

Minimum autocatalytic networks cannot be approximated within any constant factor

Anonymous manuscript

September 7, 2026

Abstract

A reflexively autocatalytic and food-generated set (RAF) is a nonempty set of reactions in a catalytic reaction system whose reactants can all be built up from a designated food set and each of whose reactions is catalysed by a food molecule or by a product of the set. Steel, Hordijk and Smith (*J. Theor. Biol.*, 2013) proved that finding a RAF of minimum size is NP-hard and asked whether the smallest RAF can be approximated in polynomial time within a constant factor of optimal. We answer this question negatively for general finite catalytic reaction systems: unless $P = NP$, no polynomial-time algorithm returns, on every system that contains a RAF, a RAF whose size is within any fixed constant factor of the minimum. The proof is an approximation-preserving reduction from SET COVER. For an instance with m elements and n sets and an amplification parameter $M \geq 1$, we construct a system with one food molecule and $m + nM$ reactions, each with at most two reactants and exactly one product, and prove an exact structural theorem: its RAFs are precisely the sets $\mathcal{R}(D)$ indexed by the set covers D , and $|\mathcal{R}(D)| = m + M|D|$. With $M \geq m$, a RAF within factor c of optimal decodes in linear time to a cover within factor $2c - 1$ of optimal, so the reduction is an AP-reduction with constant 2. The same construction shows that estimating the minimum RAF size, even without producing a RAF, is NP-hard within any constant factor, that irreducible RAFs of the constructed systems are exactly the inclusion-minimal covers, and that a single system can contain two irreducible RAFs whose sizes differ by an arbitrarily large factor. The structural theorem, the objective identity, the decoding arithmetic and the conditional hardness transfer have been formalized and checked in the Lean 4 proof assistant with warnings promoted to errors.

Keywords. autocatalytic sets; RAF theory; catalytic reaction systems; approximation hardness; set cover; approximation-preserving reductions; irreducible RAFs; formal verification.

MSC 2020. 68Q17, 68W25, 92E20, 05C65, 68V20.

1 Introduction

Reflexively autocatalytic and food-generated sets (RAFTs) are a finite combinatorial model of collectively self-sustaining reaction networks. The concept originates in Kauffman’s autocatalytic sets of proteins [19, 20], was given its present form by Hordijk and Steel [13], and has grown into a structural theory with efficient algorithms [16, 12, 29, 30] and applications to the origin of metabolism and of life [14, 15, 32]. A catalytic reaction system (CRS) consists of molecule types, reactions with reactants and products, a food set of freely available molecules, and a catalysis relation. A nonempty set of reactions is a RAF if its reactions can be carried out in some order starting from the food set and every one of them is catalysed by a molecule that is food or is produced by the set.

Because unions of RAFs are RAFs, a CRS that contains a RAF contains a unique largest one, the *maxRAF*, and the RAF algorithm of Hordijk and Steel [13] computes it in polynomial time. The optimization problem points the other way. Steel, Hordijk and Smith [28] studied

minimal autocatalytic networks: the irreducible RAFs (iRAFs), which contain no proper RAF, and the smallest RAF overall. They proved that finding a RAF of minimum size is NP-hard, gave efficiently computable lower bounds on its size, showed how to certify for a fixed number k of listed iRAFs that the list is complete, and closed by asking whether the size of a smallest RAF can be approximated within a constant factor of optimal [28]. In the language of approximation algorithms this asks whether MIN-RAF is in the class APX.

This paper answers the question in the negative for general finite CRSs.

Theorem 1.1 (Main theorem, informal). *Unless $P = NP$, there is no polynomial-time algorithm and no constant $\alpha \geq 1$ such that, on every finite catalytic reaction system that contains a RAF, the algorithm returns a RAF of size at most α times the minimum. This holds even for systems with a single food molecule in which every reaction has at most two reactants and exactly one product, and in which the whole reaction set is itself a RAF.*

The construction

The proof is a reduction from SET COVER, the problem of covering a universe $U = \{0, \dots, m-1\}$ by as few as possible of given sets $A_0, \dots, A_{n-1} \subseteq U$. SET COVER is one of the original NP-complete problems [18, 10], the greedy algorithm approximates it within $1 + \ln m$ [17, 21, 3], and a sequence of results [22, 26, 8, 1, 23, 6] culminating in Dinur and Steurer’s theorem that $(1 - \varepsilon) \ln m$ -approximation is NP-hard for every fixed $\varepsilon > 0$ shows that no polynomial-time algorithm achieves any constant factor unless $P = NP$. We use only this last, weakest consequence.

Given a SET COVER instance $\mathcal{I}$ and an amplification parameter $M \geq 1$, Construction 3.1 builds a CRS $\mathcal{Q}_M(\mathcal{I})$ from two reaction families. A chain of m gate reactions $g_0, \dots, g_{m-1}$ represents the universe: g_0 consumes the single food molecule f and produces y_0 , and g_i consumes f and y_{i-1} and produces y_i . Each set A_j is represented by a block of M chained reactions $b_{j,0}, \dots, b_{j,M-1}$ that starts from y_{m-1} and produces $z_{j,0}, \dots, z_{j,M-1}$ in turn. The final product $h_j = z_{j,M-1}$ of block j is its hub: it catalyses every reaction of block j , and it catalyses exactly those gates g_i with $i \in A_j$. Nothing else catalyses anything.

The construction is governed by an exact structural theorem rather than by inequalities. Write $\mathcal{R}(D)$ for the set consisting of all m gates together with the complete blocks of the sets indexed by $D \subseteq \{0, \dots, n-1\}$.

Theorem 1.2 (Exact correspondence, Theorem 4.4 below). *Let $m, M \geq 1$. A set S of reactions of $\mathcal{Q}_M(\mathcal{I})$ is a RAF if and only if $S = \mathcal{R}(D)$ for a set cover D of $\mathcal{I}$. Moreover $|\mathcal{R}(D)| = m + M|D|$.*

The two directions have different flavours. If D is a cover then $\mathcal{R}(D)$ can be built from f in $m + M$ closure stages, its block reactions are catalysed by their own hubs, and its gates are catalysed by the hubs of the covering sets. Conversely, let S be a RAF. Every catalyst in $\mathcal{Q}_M(\mathcal{I})$ is a hub h_j , and the only reaction producing h_j is $b_{j,M-1}$; so any block reaction in S forces the last reaction of its block into S , and food generation together with the uniqueness of the reaction producing each molecule then forces the whole block and, through y_{m-1} , the whole gate chain. Finally, a gate $g_i \in S$ must be catalysed by a hub $h_j \in \text{cl}_S(\{f\})$ with $i \in A_j$, which says exactly that the selected blocks cover i .

Consequently $\text{OPT}_{\text{RAF}}(\mathcal{Q}_M(\mathcal{I})) = m + M \tau(\mathcal{I})$, where $\tau(\mathcal{I})$ is the minimum cover size, and minimum covers correspond exactly to minimum RAFs. The additive term m is the price of representing the universe; the amplification M makes the variable term dominate. With $M \geq m$ a RAF S with $|S| \leq c \cdot \text{OPT}_{\text{RAF}}$ decodes, by reading off which final block reactions it contains, to a cover D with $|D| \leq (2c - 1)\tau(\mathcal{I})$ (Theorem 5.3). The reduction is therefore an approximation-preserving reduction in the sense of Crescenzi [4] and Ausiello et al. [2], with constant $\alpha = 2$, and Theorem 1.1 follows.

Why NP-hardness alone was not enough

The exact correspondence already holds for $M = 1$, and it gives a short new proof of the NP-hardness of MIN-RAF in a highly restricted class of systems (Corollary 4.6). But with $M = 1$ an approximate RAF of size $c(m + \tau(\mathcal{I}))$ decodes only to a cover of size at most $(c - 1)m + c\tau(\mathcal{I})$, which is useless when $\tau(\mathcal{I}) \ll m$. Constant-factor hardness needs a reduction that controls the ratio, not only the optimum, and this is what the amplification and the exact classification of *all* RAFs, not just the optimal ones, provide.

