

Hardness of computing and approximating predicates and functions with leaderless population protocols*

Amanda Belleville¹, David Doty¹, and David Soloveichik²

¹Department of Computer Science, University of California, Davis

²Department of Electrical and Computer Engineering, University of Texas, Austin

Abstract

Population protocols are a distributed computing model appropriate for describing massive numbers of agents with very limited computational power (finite automata in this paper), such as sensor networks or programmable chemical reaction networks in synthetic biology. A population protocol is said to require a leader if every valid initial configuration contains a single agent in a special “leader” state that helps to coordinate the computation. Although the class of predicates and functions computable with probability 1 (stable computation) is the same whether a leader is required or not (semilinear functions and predicates), it is not known whether a leader is necessary for fast computation. Due to the large number of agents n (synthetic molecular systems routinely have trillions of molecules), efficient population protocols are generally defined as those computing in polylogarithmic in n (parallel) time. We consider population protocols that start in leaderless initial configurations, and the computation is regarded finished when the population protocol reaches a configuration from which a different output is no longer reachable.

In this setting we show that a wide class of functions and predicates computable by population protocols are not *efficiently* computable (they require at least linear time to stabilize on a correct answer), nor are some linear functions even *efficiently approximable*. For example, our results for predicates immediately imply that the widely studied parity, majority, and equality predicates cannot be computed in sublinear time. (Existing arguments specific to majority were already known). Moreover, it requires at least linear time for a population protocol even to approximate division by a constant or subtraction (or any linear function with a coefficient outside of \mathbb{N}), in the sense that for sufficiently small $\gamma > 0$, the output of a sublinear time protocol can stabilize outside the interval $f(m)(1 \pm \gamma)$ on infinitely many inputs m . We also show that it requires linear time to exactly compute a wide range of semilinear functions (e.g., $f(m) = m$ if m is even and $2m$ if m is odd).

In a complementary positive result, we show that with a sufficiently large value of γ , a population protocol *can* approximate any linear f with nonnegative rational coefficients, within approximation factor γ , in $O(\log n)$ time.

1 Introduction

Population protocols were introduced by Angluin, Aspnes, Diamadi, Fischer, and Peralta [4] as a model of distributed computing in which the agents have very little computational power and no control over their schedule of interaction with other agents. They can be thought of as a special case of a model of concurrent processing introduced in the 1960s, known alternately as

*The first and second authors were supported by NSF grant 1619343, and the third author by NSF grant 1618895.

vector addition systems [26], Petri nets [30], or commutative semi-Thue systems (or, when all transitions are reversible, “commutative semigroups”) [12, 28]. As well as being an appropriate model for electronic computing scenarios such as sensor networks, they are a useful abstraction of “fast-mixing” physical systems such as animal populations [33], gene regulatory networks [9], and chemical reactions.

The latter application is especially germane: several recent wet-lab experiments demonstrate the systematic engineering of custom-designed chemical reactions [8, 16, 31, 34], unfortunately in all cases having a cost that scales linearly with the number of unique chemical species (states). (The cost can even be quadratic if certain error-tolerance mechanisms are employed [32].) Thus, it is imperative in implementing a molecular computational system to keep the number of distinct chemical species at a minimum. On the other hand, it is common (and relatively cheap) for the total number of such molecules (agents) to number in the trillions in a single test tube. It is thus important to understand the computational power enabled by a large number of agents n , where each agent has only a constant number of states (each agent is a finite state machine).

A population protocol is said to require a leader if every valid initial configuration contains a single agent in a special “leader” state that helps to coordinate the computation. Studying computation without a leader is important for understanding essentially distributed systems where symmetry breaking is difficult. Further, in the chemical setting obtaining single-molecule precision in the initial configuration is difficult. Thus, it would be highly desirable if the population protocol did not require an exquisitely tuned initial configuration.

1.1 Introduction to the model

A population protocol is defined by a finite set Λ of *states* that each agent may have, together with a *transition function*¹ $\delta : \Lambda^2 \rightarrow \Lambda^2$. A *configuration* is a nonzero vector $\mathbf{c} \in \mathbb{N}^\Lambda$ describing, for each $\bar{s} \in \Lambda$, the *count* $\mathbf{c}(\bar{s})$ of how many agents are in state \bar{s} . By convention we denote the number of agents by $n = \|\mathbf{c}\| = \sum_{\bar{s} \in \Lambda} \mathbf{c}(\bar{s})$. Given states $\bar{r}_1, \bar{r}_2, \bar{p}_1, \bar{p}_2 \in \Lambda$, if $\delta(\bar{r}_1, \bar{r}_2) = (\bar{p}_1, \bar{p}_2)$ (denoted $\bar{r}_1, \bar{r}_2 \rightarrow \bar{p}_1, \bar{p}_2$), and if a pair of agents in respective states \bar{r}_1 and \bar{r}_2 interact, then their states become \bar{p}_1 and \bar{p}_2 .² The next pair of agents to interact is chosen uniformly at random. The expected (parallel) time for any event to occur is the expected number of interactions, divided by the number of agents n . This measure of time is based on the natural parallel model where each agent participates in a constant number of interactions in one unit of time; hence $\Theta(n)$ total interactions are expected per unit time [6].

The most well-studied population protocol task is computing Boolean-valued *predicates*. It is known that a protocol *stably decides* a predicate $\phi : \mathbb{N}^k \rightarrow \{0, 1\}$ (meaning computes the correct answer with probability 1; see Section 4 for a formal definition) if [4] and only if [5] ϕ is semilinear.

Population protocols can also compute integer-valued *functions* $f : \mathbb{N}^k \rightarrow \mathbb{N}$. Suppose we start with $m \leq n/2$ agents in “input” state \bar{x} and the remaining agents in a “quiescent” state \bar{q} . Consider the protocol with a single transition rule $\bar{x}, \bar{q} \rightarrow \bar{y}, \bar{y}$. Eventually exactly $2m$ agents are in the “output” state \bar{y} , so this protocol computes the function $f(m) = 2m$. Furthermore (letting $\#\bar{s}$ = count of state \bar{s}), if $\#\bar{q} - 2m = \Omega(n)$ initially (e.g., $\#\bar{q} = 3m$), then it takes $\Theta(\log n)$ expected time until $\#\bar{y} = 2m$. Similarly, the transition rule $\bar{x}, \bar{x} \rightarrow \bar{y}, \bar{q}$ computes the function $f(m) = \lfloor m/2 \rfloor$, but exponentially slower, in expected time $\Theta(n)$. The transitions $\bar{x}_1, \bar{q} \rightarrow \bar{y}, \bar{q}$ and $\bar{x}_2, \bar{y} \rightarrow \bar{q}, \bar{q}$ compute

¹Some work allows nondeterministic transitions, in which the transition function maps to subsets of $\Lambda \times \Lambda$. Our results are independent of whether transitions are nondeterministic, and we choose a deterministic, symmetric transition function, rather than a more general relation $\delta \subseteq \Lambda^4$, merely for notational convenience.

²In the most generic model, there is no restriction on which agents are permitted to interact. If one prefers to think of the agents as existing on nodes of a graph, then it is the complete graph K_n for a population of n agents.

$f(m_1, m_2) = m_1 - m_2$ (assuming $m_1 \geq m_2$), also in time $\Theta(n)$ if $m_1 = m_2 + O(1)$.

