

Optimal Information Encoding in Chemical Reaction Networks

Austin Luchsinger   

Electrical and Computer Engineering, University of Texas at Austin, TX, USA

David Doty   

Computer Science, University of California–Davis, CA, USA

David Soloveichik   

Electrical and Computer Engineering, University of Texas at Austin, TX, USA

Abstract

Discrete chemical reaction networks formalize the interactions of molecular species in a well-mixed solution as stochastic events. Given their basic mathematical and physical role, the computational power of chemical reaction networks has been widely studied in the molecular programming and distributed computing communities. While for Turing-universal systems there is a universal measure of optimal information encoding based on Kolmogorov complexity, chemical reaction networks are not Turing universal unless error and unbounded molecular counts are permitted. Nonetheless, here we show that the optimal number of reactions to generate a specific count $x \in \mathbb{N}$ with probability 1 is asymptotically equal to a “space-aware” version of the Kolmogorov complexity of x , defined as $\tilde{K}_s(x) = \min_p \{ |p|/\log|p| + \log(\text{space}(\mathcal{U}(p))) : \mathcal{U}(p) = x \}$, where p is a program for universal Turing machine \mathcal{U} . This version of Kolmogorov complexity incorporates not just the length of the shortest program for generating x , but also the space usage of that program. Probability 1 computation is captured by the standard notion of stable computation from distributed computing, but we limit our consideration to chemical reaction networks obeying a stronger constraint: they “know when they are done” in the sense that they produce a special species to indicate completion. As part of our results, we develop a module for encoding and unpacking any b bits of information via $O(b/\log b)$ reactions, which is information-theoretically optimal for incompressible information. Our work provides one answer to the question of how succinctly chemical self-organization can be encoded – in the sense of generating precise molecular counts of species as the desired state.

2012 ACM Subject Classification Theory of computation → Models of computation

Keywords and phrases chemical reaction networks, Kolmogorov complexity, stable computation

Digital Object Identifier 10.4230/LIPIcs.DNA.29.9

Funding *Austin Luchsinger*: Carroll H. Dunn Endowed Graduate Fellowship in Engineering.

David Doty: NSF grants 2211793, 1900931, and CAREER-1844976.

David Soloveichik: NSF grant 1901025, Sloan Foundation Research Fellowship.

Acknowledgements The authors thank Eric Severson for the invaluable discussions on chemical reaction network computation that helped lead to the results presented here.

1 Introduction

In potential biochemical, nanotechnological, or medical applications, synthetic chemical computation could allow for the re-programming of biological regulatory networks and the insertion of control modules where traditional electronic controllers are not feasible. Understanding the design principles of chemical information processing also may achieve better understanding of the complex information processing that occurs in biological chemical interactions.



© Austin Luchsinger, David Doty, and David Soloveichik;

licensed under Creative Commons License CC-BY 4.0

29th International Conference on DNA Computing and Molecular Programming (DNA 29).

Editors: Ho-Lin Chen and Constantine G. Evans; Article No. 9; pp. 9:1–9:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Discrete chemical reaction networks, also called stochastic chemical reaction networks, is a formal model of chemical kinetics in a well-mixed solution. While in continuous chemical kinetics, continuous concentrations change in time governed by ordinary differential equations, here the state consists of non-negative integer molecular counts of the species, and reaction events occur stochastically as a continuous time Markov process. Closely related models include population protocols in distributed computing [4], as well as models without stochastic kinetics such as Petri nets [20], vector addition systems [22] and commutative semigroups [11]. The model is particularly relevant when some species are present in small molecular counts, which are not well-approximated by continuous concentrations [21]; this regime is germane for small volumes such as that of a cell, natural or artificial. For the rest of this paper, the acronym CRNs (Chemical Reaction Networks) refers to the discrete model.

Typically the ensuing sequence of reactions can be predicted only stochastically since multiple reactions compete with each other. Nonetheless certain behaviors are independent of the order in which reactions happen to occur. Such probability 1 behavior is formalized using the notion of stable computation. For example the reactions $X_1 \rightarrow 2Y$ and $X_2 + Y \rightarrow \emptyset$ compute the function $f(x_1, x_2) = \max(2x_1 - x_2, 0)$ regardless of the order in which reactions happen. Below when we say that a CRN computes something, we mean it in the sense of stable computation. It is known that stably computing CRNs are not Turing-universal [36], but instead are limited to computing semilinear predicates and functions [5, 13]. However, the scaling of the computational power of CRNs with the number of reactions and species still lacks a tight and general characterization.

Prior approaches to answering the question of reaction or species complexity – in the equivalent language of population protocols – have focused largely on predicate computation and can be divided into two groups. (We should point out that the literature makes the important distinction between population protocols with and without a “leader,” which is equivalent to starting with a single copy of a distinguished species in the initial state. The prior results described here as well as our work correspond to protocols *with* a leader.) The first line of work focuses on specific predicates – with the prototypical choice being the so-called “counting predicates” in which the task is to decide whether the count of the input species is at least some threshold $x \in \mathbb{N}$ [8, 17, 26]. In particular, close upper and lower bounds were developed: for infinitely many x , the predicate can be stably decided with $\mathcal{O}(\log \log x)$ species [8], and $\mathcal{O}((\log \log x)^{1/2-\epsilon})$ species are required [26].

Other work has focused on the more general characterization of predicate computation. It is well-known that semilinear predicates can be characterized in terms of Presburger arithmetic, the first-order theory of addition. It was subsequently shown that a CRN can decide a semilinear predicate with the number of species scaling polynomially with the size of the corresponding Presburger formula [7, 18]. There are also provable tradeoffs between the speed of computation and the number of species (e.g., [2, 6, 19]). We do not consider the time-complexity of CRNs further in this paper.

While the prior work described above involves stably deciding a counting predicate where the system recognizes if the count of some species is at least x , we investigate the problem of generating exactly x copies of a particular species Y , starting from a single copy of another species L . This idea of generation is natural for engineers of these systems who may wish to prepare a particular configuration to be used in a downstream process, and captures a certain form of chemical self-organization. (We note the conceptual connection to another type of self-organization: leader-election, in which we want to end up with exactly one molecule of a species, starting from many [6].) Our constructions can be adopted to deciding the counting predicates with only a constant more reactions – giving a novel upper bound on the number

of reactions (see Open Questions). It is also worth noting that other complexity questions have been investigated for CRNs, such as “the size of the smallest chemical reaction network that approximates a desired distribution” [10].

The goal of this paper is to connect the complexity of the most compact CRN for generating x to the well-known measures of the optimal “description length” of x . Kolmogorov complexity, a widely recognized concept across various disciplines in computer science and information theory, serves as a universal, broadly accepted measure of description length [27]. This notion quantifies the complexity of an object, such as a string or a number, by the length of the shortest program that produces it. While the minimal number of species or reactions to generate count x cannot be connected to the canonical Kolmogorov complexity, we provide tight asymptotic bounds to a modification of Kolmogorov complexity \tilde{K}_s (Equation (1)). As this quantity incorporates not only the length of the shortest program to produce x , but also the space (memory) usage of the program, it can be called “space-aware.” Unlike the canonical Kolmogorov complexity, \tilde{K}_s is computable.