Further consequences

Three further results come out of the same construction.

- *Value estimation.* The reduction is gap-preserving. Unless $P = NP$, no polynomial-time algorithm outputs a number within a constant factor of the minimum RAF size, whether as an upper estimate, a lower estimate, or a two-sided estimate, even without producing a RAF (Proposition 7.3).
- *Irreducible RAFs.* The iRAF of $\mathcal{Q}_M(\mathcal{I})$ are exactly the sets $\mathcal{R}(D)$ with D an inclusion-minimal cover (Proposition 8.2). Taking $M = 1$ and the incidence instance of a hypergraph, iRAFs are exactly minimal transversals, so enumerating the iRAFs of a CRS is at least as hard as hypergraph dualization [7, 9].
- *No universal iRAF size ratio.* For every c there is a CRS with two iRAFs whose sizes differ by a factor greater than c (Proposition 8.4). In particular, no procedure that merely returns *some* iRAF, such as the random reaction-deletion heuristic of [12], can be a constant-factor approximation algorithm for MIN-RAF.

Formal verification

The exact correspondence, the cardinality identity, the optimum correspondence, the decoding arithmetic, the encoding-size bounds, the gap arithmetic, the iRAF results and the conditional hardness transfer have been formalized in Lean 4 [5] on top of Mathlib [31] and compiled with warnings promoted to errors; details are in Section 10. The formalization treats polynomial time abstractly, as a predicate on algorithms closed under the explicit construction and decoder, and takes SET COVER inapproximability as a hypothesis; Section 6 supplies the routine encoding argument. The proofs in this paper are complete and self-contained, and the formalization is an additional guarantee rather than a substitute for them. Before the formalization, the construction was tested exhaustively on all coverable instances with at most three elements and three sets and $M \leq 3$, comparing the literal RAF semantics with the predicted normal form on 1,747,004 reaction subsets without a discrepancy; those computations are reported in Section 10 for reproducibility and play no role in the proofs.

Organization

Section 2 fixes definitions and the hardness premise. Section 3 gives the construction and its provenance lemmas. Section 4 proves the exact correspondence. Section 5 proves the objective identity and the decoding theorem. Section 6 handles encoding size and proves the main theorem. Section 7 treats value estimation. Section 8 treats irreducible RAFs. Section 9 relates the results to the bounds of [28]. Section 10 describes the Lean formalization and the computational checks, and Section 11 lists open problems.

2 Preliminaries

All sets are finite. For a finite set X we write 2^X for its power set and $[k] = \{0, 1, \dots, k-1\}$.

2.1 Catalytic reaction systems and RAFs

Definition 2.1 (CRS). A *catalytic reaction system* (CRS) is a tuple $\mathcal{Q} = (X, R, \rho, \pi, F, C)$ where X is a finite set of *molecule types*, R is a finite set of *reactions*, $\rho, \pi: R \rightarrow 2^X$ assign to each reaction its finite sets of *reactants* and *products*, $F \subseteq X$ is the *food set*, and $C \subseteq X \times R$ is the *catalysis relation*; $(x, r) \in C$ is read “ x catalyses r ”. For $S \subseteq R$ we write $\pi(S) = \bigcup_{r \in S} \pi(r)$.

Multiplicities of reactants and products play no role in what follows and are ignored. We use the closure-based semantics of food generation and catalysis; Lemma 2.4 shows that it agrees with the ordering-based definitions of [13, 28].

Definition 2.2 (Closure). Let $\mathcal{Q}$ be a CRS and $S \subseteq R$. Define $W_0(S) = F$ and, for $i \geq 0$,

$$W_{i+1}(S) = W_i(S) \cup \bigcup \{\pi(r) : r \in S, \rho(r) \subseteq W_i(S)\}.$$

The sequence is increasing and stabilizes after at most $|X|$ steps; its union is the *closure* $\text{cl}_S(F)$. A reaction r is *enabled* at stage i if $\rho(r) \subseteq W_i(S)$.

Definition 2.3 (Food-generated sets and RAFs). Let $\mathcal{Q}$ be a CRS and $S \subseteq R$.

- (i) S is *food-generated* if $\rho(r) \subseteq \text{cl}_S(F)$ for every $r \in S$.
- (ii) S is *reflexively autocatalytic* if for every $r \in S$ there is $x \in \text{cl}_S(F)$ with $(x, r) \in C$.
- (iii) S is a *RAF* if it is nonempty, food-generated and reflexively autocatalytic.

An *irreducible RAF* (iRAF) is a RAF that contains no RAF as a proper subset. When $\mathcal{Q}$ has a RAF, $\text{OPT}_{\text{RAF}}(\mathcal{Q})$ denotes the minimum cardinality of a RAF of $\mathcal{Q}$.

The optimization problem MIN-RAF is: given a CRS that contains a RAF, find a RAF of minimum cardinality [28]. The existence of a RAF can be decided in polynomial time by the RAF algorithm [13], so the promise is harmless.

Lemma 2.4 (Provenance). Let $\mathcal{Q}$ be a CRS and $S \subseteq R$.

- (i) Every $x \in \text{cl}_S(F)$ is a food molecule or a product of some reaction in S .
- (ii) If S is food-generated then $\text{cl}_S(F) = F \cup \pi(S)$. Consequently a food-generated set S is reflexively autocatalytic if and only if every $r \in S$ is catalysed by a food molecule or by a product of some reaction in S .

Proof. (i) is immediate by induction on the stage at which x enters the closure. For (ii), the inclusion $\subseteq$ is (i). Conversely $F = W_0(S) \subseteq \text{cl}_S(F)$, and if $r \in S$ then $\rho(r) \subseteq \text{cl}_S(F)$, so r is enabled at some stage i and $\pi(r) \subseteq W_{i+1}(S)$. The last sentence follows by substituting $F \cup \pi(S)$ for the closure in Definition 2.3(ii). $\square$

Hordijk and Steel [13] define S to be food-generated when its reactions admit an ordering $r_1, \dots, r_k$ with $\rho(r_j) \subseteq F \cup \pi(\{r_1, \dots, r_{j-1}\})$, and reflexively autocatalytic when every $r \in S$ has a catalyst in $F \cup \pi(S)$. Ordering the reactions of S by the stage at which they are first enabled shows that the ordering condition is equivalent to Definition 2.3(i), and Lemma 2.4(ii) shows that the two catalysis conditions agree on food-generated sets. The RAFs of Definition 2.3 are therefore exactly the RAFs of [13, 28]. The Lean development proves the second half of Lemma 2.4(ii) in the form `isRAF_iff_foodGenerated_and_productGraph` and uses it throughout.

2.2 Set cover and the hardness premise

Definition 2.5 (SET COVER). A SET COVER *instance* is a pair $\mathcal{I} = (m, (A_j)_{j \in [n]})$ with $m \geq 1$ and $A_j \subseteq [m]$ for each $j \in [n]$, subject to the promise $\bigcup_{j \in [n]} A_j = [m]$. A *cover* is a set $D \subseteq [n]$ with $\bigcup_{j \in D} A_j = [m]$, and $\tau(\mathcal{I})$ denotes the minimum cardinality of a cover. The *incidence count* is $L(\mathcal{I}) = \sum_{j \in [n]} |A_j|$, and the *size* of $\mathcal{I}$ is $m + n + L(\mathcal{I})$.

Sets may be empty or repeated; this only makes the hardness statements stronger. Since $m \geq 1$ every cover is nonempty, so $\tau(\mathcal{I}) \geq 1$.

Definition 2.6 (Solution approximation). Let $\alpha \geq 1$. A *factor- α solution approximation* for a minimization problem is an algorithm that, on every instance, returns a feasible solution of cost at most α times the optimum. For SET COVER the feasible solutions are covers with their cardinality; for MIN-RAF they are RAFs with their cardinality.

Hypothesis 2.7 (SET COVER inapproximability). If $P \neq NP$ then for no constant $\beta \geq 1$ does SET COVER have a polynomial-time factor- β solution approximation.