Formally, we say a population protocol *stably computes* a function $f : \mathbb{N}^k \rightarrow \mathbb{N}$ if, for every “valid” initial configuration $\mathbf{i} \in \mathbb{N}^\Lambda$ representing input $\mathbf{m} \in \mathbb{N}^k$ (via counts $\mathbf{i}(\bar{x}_1), \dots, \mathbf{i}(\bar{x}_k)$ of “input” states $\Sigma = \{\bar{x}_1, \dots, \bar{x}_k\} \subseteq \Lambda$) with probability 1 the system reaches from \mathbf{i} to \mathbf{o} such that $\mathbf{o}(\bar{y}) = f(\mathbf{m})$ ($\bar{y} \in \Lambda$ is the “output” state) and $\mathbf{o}'(\bar{y}) = \mathbf{o}(\bar{y})$ for every \mathbf{o}' reachable from \mathbf{o} (i.e., \mathbf{o} is *stable*). Defining what constitutes a “valid” initial configuration (i.e., what non-input states can be present initially, and how many) is nontrivial. In this paper we focus on population protocols without a *leader*—a state present in count 1, or small count—in the initial configuration. Here, we equate “leaderless” with initial configurations in which no positive state count is sublinear in the population size n .

It is known that a function $f : \mathbb{N}^k \rightarrow \mathbb{N}$ is stably computable by a population protocol if and only if its *graph* $\{(\mathbf{m}, f(\mathbf{m})) \mid \mathbf{m} \in \mathbb{N}^k\} \subset \mathbb{N}^{k+1}$ is a semilinear set [5, 15]. This means intuitively that it is piecewise affine, with each affine piece having rational slopes.

Despite the exact characterization of predicates and functions stably computable by population protocols, we still lack a full understanding of which of the stably computable (i.e., semilinear) predicates and functions are computable *quickly* (say, in time polylogarithmic in n) and which are only computable slowly (linear in n). For positive results, much is known about time to *convergence* (time to get the correct answer). It has been known for over a decade that with an initial leader, any semilinear predicate can be stably computed with polylogarithmic convergence time [6]. Furthermore, it has recently been shown that all semilinear predicates can be computed *without* a leader with sublinear convergence time [27]. (See Section 1.4 for details.)

In this paper, however, we exclusively study time to *stabilization* without a leader (time after which the answer is *guaranteed to remain correct*). Except where explicitly marked otherwise with a variant of the word “converge”, all references to time in this paper refer to time until stabilization. Section 9 explains in more detail the distinction between the two.

1.2 Contributions

Undecidability of many predicates in sublinear time. Every semilinear predicate $\phi : \mathbb{N}^k \rightarrow \{0, 1\}$ is stably decidable in $O(n)$ time [6]. Some, such as $\phi(m_1, m_2) = 1$ iff $m_1 \geq 1$, are stably decidable in $O(\log n)$ time by a leaderless protocol, in this case by the transition $\bar{x}_1, \bar{x}_2 \rightarrow \bar{x}_1, \bar{x}_2$, where \bar{x}_1 “votes” for output 1 and \bar{x}_2 votes 0. A predicate is *eventually constant* if $\phi(\mathbf{m}) = \phi(\mathbf{m}')$ for all sufficiently large $\mathbf{m} \neq \mathbf{m}'$. We show in Theorem 4.4 that unless ϕ is eventually constant, any leaderless population protocol stably deciding a predicate ϕ requires at least linear time. Examples of non-eventually constant predicates include parity ($\phi(m) = 1$ iff m is odd), majority ($\phi(m_1, m_2) = 1$ iff $m_1 \geq m_2$), and equality ($\phi(m_1, m_2) = 1$ iff $m_1 = m_2$). It does *not* include certain semilinear predicates, such as $\phi(m_1, m_2) = 1$ iff $m_1 \geq 1$ (decidable in $O(\log n)$ time) or $\phi(m_1, m_2) = 1$ iff $m_1 \geq 2$ (decidable in $O(n)$ time, and no faster protocol is known).

Definition of function computation and approximation. We formally define computation and approximation of functions $f : \mathbb{N}^k \rightarrow \mathbb{N}$ for population protocols. This mode of computation was discussed briefly in the first population protocols paper [4, Section 3.4], which focused more on Boolean predicate computation, and it was defined formally in the more general model of chemical reaction networks [15, 19]. Some subtle issues arise that are unique to population protocols. We also formally define a notion of function *approximation* with population protocols.

Inapproximability of most linear functions with sublinear time and sublinear error. Recall that the transition rule $\bar{x}, \bar{x} \rightarrow \bar{y}, \bar{q}$ computes $f(m) = \lfloor m/2 \rfloor$ in linear time. Consider the

transitions $\bar{a}, \bar{x} \rightarrow \bar{b}, \bar{y}$ and $\bar{b}, \bar{x} \rightarrow \bar{a}, \bar{q}$, starting with $\#\bar{x} = m$, $\#\bar{a} = \gamma m$ for some $0 < \gamma < 1$, and $\#\bar{y} = \#\bar{q} = 0$ (so $n = m + \gamma m$ total agents). Then eventually $\#\bar{y} \in \{m/2, \dots, m/2 + \gamma m\}$ and $\#\bar{x} = 0$ (stabilizing $\#\bar{y}$), after $O(\frac{1}{\gamma} \log n)$ expected time. (This is analyzed in more detail in Section 7.) Thus, if we tolerate an error linear in n , then f can be approximated in logarithmic time. However, Theorem 6.1 shows this error bound to be tight: *any* leaderless population protocol that approximates $f(m) = \lfloor m/2 \rfloor$, or any other linear function with a coefficient outside of \mathbb{N} (such as $\lfloor 4m/3 \rfloor$ or $m_1 - m_2$), requires at least linear time to achieve sublinear error.

As a corollary, such functions cannot be stably computed in sublinear time (since computing exactly is the same as approximating with zero error). Conversely, it is simple to show that any linear function with all coefficients in \mathbb{N} is stably computable in logarithmic time (Observation 7.1). Thus we have a dichotomy theorem for the efficiency (with regard to stabilization) of computing linear functions f by leaderless population protocols: if all of f 's coefficients are in \mathbb{N} , then it is computable in logarithmic time, and otherwise it requires linear time.

Approximability of nonnegative rational-coefficient linear functions with logarithmic time and linear error. Theorem 6.1 says that no linear function with a coefficient outside of \mathbb{N} can be stably computed with sublinear time and sublinear error. In a complementary positive result, Theorem 7.2, by relaxing the error to linear, and restricting the coefficients to be *nonnegative* rationals (but not necessarily integers), we show how to approximate any such linear function in logarithmic time. (It is open if $m_1 - m_2$ can be approximated with linear error in logarithmic time.)

Uncomputability of many nonlinear functions in sublinear time. What about non-linear functions? Theorem 8.5 states that sublinear time computation cannot go much beyond linear functions with coefficients in \mathbb{N} : unless f is *eventually* \mathbb{N} -linear, meaning linear with nonnegative integer coefficients on all sufficiently large inputs, any protocol stably computing f requires at least linear time. Examples of non-eventually- \mathbb{N} -linear functions, that provably cannot be computed in sublinear time, include $f(m_1, m_2) = \min(m_1, m_2)$ (computable slowly via $\bar{x}_1, \bar{x}_2 \rightarrow \bar{y}, \bar{q}$), and $f(m) = m - 1$ (computable slowly via $\bar{x}, \bar{x} \rightarrow \bar{x}, \bar{y}$).

