentially more efficient, than the carbon fixation processes of SDAS-1a.

To give one final example of the new potential benefit that can arise with a higher tier, it is worth noting that SDAS-2b includes some relatively long-chain amphiphiles such as hexadecanoic acid (Table S17), which might be able to form liposomes. Were this to be the case, the addition of SDAS-2b might open-up the possibility of producing membranes that altered ecosystem stability by providing protection against mechanical perturbation or by fostering spatial organization. Liposomes potentially generated by SDAS-2b might even enable long distance co-dispersal of interdependent seed chemicals, giving the whole ecosystem a selective advantage.

These results support the view that, as higher-tier SDASs become activated, a chemical ecosystem can become more productive and more resistant to environmental perturbations. As a result, if multiple systems consisting of SDASs coexisted, we might expect that chemical ecosystems with higher efficiency of utilizing environmental resources would gradually dominate. This closely resembles the ecological principle that a community’s efficiency of

resource utilization tends to increase over the course of ecological succession (24, 25). Thus, it seems plausible that chemical ecosystems may, complexify over time not only in the trivial sense of coming to support the persistence of a greater number of chemical species, but also in the sense of acquiring and optimizing new, emergent ecosystem-level properties.

3. DISCUSSION

3.1 Limitations of the analysis

In this paper, we extend the framework of chemical ecosystem ecology (5) to look for autocatalytic motifs within real chemical reaction networks. Using a stoichiometric criterion, we detected multiple potential autocatalytic cycles in both abiotic and biotic reaction networks, providing evidence that both networks share many basic features related to abiogenesis. Nonetheless, a few limitations of this analysis should be noted.

First, both reaction systems are small fragments of the much larger network of reactions that would have been available to early life, namely the total chemical reaction network possible in the physical conditions present somewhere on Earth at that time. Thus, it would be premature to conclude that, because we only found 2 SDAS tiers and a few SDASs, true chemistry would be similarly limited. Likewise, it is possible that in larger networks, the order in which seed molecules are introduced could result in the exploration of different regions of chemical space, and such exploration could allow for systems to diverge over time in an historically contingent manner, something that was not apparent in the subnetworks we studied.

Our analyses of the biotic network have additional limitations, since almost all reactions are catalyzed by evolved enzymes. One inference from this might be that only a tiny subset of the reactions in the network we studied would occur at an appreciable rate in the absence of enzymes, implying that the true network might be less interconnected than we supposed. However, the converse argument could also be made. In the absence of catalysts optimized for particular reactions, prebiotic reaction networks might contain many more reactions that have similar rate constants resulting in greater connectedness than we assumed. Given, this uncertainty, we really do not know how representative the biotic reaction database is. Nonetheless, since the results from the biotic and abiotic networks were qualitatively similar

despite differing in size and reaction density, we suspect that our conclusions regarding the overall structure of the biotic network are robust to the presence/absence of enzymes.

The criterion of autocatalysis we used is based on network topology and stoichiometry, while thermodynamics and kinetics were ignored. Because we are imagining a chemical system driven by the constant provision of input food, which can potentially drive reactions that are individually disfavored in standard conditions, the lack of thermodynamic and kinetic factors may not affect the existence of an autocatalytic motif detected by our method. However, for our theoretical framework to simulate how a reaction network may evolve over time and to further direct experimental studies by estimating appropriate reaction conditions, it would ultimately be necessary to conduct simulations that include realistic thermodynamic and kinetic parameters.

One consequence of ignoring thermodynamics and kinetics is that some of the detected autocatalytic cores might lack viability in practice. Indeed, careful analysis of the carbon fixation reactions within the biotic SDAS-1a shows that it depends on a module consisting of two thermodynamically unfavorable reactions, namely water reducing formaldehyde to methanol ($\text{H}_2\text{CO} + 2 \text{H}_2\text{O} \rightarrow \text{CH}_3\text{OH} + \text{H}_2\text{O}_2$) and a reverse Cannizzaro reaction where methanol and formic acid react to generate two molecules of formaldehyde ($\text{CH}_3\text{OH} + \text{HCOOH} \rightarrow 2 \text{H}_2\text{CO} + \text{H}_2\text{O}$) (Fig. 3). It seems unlikely, therefore, that this cycle could sustain itself under natural conditions. On the other hand, we should remember that the networks we studied do not include all possible chemical reactions. For example, the tier-1 autocatalytic core could be viable if there were other prebiotic pathways reducing formic acid to formaldehyde, which is likely (26).

It has been argued that kinetic factors make prebiotic autocatalysis an unlikely basis for the emergence of life (27). The problem is that autocatalytic cycles are only viable when the rate of production of all members of the cycle is higher than their rate of loss from the system through dilution and/or side-reactions. For this reason, Orgel (2008) (27) argued that catalysts are the potential gatekeepers of any viable metabolism and that therefore, without enzymes, and hence some kind of genetic encoding, prebiotic chemical reaction networks would not be viable or evolvable.

While Orgel (2008) (27) raised some important issues (and did so eloquently and convincingly), it seems premature to rule out a metabolism-first model for the origin of life. First, recent work has shown that many core metabolic reactions dependent on enzymes in modern life can proceed

at reasonable rates through the actions of metal or mineral catalysts (28–36). Second, it should be noted that we detected many overlapping autocatalytic cycles in both databases, which increases the probability that at least one cycle within each SDAS would be viable in a given environment. Third, it should be borne in mind that the total small-molecule chemistry available to prebiotic reaction systems is vastly richer and more diverse than the two small subnetworks of chemistry studied here. Thus, we believe that, while empirical data are urgently needed (Sect. 3.3), it would be premature to rule out the possibility of evolutionary processes in the absence of enzymes (and hence genes).

3.2 Chemical ecosystem theory as a general framework for understanding the emergence of evolution before genes

Our finding of seed-induced autocatalysis in real reaction networks is an important advance because it provides a mechanism for memorizing past environments in the dynamical state of a reaction network. This may provide a basic mechanism of heritability in the absence of digital genetic machinery. We also observed a trophic hierarchy that allows for the stepwise accretion of complexity and the gradual acquisition of motifs that enhance overall ecosystem function. Combined with pre-existing models that allow for analog or compositional inheritance (37–40), this finding supports a metabolism-first model for the origin of life entailing the following steps: (a) Planetary processes, such as solar irradiation, local redox disequilibria driven by tectonics, or global redox disequilibria driven by the loss of hydrogen to space (41), generate a steady flux of food molecules/entities into local environments. (b) Chemical seeds are intermittently introduced by rare reactions or stochastic events (e.g., delivery from space), triggering low-tier SDASs that support the persistent production of new, more complex chemicals in the environment. (c) Each additional SDAS provides a larger pool of food available to potential SDASs at yet higher tiers. (d) Some newly seeded SDASs cause the productivity of the network as a whole to increase, for example when complex molecules catalyze more efficient pathways for extracting energy from the food set. (e) Such chemical succession continues until the ecosystem hits upon catalytic and template-guided replicated polymers, which allows for even more complexity to accrete through conventional Darwinian evolution.