Hypothesis 2.7 is a theorem. Lund and Yannakakis [22] proved $\Omega(\log m)$ hardness under the assumption that NP has no quasi-polynomial-time algorithms; Raz and Safra [26] obtained $c \log m$ hardness for some $c > 0$ under $P \neq NP$; Feige [8] established the sharp threshold $(1 - \varepsilon) \ln m$ under the assumption $NP \not\subseteq DTIME(n^{O(\log \log n)})$; Alon, Moshkovitz and Safra [1] and Moshkovitz [23] narrowed the gap; and Dinur and Steurer [6] proved that $(1 - \varepsilon) \ln m$ -approximation is NP-hard for every fixed $\varepsilon > 0$. Any of the statements under $P \neq NP$ implies the hypothesis, since $\log m$ exceeds every constant on all but finitely many instances, and those finitely many can be solved exactly. We isolate the hypothesis because it is the only complexity-theoretic input to the paper and because the Lean formalization takes it as a named premise.

Definition 2.8 (AP-reductions [4, 2]). Let A and B be minimization problems. An *AP-reduction* from A to B with constant $\alpha \geq 1$ consists of polynomial-time computable maps f , taking instances x of A to instances $f(x)$ of B, and g , taking a pair (x, y) with y a feasible solution of $f(x)$ to a feasible solution $g(x, y)$ of x , such that for every $c \geq 1$: if y is within factor c of optimal for $f(x)$ then $g(x, y)$ is within factor $1 + \alpha(c - 1)$ of optimal for x .

If A AP-reduces to B and B has a polynomial-time factor- c solution approximation, then A has one with factor $1 + \alpha(c - 1)$. The reduction of this paper is an AP-reduction from SET COVER to MIN-RAF with $\alpha = 2$, since $1 + 2(c - 1) = 2c - 1$.

3 The amplified construction

Throughout, $\mathcal{I} = (m, (A_j)_{j \in [n]})$ is a SET COVER instance with $m \geq 1$ and $M \geq 1$ is an integer, the *block length* or *amplification*.

Construction 3.1 (The source $\mathcal{Q}_M(\mathcal{I})$). The CRS $\mathcal{Q}_M(\mathcal{I}) = (X, R, \rho, \pi, F, C)$ has

- molecule types $X = \{f\} \cup \{y_i : i \in [m]\} \cup \{z_{j,k} : j \in [n], k \in [M]\}$ and food set $F = \{f\}$;
- reactions $R = \{g_i : i \in [m]\} \cup \{b_{j,k} : j \in [n], k \in [M]\}$, the *gates* and the *block reactions*, with reactants and products as in Table 1;
- catalysis relation C consisting of the pairs $(h_j, b_{j,k})$ for all $j \in [n]$ and $k \in [M]$, and the pairs (h_j, g_i) for all $j \in [n]$ and $i \in A_j$, where h_j abbreviates the *hub* $z_{j,M-1}$ of block j .

For $D \subseteq [n]$ the *canonical reaction set* of D is

$$\mathcal{R}(D) = \{g_i : i \in [m]\} \cup \{b_{j,k} : j \in D, k \in [M]\}.$$

Reaction	Reactants ρ	Product π	Catalysts
g_0	$\{f\}$	y_0	h_j for every j with $0 \in A_j$
$g_i, 1 \leq i < m$	$\{f, y_{i-1}\}$	y_i	h_j for every j with $i \in A_j$
$b_{j,0}$	$\{f, y_{m-1}\}$	$z_{j,0}$	h_j
$b_{j,k}, 1 \leq k < M$	$\{f, z_{j,k-1}\}$	$z_{j,k}$	h_j

Table 1: The reactions of $\mathcal{Q}_M(\mathcal{I})$. The hub of block j is $h_j = z_{j,M-1}$. When $M = 1$ the hub is $z_{j,0}$ and the last row is empty.

The system has $m + nM$ reactions and $1 + m + nM$ molecule types. Every reaction has at most two reactants and exactly one product, and distinct reactions have distinct products. The only catalysts are hubs, and no food molecule catalyses anything. The reactions form a single chain $g_0, \dots, g_{m-1}$ from which n chains of length M branch off; the hub of each branch catalyses the whole branch and points back to the gates of the elements its set covers. Figure 1 shows a small example.

Instance: $m = 3$, $A_0 = \{0, 1\}$, $A_1 = \{1, 2\}$, $A_2 = \{2\}$; block length $M = 2$.

Gates: $g_0: f \rightarrow y_0$ $g_1: f + y_0 \rightarrow y_1$ $g_2: f + y_1 \rightarrow y_2$

Blocks: $b_{j,0}: f + y_2 \rightarrow z_{j,0}$, $b_{j,1}: f + z_{j,0} \rightarrow z_{j,1} = h_j$ ($j = 0, 1, 2$)

Catalysis: $h_0 \rightarrow \{g_0, g_1, b_{0,0}, b_{0,1}\}$, $h_1 \rightarrow \{g_1, g_2, b_{1,0}, b_{1,1}\}$, $h_2 \rightarrow \{g_2, b_{2,0}, b_{2,1}\}$

Covers: $\{0, 1\}$, $\{0, 2\}$, $\{0, 1, 2\}$. RAFs: $\mathcal{R}(\{0, 1\})$, $\mathcal{R}(\{0, 2\})$, $\mathcal{R}(\{0, 1, 2\})$, of sizes 7, 7, 9.

Figure 1: A small instance of Construction 3.1. The nine reactions have exactly three RAFs, one for each cover, and $\text{OPT}_{\text{RAF}} = 3 + 2 \cdot 2 = 7$.

The next lemma records the unique-provenance property that drives the reverse direction of the correspondence. We write cl_S for $\text{cl}_S(\{f\})$ in $\mathcal{Q}_M(\mathcal{I})$.

Lemma 3.2 (Unique provenance). *Let $S \subseteq R$ be any set of reactions of $\mathcal{Q}_M(\mathcal{I})$.*

(i) *If $y_i \in \text{cl}_S$ then $g_i \in S$.*

(ii) *If $z_{j,k} \in \text{cl}_S$ then $b_{j,k} \in S$.*

Proof. By Lemma 2.4(i), a molecule in cl_S other than f is a product of a reaction in S . The only reaction with product y_i is g_i , and the only reaction with product $z_{j,k}$ is $b_{j,k}$. $\square$

4 Exact characterization of the RAFs

We first show that a food-generated set is closed under predecessors along the chains.

Lemma 4.1 (Prefix closure). *Let S be a food-generated set of reactions of $\mathcal{Q}_M(\mathcal{I})$.*

(i) *If $g_i \in S$ then $g_q \in S$ for every $q \leq i$.*

(ii) *If $b_{j,k} \in S$ then $b_{j,l} \in S$ for every $l \leq k$.*

(iii) *If $b_{j,0} \in S$ then $g_i \in S$ for every $i \in [m]$.*

Proof. (i) By downward induction on i . If $g_i \in S$ with $i \geq 1$ then $y_{i-1} \in \rho(g_i) \subseteq \text{cl}_S$ because S is food-generated, so $g_{i-1} \in S$ by Lemma 3.2(i). (ii) Identically, using $z_{j,k-1} \in \rho(b_{j,k})$ and Lemma 3.2(ii). (iii) If $b_{j,0} \in S$ then $y_{m-1} \in \rho(b_{j,0}) \subseteq \text{cl}_S$, so $g_{m-1} \in S$ by Lemma 3.2(i), and (i) gives all gates. $\square$

Next, the catalysis condition forces the hubs.

Lemma 4.2 (Forced hubs). *Let S be a RAF of $\mathcal{Q}_M(\mathcal{I})$.*

(i) *If $b_{j,k} \in S$ then $b_{j,M-1} \in S$, and hence, by Lemma 4.1(ii), the whole block $\{b_{j,l} : l \in [M]\}$ is contained in S .*

(ii) *If $g_i \in S$ then there is j with $i \in A_j$ and $b_{j,M-1} \in S$.*

Proof. (i) The reaction $b_{j,k}$ has a catalyst in cl_S . Its only catalyst is $h_j = z_{j,M-1}$, so $z_{j,M-1} \in \text{cl}_S$ and $b_{j,M-1} \in S$ by Lemma 3.2(ii). (ii) The gate g_i has a catalyst in cl_S . Its catalysts are the hubs h_j with $i \in A_j$, so some $z_{j,M-1}$ with $i \in A_j$ lies in cl_S , and $b_{j,M-1} \in S$ by Lemma 3.2(ii). $\square$

For the forward direction we need explicit closure stages.