The only remaining semilinear functions whose asymptotic time complexity remains unknown are those “piecewise linear” functions that switch between pieces only near the boundary of \mathbb{N}^k ; for example, $f(m) = 0$ if $m \leq 3$ and $f(m) = m$ otherwise.

Note that there is a fundamental difficulty in extending the negative results to functions and predicates that “do something different only near the boundary of \mathbb{N}^k ”. This is because for inputs where one state is present in small count, the population protocol could in principle use that input as a “leader state”—and no longer be leaderless. However, this does not directly lead to a positive result for such inputs, because it is not obvious how to use (for instance) state \bar{x}_1 as a leader when its count is 1 while still maintaining correctness for larger counts of \bar{x}_1 .

Our results leave open the possibility that non-eventually constant predicates and non-eventually- \mathbb{N} -linear functions, which cannot be computed in sublinear time in our setting, *could* be efficiently computed in the following ways:

1. *With an initial leader* stabilizing to the correct answer in sublinear time,
2. Stabilizing to an output in expected sublinear time but *allowing a small probability of incorrect output* (with or without a leader), or
3. Without an initial leader but *converging* to the correct output in sublinear time. Recently Kosowski and Uznański [27] showed that this task is indeed possible. (See Section 1.4.)

1.3 Essential proof techniques

Techniques developed in previous work for proving time lower bounds [1, 20] can certainly generalize beyond leader election and majority, although it was not clear what precise category of computation they cover. However, to extend the impossibility results to all non-eventually- \mathbb{N} -linear functions, we needed to develop new tools.

Compared to our previous work showing the impossibility of sublinear time leader election [20], we achieve three main advances in proof technique. First, the previous machinery did not give us a way to affect *large*-count states predictably to change the answer, but rather focused on using surgery to remove a single leader state. Second, we need much additional reasoning to argue if a predicate is not eventually constant, then we can find infinitely many α -dense inputs that differ on their output but are close together. This leads to a contradiction when we use transition manipulation arguments to show how to absorb the small extra difference between the inputs without changing the output. Third, we need entirely different reasoning to argue that if a semilinear function is not eventually \mathbb{N} -linear, then we can find infinitely many α -dense inputs \mathbf{m} that do not appear “locally affine”: pushing a small distance \mathbf{v} from \mathbf{m} changes the function $f(\mathbf{m})$ to $f(\mathbf{m} + \mathbf{v}) = f(\mathbf{m}) + \epsilon$, but pushing by the same distance again changes it a different amount, i.e., $f(\mathbf{m} + 2\mathbf{v}) = f(\mathbf{m}) + \epsilon + \delta$, where $\epsilon \neq \delta$. This leads to a contradiction when we use transition manipulation arguments to show how, from input $\mathbf{m} + 2\mathbf{v}$, to stabilize the count of the output to the incorrect value $f(\mathbf{m}) + 2\epsilon$.³

Both in prior and current work, the high level intuition of the proof technique is as follows. The overall argument is a proof by contradiction: if sublinear time computation is possible, then we find a nefarious execution sequence that stabilizes to an incorrect output. In more detail, sublinear time computation requires avoiding “bottlenecks”—having to go through a transition in which both states are present in small count (constant independent of the number of agents n). Traversing even a single such transition requires linear time. Technical lemmas show that bottleneck-free execution sequences from α -dense initial configurations (i.e., where every state that is present is present in at least αn count) are amenable to predictable “surgery” [1, 20]. At the high level, the surgery lemmas show how states that are present in “low” count when the population protocol stabilizes, can be manipulated (added or removed) such that only “high” count other states are affected. Since it can also be shown that changing high count states in a stable configuration does not affect its stability, this means that the population protocol cannot “notice” the surgery, and remains stabilized to the previous output. For leader election, the surgery allows one to remove an additional leader state (leaving us with no leaders). For majority computation [1], the minority input must be present in low count (or absent) at the end. This allows one to add enough of the minority input to turn it into the majority, while the protocol continues to output the wrong answer.

However, applying the previously developed surgery lemmas to fool a *function* computing population protocol is more difficult. The surgery to consume additional input states affects the count of the output state, which could be present in “large count” at the end. How do we know that the effect of the surgery on the output is not consistent with the desired output of the function? In order to arrive at a contradiction we develop two new techniques, both of which are necessary to cover all cases. The first involves showing that the slope of the change in the count of the output state as a function of the input states is inconsistent. The second involves exposing the semilinear structure of the graph of the function being computed, and forcing it to enter the “wrong piece” (i.e., periodic coset).

³ These arguments are easier to understand for the special case when we can assume f is linear. Thus Section 6 concentrates on this special case, obtaining an exact characterization of the efficiently computable linear functions. Section 8 reasons about the more difficult case of arbitrary semilinear functions.

1.4 Related work

Positive results. Angluin, Aspnes, Diamadi, Fischer, and Peralta [4] showed that any semilinear predicate can be decided in expected parallel time $O(n \log n)$, later improved to $O(n)$ by Angluin, Aspnes, and Eisenstat [6]. More strikingly, the latter paper showed that if an initial leader is present (a state assigned to only a single agent in every valid initial configuration), then there is a protocol for ϕ that *converges* to the correct answer in expected time $O(\log^5 n)$. However, this protocol’s expected time to *stabilize* is still provably $\Omega(n)$. Section 9 explains this distinction in more detail. Chen, Doty, and Soloveichik [15] showed in the related model of chemical reaction networks (borrowing techniques from the related predicate results [4,5]) that any semilinear *function* (integer-output $f : \mathbb{N}^k \rightarrow \mathbb{N}$) can similarly be computed with expected convergence time $O(\log^5 n)$ if an initial leader is present, but again with much slower stabilization time $O(n \log n)$. Doty and Hajiaghayi [19] showed that any semilinear function can be computed by a chemical reaction network without a leader with expected convergence and stabilization time $O(n)$. Although the chemical reaction network model is more general, these results hold for population protocols.

Kosowski and Uznański [27] show that all semilinear predicates can be computed without an initial leader, converging in $O(\text{polylog } n)$ time if a small probability of error is allowed, and converging in $O(n^\epsilon)$ time with probability 1, where ϵ can be made arbitrarily close to 0 by changing the protocol. They also showed leader election protocols (which can be thought of as computing the constant function $f(m) = 1$) with the same properties.

Since efficient computation seems to be helped by a leader, the computational task of leader election has received significant recent attention. In particular, Alistarh and Gelashvili [3] showed that in a variant of the model allowing the number of states λ_n to grow with the population size n , a protocol with $\lambda_n = O(\log^3 n)$ states can elect a leader with high probability in $O(\log^3 n)$ expected time. Alistarh, Aspnes, Eisenstat, Gelashvili, and Rivest [1] later showed how to reduce the number of states to $\lambda_n = O(\log^2 n)$, at the cost of increasing the expected time to $O(\log^{5.3} n \log \log n)$. Gasieniec and Stachowiak [21] showed that there is a protocol with $O(\log \log n)$ states electing a leader in $O(\log^2 n)$ time in expectation and with high probability, recently improved to $O(\log n \log \log n)$ time by Gasieniec, Stachowiak, and Uznański [22]. This asymptotically matches the $\Omega(\log \log n)$ states provably required for sublinear time leader election (see negative results below).