The idea that catalytic control could have been acquired gradually during abiogenesis gains support from the structure of SDAS-2b that we identified within the biotic network. Whereas

SDAS-1a does not synthesize critical cofactors, SDAS-2b includes some larger and more complex molecules with catalytic activity, such as NADH and FAD, which can be seen as oligomers of two nucleotides (NADH) or one nucleotide and one flavin moiety (FAD). Moreover, these oligomers show potential catalytic feedbacks on lower tiers, possibly enhancing efficiency and, thus, the expected standing mass of the multi-tier system as a whole. As a result, the relationship between SDAS-2b and SDAS-1a is, abstractly, similar to the relationship between genetic/catalytic polymers (e.g., DNAs, RNAs, proteins) and metabolism in cellular life (Fig. 8). Thus, the earliest appearance of short catalytic nucleic acids and peptides might have similarly began with a fortuitous feedback on ecosystem-level fitness. However, such feedback alone cannot solve the combinatorial problem of producing sufficient functional polymers beyond a trivial length. Adding mechanisms to chemical ecology that can overcome the problematic combinatorics of polymerization remains an important focus of future work.

3.3 Future theoretical and experimental research

The two most obvious extensions of the analyses we conducted would be to add realistic kinetics so that full dynamical simulations could be explored, and to expand the network to better represent the full space of chemical possibilities. Unfortunately, it is almost impossible to build these extensions on empirical data. Only a limited number of reactions have been investigated experimentally, and those represent a non-random sample of chemistry as a whole, for example being focused on reactions relevant to biochemistry, combustion, or organic synthesis. Thus, the best hope for extrapolating our theory to chemistry as a whole would be to apply empirically-grounded rule-based algorithms to generate reaction networks *in silico*. Recent advances in computational systems chemistry are very promising (42, 43), but much more work is needed, especially to add realistic rate constants to these reactions.

The current SDAS framework allows for seeding but does not directly explain the source of seeds. Thus, while the model can identify autocatalytic modules that would be triggered if a particular molecule or a set of molecules magically appeared, it says nothing about the likelihood of seeding. This is an oversimplification. We know that a seed molecule can only arise by a chemical reaction happening somewhere, so shouldn't the reactions that generate the seed be included in the reaction network? One solution would be to include kinetics and use models, such as the Gillespie algorithm (44), that can translate low reaction rates into discrete stochastic

events. Another approach might be to utilize two or more distinct networks of allowed reactions, each interpreted as the set of feasible reactions in a particular environmental context, and then allow rare transfer of materials from one environment into another. Finally, the ultimately preferable approach would be to combine explicit kinetics and explicit spatial structure so that improbable local reactions and rare dispersal events can be allowed to simultaneously affect chemical ecosystem dynamics.

In addition to further theoretical research, we hope that our work will stimulate laboratory experiments to test our main conclusions. If our model is correct, SDASs might be a common feature of driven, “messy” chemical reaction systems. The obvious strategy to detect them in the laboratory would be to drive a chemical system with a flux of simple food chemicals using a continuously-stirred tank reactor (CSTR), or its poor-man’s equivalent, a serial batch-transfer-with-dilution experiment (45–47). We predict that it will be possible to find chemicals or sets of chemicals that are more complex than any yet in the system that, when added transiently to the vessel as seeds, cause the system to transition to a new steady-state composition. Moreover, we expect the new steady state to often include the seed chemicals (or the product of the seed chemicals in the case of pseudo-seeds). Finally, by maintaining such experiments for long periods, we believe it might be possible to detect the spontaneous emergence of evolution-like dynamics in the laboratory (47, 48). Such work would significantly advance our understanding of the origin of life by explaining how the stepwise accretion of complexity can occur even prior to the appearance of gene-based inheritance.

Our SDAS theory only requires that the system to which it is applied can be expressed in the form of a system of reaction equations with clear stoichiometry, with certain entities designated as the ultimate food. As a result, it may have significance beyond the origin-of-life field by providing a general framework to describe and analyze many complex systems that contain autocatalytic motifs, such as biological ecosystems, social media networks, and economies. While work would be needed to map concepts such as food, seeds, and stoichiometry between these cases, it seems possible that SDAS theory may ultimately have broad applicability in explaining the dynamics in diverse kinds of complex systems.

4. MATERIALS AND METHODS

4.1 Preprocessing databases of reactions

Abiotic reaction database

The reaction network assembled by Adam et al. (2021) (3) includes the following categories: free radical reactions, mineral geochemical reactions, amino acid production, chloride radical and polar reactions, nitrile radical and polar reactions, RNA nucleotide assembly, nuclear decay, and physicochemical reactions. We processed this database by the following steps.

First, we excluded the nuclear decay reactions because we did not plan to put radioactive atoms into the ultimate food set.

Second, with kind help from Dr. Zachary R. Adam and Dr. Albert C. Fahrenbach, we deleted duplicate reactions, added a few new reactions that were not in the original database, balanced some reaction equations, and excluded the reactions without clear stoichiometry. This is because our method requires stoichiometry of reactions.

Third, we added the formose reaction into the database. According to Breslow's mechanism (6), the formose reaction is driven by aldol and retro-aldol reactions and aldose-ketose isomerization. In combination these reactions allow low-carbon-number monosaccharides to generate high-carbon-number monosaccharides. Therefore, we added reversible aldol reactions and reversible aldose-ketose isomerization among formaldehyde, glycolaldehyde, and monosaccharides with no more than 8 carbon atoms. Optical isomers were not distinguished from each other. Formaldehyde dimerization was not added because it is very slow and its reaction mechanism is unclear but surely neither aldol/retro-aldol reaction nor aldose-ketose isomerization.

Fourth, every reaction labeled reversible was split into two unidirectional reactions.

Biotic reaction database

We processed the reaction database curated by Xavier et al. (2020) (4) to obtain the biotic reaction database by the following steps.

First, we removed all reactions involving chemical species that do not have specific molecular mass, such as reduced ferredoxin (KEGG: C00138), acyl-carrier protein (KEGG: C00229),

starch (KEGG: C00369), and long-chain aldehyde (KEGG: C00609), because they sometimes result in “fake” stoichiometric relationships. For example, the reaction: starch + H₂O ↔ dextrin + starch (KEGG: R02108) would make starch an infinite source of starch as long as H₂O is provided. Considering that glycans have both “G”-started (meaning “Glycan”) and “C”-started (meaning “Compound”) KEGG entries, the reactions involving “G”-started entries were also removed because these reactions are redundant.

Second, we added some obviously spontaneous reactions that were missing, such as H₂O ↔ H⁺ + OH⁻ and H₂CO₃ ↔ H⁺ + HCO₃⁻.