Our quantity \tilde{K}_s characterizes the CRN complexity of generating x in the range from $\mathcal{O}(\log \log x)$ for highly “compressible” x to $\mathcal{O}(\log x / \log \log x)$ for “incompressible” x . The module we develop for optimally encoding b bits of information with $\mathcal{O}(b / \log b)$ reactions via a permutation code may be of independent interest. The encoded information could be used for other purposes than for generating a desired amount of some species, which justifies a more general interpretation of our work as studying the encoding information in CRNs.

2 Preliminaries

We use notation from [12, 35] and stable computation definitions from [5, 14] for (discrete) chemical reaction networks. Let \mathbb{N} denote the nonnegative integers. For any finite set \mathcal{S} (of species), we write $\mathbb{N}^{\mathcal{S}}$ to mean the set of functions $f : \mathcal{S} \rightarrow \mathbb{N}$. Equivalently, $\mathbb{N}^{\mathcal{S}}$ can be interpreted as the set of vectors indexed by the elements of \mathcal{S} , and so $\mathbf{c} \in \mathbb{N}^{\mathcal{S}}$ specifies nonnegative integer counts for all elements of \mathcal{S} . For $\mathbf{a}, \mathbf{b} \in \mathbb{N}^{\mathcal{S}}$, we write $\mathbf{a} \leq \mathbf{b}$ if $\mathbf{a}(i) \leq \mathbf{b}(i), \forall i$.

2.1 Chemical Reaction Networks

A *chemical reaction network* (CRN) $\mathcal{C} = (\mathcal{S}, \mathcal{R})$ is defined by a finite set \mathcal{S} of species, and a finite set \mathcal{R} of reactions where each reaction is a pair $\langle \mathbf{r}, \mathbf{p} \rangle \in \mathbb{N}^{\mathcal{S}} \times \mathbb{N}^{\mathcal{S}}$ that denotes the *reactant* species consumed by the reaction and the *product* species generated by the reaction. For example, given $\mathcal{S} = \{A, B, C\}$, the reaction $\langle (2, 0, 0), (0, 1, 1) \rangle$ represents $2A \rightarrow B + C$. Although the definition allows for more general stoichiometry, in this paper we only consider third-order reactions (with at most three reactants and three products). For reversible reactions, we will use the notation $A + B \rightleftharpoons C + D$ to mean $A + B \rightarrow C + D$ and $C + D \rightarrow A + B$. We say that the *size* of a CRN (denoted $|\mathcal{C}|$) is simply the number of reactions in \mathcal{R} .¹

A *configuration* $\mathbf{c} \in \mathbb{N}^{\mathcal{S}}$ of a CRN assigns integer counts to every species $s \in \mathcal{S}$. When convenient, we use the notation $\{n_1 S_1, n_2 S_2, \dots, n_k S_k\}$ to describe a configuration with $n_i \in \mathbb{N}$ copies of species $S_i, \forall i \in [1, k]$. When using this notation, any species $S_j \in \mathcal{S}$ that is not listed is assumed to have a zero count (e.g., given $\mathcal{S} = \{A, B, C\}$, the configuration $\{3A, 2B\}$ has three copies of species A , two copies of species B , and zero of species C). For two configurations $\mathbf{a}, \mathbf{b} \in \mathbb{N}^{\mathcal{S}}$, we say \mathbf{b} *covers* \mathbf{a} if $\mathbf{a} \leq \mathbf{b}$; in other words, for all species, \mathbf{b} has at least as many copies as \mathbf{a} .

¹ When considering systems with third-order reactions it is clear that $|\mathcal{R}|^{1/6} \leq |\mathcal{S}| \leq 6|\mathcal{R}|$.

A reaction $\langle \mathbf{r}, \mathbf{p} \rangle$ is said to be *applicable* in configuration \mathbf{c} if $\mathbf{r} \leq \mathbf{c}$. If the reaction $\langle \mathbf{r}, \mathbf{p} \rangle$ is applicable, it results in configuration $\mathbf{c}' = \mathbf{c} - \mathbf{r} + \mathbf{p}$ if it occurs, and we write $\mathbf{c} \rightarrow \mathbf{c}'$. If there exists a finite sequence of configurations such that $\mathbf{c} \rightarrow \mathbf{c}_1 \rightarrow \dots \rightarrow \mathbf{c}_n \rightarrow \mathbf{d}$, then we say that \mathbf{d} is *reachable* from \mathbf{c} and we write $\mathbf{c} \rightsquigarrow \mathbf{d}$.

In keeping with established definitions for stable computation, we specify an *output species* $Y \in \mathcal{S}$ and a *leader species* $L \in \mathcal{S}$ for stable integer computation.² We start from an initial configuration $\mathbf{i} = \{1L\}$. A configuration \mathbf{c} is *output-stable* if $\forall \mathbf{d}$ such that $\mathbf{c} \rightsquigarrow \mathbf{d}$, $\mathbf{c}(Y) = \mathbf{d}(Y)$. CRN \mathcal{C} *stably computes* integer x if, from any configuration \mathbf{c} that is reachable from input configuration \mathbf{i} , there is an output-stable configuration \mathbf{o} reachable from \mathbf{c} with $\mathbf{o}(Y) = x$. Note that when considering systems with bounded state spaces like those discussed in this paper, stable computation is equivalent to probability 1 computing.

We also consider a much stronger constraint on CRN computation that specifies a special halting species. A species $H \in \mathcal{S}$ is a *halting species* if $\forall \mathbf{c}$ such that $\mathbf{c}(H) \geq 1$, \mathbf{c} is output stable and $\forall \mathbf{d}$ where $\mathbf{c} \rightsquigarrow \mathbf{d}$, $\mathbf{d}(H) \geq 1$. We say that a CRN \mathcal{C} *haltingly computes* an integer x if (1) \mathcal{C} stably computes x and (2) \mathcal{C} has a halting species H . Intuitively, a halting CRN knows when it is done – the halting species can initiate some downstream process that is only meant to occur when the computation is finished.

2.2 Kolmogorov Complexity

A focus of this paper is the “optimal description” of integers. As such, we often refer to the traditional notion of Kolmogorov complexity which we define here.

Let \mathcal{U} be a universal Turing machine. The Kolmogorov complexity for an integer x is the value $K(x) = \min\{|p| : \mathcal{U}(p) = x\}$. In other words, the Kolmogorov complexity of x is the size of the smallest Turing machine program p that outputs x . This captures the descriptonal complexity of x in the sense that a (smaller) description of x can be given to some machine that generates x based on the given description.

We use a “space-aware” variant of this quantity which we later connect to the size of the smallest CRN stably computing x :

$$\tilde{K}s(x) = \min \left\{ \frac{|p|}{\log|p|} + \log(\text{space}(\mathcal{U}(p))) : \mathcal{U}(p) = x \right\}. \quad (1)$$

Note that $\tilde{K}s(x)$ does not refer to CRNs in any direct way, so the tight asymptotic connection (Theorem 12) we establish may be surprising.