Negative results. The first attempt to show the limitations of sublinear time population protocols, using the more general model of chemical reaction networks, was made by Chen, Cummings, Doty, and Soloveichik [14]. They studied a variant of the problem in which negative results are easier to prove, an “adversarial worst-case” notion of sublinear time: the protocol is required to be sublinear time not only from the initial configuration, but also from any reachable configuration. They showed that the predicates computable in this manner are precisely those whose output depends only on the presence or absence of states (and not on their exact positive counts). Doty and Soloveichik [20] showed the first $\Omega(n)$ lower bound on expected time from valid initial configurations, proving that any protocol electing a leader with probability 1 takes $\Omega(n)$ time.

These techniques were improved by Alistarh, Aspnes, Eisenstat, Gelashvili, and Rivest [1], who showed that even with up to $\lambda_n = O(\log \log n)$ states, any protocol electing a leader with probability 1 requires nearly linear time: $\Omega(n/\text{polylog } n)$. They used these tools to prove time lower bounds for another important computational task: majority (detecting whether state \bar{x}_1 or \bar{x}_2 is more numerous in the initial population, by stabilizing on a configuration in which the state with the larger initial count occupies the whole population). Alistarh, Aspnes, and Gelashvili [2] strengthened the state lower bound, showing that $\Omega(\log n)$ states are required to compute majority in $O(n^{1-c})$ time for some $c > 0$, when a certain “natural” condition is imposed on the protocol that

holds for all known protocols.

In contrast to these previous results on the specific tasks of leader election and majority, we obtain time lower bounds for a broad class of functions and predicates, showing “most” of those computable at all by population protocols, cannot be computed in sublinear time. Since they all *can* be computed in linear time, this settles their asymptotic population protocol time complexity.

Informally, one explanation for our result could be that some computation requires electing “leaders” as part of the computation, and other computation does not. Since leader election itself requires linear time as shown in [20], the computation that requires it is necessarily inefficient. It is not clear, however, how to define the notion of a predicate or function computation requiring electing a leader somewhere in the computation, but recent work by Michail and Spirakis helps to clarify the picture [29].

1.5 Organization of this paper

Section 2 defines population protocol model and notation. Section 3 proves the technical lemmas that are used in all the time lower bound proofs. Section 4 shows that a wide class of predicates requires $\Omega(n)$ time to compute. Section 5 explains our definitions of function computation and approximation. Section 6 shows that linear functions with either a negative (e.g., $m_1 - m_2$) or non-integer (e.g., $\lfloor m/2 \rfloor$) coefficient cannot be stably approximated with $o(n)$ error in $o(n)$ time. Section 7 shows our positive result, Theorem 7.2, that linear functions with all nonnegative rational coefficients (e.g., $\lfloor 2m_1/3 \rfloor + 3m_2$) *can* be stably approximated with $O(n)$ error in $O(\log n)$ time. Section 8 studies non-linear functions, showing that a large class of those computable by population protocols require $\Omega(n)$ time to compute. Section 9 states conclusions and open questions.

2 Preliminaries

If Λ is a finite set (in this paper, of *states*, which will be denoted as lowercase Roman letters with an overbar such as \bar{s}), we write \mathbb{N}^Λ to denote the set of functions $\mathbf{c} : \Lambda \rightarrow \mathbb{N}$. Equivalently, we view an element $\mathbf{c} \in \mathbb{N}^\Lambda$ as a vector of $|\Lambda|$ nonnegative integers, with each coordinate “labeled” by an element of Λ . (By assuming some canonical ordering $\bar{s}_1, \dots, \bar{s}_k$ of Λ , we also interpret $\mathbf{c} \in \mathbb{N}^\Lambda$ as a vector $\mathbf{c} \in \mathbb{N}^k$.) Given $\bar{s} \in \Lambda$ and $\mathbf{c} \in \mathbb{N}^\Lambda$, we refer to $\mathbf{c}(\bar{s})$ as the *count of \bar{s} in \mathbf{c}* . Let $\|\mathbf{c}\| = \|\mathbf{c}\|_1 = \sum_{\bar{s} \in \Lambda} \mathbf{c}(\bar{s})$. We write $\mathbf{c} \leq \mathbf{c}'$ to denote that $\mathbf{c}(\bar{s}) \leq \mathbf{c}'(\bar{s})$ for all $\bar{s} \in \Lambda$. Since we view vectors $\mathbf{c} \in \mathbb{N}^\Lambda$ equivalently as multisets of elements from Λ , if $\mathbf{c} \leq \mathbf{c}'$ we say \mathbf{c} is a *subset* of \mathbf{c}' . For $\alpha > 0$, we say that $\mathbf{c} \in \mathbb{N}^k$ is α -dense if, for all $i \in \{1, \dots, k\}$, if $\mathbf{c}(i) > 0$, then $\mathbf{c}(i) \geq \alpha \|\mathbf{c}\|$.

It is sometimes convenient to use multiset notation to denote vectors, e.g., $\{\bar{x}, \bar{x}, \bar{y}\}$ and $\{2\bar{x}, \bar{y}\}$ both denote the vector \mathbf{c} defined by $\mathbf{c}(\bar{x}) = 2$, $\mathbf{c}(\bar{y}) = 1$, and $\mathbf{c}(\bar{s}) = 0$ for all $\bar{s} \notin \{\bar{x}, \bar{y}\}$. Given $\mathbf{c}, \mathbf{c}' \in \mathbb{N}^\Lambda$, we define the vector component-wise operations of addition $\mathbf{c} + \mathbf{c}'$, subtraction $\mathbf{c} - \mathbf{c}'$, and scalar multiplication $m\mathbf{c}$ for $m \in \mathbb{N}$. For a set $\Delta \subset \Lambda$, we view a vector $\mathbf{c} \in \mathbb{N}^\Lambda$ equivalently as a vector $\mathbf{c} \in \mathbb{N}^\Delta$ by assuming $\mathbf{c}(\bar{s}) = 0$ for all $\bar{s} \in \Lambda \setminus \Delta$. Write $\mathbf{c} \upharpoonright \Delta$ to denote the vector $\mathbf{d} \in \mathbb{N}^\Delta$ such that $\mathbf{c}(\bar{s}) = \mathbf{d}(\bar{s})$ for all $\bar{s} \in \Delta$. For any vector or matrix \mathbf{c} , let $\text{amax}(\mathbf{c})$ denote the largest absolute value of any component of \mathbf{c} . Also, given $b \in \mathbb{N}$ and $\mathbf{m} \in \mathbb{N}^k$, $\max(b, \mathbf{m})$ is a shorthand for $\max\left(\{b\} \cup \bigcup_{i=1}^k \{\mathbf{m}(i)\}\right)$, and similar for $\text{amax}(b, \mathbf{m})$.

In this paper, the floor function $\lfloor \cdot \rfloor : \mathbb{R} \rightarrow \mathbb{Z}$ is defined to be the integer *closest to 0* that is distance < 1 from the input, e.g., $\lfloor -3.4 \rfloor = -3$ and $\lfloor 3.4 \rfloor = 3$. For an (infinite) set/sequence of configurations C , let $\text{bdd}(C) = \{\bar{s} \in \Lambda \mid (\exists b \in \mathbb{N})(\forall \mathbf{c} \in C) \mathbf{c}(\bar{s}) < b\}$ be the set of states whose counts are bounded by a constant in C . Let $\text{unbdd}(C) = \Lambda \setminus \text{bdd}(C)$. For $m \in \mathbb{N}$, let

$\mathbb{N}_{\geq m_0}^k = \{\mathbf{m} \in \mathbb{N}^k \mid (\forall i \in \{1, \dots, k\}) \mathbf{m}(i) \geq m_0\}$, denote the set of vectors in which each coordinate is at least m_0 .