Third, we added the reaction KEGG R06974 into the biotic reaction database. This reaction is very similar to the reaction R06975 (Fig. S5): both reactions use HCOOH as the carbon donor to add a -CHO to -NH₂ and form a -NH-CHO with ATP hydrolysis providing energy for the reaction. However, R06975 is in the network curated by Xavier et al. (2020) (4) while R06974 is not, presumably because the annotations of R06974 in the KEGG database are not as detailed as those of R06975, and thus R06974 was filtered out. We also conducted analyses without R06974. In that case, SDAS-1a, SDAS-1b, and SDAS-2a were not affected but SDAS-1c and SDAS-2b were missing. However, we opted to present results that included the plausible reaction R06974 so as to better illustrate the potential for ecosystem-level feedback without resorting to analyzing the entirety of KEGG.

Fourth, as all reactions in the KEGG biochemical reaction database are labeled reversible, every reaction was split into two unidirectional reactions. The reaction following the forward direction specified in the KEGG database has a suffix “.a” to its entry, and that of the reverse direction has a suffix “.b”.

Fifth, the reactions that are labeled as multi-step were removed because each step is already a reaction in the database. Although keeping these multi-step reactions may not have big impact on the detection of SDAS existence, decreasing the number of reactions in the stoichiometric matrix should help accelerate the computation.

4.2 Network expansion

The set $\mathbf{R} = \{r_1, r_2, \dots, r_i, \dots, r_n\}$ is a set of multiple reactions r_i 's that are allowed. Each r_i specifies reactants and products, and the union of all reactants and products across all r_i 's is the maximum set of chemical species $\mathbf{S} = \{k_1, k_2, \dots, k_j, \dots, k_m\}$. We define an operation called full network expansion, $\Xi(\mathbf{S}_O, \mathbf{R}) = (\mathbf{S}_E, \mathbf{R}_E)$, where \mathbf{S}_O is the subset of \mathbf{S} where the expansion starts, \mathbf{S}_E the set of chemical species resulting from the expansion, and \mathbf{R}_E the expanded set of reactions resulting from the expansion. The expansion is conducted as follows:

- (i) Let $\mathbf{R}_E = \emptyset$; define a set of reactions $\mathbf{R}' = \mathbf{R}$; let $\mathbf{S}_E = \mathbf{S}_O$.
- (ii) Define a temporary set of chemical species $\mathbf{S}' = \emptyset$.
- (iii) For a reaction r_i in \mathbf{R}' , check if the reactants required by r_i are all present in \mathbf{S}_E ; if so, move r_i from \mathbf{R}' to \mathbf{R}_E , and scan through the products of r_i to add the chemical species that are not in \mathbf{S}_E to \mathbf{S}' . Do this for all reactions in \mathbf{R}' . Then add all chemical species in \mathbf{S}' to \mathbf{S}_E . If during this step, no reaction in \mathbf{R}' is moved, then the expansion is finished; otherwise, proceed to (ii).

4.3 Identifying cliques

Let us assume that \mathbf{S}_F is a set of chemical species resulting from a full expansion within the set of allowed reactions \mathbf{R} from a set of ultimate food \mathbf{S}_{UF} . Two non-empty seed sets of non-food chemical species \mathbf{S}_{P1} and \mathbf{S}_{P2} are said to be in the same clique if (a) $\Xi(\mathbf{S}_F \cup \mathbf{S}_{P1}, \mathbf{R}) = \Xi(\mathbf{S}_F \cup \mathbf{S}_{P2}, \mathbf{R})$, and (b) for any proper subset \mathbf{S}'_{P1} of \mathbf{S}_{P1} and any proper subset \mathbf{S}'_{P2} of \mathbf{S}_{P2} , $\Xi(\mathbf{S}_F \cup \mathbf{S}'_{P1}, \mathbf{R}) \neq \Xi(\mathbf{S}_F \cup \mathbf{S}_{P1}, \mathbf{R})$ and $\Xi(\mathbf{S}_F \cup \mathbf{S}'_{P2}, \mathbf{R}) \neq \Xi(\mathbf{S}_F \cup \mathbf{S}_{P2}, \mathbf{R})$.

In this paper, we only investigated the cliques consisting of seeds that are individual chemical species (i.e., singleton seeds). Nonetheless, the principle could be expanded to potential seed sets comprising more than one chemical species.

4.4 Detecting Seed-Dependent Autocatalytic Systems (SDASs) by linear programming

Let us assume that a $(p-1) \times (q-1)$ stoichiometric matrix, where each row represents a chemical species and each column represents a unidirectional reaction, results from a full expansion within the set of allowed reactions \mathbf{R} . The row labels of this $(p-1) \times (q-1)$ stoichiometric matrix form a

chemical species set $S_F = \{k_1, k_2, \dots, k_{p-1}\}$, which is defined as the external food for this detection process. Now we select a non-empty set of non-food chemical species $S_P = \{k_p, k_{p+1}, \dots, k_{p+h}\}$ to serve as the potential seed set. For example, we can treat all chemicals not in the external food as potential seeds and search through these one at a time. Then we conduct a full expansion from the food set and the seed set, $\Xi(S_F \cup S_P, R) = (S_{FP}, R_{FP})$, generating $S_{FP} = \{k_1, k_2, \dots, k_m\}$ and $R_{FP} = \{r_1, r_2, \dots, r_n\}$.

A SDAS feeding on S_F exists if there is a vector of non-negative elements $\mathbf{x} = (x_q, x_{q+1}, \dots, x_n)$ such that

$$\sum_{j=q}^n x_j s_{ij} > 0 \quad (x_j \geq 0) \quad \forall i \in [p, m], \quad [1]$$

where s_{ij} is the entry at the i th row and j th column of the stoichiometric matrix.

To determine whether such an \mathbf{x} exists, we used the linear programming tool provided by SciPy v1.6.2 (<https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.linprog.html>). In addition to the constraint set by [1], this linear programming tool requires an objective function. Since we knew that the growth of an autocatalytic system feeding on the external food is unbounded when the external food is unlimited (Fig. 1), we could set an objective function to find the maximum $\sum_{j=q}^n x_j s_{ij}$ for an internal species k_i ($i \in [p, m]$). If this objective function was found to be unbounded, we knew that a feasible region constrained by [1] must exist, indicating that a SDAS must exist. We simply let $i = p$ and set

$$\max_{x_q, \dots, x_n} \sum_{j=q}^n x_j s_{pj} \quad [2]$$

as the objective function.

We used the “highs” method (49) to confirm the existence of SDASs. Once the SDAS was confirmed to exist, we ran the integer programming process, described in the next section, to find autocatalytic cores within the SDAS, subject to further specific constraints.

4.5 Detecting minimum-reaction-number autocatalytic cores by integer programming

We can use linear programming to identify a SDAS, but we also desire to find small autocatalytic cores within each SDAS, because these are easier to visualize and could potentially guide future experimental studies. Specifically, considering that there were likely few types of catalysts in the prebiotic world, we focus on finding the autocatalytic cores with as few reactions as possible. As a result, we want to minimize the number of positive components of \mathbf{x} while the reactions corresponding to positive x_j 's still form an autocatalytic core feeding on the external food. In fact, the original SDAS may contain multiple autocatalytic cores. In this section we describe a method based on integer linear programming to enumerate small-cardinality autocatalytic cores within a given SDAS.