$\tilde{K}s(x)$ is similar to the Kolmogorov complexity variant defined as $Ks(x) = \min\{|p| + \log(\text{space}(\mathcal{U}(p))) : \mathcal{U}(p, i) = x[i]\}$ by Allender, Koucký, Ronneburger, and Roy [3] in that it additively mixes program size with the log of the space usage. There are two differences: (1) The program size component of $\tilde{K}s$ is $|p|/\log|p|$ rather than $|p|$. The intuition is that a single chemical reaction can encode more than one bit of information; thus, a Turing machine program p can be converted to a “CRN program” with a number of reactions that is asymptotically smaller than the number of bits of p . (2) $Ks(x)$ is defined with respect to programs that, given index i as input, output $x[i]$, the i 'th bit of x , while our $\tilde{K}s(x)$ is defined with respect to programs that (taking no input) directly output all of x . Thus $\tilde{K}s(x) \geq \log|x|$, since the Turing machine must at least store the output integer, while $Ks(x)$

² For stable *function* computation, an ordered subset of input species $\{X_1, X_2, \dots, X_n\} \subset \mathcal{S}$ is also included; however, stable integer computation would be something along the lines of $f(1) = x$, so a single copy of the leader species serves as the “input” here.

may be smaller in principle. Due to the ability of efficient universal Turing machines to simulate each other efficiently, $\tilde{K}s$ (like Ks) is invariant within multiplicative constants to the choice of universal Turing machine \mathcal{U} , as long as \mathcal{U} is space-efficient. Note that if $\tilde{K}s$ were not robust to the choice of \mathcal{U} , it could hardly be a universal measure.

It is worth noting that unlike $K(x)$, $\tilde{K}s(x)$ is computable. To see this, one can enumerate all programs for universal Turing machine \mathcal{U} and run them in order from smallest to largest, stopping on the first machine that outputs x . Since the space usage of $\mathcal{U}(p)$ is included in $\tilde{K}s$, we can terminate executions as soon as they start using too much space. This ensures that no execution will run forever, and so we are guaranteed to find the smallest p that outputs x .

2.3 Overview

Here, we give a high level overview for the constructions and results presented in the subsequent sections of this paper.

Our constructions rely on the ability of CRNs to “efficiently” simulate space-bounded Turing machines (in terms of program size and space usage, not time) by “efficiently” simulating bounded-count register machines. Section 3 details how to use a combination of previous results to achieve this. The first half of the section describes how to construct a CRN to faithfully simulate a bounded-count register machine. The second half of the section shows how to generate a large register machine bound (2^{2^n}) with very few species/reactions (n). While the latter result is from previous work [11], we translate their construction from a commutative semigroup presentation into a chemical reaction network.

In Section 4, we present a method for constructing a CRN \mathcal{C}_x which (optimally) haltingly computes n -bit integer x with $|\mathcal{C}_x| = \mathcal{O}(n/\log n)$ by using a permutation code (Theorem 6). The idea of the construction is to generate a specified permutation and convert that permutation to a mapped target integer x . This construction relies on the “efficient” bounded-count register machine and space-bounded Turing machine simulations.

We then show how to use our permutation construction to achieve an optimal encoding (within global multiplicative constants) for algorithmically compressible integers in Section 5. Here, we use our permutation code technique to “unpack” a Turing machine program that outputs x , resulting in a CRN that haltingly computes x with $\mathcal{O}(\tilde{K}s)$ reactions (Theorem 7). Afterwards, we use a result from Künnemann et al. [24] to show that the size of our constructed CRN is within multiplicative constants of the optimal size of a CRN that stably computes x , denoted $K_{\text{crn}}(x)$ (Lemma 11). The results of the paper culminate with us connecting $K_{\text{crn}}(x)$ and $\tilde{K}s$ in Theorem 12 (our main theorem), which is directly implied by the combination of Theorem 7 and Lemma 11.

Lastly, we present some open questions for future work in Section 6.

3 Efficient Simulation of Bounded Register Machines

3.1 Register machines


A register machine is a finite state machine along with a fixed number of registers, each with non-negative integer counts. The two fundamental instructions for a register machine are increment $inc(r_i, s_j)$ and decrement $dec(r_i, s_j, s_k)$. The first instruction increments register r_i and transitions the machine to state s_j . The second instruction decrements register r_i if it is non-zero and transitions the machine to state s_j , otherwise the machine just transitions to state s_k . We also consider the more advanced instruction of $copy(r_i, r_j, s_k)$, which adds the value of register r_i to register r_j , i.e., it is equivalent to the assignment statement $r_j := r_j + r_i$

(note that the value is preserved in r_i). It is clear that *copy* can be constructed with a constant number of register machine states. In fact, register machines are known to be Turing-universal with three registers [30].³


In [36], a simple CRN construction was shown to simulate register machines with some possibility of error (thus not directly compatible with stable computation). The source of the error is due to the zero-checking in a *dec* instruction. For the simulation, the CRN has a finite set of species (one for each register and one for each state of the register machine) and a finite set of reactions (one for each instruction in the register machine program). Each *inc*(r_i, s_j) instruction corresponds to the reaction $S_{j'} \rightarrow R_i + S_j$, and each *dec*(r_i, s_j, s_k) instruction to two reactions $S_{j'} + R_i \rightarrow S_j$ and $S_{j'} \rightarrow S_k$. In the chemical reaction network implementation of a *dec* instruction, the two reactions are competing for the state species $S_{j'}$. While in general this is an unavoidable problem, in the special case that the maximum value in our counters is bounded by a constant, we can remedy this following the idea from [28] as follows.

Let's consider bounded registers that can contain a value no greater than $b \in \mathbb{N}$. For each register r_i , we can use two species R_i^A and R_i^I as “active” and “inactive” species for register r_i , respectively. The idea is that the total sum of the counts of species R_i^A and R_i^I is always equal to b : whenever one is consumed, the other is produced. Now, an *inc*(r_i, s_j) instruction could be implemented with the reaction $S_{j'} + R_i^I \rightarrow R_i^A + S_j$, and a *dec*(r_i, s_j, s_k) instruction could be implemented with the reactions $S_{j'} + R_i^A \rightarrow R_i^I + S_j$ and $S_{j'} + bR_i^I \rightarrow bR_i^I + S_k$. With this approach, register r_i has a zero count exactly when inactive species R_i^I has a count of b , and so we can zero-check without error. Notice that this approach uses reactions with a large stoichiometric coefficient b . At this point, there are two issues to be addressed: (1) how to generate an initial b count of inactive species R_i^I , and (2) how to transform the reactions into a series of third-order reactions (avoiding the large stoichiometric coefficient b).

Let's first consider a very simple construction which addresses the above concerns, albeit suboptimally. Suppose $b = 2^n$ is a power of two. To handle (1), we can initially produce count b of R_i^I from a single copy of A_1 using $O(\log b)$ species with reactions



To handle (2), we can transform the decrement reactions into a series of n bimolecular reactions by adding reversible versions of the reactions from (1) and “counting down” to some unique zero count indicator species C_1 :



Then C_1 is producible if and only if R_i^I had count $\geq b$, so the reaction $S_j + C_1 \rightarrow S_k + C_1$ implements the “jump to state k if $r_i = 0$ ” portion of the *dec*(r_i, s_j, s_k) command. This

³ Turing-universality has also been shown for machines with two registers, but only when a nontrivial encoding of the input/output is allowed [30, 33].

construction allows an error-free simulation of a register machine with counters with bound b exponential in the number of species. Now we discuss a more sophisticated construction, based on previous results [11, 28], that achieves a counter bound b that is *doubly* exponential in the number of species.