2.1 Population Protocols

A *population protocol* is a pair $\mathcal{P} = (\Lambda, \delta)$, where Λ is a finite set of *states* and $\delta : \Lambda^2 \rightarrow \Lambda^2$ is the (symmetric) *transition function*. A *configuration* of a population protocol is a vector $\mathbf{c} \in \mathbb{N}^\Lambda$, with the interpretation that $\mathbf{c}(\bar{s})$ agents are in state $\bar{s} \in \Lambda$. If there is some “current” configuration \mathbf{c} understood from context, we write $\#\bar{s}$ to denote $\mathbf{c}(\bar{s})$. By convention, the value $n \in \mathbb{Z}_{\geq 1}$ represents the total number of agents $\|\mathbf{c}\|$. A *transition* is a 4-tuple $\tau = (\bar{r}_1, \bar{r}_2, \bar{p}_1, \bar{p}_2) \in \Lambda^4$, written $\tau : \bar{r}_1, \bar{r}_2 \rightarrow \bar{p}_1, \bar{p}_2$, such that $\delta(\bar{r}_1, \bar{r}_2) = (\bar{p}_1, \bar{p}_2)$. If an agent in state \bar{r}_1 interacts with an agent in state \bar{r}_2 , then they change states to \bar{p}_1 and \bar{p}_2 . This paper typically defines a protocol by a list of transitions, with δ implicit. There is a *null* transition $\delta(\bar{r}_1, \bar{r}_2) = (\bar{r}_1, \bar{r}_2)$ if a different output for $\delta(\bar{r}_1, \bar{r}_2)$ is not specified.

Given $\mathbf{c} \in \mathbb{N}^\Lambda$ and transition $\tau : \bar{r}_1, \bar{r}_2 \rightarrow \bar{p}_1, \bar{p}_2$, we say that τ is *applicable* to \mathbf{c} if $\mathbf{c} \geq \{\bar{r}_1, \bar{r}_2\}$, i.e., \mathbf{c} contains 2 agents, one in state \bar{r}_1 and one in state \bar{r}_2 . If τ is applicable to \mathbf{c} , then write $\tau(\mathbf{c})$ to denote the configuration $\mathbf{c} - \{\bar{r}_1, \bar{r}_2\} + \{\bar{p}_1, \bar{p}_2\}$ (i.e., that results from applying τ to \mathbf{c}); otherwise $\tau(\mathbf{c})$ is undefined. A finite or infinite sequence of transitions (τ_i) is a *transition sequence*. Given a $\mathbf{c}_0 \in \mathbb{N}^\Lambda$ and a transition sequence (τ_i) , the induced *execution sequence* (or *path*) is a finite or infinite sequence of configurations $(\mathbf{c}_0, \mathbf{c}_1, \dots)$ such that, for all $i \geq 1$, $\mathbf{c}_i = \tau_{i-1}(\mathbf{c}_{i-1})$.⁴ If a finite execution sequence, with associated transition sequence \bar{q} , starts with \mathbf{c} and ends with \mathbf{c}' , we write $\mathbf{c} \Longrightarrow_q \mathbf{c}'$. We write $\mathbf{c} \Longrightarrow_{\mathcal{P}} \mathbf{c}'$ (or $\mathbf{c} \Longrightarrow \mathbf{c}'$ when \mathcal{P} is clear from context) if such a path exists (i.e., it is possible to reach from \mathbf{c} to \mathbf{c}') and we say that \mathbf{c}' is *reachable* from \mathbf{c} . Let $\text{post}_{\mathcal{P}}(\mathbf{c}) = \{\mathbf{c}' \mid \mathbf{c} \Longrightarrow_{\mathcal{P}} \mathbf{c}'\}$ to denote the set of all configurations reachable from \mathbf{c} , writing $\text{post}(\mathbf{c})$ when \mathcal{P} is clear from context. If it is understood from context what is the initial configuration \mathbf{i} , then say \mathbf{c} is simply *reachable* if $\mathbf{i} \Longrightarrow \mathbf{c}$. If a transition $\tau : \bar{r}_1, \bar{r}_2 \rightarrow \bar{p}_1, \bar{p}_2$ has the property that for $i \in \{1, 2\}$, $\bar{r}_i \notin \{\bar{p}_1, \bar{p}_2\}$, or if $(\bar{r}_1 = \bar{r}_2 \text{ and } (\bar{r}_i \neq \bar{p}_1 \text{ or } \bar{r}_i \neq \bar{p}_2))$, then we say that τ *consumes* \bar{r}_i ; i.e., applying τ reduces the count of \bar{r}_i . We say τ *produces* \bar{p}_i if it increases the count of \bar{p}_i .

2.2 Time Complexity

The model used to analyze time complexity is a discrete-time Markov process, whose states correspond to configurations of the population protocol. In any configuration the next interaction is chosen by selecting a pair of agents uniformly at random and applying transition function δ to determine the next configuration. Since a transition may be null, self-loops are allowed. To measure time we count the expected total number of interactions (including null), and divide by the number of agents n . (In the population protocols literature, this is often called “parallel time”; i.e. n interactions among a population of n agents corresponds to one unit of time). Let $\mathbf{c} \in \mathbb{N}^\Lambda$ and $C \subseteq \mathbb{N}^\Lambda$. Denote the probability that the protocol reaches from \mathbf{c} to some configuration $\mathbf{c}' \in C$ by $\Pr[\mathbf{c} \Longrightarrow C]$. If $\Pr[\mathbf{c} \Longrightarrow C] = 1$,⁵ define the *expected time to reach from \mathbf{c} to C* , denoted $\mathbb{T}[\mathbf{c} \Longrightarrow C]$, to be the expected number of interactions to reach from \mathbf{c} to some $\mathbf{c}' \in C$, divided by the number of agents $n = \|\mathbf{c}\|$. If $\Pr[\mathbf{c} \Longrightarrow C] < 1$ then $\mathbb{T}[\mathbf{c} \Longrightarrow C] = \infty$.

⁴When the initial configuration to which a transition sequence is applied is clear from context, we may overload terminology and refer to a transition sequence and an execution sequence interchangeably.

⁵Since population protocols have a finite reachable configuration space, this is equivalent to requiring that for all $\mathbf{x} \in \text{post}(\mathbf{c})$, there is a $\mathbf{c}' \in C \cap \text{post}(\mathbf{x})$.

3 Technical tools

In this section we explain some technical results that are used in proving the time lower bounds of Theorems 4.4, 6.3, 6.4, 8.4, and 8.5. In some cases the main ideas are present in previous papers, but several must be adapted significantly to the current problem. Throughout Section 3, let $\mathcal{P} = (\Lambda, \delta)$ be a population protocol.

Although other results from this section are used in this paper, the key technical result of this section is Corollary 3.11. It gives a generic method to start with an initial configuration \mathbf{i} reaching in sublinear time to a configuration \mathbf{o} (in all our uses \mathbf{o} is a stable configuration, but this is not required by the corollary), and starting from *two* copies of \mathbf{i} , to manipulate the transitions leading from $2\mathbf{i}$ to $2\mathbf{o}$ while having a predictable effect on the counts of certain states, possibly also starting with a “small” number of extra states, denoted \mathbf{d}^Δ in Corollary 3.11. This leads to a contradiction when the effect on the counts of the states representing the output can be shown to be incorrect for the given input $2\mathbf{i} + \mathbf{d}^\Delta$.