Lemma 4.3 (Reachability). *Let $D \subseteq [n]$ and $S = \mathcal{R}(D)$. Then $y_i \in W_{i+1}(S)$ for every $i \in [m]$, and $z_{j,k} \in W_{m+k+1}(S)$ for every $j \in D$ and $k \in [M]$. In particular S is food-generated.*

Proof. By induction on i : $g_0 \in S$ is enabled at stage 0 since $\rho(g_0) = \{f\} = W_0(S)$, so $y_0 \in W_1(S)$; and if $y_{i-1} \in W_i(S)$ then $g_i \in S$ is enabled at stage i , so $y_i \in W_{i+1}(S)$. Now let $j \in D$. Since $y_{m-1} \in W_m(S)$, the reaction $b_{j,0} \in S$ is enabled at stage m and $z_{j,0} \in W_{m+1}(S)$; and if $z_{j,k-1} \in W_{m+k}(S)$ then $b_{j,k} \in S$ is enabled at stage $m+k$, so $z_{j,k} \in W_{m+k+1}(S)$. Every reactant of every reaction in S is f or one of these molecules, so S is food-generated. $\square$

Theorem 4.4 (RAF–cover correspondence). *Let $\mathcal{I}$ be a SET COVER instance with $m \geq 1$ and let $M \geq 1$. A set S of reactions of $\mathcal{Q}_M(\mathcal{I})$ is a RAF if and only if there is a cover D of $\mathcal{I}$ with $S = \mathcal{R}(D)$. The cover is unique, namely $D = \{j \in [n] : b_{j,M-1} \in S\}$.*

Proof. Suppose S is a RAF and put $D = \{j : b_{j,M-1} \in S\}$. By Lemma 4.2(i) and Lemma 4.1(ii), a block reaction $b_{j,k}$ lies in S if and only if $j \in D$, and in that case S contains the whole block j . Since S is nonempty it contains either a block reaction or a gate; in the latter case Lemma 4.2(ii) supplies a block reaction in S anyway. So $D \neq \emptyset$, and $b_{j,0} \in S$ for some j , whence S contains every gate by Lemma 4.1(iii). Thus $S = \mathcal{R}(D)$. Finally, for every $i \in [m]$ the gate g_i lies in S , so by Lemma 4.2(ii) there is $j \in D$ with $i \in A_j$; that is, D is a cover.

Conversely, let D be a cover and $S = \mathcal{R}(D)$. Then S is nonempty since $m \geq 1$, and food-generated by Lemma 4.3. For catalysis, each $b_{j,k} \in S$ has $j \in D$, and its catalyst $h_j = z_{j,M-1}$ lies in $W_{m+M}(S) \subseteq \text{cl}_S$ by Lemma 4.3. Each gate g_i has, because D is a cover, some $j \in D$ with $i \in A_j$, and $h_j \in \text{cl}_S$ catalyses g_i . Hence S is a RAF. Uniqueness of D is clear from the definition of $\mathcal{R}(D)$, since $M \geq 1$. $\square$

Remark 4.5. Theorem 4.4 classifies every RAF, not only the optimal ones, and it does so for every $M \geq 1$. The map $D \mapsto \mathcal{R}(D)$ is an isomorphism from the poset of covers of $\mathcal{I}$, ordered by inclusion, onto the poset of RAFs of $\mathcal{Q}_M(\mathcal{I})$ (Lemma 8.1 below). Since $[n]$ itself is a cover, the whole reaction set $R = \mathcal{R}([n])$ is a RAF: the maxRAF of $\mathcal{Q}_M(\mathcal{I})$ is R , and the RAF algorithm of [13] returns all of R on every instance of the family.

Corollary 4.6 (NP-hardness in a restricted class). *MIN-RAF is NP-hard, and remains so on CRSs with a single food molecule, at most two reactants and exactly one product per reaction, no food catalysis, and in which the entire reaction set is a RAF.*

Proof. Take $M = 1$. By Theorem 4.4 and Lemma 5.1 below, $\text{OPT}_{\text{RAF}}(\mathcal{Q}_1(\mathcal{I})) = m + \tau(\mathcal{I})$, and a minimum RAF yields a minimum cover by reading off its block reactions. The construction is computable in polynomial time (Lemma 6.1). $\square$

Corollary 4.6 recovers the NP-hardness theorem of [28] by a different reduction. Its interest here is that it shows the two ingredients that the approximation result adds: the amplification M , and the fact that the classification in Theorem 4.4 covers every RAF.

5 Objective identity and approximation transfer

Lemma 5.1 (Affine cardinality). *For every $D \subseteq [n]$, $|\mathcal{R}(D)| = m + M|D|$.*

Proof. The map sending g_i to i and $b_{j,k}$ to (j, k) is a bijection from $\mathcal{R}(D)$ onto the disjoint union $[m] \sqcup (D \times [M])$. $\square$

Corollary 5.2 (Optimum correspondence). *Let $m, M \geq 1$. Then $\text{OPT}_{\text{RAF}}(\mathcal{Q}_M(\mathcal{I})) = m + M\tau(\mathcal{I})$, and a cover D has minimum cardinality among covers if and only if $\mathcal{R}(D)$ has minimum cardinality among RAFs.*

Proof. By Theorem 4.4 the RAFs are exactly the sets $\mathcal{R}(D)$ with D a cover, and by Lemma 5.1 the map $D \mapsto |\mathcal{R}(D)| = m + M|D|$ is strictly increasing in $|D|$. $\square$

Theorem 5.3 (Decoding). *Let $\mathcal{I}$ be a SET COVER instance with $m \geq 1$, let $M \geq m$, and let $c \geq 1$ be real. If S is a RAF of $\mathcal{Q}_M(\mathcal{I})$ with*

$$|S| \leq c \cdot \text{OPT}_{\text{RAF}}(\mathcal{Q}_M(\mathcal{I})),$$

then $D = \{j \in [n] : b_{j,M-1} \in S\}$ is a cover of $\mathcal{I}$ with $|D| \leq (2c - 1)\tau(\mathcal{I})$.

Proof. By Theorem 4.4, D is a cover and $S = \mathcal{R}(D)$. Write $\tau = \tau(\mathcal{I})$ and $d = |D|$. By Lemma 5.1 and Corollary 5.2 the hypothesis reads

$$m + Md \leq c(m + M\tau).$$

Subtracting m and using $m \leq M$ and $c - 1 \geq 0$,

$$Md \leq (c - 1)m + cM\tau \leq (c - 1)M + cM\tau = M(c\tau + (c - 1)),$$

and dividing by $M > 0$ gives $d \leq c\tau + (c - 1)$. Since $\tau \geq 1$ we have $c - 1 \leq (c - 1)\tau$, so $d \leq c\tau + (c - 1)\tau = (2c - 1)\tau$. $\square$

Remark 5.4. Together with Lemma 6.1 below, Theorem 5.3 says that $f(\mathcal{I}) = \mathcal{Q}_{\max\{2, m\}}(\mathcal{I})$ and $g(\mathcal{I}, S) = \{j : b_{j,M-1} \in S\}$ form an AP-reduction from SET COVER to MIN-RAF with constant $\alpha = 2$ (Definition 2.8). Any $M \geq m$ would do; the Lean formalization fixes $M = \max\{2, m\}$. The reduction is not an L-reduction in the sense of Papadimitriou and Yannakakis [25], because $\text{OPT}_{\text{RAF}} = m + M\tau$ is not bounded by a constant multiple of τ ; it is the ratio, not the optimum, that is preserved. The bound $2c - 1$ is not claimed to be tight; the value of the constant does not matter for the main theorem.