3.2 Counting to 2^{2^n} with n species

The CRN constructions in this paper simulate bounded register machines in the manner discussed previously. Since we are focused on reducing the size of our CRN, we want to do this simulation with as few species (reactions) as possible. Fortunately, we can rely on established results from prior work to do this. Lipton provided a construction for which the largest producible amount of a species is a doubly exponential count [28]. However, this amount is only produced non-deterministically and (most) paths produce less. Cardoza et al. went on to present a fully reversible system that can achieve this doubly exponential count as well [11]. Further, their system is halting in the sense that a new species is produced precisely when the maximum amount is reached.

While Cardoza et al. [11] describe their construction in the language of commutative semigroup presentations, we present a modified construction in Figure 1 articulated as a CRN. In the figure and in the text below, we use the “box” notation to indicate *meta-reactions*, which correspond to a set of reactions. Note that in Lemma 4 we will see that the combined behavior of the reactions in a meta-reaction module faithfully implement the meta-reaction semantics. By construction, the sets of reactions that meta-reactions expand to overlap, and we include only one copy of any repeated reaction. Each layer of the construction introduces $\mathcal{O}(1)$ more reactions and species – 9 reactions ((1)–(9)) and 9 species ($S_i^k, H_i^k, X_i^k, T1_i^k, T2_i^k, C1_i^k, C2_i^k, C3_i^k, C4_i^k$) for each $i \in \{1, 2, 3, 4\}$.

The idea of the construction is to produce (or consume) a doubly exponential count of species X by recursively producing (or consuming) quadratically more X 's than the previous layer. Each species type performs a different role. X is the counting species to be generated or consumed. S starts the process to generate/consume many molecules of species X . T transforms different types of X species into one another. H indicates that the generation/consumption process has completed. C “cleans up” the H species. Reaction (5) in the meta-reaction implementation (which converts X_2^{k-1} into X_3^{k-1}) changes based on i . If $i \in \{1, 2\}$, then X_i appears as a product and is generated by this reaction. If $i \in \{3, 4\}$, the X_i appears as a reactant and is consumed by this reaction. A high level diagram of a layer- k meta-reaction is shown in Figure 2, which is helpful in understanding the behavior of the system.


► **Definition 1.** Let \mathbf{c} be a configuration of CRN \mathcal{C} given above. We say \mathbf{c} is well-led if $\mathbf{c}(S_*^*) + \mathbf{c}(H_*^*) + \mathbf{c}(T_*^*) = 1$ where the notation S_*^* denotes any species with label S , regardless of the subscript or superscript. In other words, there is only a single leader in the system and it either has the label S , H , or T . We call species S_*^* , H_*^* , and T_*^* leader species.

► **Observation 2.** Every reaction has exactly one leader species as a reactant, and exactly one leader species as a product.


The following is immediate from Observation 2:

► **Corollary 3.** Let \mathbf{c} be a well-led configuration of CRN \mathcal{C} given above. Then any configuration \mathbf{d} such that $\mathbf{c} \rightsquigarrow \mathbf{d}$ is also well-led. In other words, the well-led property is forward invariant.

Informally, the observation above together with the well-led condition implies that we can reason about the meta-reactions in isolation, without fear of cross-talk – because while one meta-reaction is executing, no reactions outside of it are applicable. This allows us to inductively prove the main result of this section:



■ **Figure 1** Doubly exponential counting construction. (Left) The base reactions and layer- k meta-reactions. We use the box notation to indicate meta-reactions, which correspond to a set of reactions. Note that the reactions corresponding to the different meta-reactions overlap; when all the meta-reactions are expanded we include only one reaction copy. (Right) Explicit reactions for the layer- k meta-reaction in terms of reactions and other meta-reactions. The core functionality is the same for any i , but reaction (5) either generates X_i^k 's (if $i \in \{1, 2\}$) or consumes X_i^k 's (if $i \in \{3, 4\}$).



■ **Figure 2** A visualization of the states for a layer- k meta-reaction. The state transitions effectively execute a nested loop. In order to iterate the outer loop (transition from V \rightarrow II), the inner loop IV must be executed $2^{2^{k-1}}$ times. And in order to leave state VI and produce an H , the outer loop must be executed $2^{2^{k-1}}$. So, state IV must be executed a total of 2^{2^k} times, which either produces or consumes that many X_i^k 's depending on the type of meta-reaction.

► **Lemma 4** (Production). Consider the CRN implementing $S_i^k \rightleftharpoons H_i^k + 2^{2^k} X_i^k$. For any $n \in \mathbb{N}$, let $\mathbf{s} = \{nX_i^k, 1S_i^k\}$ and $\mathbf{h} = \{(2^{2^k} + n)X_i^k, 1H_i^k\}$, and let \mathbf{c} be any configuration reachable from \mathbf{s} or \mathbf{h} . Then: (a) Both \mathbf{s} and \mathbf{h} are reachable from \mathbf{c} . (b) If \mathbf{c} contains S_i^k then $\mathbf{c} = \mathbf{s}$, and if \mathbf{c} contains H_i^k then $\mathbf{c} = \mathbf{h}$.

► **Lemma 5** (Consumption). Consider the CRN implementing $2^{2^k} X_i^k + S_i^k \rightleftharpoons H_i^k$. For any $n \in \mathbb{N}$, let $\mathbf{s} = \{(2^{2^k} + n)X_i^k, 1S_i^k\}$ and $\mathbf{h} = \{nX_i^k, 1H_i^k\}$, and let \mathbf{c} be any configuration reachable from \mathbf{s} or \mathbf{h} . Then: (a) Both \mathbf{s} and \mathbf{h} are reachable from \mathbf{c} . (b) If \mathbf{c} contains S_i^k then $\mathbf{c} = \mathbf{s}$, and if \mathbf{c} contains H_i^k then $\mathbf{c} = \mathbf{h}$.

Proof. (Of Lemma 4 and Lemma 5, Sketch) Both lemmas are proven by induction over the layers of the construction. The base case ($k = 1$) can be checked by inspection. Now assume the lemmas are true for $k - 1$ layers, and we want to prove them true for k layers.

First we argue that the construction is correct if the $k - 1$ layer meta-reactions are “atomic” and occur in one step. As visualized in Figure 2, the CRN iterates through a nested loop process. Each state transition (states I^k through VI^k) is coupled to a conversion of the leader species; the well-led condition ensures that the CRN is in exactly one state at any given time. Each net forward traversal of the outer loop converts a X_1^{k-1} to X_4^{k-1} , and each forward traversal of the inner loop converts a X_2^{k-1} to X_3^{k-1} . Step I^k makes $2^{2^{k-1}} X_1^{k-1}$, bounding the net maximum number of times that the outer loop can happen in the forward direction. Step III^k makes $2^{2^{k-1}} X_2^{k-1}$, bounding the net maximum number of times that the inner loop can happen in the forward direction for every net forward traversal of the outer loop. This implies that reaction (5) can fire at most a net total 2^{2^k} times (producing at most a net total $2^{2^k} X_i^k$'s).