We often deal with infinite sequences of configurations.⁶ The following lemma, used frequently in reasoning about population protocols, shows that we can always take a nondecreasing subsequence.

Lemma 3.1 (Dickson’s Lemma [17]). *Any infinite sequence $\mathbf{c}_0, \mathbf{c}_1, \dots \in \mathbb{N}^k$ has an infinite nondecreasing subsequence $\mathbf{c}_{i_0} \leq \mathbf{c}_{i_1} \leq \dots$, where $i_0 < i_1 < \dots \in \mathbb{N}$.*

3.1 Bottleneck transitions take linear time

Let $b \in \mathbb{N}$. A transition $\bar{r}_1, \bar{r}_2 \rightarrow \bar{p}_1, \bar{p}_2$ is a *b-bottleneck* for configuration \mathbf{c} if $\mathbf{c}(\bar{r}_1) \leq b$ and $\mathbf{c}(\bar{r}_2) \leq b$.

The next observation, proved in [20], states that, if to get from a configuration $\mathbf{x} \in \mathbb{N}^\Lambda$ to some configuration in a set $S \subseteq \mathbb{N}^\Lambda$, it is necessary to execute a transition $\bar{r}_1, \bar{r}_2 \rightarrow \bar{p}_1, \bar{p}_2$ in which the counts of \bar{r}_1 and \bar{r}_2 are both at most some number b , then the expected time to reach from \mathbf{x} to some configuration in S is $\Omega(n/b^2)$.

Observation 3.2 ([20]). *Let $b \in \mathbb{N}$, $\mathbf{x} \in \mathbb{N}^\Lambda$, and $S \subseteq \mathbb{N}^\Lambda$ such that $\Pr[\mathbf{x} \Longrightarrow S] = 1$. If every path taking \mathbf{x} to a configuration $\mathbf{o} \in S$ has a *b-bottleneck*, then $\mathbb{T}[\mathbf{x} \Longrightarrow S] \geq \frac{n-1}{2(b|\Lambda|)^2} > \frac{n}{3b^2|\Lambda|^2}$.*

The next corollary is useful.

Observation 3.3 ([20]). *Let $\gamma > 0$, $b \in \mathbb{N}$, $\mathbf{c} \in \mathbb{N}^\Lambda$, and $X, S \subseteq \mathbb{N}^\Lambda$ such that $\Pr[\mathbf{c} \Longrightarrow X] \geq \gamma$, $\Pr[\mathbf{c} \Longrightarrow S] = 1$, and every path from every $\mathbf{x} \in X$ to some $\mathbf{o} \in S$ has a *b-bottleneck*. Then $\mathbb{T}[\mathbf{c} \Longrightarrow S] > \gamma \frac{n}{3b^2|\Lambda|^2}$.*

3.2 Transition ordering lemma

The following lemma was originally proved in [14] and was restated in the language of population protocols as Lemma 4.5 in [20]. Intuitively, the lemma states that a “fast” transition sequence (meaning one without a bottleneck transition) that decreases certain states from large counts to small counts must contain transitions of a certain restricted form. In particular the form is as follows: if Δ is the set of states whose counts decrease from large to small, then we can write the states in Δ in some order $\bar{d}_1, \bar{d}_2, \dots, \bar{d}_d$, such that for each $1 \leq i \leq d$, there is a transition τ_i that consumes \bar{d}_i , and every other state involved in τ_i is either not in Δ , or comes later in the ordering. These transitions will later be used to do controlled “surgery” on fast transition sequences, because

⁶In general these will not be *execution* sequences. Typically none of the configurations are reachable from any others because they are configurations with increasing numbers of agents.

they give a way to alter the count of \bar{d}_i , by inserting or removing the transitions τ_i , knowing that this will not affect the counts of $\bar{d}_1, \dots, \bar{d}_{i-1}$.

Let $\Delta \subseteq \Lambda$, with $d = |\Delta|$. We say that \mathcal{P} is Δ -ordered (via τ_1, \dots, τ_d) if there is an order on Δ , so that we may write $\Delta = \{\bar{d}_1, \dots, \bar{d}_d\}$, such that, for all $i \in \{1, \dots, d\}$, there is a transition $\tau_i : \bar{d}_i, \bar{s}_i \rightarrow \bar{o}_i, \bar{o}'_i$, such that $\bar{s}_i, \bar{o}_i, \bar{o}'_i \notin \{\bar{d}_1, \dots, \bar{d}_i\}$. In other words, for each i there is a transition consuming exactly one \bar{d}_i without affecting $\bar{d}_1, \dots, \bar{d}_{i-1}$.

Lemma 3.4 (adapted from [14]). *Let $b_1, b_2 \in \mathbb{N}$ such that $b_2 > |\Lambda| \cdot b_1$. Let $\mathbf{x}, \mathbf{o} \in \mathbb{N}^\Lambda$ such that $\mathbf{x} \xrightarrow{p} \mathbf{o}$ via a path p without a b_2 -bottleneck. Define $\Delta = \{\bar{d} \in \Lambda \mid \mathbf{x}(\bar{d}) \geq b_2 \text{ and } \mathbf{o}(\bar{d}) \leq b_1\}$. Then \mathcal{P} is Δ -ordered via τ_1, \dots, τ_d , and each τ_i occurs at least $(b_2 - |\Lambda| \cdot b_1)/|\Lambda|^2$ times in p .*

3.3 Sublinear time from dense configuration implies bottleneck free path from dense configuration with every state present

Say that $\mathbf{c} \in \mathbb{N}^\Lambda$ is *full* if $(\forall \bar{s} \in \Lambda) \mathbf{c}(\bar{s}) > 0$, i.e., every state is present. The following theorem states that with high probability, a population protocol will reach from an α -dense configuration to a configuration in which all states are present (full) in “large” count (β -dense, for some $0 < \beta < \alpha$).⁷ It was proven in [18] in the more general model of chemical reaction networks, for a subclass of such networks that includes all population protocols.

Theorem 3.1 (adapted from [18]). *Let $\alpha > 0$. Then there are constants $\epsilon, \beta > 0$ such that, letting $X_\beta = \{\mathbf{x} \in \mathbb{N}^\Lambda \mid \mathbf{x} \text{ is full and } \beta\text{-dense}\}$, for all sufficiently large α -dense $\mathbf{i} \in \mathbb{N}^\Lambda$, $\Pr[\mathbf{i} \Longrightarrow X_\beta] \geq 1 - 2^{-\epsilon \|\mathbf{i}\|}$.*

The following was originally proved as Lemma 4.4 in [20]. The result was stated with S being the set of what was called “ Q -stable configs,” but we have adapted it to make the statement more general and quantitatively relate the bound $b(n)$ to the expected time $t(n)$. It states that if a protocol goes from an α -dense configuration to a set of states S in expected time $\leq t(n)$, then there is a full β -dense (for $0 < \beta < \alpha$) reachable configuration \mathbf{x} and a path from \mathbf{x} to a state in S with no $b(n)$ -bottleneck transition, where $b(n) = O\left(\sqrt{\frac{n}{t(n)}}\right)$. If $t(n) = o(n)$, then $b(n) = \omega(1)$, which suffices for our subsequent results.