6 Encoding size and the main theorem

A CRS is encoded in the natural sparse way: a list of molecule identifiers, a list of food identifiers, for each reaction the lists of its reactants and products, and the list of catalysis pairs. Its *encoding size* is the total number of identifiers in these lists; the bit length is larger by a logarithmic factor, which is irrelevant to polynomiality.

Lemma 6.1 (Encoding size). *Let $\mathcal{I}$ be a SET COVER instance with $m \geq 1$ and let $M \geq 1$. The CRS $\mathcal{Q}_M(\mathcal{I})$ has $m + nM$ reactions and $1 + m + nM$ molecule types, and its encoding size is*

$$1 + 5(m + nM) + nM + L(\mathcal{I}) \leq 1 + 6(m + nM) + nm.$$

It can be written down in time polynomial in m, n, M and $L(\mathcal{I})$. With $M = \max\{2, m\}$ the encoding size is at most $1 + 6(m + n(m + 1)) + nm$, which is polynomial in the size $m + n + L(\mathcal{I})$ of $\mathcal{I}$. The decoder $S \mapsto \{j : b_{j,M-1} \in S\}$ runs in time linear in the encoding of S .

Proof. The reaction and molecule counts are read off from Construction 3.1. The lists contain $1 + m + nM$ molecule identifiers, one food identifier, $m + nM$ reaction identifiers, $2(m + nM) - 1$ reactant identifiers (g_0 has one reactant and every other reaction has two), $m + nM$ product identifiers, and $nM + L(\mathcal{I})$ catalysis pairs, each counted as one identifier pair; summing gives the stated expression, and $L(\mathcal{I}) \leq nm$ gives the bound. For $M = \max\{2, m\} \leq m + 1$ substitute $nM \leq n(m + 1)$. Each entry is produced by a constant number of arithmetic operations on indices, and the decoder scans the n reactions $b_{j, M-1}$. $\square$

Theorem 6.2 (Main theorem). *Assume $P \neq NP$. Then for no constant $\alpha \geq 1$ is there a polynomial-time factor- α solution approximation for MIN-RAF. The conclusion holds already for MIN-RAF restricted to the CRSs $\mathcal{Q}_{\max\{2, m\}}(\mathcal{I})$, all of which have a single food molecule, at most two reactants and exactly one product per reaction, no food catalysis, and reaction set equal to their maxRAF.*

Proof. Suppose $\mathcal{A}$ is a polynomial-time algorithm that, on every CRS containing a RAF, returns a RAF of size at most α times the minimum. Consider the following algorithm for SET COVER. Given an instance $\mathcal{I}$ with $m \geq 1$, construct $\mathcal{Q} = \mathcal{Q}_{\max\{2, m\}}(\mathcal{I})$, which contains a RAF by Remark 4.5; run $\mathcal{A}$ on $\mathcal{Q}$ to obtain a RAF S with $|S| \leq \alpha \text{OPT}_{\text{RAF}}(\mathcal{Q})$; and return $D = \{j : b_{j, M-1} \in S\}$. By Lemma 6.1 the whole procedure runs in polynomial time, and by Theorem 5.3 with $M = \max\{2, m\} \geq m$ and $c = \alpha$, the set D is a cover with $|D| \leq (2\alpha - 1)\tau(\mathcal{I})$. This is a polynomial-time factor- $(2\alpha - 1)$ solution approximation for SET COVER, contradicting Hypothesis 2.7. The restriction to the family $\mathcal{Q}_{\max\{2, m\}}(\mathcal{I})$ is built into the argument, and the listed syntactic properties are those noted after Construction 3.1 and in Remark 4.5. $\square$

Remark 6.3. Theorem 5.3 is stated for real $c \geq 1$. The Lean formalization proves it for natural numbers c in natural-number arithmetic, without division; for the main theorem this loses nothing, since a factor- α guarantee is a factor- $\lceil \alpha \rceil$ guarantee.

7 Estimating the minimum without producing a RAF

Theorem 6.2 concerns algorithms that return a RAF. A weaker task is to output a number that estimates OPT_{RAF} . Three conventions are in use: a *factor- c upper estimate* is a number V with $\text{OPT}_{\text{RAF}} \leq V \leq c \text{OPT}_{\text{RAF}}$; a *factor- c lower estimate* is V with $V \leq \text{OPT}_{\text{RAF}} \leq cV$; and a *two-sided factor- c estimate* is V with $\text{OPT}_{\text{RAF}} \leq cV$ and $V \leq c \text{OPT}_{\text{RAF}}$. The construction transfers gaps as well as ratios, so all three are hard. This uses the gap form of the SET COVER hardness theorems, which is what [26, 8, 6] actually prove.

Hypothesis 7.1 (Gap SET COVER). *If $P \neq NP$ then for every constant $\kappa > 1$ the following promise problem is not solvable in polynomial time: given a SET COVER instance $\mathcal{I}$ and an integer $a \geq 1$ with the promise that either $\tau(\mathcal{I}) \leq a$ or $\tau(\mathcal{I}) > \kappa a$, decide which.*

Lemma 7.2 (Gap arithmetic). *Let $m \geq 1$, $M \geq m$, $a \geq 1$ and $c \geq 1$, and let $\kappa > 2c - 1$. Then*

$$c(m + Ma) < m + M\kappa a.$$

If moreover $\kappa > 2c^2 - 1$ then $c^2(m + Ma) < m + M\kappa a$.

Proof. We have $c(m + Ma) - m = (c - 1)m + cMa \leq (c - 1)M + cMa \leq (c - 1)Ma + cMa = (2c - 1)Ma < \kappa Ma$, using $m \leq M$, $a \geq 1$ and $c - 1 \geq 0$. The second statement is the first with c^2 in place of c . $\square$

Proposition 7.3 (Value estimation is hard). *Assume $P \neq NP$ and Hypothesis 7.1. Then for no constant $c \geq 1$ is there a polynomial-time algorithm that, on every CRS containing a RAF, outputs a factor- c upper estimate of OPT_{RAF} ; nor one that outputs a factor- c lower estimate; nor one that outputs a two-sided factor- c estimate.*

Proof. Fix $c \geq 1$ and put $\kappa = 2c^2$, so that $\kappa > 2c - 1$ and $\kappa > 2c^2 - 1$. Given $(\mathcal{I}, a)$ as in Hypothesis 7.1, build $\mathcal{Q} = \mathcal{Q}_{\max\{2, m\}}(\mathcal{I})$ in polynomial time (Lemma 6.1) and put $M = \max\{2, m\} \geq m$. By Corollary 5.2, $\text{OPT}_{\text{RAF}}(\mathcal{Q}) = m + M\tau(\mathcal{I})$. In the YES case $\text{OPT}_{\text{RAF}}(\mathcal{Q}) \leq m + Ma$; in the NO case $\text{OPT}_{\text{RAF}}(\mathcal{Q}) > m + M\kappa a$. By Lemma 7.2,

$$c^2(m + Ma) < m + M\kappa a. \quad (*)$$

Let V be the output of a putative estimator on $\mathcal{Q}$.

Upper estimate. In the YES case $V \leq c\text{OPT}_{\text{RAF}} \leq c(m + Ma)$; in the NO case $V \geq \text{OPT}_{\text{RAF}} > m + M\kappa a$. By $(*)$ the two ranges are disjoint, so comparing V with $c(m + Ma)$ decides the promise problem.

Lower estimate. In the YES case $V \leq \text{OPT}_{\text{RAF}} \leq m + Ma$; in the NO case $V \geq \text{OPT}_{\text{RAF}}/c > (m + M\kappa a)/c > m + Ma$ by $(*)$. Compare V with $m + Ma$.

Two-sided estimate. In the YES case $V \leq c\text{OPT}_{\text{RAF}} \leq c(m + Ma)$; in the NO case $V \geq \text{OPT}_{\text{RAF}}/c > (m + M\kappa a)/c > c(m + Ma)$ by $(*)$. Compare V with $c(m + Ma)$.