Step V^k consumes $2^{2^{k-1}} X_3^{k-1}$, requiring the net total number of forward traversals of the inner loop to be at least $2^{2^{k-1}}$ for every net forward traversal of the outer loop. Step VI^k consumes $2^{2^{k-1}} X_4^{k-1}$, requiring the net total number of forward traversals of the outer loop to be at least $2^{2^{k-1}}$. This implies that reaction (5) must fire at least a net total 2^{2^k} times (producing at least a net total $2^{2^k} X_i^k$'s).

Thus, reaction (5) must be executed exactly 2^{2^k} times (producing exactly $2^{2^k} X_i^k$'s). Notice that an excess of X_i^k (as allowed by the statement of the lemma) does not affect the net total number of times reaction (5) can fire (forward or backward) since X_2^{k-1} and X_3^{k-1} are the limiting factors.

Now we need to make sure that this behavior is preserved once the meta-reactions are expanded to their constituent reactions. Each meta-reaction i in Figure 1 expands to some set R_i of reactions. First we note that for each meta-reaction, R_i overlaps with reactions not in R_i only over species S_i^{k-1} , H_i^{k-1} , and X_i^{k-1} . We are not worried about cross-talk in species S_i^{k-1} and H_i^{k-1} because of the well-led property. We may still be concerned, however, that external consumption of X_i^{k-1} might somehow interfere with the meta-reaction. Luckily, the well-led property and Observation 2 enforce that unless we have S_i^{k-1} or H_i^{k-1} (i.e., we are at the beginning or end of the meta-reaction), it is never the case that a reaction in R_i and a reaction not in R_i are applicable at the same time. Thus nothing outside the meta-reaction can change X_i^{k-1} while the meta-reaction is executing. ◀

Note that although we chose to write Lemma 4 and Lemma 5 separately, we could have just one kind of meta-reaction (production or consumption) and obtain the other kind by running the meta-reaction backward switching the roles of S and H . We include the two different versions because it is conceptually easier to just think about the intended execution being in the forward direction.

4 Optimal Encoding

4.1 Encoding Information in CRNs

In this section we discuss the encoding of an integer in a chemical reaction network. In the same sense as Kolmogorov-optimal programs for Turing machines, we consider a similar measure of optimality for chemical reaction networks. In particular, we ask the question, “what is the smallest chemical reaction network that can produce a desired count of a particular chemical species?”

A simple construction shows that x copies of some species can be produced using $\mathcal{O}(\log x)$ reactions. The idea is to have a reaction for each bit b_i of the binary expansion of x , and produce a copy of your output species in each reaction where $b_i = 1$. More concretely, consider $\log x$ reactions of the form $X_i \rightarrow 2X_{i+1}$ and $X_i \rightarrow 2X_{i+1} + Y$. For each bit b_i in the binary expansion of x , use the first reaction if $b_i = 0$ and use the second reaction if $b_i = 1$. Each species X_i will have a count equal to 2^i , and species Y will have a count equal to the sum of the powers of two that were chosen (which is x). While this simple construction generates x with $\log x$ reactions, it is not immediately clear how to improve upon it.

Our first result shows how to construct a CRN that can generate x copies of an output species (from an initial configuration with only a single molecule) yet uses only $\mathcal{O}(\log x / \log \log x)$ many reactions. This matches the lower bound dictated by Kolmogorov complexity (see end of Section 4), which suggests that the full power of CRNs is really being used in our construction. Our construction is achieved through the simulation of (space-bounded) Turing machines via the simulation of (space-bounded) register machines. A key aspect in this process is the ability of CRNs to use the previously discussed recursive counting technique to count very high with very few species (counting to 2^{2^k} with k species).

4.2 Our Construction

Now, we present an encoding scheme to produce count x of a particular species with $\mathcal{O}(n/\log n)$ CRN reactions, where $n = \log x$. In the simple CRN given in Section 4.1, each reaction encodes a single bit of x . In the optimized construction with k reactions, each reaction will encode $\log k$ bits instead.⁴ A sketch of our construction is as follows:

Sketch: We start with a CRN in configuration $\mathbf{c}_1 = \{1L\}$ and create a configuration $\mathbf{c}_2 = \{1S_i, m_1R_1, m_2R_2, \dots, m_kR_k\}$ that represents a particular permutation of k distinct elements. We encode this permutation in the count of a species I , transforming configuration \mathbf{c}_2 into a configuration $\mathbf{c}_3 = \{1S_j, mI\}$. The count of species I can be interpreted as the input to a Turing machine, so we simulate a Turing machine that maps the permutation to a unique integer via Lehmer code/factorial number system [25, 34] (by choosing the right value of k , we can ensure there are sufficiently many permutations to let us map to x). This Turing machine simulation transforms configuration \mathbf{c}_3 into configuration $\mathbf{c}_4 = \{1H, xY\}$.

► **Theorem 6.** For any $n \in \mathbb{N}$ and any n -bit integer x , there exists a chemical reaction network \mathcal{C}_x that haltingly computes x from initial configuration $\{1L\}$ with $|\mathcal{C}_x| = \mathcal{O}(n/\log n)$.

Proof. First, we describe how to construct CRN \mathcal{C}_x that haltingly computes x from starting configuration $\{1L\}$, then we describe the size of $|\mathcal{C}_x|$. Let $k = \lceil n/\log n \rceil$. We will map a permutation of k distinct elements to the integer x , and this value of k ensures there are at least x permutations. We break the construction into three primary steps.

⁴ Adleman et al. [1] provided a clever base conversion trick for tile assembly programs. Here, we employ a permutation encoding trick to yield the same effect.

Step 1: $\{1L\} \rightsquigarrow \{1S_i, m_1R_1, m_2R_2, \dots, m_kR_k\}$. We can transform $\{1L\}$ into a configuration $\{1S_i, m_1R_1, m_2R_2, \dots, m_kR_k\}$ where (m_1, m_2, \dots, m_k) is a permutation of the integers 1 through k . This can be achieved with k registers and $2k$ register machine states. For example, to set the permutation $(2, 4, 3, 1)$, use instructions

$$\begin{aligned} s_0 &: inc(r_2, s_1) \\ s_1 &: copy(r_2, r_4, s_2) \\ s_2 &: inc(r_2, s_3) \\ s_3 &: copy(r_2, r_1, s_4) \\ s_4 &: inc(r_2, s_5) \\ s_5 &: copy(r_2, r_3, s_6) \\ s_6 &: inc(r_2, s_7) \end{aligned}$$