If there exists a vector \mathbf{x} , such that [1] is fulfilled, then by scaling \mathbf{x} , we may ensure that

$$\sum_{j=q}^n x_j s_{ij} \geq 1 \quad (x_j \geq 0) \quad \forall i \in [p, m]. \quad [3]$$

Thus, without loss of generality, we may use [3] as the constraint. There may be many possible \mathbf{x} 's that fulfill [3], and we used integer programming to seek and enumerate SDASs with desirable properties.

To find smaller systems, we seek a set of columns $\mathbf{T} \subseteq [q, n]$ such that if $j \in \mathbf{T}$ and $\exists i \in [p, m]$ which makes $s_{ij} \neq 0$, then $\exists x_j \in \mathbb{R}_+^{n-q+1}$ such that $\sum_{j=q}^n x_j s_{ij} \geq 1$. Of course, the set \mathbf{T} should have $|\mathbf{T}| \geq 1$. Finding such a set \mathbf{T} can be accomplished in a systematic manner by seeking solutions to a linear-inequality system wherein some of the variables are required to take integer values.

In the formulation, we use binary variables $z_j \in \{0, 1\}$ that take the value 1 if and only if column $j \in [q, n]$ is in the set \mathbf{T} , and we will minimize $\sum_{j=q}^n z_j$ to minimize the cardinality of the selected set. Because an autocatalytic core must have at least one reaction, it is obvious that $\sum_{j=q}^n z_j \geq 1$.

Let β_j be the upper bound on the number of the reaction r_j that can occur in an autocatalytic core. We must enforce that if column j is not selected for the autocatalytic core (i.e., $z_j = 0$), then its level of reaction x_j must also be zero, which is done with the algebraic constraints $x_j \leq \beta_j z_j$.

For the selected set of reactions \mathbf{T} to be an autocatalytic core, let us first define a set $\Omega_{\mathbf{T}}$ that contains and only contains the row indices of all chemical species that are involved in the reactions in \mathbf{T} and are not external species (i.e., $\Omega_{\mathbf{T}} \subseteq [p, m]$), then we must make sure that

$$\sum_{j=q}^n x_j s_{ij} \geq 1 \quad \forall i \in \Omega_{\mathbf{T}}. \quad [4]$$

This is done by introducing additional binary variables $y_i \in \{0, 1\}$ ($i \in [p, m]$) indicating if species k_i is involved in the autocatalytic core represented by the positive components of the vector $\mathbf{z} = (z_q, z_{q+1}, \dots, z_n)$. Then, the following set of linear inequalities accomplish the conditions in [4]

$$\sum_{j=q}^n x_j s_{ij} \geq 1 - M_i(1 - y_i) \quad \forall i \in [p, m], \quad [5]$$

where

$$M_i = 1 - \sum_{j=q: s_{ij} < 0}^n \beta_j s_{ij}. \quad [6]$$

To understand how [5] works, imagine that k_i is involved in the autocatalytic core, then $y_i = 1$ and [5] is equivalent to [4]. In contrast, if k_i is not involved in the autocatalytic core, then $y_i = 0$ and [5] is equivalent to

$$\sum_{j=q}^n x_j s_{ij} \geq \sum_{j=q: s_{ij} < 0}^n \beta_j s_{ij}. \quad [7]$$

Because β_j is the upper bound on x_j , [7] should always hold and thus it is redundant in the linear system, representing the fact that [4] does not need to be considered for a k_i that is not included in the autocatalytic core.

It is necessary to link a reaction and the chemical species that are involved in the reaction. If a reaction r_j ($j \in [q, n]$) is selected for an autocatalytic core (i.e., $z_j = 1$), any chemical species k_i that is involved in r_j must also exist in the autocatalytic core (i.e., $y_i = 1$). Therefore, for any reaction r_j ($j \in [q, n]$), we define a set Ω_j that contains and only contains the row indices of all

chemical species that are involved in r_j and are not external species (i.e., $\Omega_j \subseteq [p, m]$). Then we apply the constraint

$$y_i \geq z_j \quad \forall j \in [q, n], \forall i \in \Omega_j \quad [8]$$

which guarantees that once r_j is included in an autocatalytic core (i.e., $z_j = 1$), k_i need to be sustainably synthesized (i.e., $y_i = 1$).

This gives a full integer programming formulation for finding a minimum-cardinality autocatalytic core among the reactions $\{r_q, r_{q+1}, \dots, r_n\}$. We set the integer programming problem as to find

$$\min_{x,y,z} \sum_{j=q}^n z_j \quad [9]$$

which is constrained by

$$\left\{ \begin{array}{l} x_j \geq 0 \quad \forall j \in [q, n] \\ z_j \in \{0, 1\} \quad \forall j \in [q, n] \\ y_i \in \{0, 1\} \quad \forall i \in [p, m] \\ \sum_{j=q}^n z_j \geq 1 \\ x_j \leq \beta_j z_j \quad \forall j \in [q, n] \\ \sum_{j=q}^n x_j s_{ij} \geq 1 - M_i(1 - y_i) \quad \forall i \in [p, m] \\ y_i \geq z_j \quad \forall j \in [q, n], \forall i \in \Omega_j \end{array} \right. \quad [10]$$

This formulation can be solved computationally using a state-of-the-art integer programming software, such as Gurobi. The positive elements of the binary solution vector \mathbf{z} indicate the reaction set \mathbf{T} in a minimum-cardinality autocatalytic core chosen from the set of reactions $\{r_q, r_{q+1}, \dots, r_n\}$.

It should be noted that multiple autocatalytic cores with the same number of reactions may exist, and our integer programming can enumerate all minimum-cardinality autocatalytic cores by solving a sequence of integer programs. After a binary solution vector \mathbf{z} and its associated reaction set \mathbf{T} are identified, the constraint $\sum_{j=q:j \in \mathbf{T}} z_j \leq |\mathbf{T}| - 1$ may be added to [10], and then

the process repeats. Furthermore, if we want to find minimum-cardinality autocatalytic cores with at least D reactions, we can do it by replacing $\sum_{j=q}^n z_j \geq 1$ with $\sum_{j=q}^n z_j \geq D$ in [10].

V. ACKNOWLEDGMENTS AND FUNDING SOURCES

ACKNOWLEDGEMENTS

We thank Dr. Zachary R. Adam and Dr. Albert C. Fahrenbach for their kind help with curating the abiotic reaction database and the following for useful discussions: Alyssa Adams, Stephanie Colón-Santos, Emily Dolson, Praful Gagrani, Juan Perez Mercader, Alex Plum, Daniel Segrè, D. Eric Smith, Lena Vincent, and participants in the Evolving Chemical Systems workshop at the Santa Fe Institute (Nov. 2019).

FUNDING

This project is supported by NASA-NSF CESPOoL (Chemical Ecosystem Selection Paradigm for the Origins of Life) Ideas Lab grant (NASA-80NSSC17K0296) and University of Wisconsin Vice-Chancellor for Research and Graduate Education.