In each case a polynomial-time estimator would solve the promise problem with $\kappa = 2c^2$, contradicting Hypothesis 7.1. $\square$

Remark 7.4. The proof shows that the one-sided estimates need only the linear gap $\kappa > 2c - 1$, while the two-sided estimate needs the quadratic gap $\kappa > 2c^2 - 1$; the Lean development records both inequalities and the resulting separation statements (Section 10). Hypothesis 7.1 is, like Hypothesis 2.7, a consequence of [26] or [6]: those results exhibit polynomial-time reductions from an NP-complete problem to gap instances with $\kappa = \Omega(\log m)$, which exceeds any constant on all but finitely many instances.

8 Irreducible RAFs

Irreducible RAFs, the inclusion-minimal RAFs of a CRS, are the objects studied in [28], where it is shown that their number can be exponential and that finding a smallest one is NP-hard. The construction of this paper describes them exactly.

Lemma 8.1. *Let $M \geq 1$ and $C, D \subseteq [n]$. Then $\mathcal{R}(D) \subseteq \mathcal{R}(C)$ if and only if $D \subseteq C$.*

Proof. If $D \subseteq C$ the inclusion is immediate from the definition of $\mathcal{R}$. Conversely if $j \in D$ then $b_{j, M-1} \in \mathcal{R}(D) \subseteq \mathcal{R}(C)$, so $j \in C$. $\square$

Proposition 8.2 (iRAFs are inclusion-minimal covers). *Let $m, M \geq 1$ and $C \subseteq [n]$. Then $\mathcal{R}(C)$ is an irreducible RAF of $\mathcal{Q}_M(\mathcal{I})$ if and only if C is an inclusion-minimal cover of $\mathcal{I}$, that is, a cover none of whose proper subsets is a cover. Every iRAF of $\mathcal{Q}_M(\mathcal{I})$ is of this form.*

Proof. By Theorem 4.4 every RAF is $\mathcal{R}(D)$ for a cover D , and by Lemma 8.1 the map $D \mapsto \mathcal{R}(D)$ is an order isomorphism from the covers onto the RAFs. Minimal elements correspond under an order isomorphism. $\square$

Corollary 8.3 (Minimal transversals). *Let $\mathcal{H}$ be a hypergraph with vertex set $[v]$ and nonempty edges $E_0, \dots, E_{e-1}$, $e \geq 1$. Let $\mathcal{I}_{\mathcal{H}}$ be the SET COVER instance with universe $[e]$ and sets $A_x = \{t : x \in E_t\}$ for $x \in [v]$. Then the iRAFs of $\mathcal{Q}_1(\mathcal{I}_{\mathcal{H}})$ are exactly the sets $\mathcal{R}(T)$ with T a minimal transversal of $\mathcal{H}$. In particular, enumerating the iRAFs of a CRS is at least as hard as enumerating the minimal transversals of a hypergraph.*

Proof. A set $T \subseteq [v]$ covers $\mathcal{I}_{\mathcal{H}}$ if and only if every edge meets T , that is, if and only if T is a transversal. Apply Proposition 8.2 with $M = 1$. The CRS $\mathcal{Q}_1(\mathcal{I}_{\mathcal{H}})$ has $e + v$ reactions and is computable in polynomial time from $\mathcal{H}$, and the bijection $T \mapsto \mathcal{R}(T)$ and its inverse are computable in linear time, so an enumeration algorithm for iRAFs yields one for minimal transversals with the same delay up to a polynomial factor. $\square$

Enumerating minimal transversals is the hypergraph dualization problem [7]. Whether it can be solved in time polynomial in the combined input and output size is a well-known open problem; the algorithm of Fredman and Khachiyan [9] runs in quasi-polynomial time $N^{o(\log N)}$ in that measure. Corollary 8.3 places iRAF enumeration at or above this level of difficulty.

Proposition 8.4 (No universal iRAF size ratio). *For every $c \geq 1$ there is a CRS with two irreducible RAFs S_1, S_2 such that $|S_2| > c|S_1|$. Explicitly, let $q \geq 2c$ be an integer, let $\mathcal{I}_q$ have universe $[q]$ and sets $A_0 = [q]$ and $A_{i+1} = \{i\}$ for $i \in [q]$, and let $M = q$. Then $\mathcal{R}(\{0\})$ and $\mathcal{R}(\{1, \dots, q\})$ are iRAFs of $\mathcal{Q}_q(\mathcal{I}_q)$ of sizes $2q$ and $q(q+1)$.*

Proof. Both $\{0\}$ and $\{1, \dots, q\}$ are covers, and each is inclusion-minimal: removing 0 from the first leaves nothing, and removing $i+1$ from the second uncovers i . By Proposition 8.2 both canonical sets are iRAFs, and by Lemma 5.1 their sizes are $q+q \cdot 1 = 2q$ and $q+q \cdot q = q(q+1)$. Finally $q(q+1) > c \cdot 2q$ if and only if $q+1 > 2c$, which holds for $q \geq 2c$. $\square$

Remark 8.5. Proposition 8.4 has an algorithmic reading. Several procedures return an iRAF: for instance, repeatedly deleting a reaction and recomputing the maxRAF until no deletion preserves the RAF property [12]. Any such procedure, however it chooses deletion order, may return either of the two iRAFs in Proposition 8.4, so producing an iRAF is not by itself an approximation guarantee. Theorem 6.2 shows that the failure is not an artefact of the deletion heuristic but is inherent in the problem.

9 Relation to the earlier bounds

Steel, Hordijk and Smith [28] proved that MIN-RAF is NP-hard and, in the absence of an approximation algorithm, gave two polynomial-time lower bounds on OPT_{RAF} . The *mandatory-reaction bound* counts the reactions that lie in every RAF, that is, the reactions whose deletion destroys every RAF. The *cycle bound* applies when no reaction of the RAF under consideration is catalysed by a food molecule, and bounds its size below by the length of a shortest directed cycle in a graph that records which reactions produce catalysts of which other reactions. They also gave, for fixed k , a polynomial-time test of whether a given list of k iRAFs exhausts all iRAFs, with an exponent that depends on k . These are certification and lower-bound results, not worst-case multiplicative guarantees, and in their concluding remarks the authors asked explicitly whether a constant-factor approximation of the smallest RAF is possible [28]. Theorem 6.2 answers that question negatively.

The construction also shows concretely why the two lower bounds cannot supply a ratio. In $\mathcal{Q}_M(\mathcal{I})$ the RAFs are the sets $\mathcal{R}(D)$ with D a cover, so by Theorem 4.4 a gate is mandatory always, and a block reaction $b_{j,k}$ is mandatory if and only if j lies in every cover, that is, if and only if some element of $[m]$ lies in A_j and in no other set. The mandatory-reaction bound therefore equals $m + M|\{j : A_j \text{ contains a privately covered element}\}|$, which equals m on every instance without privately covered elements, while $\text{OPT}_{\text{RAF}} = m + M\tau(\mathcal{I})$; for $M = m$ the ratio between the optimum and the bound is then $1 + \tau(\mathcal{I})$, which is unbounded. The cycle bound is applicable, since no reaction of $\mathcal{Q}_M(\mathcal{I})$ is food-catalysed, but each hub h_j is produced by the reaction $b_{j,M-1}$ that it catalyses, so the catalyst graph has a loop at every block and the bound is trivial.