Step 2: $\{1S_i, m_1R_1, m_2R_2, \dots, m_kR_k\} \rightsquigarrow \{1S_j, mI\}$. Now, we can transform configuration $\{1S_i, m_1R_1, m_2R_2, \dots, m_kR_k\}$ into configuration $\{1S_j, mI\}$, encoding the permutation as the integer count m of species I . For each register r_i for i from 1 to k in order, we can decrement the register to 0. On each decrement, we double the count of I and then add 1 to it, i.e., appending a 1 to m 's binary expansion. After the register reaches 0, before moving to the next register, we double the count of I again, appending a 0 to m 's binary expansion. For example, if the permutation configuration was $\{3R_1, 1R_2, 2R_3\}$, the resulting count of I in binary would be

$$\underbrace{111}_3 \underbrace{0}_1 \underbrace{1}_1 \underbrace{0}_1 \underbrace{11}_2.$$

Step 3: $\{1S_j, mI\} \rightsquigarrow \{1H, xY\}$. At this point, we can consider the value in register I , expressed as a binary string, to be the input tape content for a Turing machine that maps the permutation to the integer x using a standard Lehmer code/factorial number system technique [25, 34]. The output of the Turing machine will be the count of Y in configuration \mathbf{c} at the end of the computation (with $\mathbf{c}(Y) = x$). Our register machine will have a state species that corresponds to the halted state of the Turing machine – and such a species serves as our halting species H .

Now we argue the size of CRN $|\mathcal{C}_x| = \mathcal{O}(n/\log n)$, i.e., it uses $\mathcal{O}(k)$ reactions. The register machine program from Step 1 generates the permutation using k registers and $2k$ register machine states, which results in $\mathcal{O}(k)$ CRN reactions. The register machine program from Step 2 encodes the permutation as a binary number in register I using $\mathcal{O}(k)$ registers and $\mathcal{O}(k)$ register machine states, which also results in $\mathcal{O}(k)$ CRN reactions. Even a naive algorithm for the Turing machine from Step 3 maps the permutation to an integer using $\mathcal{O}(k^2 \log k)$ space ($\mathcal{O}(k^2)$ bits to store the initial permutation, $\mathcal{O}(k \log k)$ bits to store the Lehmer code, $\mathcal{O}(k^2 \log k)$ bits to store factorial bases $1!$ through $k!$, and $\mathcal{O}(k \log k)$ bits to store the integer x). Recall, a Turing machine using space $\mathcal{O}(k^2 \log k)$ can be simulated by a register machine with count bound $\mathcal{O}(2^{k^2 \log k})$ on its registers. This can in turn be simulated by a CRN via the construction of Section 3.2 with $\mathcal{O}(\log \log 2^{k^2 \log k}) = \mathcal{O}(\log k)$ reactions. Thus $\mathcal{O}(k)$ reactions suffices to simulate the register machine instructions as well as the bounded counters for our register machine to simulate this Turing machine. ◀

The above construction is optimal for almost all integers x in the following sense. Any CRN of $|\mathcal{C}|$ reactions, each with $\mathcal{O}(1)$ reactants and products, can be encoded in a string of length $\mathcal{O}(|\mathcal{C}| \log |\mathcal{C}|)$. Given an encoded CRN stably computing an integer x , a fixed-size

program can simulate it and return x . Thus $K(x) \leq \mathcal{O}(|\mathcal{C}| \log |\mathcal{C}|)$. The pigeonhole principle argument for Kolmogorov complexity implies that $K(x) < \lceil \log x \rceil - \Delta$ for at most $(1/2)^\Delta$ of all x [27]. Together these observations imply that there is a c such that for most x there does not exist a CRN \mathcal{C} of size smaller than $cn/\log n$ that stably computes x .

5 Algorithmic Compression

The construction in Section 4 is optimal for incompressible integers (integers x where $K(x) \approx |x|$, which is the case for “most” integers). Now we extend the construction to be optimal within global multiplicative constants for *all* integers. For algorithmically compressible integers x , there exists a p such that $\mathcal{U}(p) = x$ and $|p| < |x|$. We discuss the construction in Section 5.1 and we argue optimality of our construction in Section 5.2.

5.1 Our Construction

We now show how to fully exploit the encoding scheme and doubly exponential counter from Section 4 to achieve an optimal result for all integers. A sketch of our construction is as follows:

Sketch: Given a program p for a fixed Universal Turing Machine \mathcal{U} such that $\mathcal{U}(p) = x$, we construct a CRN that simulates running p on \mathcal{U} via a register machine simulation. The idea is to use $\mathcal{O}(|p|/\log|p|)$ reactions to encode p , and to use $\mathcal{O}(\log(\text{space}(\mathcal{U}(p))))$ reactions for a counter machine simulation of $\mathcal{U}(p)$.

► **Theorem 7.** *For any integer x , there exists a CRN \mathcal{C}_x that haltingly computes x from initial configuration $\{1L\}$ with $|\mathcal{C}_x| = \mathcal{O}(\tilde{K}s(x))$.*

Proof. Let p be a program for a fixed Universal Turing Machine \mathcal{U} such that $\mathcal{U}(p) = x$. We encode p in the manner provided by Theorem 6 using $\mathcal{O}(|p|/\log|p|)$ reactions. This results in p count of species Y (specifically, configuration $\{1H, pY\}$). Since haltingly-computing CRNs are composable via concatenation [12, 35], we can consider $\{1H, pY\}$ to be taken as the input for another system which simulates running $\mathcal{U}(p)$ via the previously described register machine method with bounded register count (Section 3.2). Again, we need enough species/reactions to ensure our bounded registers can count high enough. The registers must be able to store an integer that represents the current configuration of the Turing machine being simulated (at most this is $2^{\text{space}(\mathcal{U}(p))}$). Since we have doubly exponential counters, an additional $\log(\text{space}(\mathcal{U}(p)))$ species are needed to do this. So, the total size of our CRN is $\mathcal{O}(|p|/\log|p| + \log(\text{space}(\mathcal{U}(p))))$ and by choosing the program p that minimizes this expression, we see $|\mathcal{C}_x| = \mathcal{O}(\tilde{K}s(x))$. ◀

It is interesting to note the appearance of our “space-aware” version of Kolmogorov complexity. Importantly, this notion is different from space-bounded Kolmogorov complexity that puts a limit on the space usage of the program that outputs x . This alternate version allows a trade-off between compact program descriptions and the space required to run those programs, which seems natural for systems like CRNs. Perhaps it is surprising that this (computable) measure of complexity shows up here, and at first it may seem like \log of this space usage is a bit arbitrary, but we will show that this is indeed optimal (within global multiplicative constants) for CRNs.

5.2 Optimality

Here, we argue that size of CRN \mathcal{C}_x from Theorem 7 is optimal. We begin by giving a definition for the size of the optimal CRN that haltingly computes an integer x .

► **Definition 8.** *For any integer x , define $K_{\text{crn}}(x) = \min\{|\mathcal{C}| : \text{CRN } \mathcal{C} \text{ haltingly computes } x\}$. In other words, $K_{\text{crn}}(x)$ is the size of the smallest CRN that haltingly computes x .*

Our argument relies on a Turing machine that solves the coverability problem for CRNs. We give the definition for this problem in Definition 9 and discuss its space complexity in Lemma 10.