CONFLICT OF INTEREST

The authors declare that there is no conflict of interest.

VI. REFERENCES

1. J. Ladyman, J. Lambert, K. Wiesner, What is a complex system? *Euro Jnl Phil Sci* **3**, 33–67 (2013).
2. M. Mitchell, M. Newman, Complex systems theory and evolution. 5 (2001).
3. Z. R. Adam, A. C. Fahrenbach, S. M. Jacobson, B. Kacar, D. Y. Zubarev, Radiolysis generates a complex organosynthetic chemical network. *Scientific Reports* **11**, 1743 (2021).
4. J. C. Xavier, W. Hordijk, S. Kauffman, M. Steel, W. F. Martin, Autocatalytic chemical networks at the origin of metabolism. *Proc. R. Soc. B.* **287**, 20192377 (2020).
5. Z. Peng, A. M. Plum, P. Gagrani, D. A. Baum, An ecological framework for the analysis of prebiotic chemical reaction networks. *Journal of Theoretical Biology* **507**, 110451 (2020).
6. R. Breslow, On the mechanism of the formose reaction. *Tetrahedron Letters* **1**, 22–26 (1959).
7. H. J. (Jim) Cleaves, “Formose Reaction” in *Encyclopedia of Astrobiology*, M. Gargaud, *et al.*, Eds. (Springer, 2011), pp. 600–605.
8. M. Haas, S. Lamour, S. B. Christ, O. Trapp, Mineral-mediated carbohydrate synthesis by mechanical forces in a primordial geochemical setting. *Commun Chem* **3**, 1–6 (2020).
9. M. Kanehisa, Toward understanding the origin and evolution of cellular organisms. *Protein Science* **28**, 1947–1951 (2019).
10. M. Kanehisa, M. Furumichi, Y. Sato, M. Ishiguro-Watanabe, M. Tanabe, KEGG: integrating viruses and cellular organisms. *Nucleic Acids Research* **49**, D545–D551 (2021).
11. M. Kanehisa, S. Goto, KEGG: Kyoto Encyclopedia of Genes and Genomes. *Nucleic Acids Research* **28**, 27–30 (2000).
12. B. Damer, D. Deamer, Coupled Phases and Combinatorial Selection in Fluctuating Hydrothermal Pools: A Scenario to Guide Experimental Approaches to the Origin of Cellular Life. *Life* **5**, 872–887 (2015).
13. D. J. Donaldson, H. Tervahattu, A. F. Tuck, V. Vaida, Organic Aerosols and the Origin of Life: An Hypothesis. *Orig Life Evol Biosph* **34**, 57–67 (2004).
14. R. Lathe, Fast tidal cycling and the origin of life. *Icarus* **168**, 18–22 (2004).
15. M. R. Marín-Yaseli, E. González-Toril, C. Mompeán, M. Ruiz-Bermejo, The Role of Aqueous Aerosols in the “Glyoxylate Scenario”: An Experimental Approach. *Chemistry – A European Journal* **22**, 12785–12799 (2016).

16. W. Martin, J. Baross, D. Kelley, M. J. Russell, Hydrothermal vents and the origin of life. *Nature Reviews Microbiology* **6**, 805–814 (2008).
17. S. Maruyama, *et al.*, Nine requirements for the origin of Earth's life: Not at the hydrothermal vent, but in a nuclear geyser system. *Geoscience Frontiers* **10**, 1337–1357 (2019).
18. V. Sojo, B. Herschy, A. Whicher, E. Camprubí, N. Lane, The Origin of Life in Alkaline Hydrothermal Vents. *Astrobiology* **16**, 181–197 (2016).
19. G. Wächtershäuser, Before enzymes and templates: theory of surface metabolism. *Microbiol Rev* **52**, 452–484 (1988).
20. F. Westall, *et al.*, A Hydrothermal-Sedimentary Context for the Origin of Life. *Astrobiology* **18**, 259–293 (2018).
21. A. Blokhuis, D. Lacoste, P. Nghe, Universal motifs and the diversity of autocatalytic systems. *PNAS* **117**, 25230–25236 (2020).
22. M. Steel, J. C. Xavier, D. H. Huson, The structure of autocatalytic networks, with application to early biochemistry. *Journal of The Royal Society Interface* **17**, 20200488 (2020).
23. T. J. Smith, K. G. Boto, S. D. Frusher, R. L. Giddins, Keystone species and mangrove forest dynamics: the influence of burrowing by crabs on soil nutrient status and forest productivity. *Estuarine, Coastal and Shelf Science* **33**, 419–432 (1991).
24. R. MacArthur, Species packing and competitive equilibrium for many species. *Theoretical Population Biology* **1**, 1–11 (1970).
25. R. MacArthur, Species Packing, and What Competition Minimizes. *PNAS* **64**, 1369–1371 (1969).
26. H. J. Cleaves, The prebiotic geochemistry of formaldehyde. *Precambrian Research* **164**, 111–118 (2008).
27. L. E. Orgel, The Implausibility of Metabolic Cycles on the Prebiotic Earth. *PLOS Biology* **6**, e18 (2008).
28. K. B. Muchowska, *et al.*, Metals promote sequences of the reverse Krebs cycle. *Nat Ecol Evol* **1**, 1716–1721 (2017).
29. S. J. Varma, K. B. Muchowska, P. Chatelain, J. Moran, Native iron reduces CO₂ to intermediates and end-products of the acetyl-CoA pathway. *Nat Ecol Evol* **2**, 1019–1024 (2018).
30. K. B. Muchowska, S. J. Varma, J. Moran, Synthesis and breakdown of universal metabolic precursors promoted by iron. *Nature* **569**, 104–107 (2019).