10 Formal verification and computation

10.1 Lean formalization

The definitions and results of Sections 2–8 have been formalized in Lean 4 (toolchain v4.30.0) against a pinned version of Mathlib (tag v4.30.0, commit c5ea003). The formalization follows

the closure semantics of Definitions 2.2 and 2.3 verbatim: a CRS is a structure with fields `inputs`, `outputs` and `food`, the closure $W_i(S)$ is `closureAt Q S i`, defined by recursion on i , and `FoodGenerated`, `ReflexivelyAutocatalytic` and `IsRAF` are the predicates of Definition 2.3. A SET COVER instance is a structure `SetCoverInstance (Fin m) (Fin n)` consisting of the map $j \mapsto A_j$ and a proof of the coverability promise. Reactions of the source are the inductive type with constructors `gate i` and `block j k`, molecules are `food`, `y i` and `z j k`, and Construction 3.1 is the pair `crs m n M`, `catalysis I`. The canonical set $\mathcal{R}(D)$ is `canonicalRAF D`. Theorem 4.4 is, in ASCII rendering,

```
theorem MinRAFAprox.SetCoverSource.setCoverSource_raf_iff
  (I : SetCoverInstance (Fin m) (Fin n)) (hm : 0 < m) (hM : 0 < M)
  (S : Finset (Reaction (Fin m) (Fin n) (Fin M))) :
  IsRAF (crs m n M) (catalysis I) S <->
  exists C : Finset (Fin n), I.Covers C /\ S = canonicalRAF C
```

and Lemma 5.1 is `canonicalRAF_card`, which states

```
(canonicalRAF C).card = m + M * C.card
```

and is proved through an explicit equivalence with the type $[m] \sqcup (C \times [M])$ rather than by enumeration. Corollary 5.2 is `minimumCover_iff_minimumRAF`. Theorem 5.3 is `approximation_decode`, whose arithmetic core `approximation_decode_arithmetic` is the natural-number inequality

$$1 \leq c, 1 \leq \tau, 0 < M, m \leq M, m + Mk \leq c(m + M\tau) \implies k \leq (2c - 1)\tau.$$

The complexity-theoretic conclusion is formalized relative to abstract predicates. A `CoverAlgorithm` is a function from instances to index sets, a `SourceRAFAAlgorithm` is a function from instances and block lengths to reaction sets, `IsCoverApprox c A` and `IsSourceRAFAApprox c A` are the factor- c solution-approximation properties of Definition 2.6 on the respective families, and `inducedCover A` composes a source RAF algorithm with the decoder $S \mapsto \{j : b_{j,M-1} \in S\}$ at $M = \max\{2, m\}$. The theorem `sourceRAF_approx_implies_cover_approx` states that `IsSourceRAFAApprox c A` implies `IsCoverApprox (2*c-1) (inducedCover A)`, and the publication entry point is

```
theorem MinRAFAprox.SetCoverSource.publication_constant_factor_resolution
  (P_eq_NP : Prop)
  (PolynomialCover : CoverAlgorithm -> Prop)
  (PolynomialRAF : SourceRAFAAlgorithm -> Prop)
  (setCoverHardness :
    Not P_eq_NP -> NoPolynomialCoverConstApprox PolynomialCover)
  (reductionPolynomial :
    forall A, PolynomialRAF A -> PolynomialCover (inducedCover A)) :
  Not P_eq_NP -> NoPolynomialSourceRAFConstApprox PolynomialRAF
```

Here `NoPolynomialCoverConstApprox P` says that no algorithm satisfying P is a factor- c solution approximation for any $c \geq 1$, and similarly for RAFs. The premise `setCoverHardness` is Hypothesis 2.7; the premise `reductionPolynomial` is the closure of the polynomial-time predicate under the construction and decoder, which Lemma 6.1 establishes for the ordinary encoding. Polynomial time is thus treated as a parameter rather than through a formalized machine model, which is the honest scope of the formal result. The formal statement is restricted to the source family, as Theorem 6.2 permits.

The remaining results are formalized as the following declarations.

- Lemma 6.1: `source_reaction_count` (the reaction count), `activeMolecules_card` (the molecule count $1 + m + nM$, together with a proof that every closure stage stays inside this finite carrier), and `sourceEncodingSize_le` and `amplified_sourceEncodingSize_le` (the encoding bounds).

- The syntactic restrictions of Theorem 6.2: `source_input_card_le_two`, `source_output_card_eq_one` and `source_food_card_eq_one`.
- Lemma 7.2: `linear_gap_condition` and `squared_gap_condition`. The three separation arguments in the proof of Proposition 7.3: `no_common_lower_estimate`, `no_common_upper_estimate` and `no_common_symmetric_estimate`.
- Proposition 8.2: `irreducibleRAF_iff_inclusionMinimalCover`. Corollary 8.3: `irreducibleRAF_iff_minimalTransversal`. Proposition 8.4: `irreducibleRAF_size_ratio_unbounded`, with $q = 2c + 2$.

Table 2 lists the modules. Not formalized are Hypotheses 2.7 and 7.1 themselves, the polynomial-time model, and the equivalence with the ordering-based definitions of [13] discussed after Lemma 2.4.

Lean module	Content
<code>RAF/Core/CRS.lean</code> , <code>RAF/Core/Closure.lean</code> , <code>RAF/Core/RAF.lean</code>	CRS, closure, food generation, RAFs (Section 2.1)
<code>RAF/Frankl/Antimatroid.lean</code>	Lemma 2.4
<code>Basic.lean</code>	SET COVER instances, covers, reaction labels
<code>Gadgets/SetCoverSource.lean</code>	Construction 3.1, Lemma 3.2
<code>Gadgets/BlockCoverNormalForm.lean</code>	abstract block-cover normal form and its equivalence with covers
<code>Gadgets/SetCoverProvenance.lean</code>	Lemmas 4.1 and 4.2
<code>Gadgets/SetCoverReachability.lean</code>	Lemma 4.3
<code>Gadgets/SetCoverBijection.lean</code>	Theorem 4.4
<code>Gadgets/SetCoverCardinality.lean</code>	Lemma 5.1
<code>Gadgets/FiniteSupport.lean</code>	finite molecule carrier, molecule count (Lemma 6.1)
<code>ApproximationTransfer.lean</code>	Corollary 5.2, Theorem 5.3
<code>ComplexityTransfer.lean</code>	reaction count, decoder, AP transfer, conditional hardness (Theorem 6.2)
<code>EncodingSize.lean</code>	encoding bounds (Lemma 6.1)
<code>NumericalEstimates.lean</code>	Lemma 7.2, separations for Proposition 7.3
<code>Irreducible.lean</code> ,	Propositions 8.2 and 8.4, Corollary 8.3
<code>UnboundedIrreducible.lean</code>	
<code>Resolution.lean</code> ,	syntactic restrictions; terminal and aggregate statements
<code>PublicationResolution.lean</code>	

Table 2: Correspondence between the paper and the Lean development. Modules are in the directory `proofs/MinRAFApprox` unless a path is indicated, in which case they are relative to `proofs/`.

The aggregate module `PublicationResolution.lean`, which imports every module in Table 2, was compiled with the option `-DwarningAsError=true`, which promotes warnings to errors. The build reports exit code 0, empty standard output and standard error, and no use of `sorry`. The axioms on which `publication_constant_factor_resolution` depends are exactly `propext`, `Classical.choice` and `Quot.sound`, the standard axioms of Lean’s classical logic. The verified source of the aggregate module has SHA-256 hash

2c9960b72523aa7f8d589edd945a8bf1ad721fcdc455b3aa672ba78e6dcc5060,

and the hash of its elaborated declaration interface is

de349b3a9325bffff505978f9c8e0cc8a981a9ef5619b18958b32d7691bc01572.

The terminal module `Resolution.lean`, containing the three syntactic-restriction theorems and the resolution theorem, was verified separately in the same way with the same axiom set; its source hash is 2a812b14...b1413 and its interface hash is 104f88bc...8730, recorded in full in the repository receipts.

10.2 Exhaustive finite checks

Before the formalization, Construction 3.1 was tested exhaustively by a short program that implements Definitions 2.2 and 2.3 literally. For every $m, n \in \{1, 2, 3\}$, every ordered family of n subsets of $[m]$ whose union is $[m]$ (empty and repeated sets allowed), and every $M \in \{1, 2, 3\}$, the program built $\mathcal{Q}_M(\mathcal{I})$, and for every one of the 2^{m+nM} reaction subsets S it computed the closure $\text{cl}_S(\{f\})$ by iteration, tested the RAF conditions, and compared the outcome with the normal form predicted by Theorem 4.4, namely “ $S = \mathcal{R}(D)$ for a cover D ”. It also compared the minimum RAF size found by enumeration with $m + M\tau(\mathcal{I})$. This covered 1,323 systems and 1,747,004 reaction subsets, of which 4,485 were RAFs, with no discrepancy in either comparison. An earlier variant of the construction, in which the block catalysts formed a directed cycle rather than pointing to a single hub, passed the same test on 882 systems and 1,722,312 subsets before being simplified to the present form. These computations informed the proofs and are reported for reproducibility; they are not used as evidence in any proof.