► **Definition 9 (Coverability).** *Given a CRN \mathcal{C} , initial configuration \mathbf{s} , and target configuration \mathbf{u} , does there exist a configuration $\mathbf{t} \geq \mathbf{u}$ such that $\mathbf{s} \rightsquigarrow \mathbf{t}$?*

Using the natural notion of problem size n for the specification of a coverability problem, Lipton provided a $2^{\Omega(\sqrt{n})}$ space lower bound for coverability [28], which was later improved to $2^{\Omega(n)}$ by Mayr and Meyer [29]. As for upper bounds, Rackoff provided an algorithm to decide coverability that uses $2^{\mathcal{O}(n \log n)}$ space [31]. Following this, Koppenhagen and Mayr gave an algorithm that decided coverability in $2^{\mathcal{O}(n)}$ space for *reversible* systems, closing the gap for this class of systems [23]. A recent result by Künnemann et al. also closes this gap [24] for Vector Addition Systems with States. Our work uses this latest result.

► **Lemma 10** (Implied by Theorem 3.3 from [24]). *Let CRN $\mathcal{C} = (S, \mathcal{R})$ be a CRN that haltingly computes x . Then there exists an algorithm which solves coverability for CRN \mathcal{C} for initial configuration $\{1L\}$ and target configuration $\{1H\}$ which uses $2^{\mathcal{O}(|\mathcal{C}|)}$ space.*

Proof. This result follows from Theorem 3.3 from the recent work by Künnemann et al. [24]. There, the authors consider the problem of coverability in Vector Addition Systems with States (VASS). They show that if the answer to coverability is yes, the length of the longest path is $n^{2^{\mathcal{O}(d)}}$ where n is the maximum value change of any transition and d is the dimension of the vector. For us, $n = 2$ since each reaction has at most two reactants/products and $d = |S|$. While vector addition systems are not capable of “catalytic” transitions, it is known that the same effect can be achieved by decomposing transitions into two vector additions. So the path length would at most double for our systems.

With this bound on the path length, we can consider an algorithm that non-deterministically explores the state space of \mathcal{C} (from starting configuration $\{1L\}$) by simulating reactions on a current configuration of the system until a configuration that covers $\{1H\}$ is found. By Savitch’s Theorem [32], this can be converted to a deterministic algorithm using the same space: $\mathcal{O}(|\mathcal{R}| \log |\mathcal{R}|)$ bits to hold a description of \mathcal{C} , $2^{\mathcal{O}(|S|)}$ bits for a path length counter, and $2^{\mathcal{O}(|S|)} \cdot \log(|S|)$ bits to store the current configuration of \mathcal{C} . All of these values are absorbed under a $2^{\mathcal{O}(|\mathcal{C}|)}$ bound. ◀

With this space bound on the coverability problem established, we can now argue that the size of our constructed CRN from Theorem 7 is asymptotically equal to the size of the the smallest CRN that haltingly computes x .

► **Lemma 11.** *For all $x \in \mathbb{N}$, letting \mathcal{C}_x be the CRN from Theorem 7, $|\mathcal{C}_x| = \Theta(K_{\text{crn}}(x))$.*

Proof. Clearly $|\mathcal{C}_x| = \Omega(K_{\text{crn}}(x))$, by definition of $K_{\text{crn}}(x)$ and since \mathcal{C}_x from Section 5 is an instance of a CRN that haltingly computes x . Now, we argue that $|\mathcal{C}_x| = \mathcal{O}(K_{\text{crn}}(x))$. The big picture is that one of the programs over which $\tilde{K}_s(x)$ is minimized in the construction of \mathcal{C}_x is the program solving coverability for the optimal CRN for generating x .

Start with the CRN $\mathcal{K} = (\mathcal{S}_{\mathcal{K}}, \mathcal{R}_{\mathcal{K}})$ that haltingly computes x with optimal size $|\mathcal{K}| = K_{\text{crn}}(x) = n$. Consider program $p_{\mathcal{K}}$ that solves coverability for \mathcal{K} with initial configuration $\{1L\}$ and target configuration $\{1H\}$, and outputs x . Now, with that program $p_{\mathcal{K}}$, build a CRN \mathcal{K}' by following our construction for Theorem 7. We know $|p_{\mathcal{K}}| = \mathcal{O}(n \log n)$ so our final CRN needs $\mathcal{O}(n)$ reactions to encode $p_{\mathcal{K}}$ (by Theorem 6). By Lemma 10, we know that the space usage of $\mathcal{U}(p_{\mathcal{K}})$ is $2^{\mathcal{O}(|\mathcal{K}|)}$, so our final CRN needs $\mathcal{O}(|\mathcal{K}|)$ additional reactions to have large enough registers for the simulation of $\mathcal{U}(p_{\mathcal{K}})$. Thus, the total size of our final CRN is $\mathcal{O}(|\mathcal{K}| + |\mathcal{K}'|) = \mathcal{O}(K_{\text{crn}}(x))$. ◀

The following theorem, which is the main result of our paper, follows immediately from Theorem 7 and Lemma 11. It characterizes the optimal number of reactions haltingly computing a number x using the space-aware Kolmogorov complexity measure \tilde{K}_s defined in Section 2.2.

► **Theorem 12.** *For all $x \in \mathbb{N}$, $K_{\text{crn}}(x) = \Theta(\tilde{K}_s(x))$.*

Although, as mentioned above, CRN stable computation is not Turing universal, the theorem underlines its essential connection to space-bounded Turing machine computation.

6 Open Questions

Our results rely on the fact that we consider CRNs that perform halting computation – the end of the computation is indicated by the production of a designated halting species. This constraint, intuitively that the systems know when they have finished a computation, is rather strong. It is known that a much larger class of functions can be stably computed than can be haltingly computed [15]. It remains an open question if lifting this halting requirement (and allowing just stable computation) reduces the reaction complexity.

It is also worth noting that our approach starts with exactly one copy of a special leader species. Recently, Czerner showed that leaderless protocols are capable of deciding doubly exponential thresholds [16]. While starting in some uniform state and converging to a specific state would be a better expression of “chemical self-organization,” their construction seems incompatible with our register machine simulation. Leaderless stable integer computation remains an area for future work.

Making a tight connection between stable integer computation and counting predicate computation commonly studied in population protocols [16] also remains open. We can easily follow the halting generation of a specific amount of x by running the “less-than-or-equal-to” predicate, thereby converting our constructions to compute a counting predicate with only a constant more reactions. This gives a new general upper bound on the complexity of counting in terms of $\tilde{K}_s(x)$. However, it is unclear whether counting predicate constructions carry over to the generation problem, leaving it open whether counting may be easier.

Our notion of “space-aware” Kolmogorov complexity \tilde{K}_s is interesting in its own right. While the similar quantity K_s has been previously studied in the context of computational complexity theory [3] (see also Section 2.2), it is not clear which properties proven of K_s carry over to \tilde{K}_s . Although the robustness to the choice of \mathcal{U} carries over, other properties may not. For example, it is not obvious whether our results still hold if we consider programs that output a single bit of x at a time (like K_s does).