Lemma 3.5 (adapted from [20]). *For all $\alpha > 0$, there is a $\beta > 0$ such that the following holds. Suppose that for some $t : \mathbb{N} \rightarrow \mathbb{N}$, some set $S \subseteq \mathbb{N}^\Lambda$ and some set of α -dense initial configurations I , for all $\mathbf{i} \in I$, $\mathbb{T}[\mathbf{i} \Longrightarrow S] \leq t(n)$. Define $b(n) = \frac{1}{|\Lambda|} \sqrt{\frac{n}{6t(n)}}$. There is an $n_0 \in \mathbb{N}$ such that for all $\mathbf{i} \in I$ with $\|\mathbf{i}\| = n \geq n_0$, there is $\mathbf{x} \in \text{post}(\mathbf{i})$ and path p such that:*

1. $\mathbf{x}(\bar{s}) \geq \beta n$ for all $\bar{s} \in \Lambda$,
2. $\mathbf{x} \xrightarrow{p} \mathbf{o}$, where $\mathbf{o} \in S$, and
3. p has no $b(n)$ -bottleneck transition.

Proof. Intuitively, the lemma follows from the fact that state \mathbf{x} is reached with high probability by Theorem 3.1, and if no paths such as p existed, then all paths from \mathbf{x} to a stable configuration would have a bottleneck and require more than the stated time by Observation 3.3. Since \mathbf{x} is reached with high probability, this would imply the entire expected time is linear.

⁷With the same probability, this happens in time $O(1)$, although this fact is not needed in this paper.

For any configuration \mathbf{x} reachable from some configuration $\mathbf{i} \in I$, there is a transition sequence p satisfying condition (2) by the fact that $\Pr[\mathbf{i} \implies S] = 1$. It remains to show we can find \mathbf{x} and p satisfying conditions (1) and (3).

By Theorem 3.1 there exist ϵ, β (which depend only on \mathcal{A} and α) such that, starting in any sufficiently large initial configuration \mathbf{i} , with probability at least $1 - 2^{-\epsilon n}$, \mathcal{A} reaches a configuration \mathbf{x} where all states $\bar{s} \in \Lambda$ have count at least βn , where $n = \|\mathbf{i}\|$. For all \mathbf{i} , let $X_{\mathbf{i}} = \text{post}(\mathbf{i}) \cap X_{\beta} = \{\mathbf{x} \mid \mathbf{i} \implies \mathbf{x} \text{ and } (\forall \bar{s} \in \Lambda) \mathbf{x}(\bar{s}) \geq \beta \|\mathbf{i}\|\}$. Let n_0 be a lower bound on n such that Theorem 3.1 applies for all $n \geq n_0$ and $1 - 2^{-\epsilon n_0} \geq \frac{1}{2}$. Then for all $\mathbf{i} \in I$ such that $\|\mathbf{i}\| = n \geq n_0$, $\Pr[\mathbf{i} \implies X_{\mathbf{i}}] \geq \frac{1}{2}$. Choose any $n \geq n_0$ for which there is $\mathbf{i} \in I$ with $\|\mathbf{i}\| = n$. Then any $\mathbf{x} \in X_{\mathbf{i}}$ satisfies condition (1): $\mathbf{x}(\bar{s}) \geq \beta n$ for all $\bar{s} \in \Lambda$. We now show that by choosing \mathbf{x} from $X_{\mathbf{i}}$ for a large enough n , we can find a corresponding p satisfying condition (3) as well.

Suppose for the sake of contradiction that, we cannot satisfy condition (3) when choosing \mathbf{x} as above, no matter how large we make n . This means that for infinitely many $\mathbf{i} \in I$, (and therefore infinitely many population sizes $n = \|\mathbf{i}\|$), all transition sequences from $X_{\mathbf{i}}$ to S have a $b(n)$ -bottleneck. Applying Observation 3.3, letting $\mathbf{c} = \mathbf{i}$, $\gamma = \frac{1}{2}$, $X = X_{\mathbf{i}}$, tells us that $t(n) = \mathbb{T}[\mathbf{i} \implies S] > \frac{1}{2} \frac{n}{3b(n)^2 |\Lambda|^2}$, so $b(n) > \frac{1}{|\Lambda|} \sqrt{\frac{n}{6t(n)}}$, a contradiction. \square

In the following lemma, note that the indexing is over a subset $N \subseteq \mathbb{N}$; for example, the sequence might be indexed $\mathbf{i}_3, \mathbf{i}_6, \mathbf{i}_7, \mathbf{i}_{12}, \dots$ if $N = \{3, 6, 7, 12, \dots\}$, allowing us to retain the convention that the population size $\|\mathbf{i}_n\|$ is represented by n . Lemma 3.6 essentially states that, if infinitely many configurations \mathbf{i} satisfy the hypothesis of Lemma 3.5, then we can find three infinite sequences satisfying the conclusion of Lemma 3.5: initial configurations \mathbf{i}_n , intermediate full configurations \mathbf{x}_n , and ‘‘final’’ configurations \mathbf{o}_n (in our applications all \mathbf{o}_n will be stable), which by Dickson’s lemma can all be assumed nondecreasing.

Lemma 3.6. *For all $\alpha > 0$, there is a $\beta > 0$ such that the following holds. Suppose that for some set $S \subseteq \mathbb{N}^{\Lambda}$ and infinite set of α -dense initial configurations I , for all $\mathbf{i} \in I$, $\mathbb{T}[\mathbf{i} \implies S] \leq t(n)$. Define $b(n) = \frac{1}{4|\Lambda|} \sqrt{\frac{n}{t(n)}}$. There is an infinite set $N \subseteq \mathbb{N}$ and infinite sequences of configurations $(\mathbf{i}_n \in I)_{n \in N}$, $(\mathbf{x}_n \in \mathbb{N}^{\Lambda})_{n \in N}$, $(\mathbf{o}_n \in \mathbb{N}^{\Lambda})_{n \in N}$, where $(\mathbf{x}_n)_{n \in N}$ and $(\mathbf{o}_n)_{n \in N}$ are nondecreasing, and an infinite sequence of paths $(p_n)_{n \in N}$ such that, for all $n \in N$,*

1. $\|\mathbf{i}_n\| = \|\mathbf{x}_n\| = \|\mathbf{o}_n\| = n$,
2. $\mathbf{i}_n \implies \mathbf{x}_n$,
3. $\mathbf{x}_n(\bar{s}) \geq \beta n$ for all $\bar{s} \in \Lambda$,
4. $\mathbf{x}_n \implies_{p_n} \mathbf{o}_n$, where $\mathbf{o}_n \in S$, and
5. p_n has no $b(n)$ -bottleneck transition.

Proof. Since I is infinite, the set $I_{n_0} = \{\mathbf{i} \in I \mid \|\mathbf{i}\| \geq n_0\}$ is infinite. Pick an infinite sequence (\mathbf{i}_n) from I_{n_0} , where $\|\mathbf{i}_n\| = n$ (n may range over a subset of \mathbb{N} here, but for each $n \in \mathbb{N}$, at most one configuration in the sequence has size n). For each \mathbf{i}_n in the sequence, pick \mathbf{x}_n , p_n and \mathbf{o}_n for \mathbf{i}_n as in Lemma 3.5. By Dickson’s Lemma (Lemma 3.1) there is an infinite subset $N \subseteq \mathbb{N}$ such that $(\mathbf{x}_n)_{n \in N}$ and $(\mathbf{o}_n)_{n \in N}$ are nondecreasing on the respective subsequences of $(\mathbf{x}_n)_{n \in \mathbb{N}}$ and $(\mathbf{o}_n)_{n \in \mathbb{N}}$ corresponding to N . Lemma 3.5 ensures that properties (1)-(5) are satisfied. \square

The conclusion of Lemma 3.6, with its various infinite sequences, is quite complex. The hypothesis of Lemma 3.9 is equally complex; they are used in tandem to prove Lemma 3.10 and

Corollary 3.11, the latter being our main technical tool for proving the time lower bounds of Theorems 4.4, 6.3, 6.4, 8.4, and 8.5.