31. K. B. Muchowska, S. J. Varma, J. Moran, Nonenzymatic Metabolic Reactions and Life's Origins. *Chem. Rev.* **120**, 7708–7744 (2020).
32. M. A. Keller, D. Kampjut, S. A. Harrison, M. Ralser, Sulfate radicals enable a non-enzymatic Krebs cycle precursor. *Nature Ecology & Evolution* **1**, 1–9 (2017).
33. M. Ralser, An appeal to magic? The discovery of a non-enzymatic metabolism and its role in the origins of life. *Biochemical Journal* **475**, 2577–2592 (2018).
34. M. A. Keller, A. V. Turchyn, M. Ralser, Non-enzymatic glycolysis and pentose phosphate pathway-like reactions in a plausible Archean ocean. *Mol Syst Biol* **10**, 725 (2014).
35. M. A. Keller, *et al.*, Conditional iron and pH-dependent activity of a non-enzymatic glycolysis and pentose phosphate pathway. *Sci. Adv.* **2**, e1501235 (2016).
36. J. Sobotta, *et al.*, A Possible Primordial Acetyleno/Carboxydutrophic Core Metabolism. *Life* **10**, 35 (2020).
37. D. A. Baum, The origin and early evolution of life in chemical composition space. *Journal of Theoretical Biology* **456**, 295–304 (2018).
38. V. Vasas, C. Fernando, M. Santos, S. Kauffman, E. Szathmáry, Evolution before genes. *Biology Direct* **7**, 1 (2012).
39. D. Segré, D. Ben-Eli, D. Lancet, Compositional genomes: Prebiotic information transfer in mutually catalytic noncovalent assemblies. *PNAS* **97**, 4112–4117 (2000).
40. D. Segré, B. Shenhav, R. Kafri, D. Lancet, The Molecular Roots of Compositional Inheritance. *Journal of Theoretical Biology* **213**, 481–491 (2001).
41. E. Smith, H. J. Morowitz, *The Origin and Nature of Life on Earth: The Emergence of the Fourth Geosphere* (Cambridge University Press, 2016)
<https://doi.org/10.1017/CBO9781316348772> (July 17, 2021).
42. A. Wołos, *et al.*, Synthetic connectivity, emergence, and self-regeneration in the network of prebiotic chemistry. *Science* **369** (2020).
43. J. L. Andersen, C. Flamm, D. Merkle, P. F. Stadler, Chemical Transformation Motifs—Modelling Pathways as Integer Hyperflows. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* **16**, 510–523 (2019).
44. D. T. Gillespie, Stochastic Simulation of Chemical Kinetics. *Annual Review of Physical Chemistry* **58**, 35–55 (2007).
45. S. Colón-Santos, G. J. T. Cooper, L. Cronin, Taming the Combinatorial Explosion of the Formose Reaction via Recursion within Mineral Environments. *ChemSystemsChem* **1**, e1900033 (2019).

46. A. J. Surman, *et al.*, Environmental control programs the emergence of distinct functional ensembles from unconstrained chemical reactions. *PNAS* **116**, 5387–5392 (2019).
47. Vincent, *et al.*, Chemical Ecosystem Selection on Mineral Surfaces Reveals Long-Term Dynamics Consistent with the Spontaneous Emergence of Mutual Catalysis. *Life* **9**, 80 (2019).
48. D. A. Baum, K. Vetsigian, An Experimental Framework for Generating Evolvable Chemical Systems in the Laboratory. *Orig Life Evol Biosph* **47**, 481–497 (2017).
49. Q. Huangfu, J. A. J. Hall, Parallelizing the dual revised simplex method. *Math. Prog. Comp.* **10**, 119–142 (2018).

VII. FIGURE LEGENDS




Fig. 1. Detection of SDASs. A stoichiometric matrix resulting from the network expansion starting with an external food set and a seed set can be split into four submatrices: the upper left one with rows representing external food and columns representing reactions involving only the external food, the upper right one with rows representing external food and columns representing reactions induced by seeding, the lower left one with rows representing internal chemical species induced by seeding and columns representing reactions involving only the external food, and the lower right one with rows representing internal chemical species induced by seeding and columns representing reactions induced by seeding. The four submatrices have different importance for detecting a SDAS. The two left ones can be ignored because none of their reactions involve chemical species outside the tier-0 system. The two right submatrices have reactions induced by seeding. Of these, the top submatrix can also be ignored, because these chemicals are provided for free by the environment, either as ultimate food or as products of reactions in the tier-0 system. So, our goal is to determine if all chemical species in the lower right submatrix can be sustainably produced by consuming the external food. Therefore, we want to find a linear combination of the reactions in this submatrix such that every internal chemical species can have a positive net change after this combination of reactions happen.


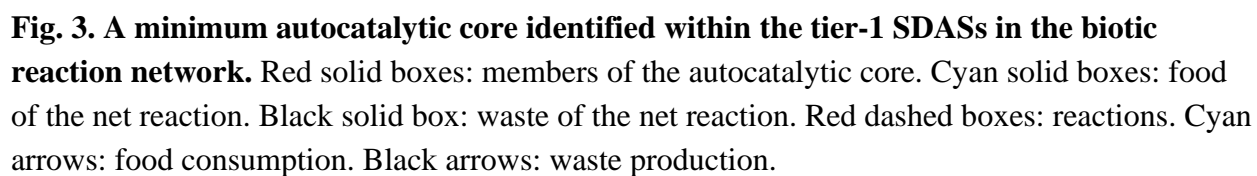


Fig. 2. A minimum autocatalytic core identified within the tier-1 SDAS in the abiotic reaction network. Red solid boxes: members of the autocatalytic core. Cyan solid boxes: food of the net reaction. Black solid box: waste of the net reaction. Red dashed boxes: reactions. Cyan arrows: food consumption. Black arrows: waste production. VL: visible light. IR: infrared.