A core piece of this work is simulating space-bounded Turing machines, so it is very natural to extend the discussion to Boolean circuits (computing functions $\phi : \{0, 1\}^n \rightarrow \{0, 1\}$). When attempting to compute Boolean functions with CRNs, one may be tempted to directly implement a Boolean circuit by creating $\mathcal{O}(1)$ reactions per gate in the circuit. However,

our results imply that reaction complexity can be improved by doing a space-bounded Turing machine simulation instead – when the circuit is algorithmically “compressible.” An important class of such compressible circuits are uniform circuits, i.e., those constructable by a fixed Turing machine given an input size. Prior work established a quadratically tight connection between the depth of uniform circuits and Turing machine space [9]. Further investigation into optimal Boolean function computation is warranted.

References

- 1 Leonard Adleman, Qi Cheng, Ashish Goel, and Ming-Deh Huang. Running time and program size for self-assembled squares. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 740–748, 2001.
- 2 Dan Alistarh, James Aspnes, and Rati Gelashvili. Space-optimal majority in population protocols. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2221–2239. SIAM, 2018.
- 3 Eric Allender, Michal Koucký, Detlef Ronneburger, and Sambuddha Roy. The pervasive reach of resource-bounded Kolmogorov complexity in computational complexity theory. *Journal of Computer and System Sciences*, 77(1):14–40, 2011.
- 4 Dana Angluin, James Aspnes, Zoë Diamadi, Michael J Fischer, and René Peralta. Computation in networks of passively mobile finite-state sensors. In *Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing*, pages 290–299, 2004.
- 5 Dana Angluin, James Aspnes, and David Eisenstat. Stably computable predicates are semilinear. In *Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*, pages 292–299, 2006.
- 6 Petra Berenbrink, George Giakkoupis, and Peter Kling. Optimal time and space leader election in population protocols. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 119–129, 2020.
- 7 Michael Blondin, Javier Esparza, Blaise Genest, Martin Helfrich, and Stefan Jaax. Succinct population protocols for presburger arithmetic. *arXiv preprint arXiv:1910.04600*, 2019.
- 8 Michael Blondin, Javier Esparza, and Stefan Jaax. Large flocks of small birds: on the minimal size of population protocols. *arXiv preprint arXiv:1801.00742*, 2018.
- 9 Allan Borodin. On relating time and space to size and depth. *SIAM journal on computing*, 6(4):733–744, 1977.
- 10 Daniele Cappelletti, Andrés Ortiz-Muñoz, David F Anderson, and Erik Winfree. Stochastic chemical reaction networks for robustly approximating arbitrary probability distributions. *Theoretical Computer Science*, 801:64–95, 2020.
- 11 E Cardoza, R Lipton, and Albert R Meyer. Exponential space complete problems for Petri nets and commutative semigroups (preliminary report). In *Proceedings of the eighth annual ACM symposium on Theory of computing*, pages 50–54, 1976.
- 12 Cameron Chalk, Niels Kornerup, Wyatt Reeves, and David Soloveichik. Composable rate-independent computation in continuous chemical reaction networks. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 18(1):250–260, 2019.
- 13 Ho-Lin Chen, David Doty, and David Soloveichik. Deterministic function computation with chemical reaction networks. *Natural computing*, 13:517–534, 2014.
- 14 Ben Chugg, Hooman Hashemi, and Anne Condon. Output-oblivious stochastic chemical reaction networks. In *22nd International Conference on Principles of Distributed Systems (OPODIS 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- 15 Rachel Cummings, David Doty, and David Soloveichik. Probability 1 computation with chemical reaction networks. *Natural Computing*, 15(2):245–261, 2016. Special issue of invited papers from DNA 2014. doi:10.1007/s11047-015-9501-x.
- 16 Philipp Czerner. Leaderless population protocols decide double-exponential thresholds. *arXiv preprint arXiv:2204.02115*, 2022.

- 17 Philipp Czerner and Javier Esparza. Lower bounds on the state complexity of population protocols. In *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*, pages 45–54, 2021.
- 18 Philipp Czerner, Roland Guttenberg, Martin Helfrich, and Javier Esparza. Fast and succinct population protocols for presburger arithmetic. *arXiv preprint arXiv:2202.11601*, 2022.
- 19 David Doty, Mahsa Eftekhari, Leszek Gąsieniec, Eric Severson, Przemysław Uznański, and Grzegorz Stachowiak. A time and space optimal stable population protocol solving exact majority. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1044–1055. IEEE, 2022.
- 20 Javier Esparza and Mogens Nielsen. Decidability issues for Petri nets—a survey. *Journal of Information Processes and Cybernetics*, 3:143–160, 1994.
- 21 Daniel T Gillespie. Stochastic simulation of chemical kinetics. *Annu. Rev. Phys. Chem.*, 58:35–55, 2007.
- 22 Richard M Karp and Raymond E Miller. Parallel program schemata. *Journal of Computer and System Sciences*, 3(2):147–195, 1969.
- 23 Ulla Koppenhagen and Ernst W Mayr. Optimal algorithms for the coverability, the subword, the containment, and the equivalence problems for commutative semigroups. *Information and Computation*, 158(2):98–124, 2000.
- 24 Marvin Künnemann, Filip Mazowiecki, Lia Schütze, Henry Sinclair-Banks, and Karol Węgrzycki. Coverability in VASS revisited: Improving Rackoff’s bound to obtain conditional optimality. *arXiv preprint arXiv:2305.01581*, 2023.
- 25 Derrick H Lehmer. Teaching combinatorial tricks to a computer. In *Proc. Sympos. Appl. Math. Combinatorial Analysis*, volume 10, pages 179–193, 1960.
- 26 Jérôme Leroux. State complexity of protocols with leaders. In *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*, pages 257–264, 2022.
- 27 Ming Li, Paul Vitányi, et al. *An introduction to Kolmogorov complexity and its applications*, volume 3. Springer, 2008.
- 28 Richard Lipton. The reachability problem requires exponential space. *Research Report 62. Department of Computer Science, Yale University*, 1976.
- 29 Ernst W Mayr and Albert R Meyer. The complexity of the word problems for commutative semigroups and polynomial ideals. *Advances in mathematics*, 46(3):305–329, 1982.
- 30 Marvin Lee Minsky. *Computation*. Prentice-Hall Englewood Cliffs, 1967.
- 31 Charles Rackoff. The covering and boundedness problems for vector addition systems. *Theoretical Computer Science*, 6(2):223–231, 1978.
- 32 Walter J Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of computer and system sciences*, 4(2):177–192, 1970.
- 33 Rich Schroepfel. A two counter machine cannot calculate 2^N . Technical report, Massachusetts Institute Of Technology, Artificial Intelligence Lab, 1972.
- 34 Robert Sedgewick. Permutation generation methods. *ACM Computing Surveys (CSUR)*, 9(2):137–164, 1977.
- 35 Eric E Severson, David Haley, and David Doty. Composable computation in discrete chemical reaction networks. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, pages 14–23, 2019.
- 36 David Soloveichik, Matthew Cook, Erik Winfree, and Jehoshua Bruck. Computation with finite stochastic chemical reaction networks. *natural computing*, 7:615–633, 2008.