11 Discussion

The smallest RAF is not merely hard to compute exactly. In unrestricted finite catalytic reaction systems it cannot, unless $\text{P} = \text{NP}$, be approximated in polynomial time within any fixed constant factor, and the obstruction already appears in systems with a single food molecule, at most two reactants and one product per reaction, and no food catalysis. The mechanism is an exact structural correspondence: in the constructed systems, RAFs *are* set covers, and reaction counts are an affine image of cover sizes in which the variable part has been amplified past the fixed scaffold. Several questions remain.

Question 11.1. What is the approximation threshold of MIN-RAF? Theorem 5.3 transfers a factor- c guarantee for MIN-RAF to a factor- $(2c - 1)$ guarantee for SET COVER, so any logarithmic hardness for SET COVER yields logarithmic hardness for MIN-RAF measured in the universe size m ; the constant obtained in terms of the number of reactions depends on the relation between m and n in the hardness instances and we have not pursued it. On the positive side we know of no polynomial-time $O(\log |R|)$ -approximation for MIN-RAF, nor of any nontrivial approximation at all. Is MIN-RAF in Log-APX, or is it harder, for instance polylogarithmically hard like label cover?

Question 11.2. Which structured classes of CRSs admit constant-factor approximation of the smallest RAF? The binary polymer model of [19, 13, 24, 11] is the natural test case. The systems $\mathcal{Q}_M(\mathcal{I})$ are far from that model: their reaction graph is a tree of chains and their catalysis is concentrated on n hub molecules. It is open whether the hardness persists under the arity and catalysis-density constraints of the polymer model, or whether those constraints admit good approximations.

Question 11.3. Is there a polynomial-time algorithm that computes a factor- c lower estimate of OPT_{RAF} for structured classes, or that certifies optimality of a given RAF in interesting cases? The lower bounds of [28] are of this kind; Section 9 shows exactly where they lose on the present family, and Proposition 7.3 shows that no universal constant-factor estimator exists.

Question 11.4. Corollary 8.3 shows that enumerating iRAFs is at least as hard as hypergraph dualization. Is it equivalent to it? An output-polynomial or quasi-polynomial enumeration algorithm for iRAFs would be of practical interest, since iRAFs are the units of analysis in much of the applied RAF literature [16, 12, 30].

Finally, the results concern the combinatorial RAF model only. They say nothing about the kinetic or thermodynamic feasibility of the constructed systems, nor about variants of RAF theory with inhibition or with the interior-operator viewpoint of [27]; in the latter setting the

family of RAFs of $\mathcal{Q}_M(\mathcal{I})$ is the union-closed family of covers of $\mathcal{I}$ together with the empty set, which may be of independent interest.

Availability

The Lean sources, the verification receipts recording the hashes and axioms quoted in Section 10, and the enumeration program are available in the repository <https://github.com/mmislán/Erdos>, in the directories `proofs/MinRAFApprox` and `problem_workspaces/RAF_minRAF_const_factor_approx`.

References

- [1] Noga Alon, Dana Moshkovitz, and Shmuel Safra. Algorithmic construction of sets for k -restrictions. *ACM Transactions on Algorithms*, 2(2):153–177, 2006.
- [2] Giorgio Ausiello, Pierluigi Crescenzi, Giorgio Gambosi, Viggo Kann, Alberto Marchetti-Spaccamela, and Marco Protasi. *Complexity and Approximation: Combinatorial Optimization Problems and Their Approximability Properties*. Springer, Berlin, 1999.
- [3] Vašek Chvátal. A greedy heuristic for the set-covering problem. *Mathematics of Operations Research*, 4(3):233–235, 1979.
- [4] Pierluigi Crescenzi. A short guide to approximation preserving reductions. In *Proceedings of the 12th Annual IEEE Conference on Computational Complexity*, pages 262–273. IEEE, 1997.
- [5] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer, 2021.
- [6] Irit Dinur and David Steurer. Analytical approach to parallel repetition. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC 2014)*, pages 624–633. ACM, 2014.
- [7] Thomas Eiter and Georg Gottlob. Identifying the minimal transversals of a hypergraph and related problems. *SIAM Journal on Computing*, 24(6):1278–1304, 1995.
- [8] Uriel Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.
- [9] Michael L. Fredman and Leonid Khachiyan. On the complexity of dualization of monotone disjunctive normal forms. *Journal of Algorithms*, 21(3):618–628, 1996.
- [10] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, San Francisco, 1979.
- [11] Wim Hordijk, Stuart A. Kauffman, and Mike Steel. Required levels of catalysis for emergence of autocatalytic sets in models of chemical reaction systems. *International Journal of Molecular Sciences*, 12(5):3085–3101, 2011.
- [12] Wim Hordijk, Joshua I. Smith, and Mike Steel. Algorithms for detecting and analysing autocatalytic sets. *Algorithms for Molecular Biology*, 10:15, 2015.
- [13] Wim Hordijk and Mike Steel. Detecting autocatalytic, self-sustaining sets in chemical reaction systems. *Journal of Theoretical Biology*, 227(4):451–461, 2004.

- [14] Wim Hordijk and Mike Steel. Chasing the tail: The emergence of autocatalytic networks. *Biosystems*, 152:1–10, 2017.
- [15] Wim Hordijk and Mike Steel. Autocatalytic networks at the basis of life’s origin and organization. *Life*, 8(4):62, 2018.
- [16] Wim Hordijk, Mike Steel, and Stuart Kauffman. The structure of autocatalytic sets: Evolvability, enablement, and emergence. *Acta Biotheoretica*, 60(4):379–392, 2012.
- [17] David S. Johnson. Approximation algorithms for combinatorial problems. *Journal of Computer and System Sciences*, 9(3):256–278, 1974.
- [18] Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller and James W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, New York, 1972.
- [19] Stuart A. Kauffman. Autocatalytic sets of proteins. *Journal of Theoretical Biology*, 119(1):1–24, 1986.
- [20] Stuart A. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, New York, 1993.
- [21] László Lovász. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13(4):383–390, 1975.
- [22] Carsten Lund and Mihalis Yannakakis. On the hardness of approximating minimization problems. *Journal of the ACM*, 41(5):960–981, 1994.
- [23] Dana Moshkovitz. The projection games conjecture and the NP-hardness of $\ln n$ -approximating set-cover. *Theory of Computing*, 11(7):221–235, 2015.
- [24] Elchanan Mossel and Mike Steel. Random biochemical networks: the probability of self-sustaining autocatalysis. *Journal of Theoretical Biology*, 233(3):327–336, 2005.
- [25] Christos H. Papadimitriou and Mihalis Yannakakis. Optimization, approximation, and complexity classes. *Journal of Computer and System Sciences*, 43(3):425–440, 1991.
- [26] Ran Raz and Shmuel Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing (STOC 1997)*, pages 475–484. ACM, 1997.
- [27] Mike Steel. Interior operators and their relationship to autocatalytic networks. *Acta Biotheoretica*, 71(4):21, 2023.
- [28] Mike Steel, Wim Hordijk, and Joshua Smith. Minimal autocatalytic networks. *Journal of Theoretical Biology*, 332:96–107, 2013.
- [29] Mike Steel, Wim Hordijk, and Joana C. Xavier. Autocatalytic networks in biology: structural theory and algorithms. *Journal of the Royal Society Interface*, 16(151):20180808, 2019.
- [30] Mike Steel, Joana C. Xavier, and Daniel H. Huson. The structure of autocatalytic networks, with application to early biochemistry. *Journal of the Royal Society Interface*, 17(171):20200488, 2020.

- [31] The mathlib Community. The Lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020)*, pages 367–381. ACM, 2020.
- [32] Joana C. Xavier, Wim Hordijk, Stuart Kauffman, Mike Steel, and William F. Martin. Autocatalytic chemical networks at the origin of metabolism. *Proceedings of the Royal Society B*, 287(1922):20192377, 2020.