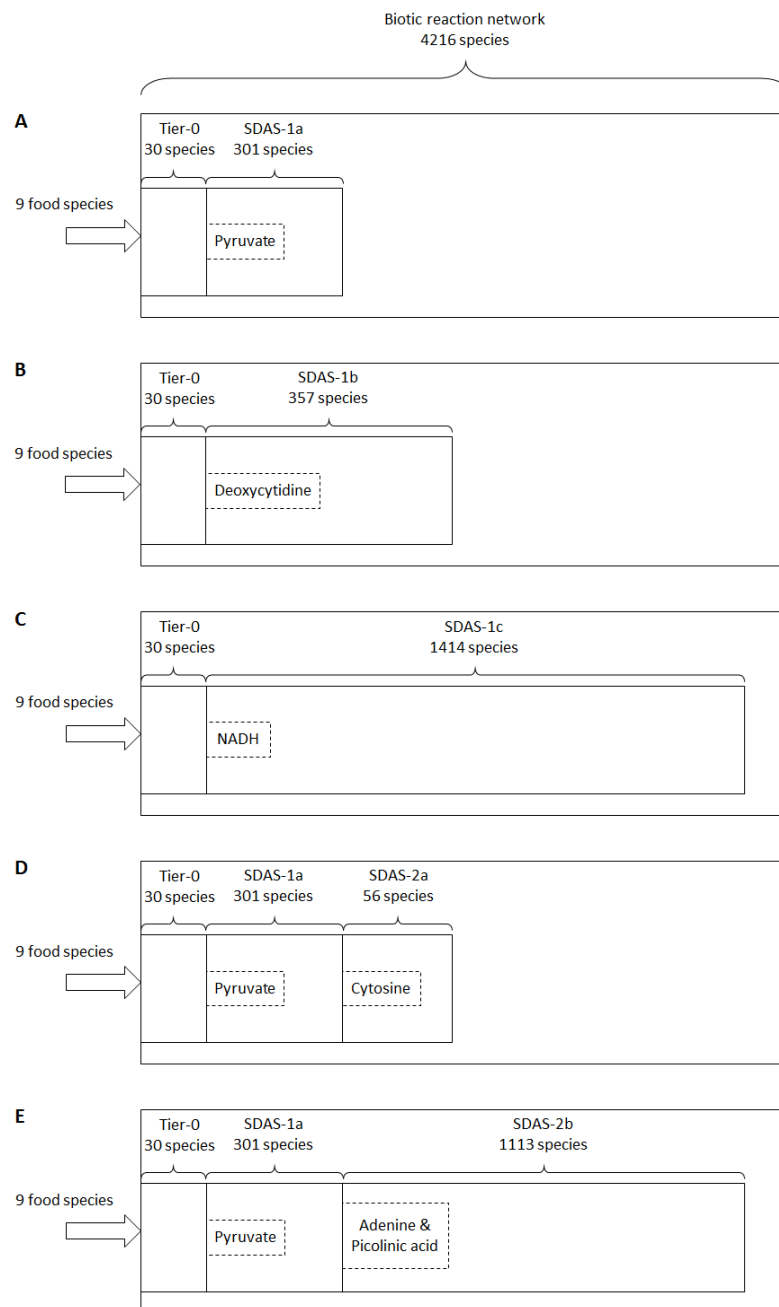


Fig. 4. Sequential seeding of SDASs at different trophic levels by simple molecules achieves a system that can be induced by a complex singleton seed. (A) SDAS-1a can be seeded by pyruvate. **(B)** SDAS-1b can be seeded by deoxycytidine. **(C)** SDAS-1c can be seeded by NADH. **(D)** Sequential seeding by pyruvate and cytosine results in a two-tier autocatalytic system, which is equivalent to SDAS-1b. **(E)** Sequential seeding by pyruvate and adenine together with picolinic acid results in a two-tier autocatalytic system, which is equivalent to SDAS-1c.

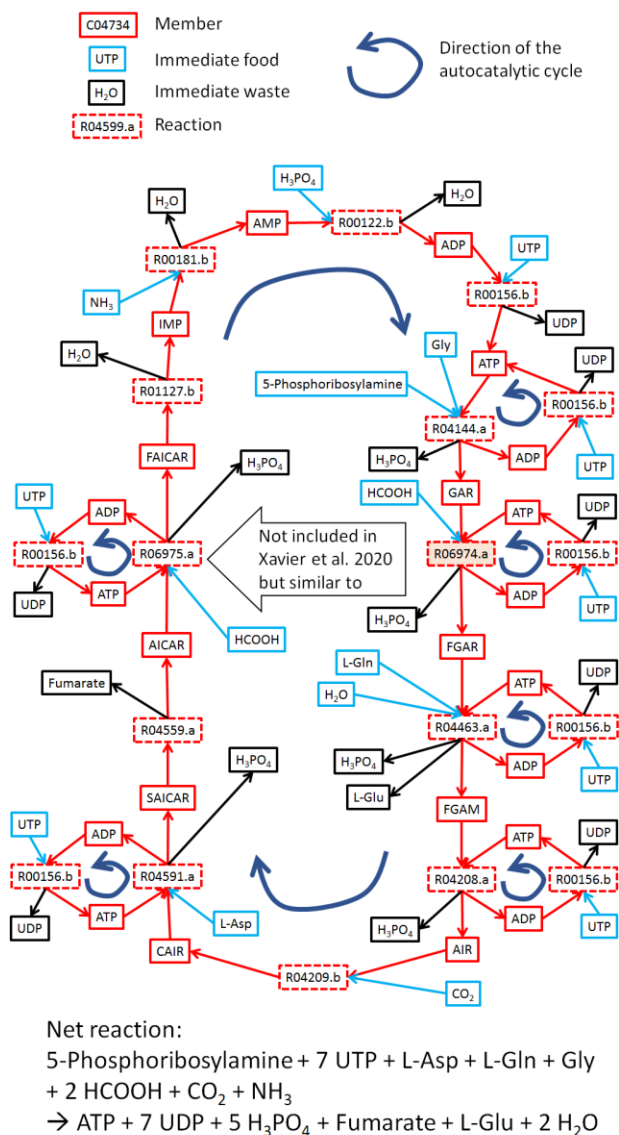


Fig. 5. An autocatalytic core synthesizing ATP within SDAS-2b detected in the biotic database. This autocatalytic core requires prior establishment of a lower-tier system able to supply specific organic molecules such as glycine and formic acid. Note that some of the immediate food chemicals, such as 5-phosphoribosylamine and UTP, are also synthesized by SDAS-2b reactions not shown in this graph. Red solid boxes: members of the autocatalytic core. Cyan solid boxes: immediate food of the autocatalytic core. Black solid box: immediate waste of the autocatalytic core. Red dashed boxes: reactions. Cyan arrows: food consumption. Black arrows: waste production.

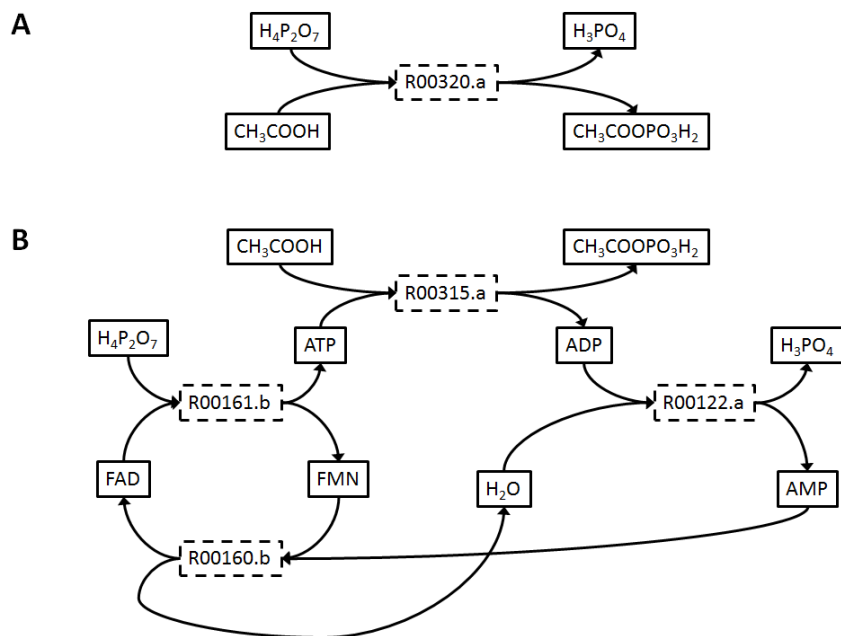


Fig. 6. Higher-tier chemicals may catalyze lower-tier reactions. An example in the biotic reaction network shows that FAD, a SDAS-2b chemical allows an alternative pathway for the phosphorylation of CH_3COOH , a tier-1 reaction, potentially increasing the rate at which the chemical ecosystem can assimilate phosphate. Solid box: chemical species. Dashed box: reaction.

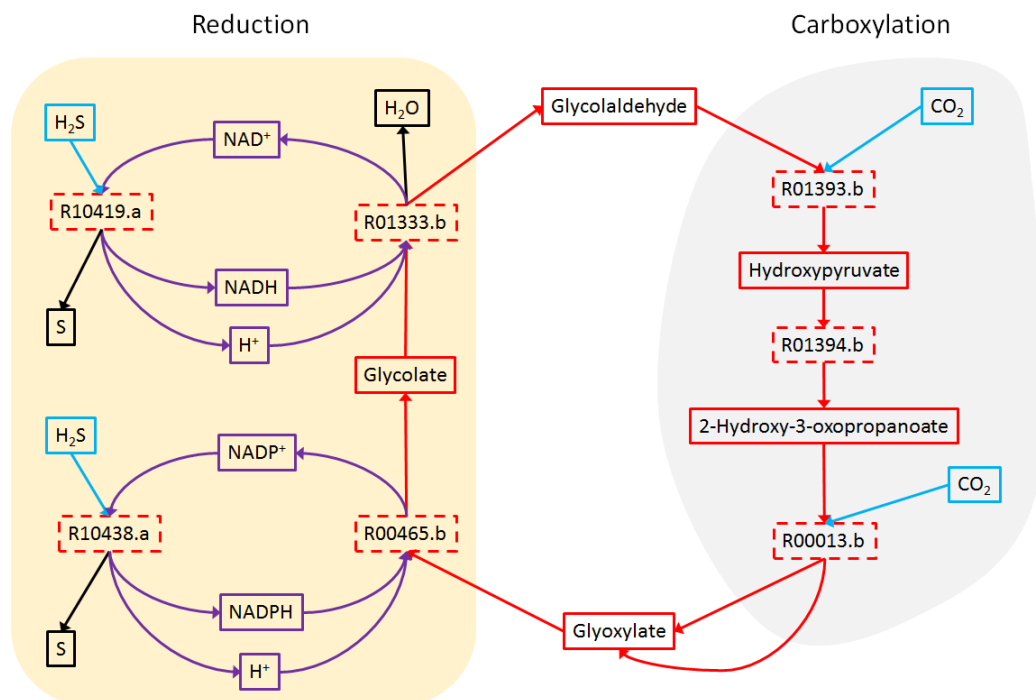


Fig. 7. A carbon-fixing autocatalytic cycle that becomes possible when NADH and NADPH are available. This autocatalytic cycle has two major modules: a reduction module leading from glyoxylate to glycolaldehyde, and a carboxylation module leading from glycolaldehyde to more glyoxylate. Because this cycle entails just seven reactions, it is likely to occur at a higher net rate than the pre-existing carbon fixation mechanisms. Red solid boxes: members of the autocatalytic cycle. Purple solid boxes: catalysts mediating the transfer of hydrogen from H_2S to organic molecules. Cyan solid boxes: immediate food of the autocatalytic cycle. Black solid box: immediate waste of the autocatalytic cycle. Red dashed boxes: reactions. Red arrows: syntheses of organic molecules. Purple arrows: catalysis. Cyan arrows: food consumption. Black arrows: waste production.

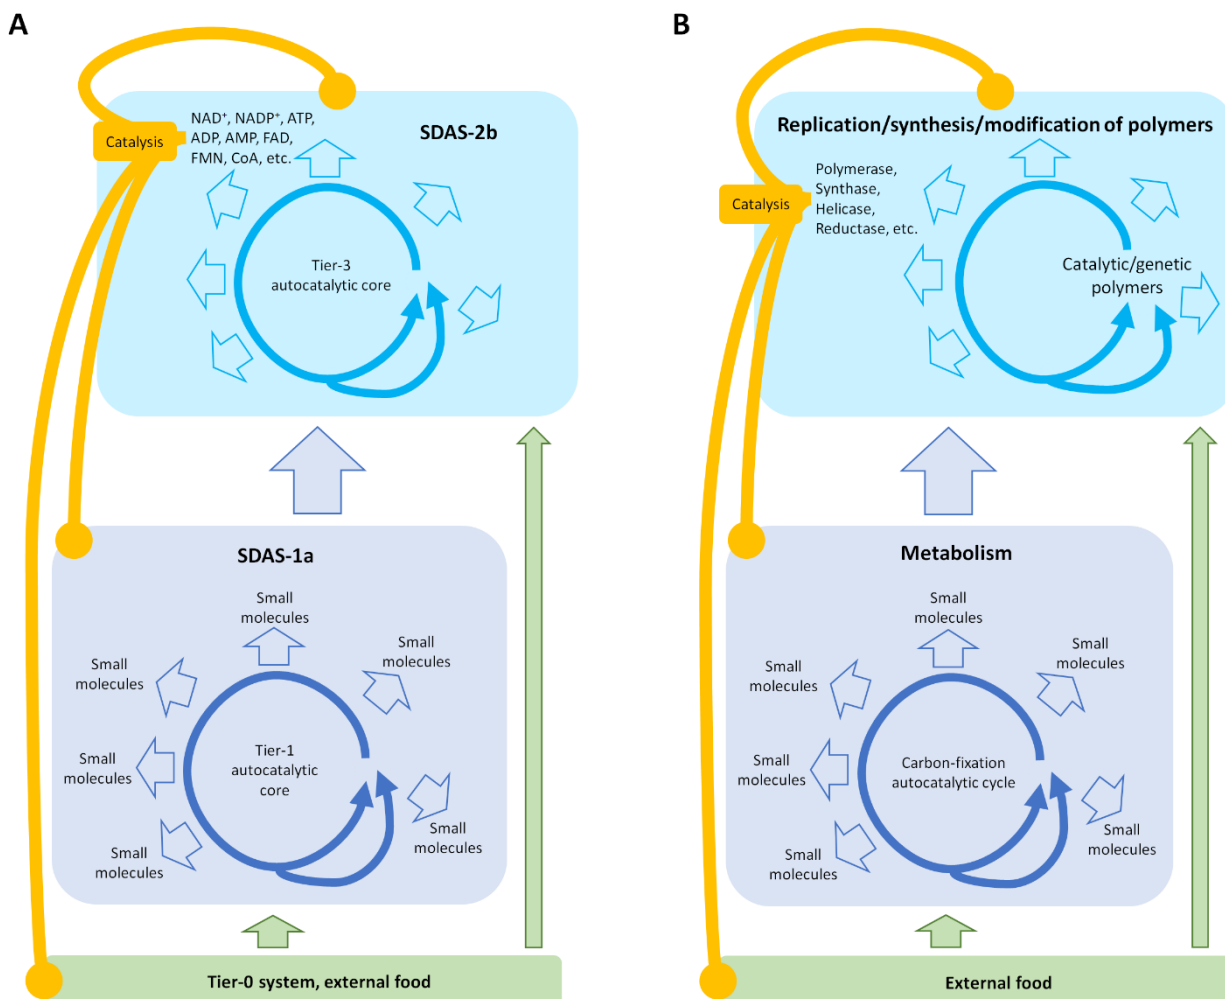


Fig. 8. The interactions between SDAs discovered in the biotic network are topologically similar to the interactions between metabolism, enzymes, and genetic polymers. (A) The summary of the interactions between SDAS-1a and SDAS-2b we found in the biotic reaction database. **(B)** A simplified model of the interactions between carbon fixation metabolism, enzymes, and genetic materials in extant life.