The idea of Lemma 3.10 is to start with a protocol satisfying the hypothesis of Lemma 3.6, which reaches in sublinear time from some set I of α -dense initial configurations to some set S (in all applications, S is the set of stable configurations reachable from I). Then, invoke Lemma 3.9 to show that it is possible from certain initial configurations to drive some states in the set $\Delta = \text{bdd}(\mathbf{o}_n)$ to 0.

The reason that the statement of Lemma 3.10 is also fairly complex, and references some of these infinite sequences, is that the set Δ appearing in the conclusion of Lemma 3.9 depends on the particular infinite sequence (\mathbf{o}_n) defined in the conclusion of Lemma 3.6. Several infinite sequences, each with their own Δ , could satisfy the hypothesis of Lemma 3.9, and it matters which one we pick. Thus, in applying these results, before reaching the conclusion of Lemma 3.9, we must explicitly define these infinite sequences to know the particular Δ to which the conclusion of Lemma 3.9 applies.

3.4 Path manipulation

This is the most technically dense subsection, with many intermediate technical lemmas that culminate in our primary technical tool for proving time lower bounds, Corollary 3.11. Each lemma statement is complex and involves many interacting variables. The first three lemmas are accompanied by an example and figures to help trace through the intuition.

The next two lemmas, Lemmas 3.7 and 3.8, apply to population protocols that have transitions as described in Lemma 3.4. Both use these transitions in order to manipulate a configuration (by manipulating a “fast” path leading to it from another configuration) until it has prescribed counts of states in Δ from Lemma 3.4.

Lemmas 3.7 and 3.8 are based on statements first proven as “Claim 1” and “Claim 2” in [20]. Since their statements in that paper were not self-contained (being claims as part of a larger proof), we have rephrased them as self-contained Lemmas 3.7 and 3.8, and we give self-contained proofs. Furthermore, we have significantly adapted both the statements and proofs to make them more generally useful for proving negative results, in particular stating the minimum conditions required to apply the lemmas, in addition to quantitatively accounting for the precise effect that the path manipulation has on the underlying configurations.

We use linear algebra to describe changes in counts of states. It is beneficial to fix some notational conventions first. Recall Λ is the set of all states, $\Delta \subset \Lambda$ where $\Delta = \{\bar{d}_1, \dots, \bar{d}_d\}$, and $\Gamma = \Lambda \setminus \Delta$ where $\Gamma = \{\bar{g}_1, \dots, \bar{g}_g\}$. A matrix $\tilde{\mathbf{C}} \in \mathbb{Z}^{\Gamma \times \Delta}$ is an integer-valued matrix with g rows and d columns, with row j corresponding to state \bar{g}_j and column i corresponding to state \bar{d}_i . Given a vector $\mathbf{c}^\Delta \in \mathbb{N}^\Delta$ representing counts of states in Δ , then $\tilde{\mathbf{C}} \cdot \mathbf{c}^\Delta = \mathbf{d}^\Gamma$ is a vector $\mathbf{d}^\Gamma \in \mathbb{Z}^\Gamma$ representing changes in counts of states in Γ .

Our notation for indexing these matrices will generally follow our usual vector convention of using the name of the state itself, rather than an integer index, so for example, $\tilde{\mathbf{C}}(\bar{g}_i, \bar{d}_j)$ refers to the entry in the column corresponding to \bar{d}_j and the row corresponding to \bar{g}_i . If necessary to identify the position, this will correspond to the i 'th row and j 'th column. Where convenient, we also use the traditional notation $\tilde{\mathbf{C}}(i, j)$ as well: for instance, a protocol being Δ -ordered implies a 1-1 correspondence between transitions τ_1, \dots, τ_d and $\Delta = \{\bar{d}_1, \dots, \bar{d}_d\}$, which can both be indexed by $i \in \{1, \dots, d\}$.

Similarly, when convenient we will abuse notation and consider a vector $\mathbf{v} \in \mathbb{N}^k$, for a predicate or function with k inputs, to equivalently represent a configuration or subconfiguration in \mathbb{N}^Σ , where $\Sigma \subseteq \Lambda$ is the set of k input states of the population protocol.

The next lemma says that for any amount $\mathbf{c}^\Delta \in \mathbb{N}^\Delta$ of states in Δ , there exists an amount $\mathbf{e} \in \mathbb{N}^\Lambda$ of states that, if present in addition to \mathbf{c}^Δ , can be used to remove \mathbf{c}^Δ and $\mathbf{e} \upharpoonright \Delta$ (the states of \mathbf{e} that are in Δ), resulting in a configuration $\mathbf{z}^\Gamma \in \mathbb{N}^\Gamma$ with no states in Δ . Furthermore, both \mathbf{e} and \mathbf{z}^Γ are linear functions of \mathbf{c}^Δ .

So when we employ Lemma 3.7 later, where will these extra agents \mathbf{e} come from? Although we talk about them as if they are somehow physically added, in actuality, we'll start with a larger initial configuration and “guide” some of the agents to the desired states that make up \mathbf{e} ; this is the work of Lemma 3.8.

Lemma 3.7 (adapted from [20]). *Let $\Delta \subseteq \Lambda$ such that \mathcal{P} is Δ -ordered, $d = |\Delta|$, and let $\Gamma = \Lambda \setminus \Delta$. Then there are matrices $\mathbf{C}_1 \in \mathbb{N}^{\Gamma \times \Delta}$ and $\mathbf{C}_2 \in \mathbb{N}^{\Lambda \times \Delta}$, with $\max(\mathbf{C}_1) < 2^{d+1}$, $\max(\mathbf{C}_2) < 2^d$, such that, for all $\mathbf{c}^\Delta \in \mathbb{N}^\Delta$, setting $\mathbf{e} = \mathbf{C}_2 \cdot \mathbf{c}^\Delta \in \mathbb{N}^\Lambda$ and $\mathbf{z}^\Gamma = \mathbf{C}_1 \cdot \mathbf{c}^\Delta \in \mathbb{N}^\Gamma$, then $\mathbf{c}^\Delta + \mathbf{e} \implies \mathbf{z}^\Gamma$.*

Proof. Intuitively, the proof works as follows. Since \mathcal{P} is Δ -ordered, for each $i \in \{1, \dots, d\}$, there is a transition $\tau_i : \bar{d}_i, \bar{s}_i \rightarrow \bar{o}_i, \bar{o}'_i$, such that for all i , $\bar{d}_i \in \Delta$ and $\bar{o}_i, \bar{o}'_i \notin \{\bar{d}_1, \dots, \bar{d}_i\}$. We will construct the path p such that $\mathbf{c}^\Delta + \mathbf{e} \implies_p \mathbf{z}$ as follows. A naïve approach would simply consume states in \mathbf{c}^Δ , by adding $\mathbf{c}^\Delta(i)$ copies of τ_i to the path p , and $\mathbf{c}^\Delta(i)$ copies of the other input \bar{s}_i to \mathbf{e} . Since \mathcal{P} is Δ -ordered this would indeed result in count 0 of \bar{d}_1 . However, although for $i \in \{2, \dots, d\}$, this consumes $\mathbf{c}^\Delta(\bar{d}_i)$ copies of \bar{d}_i , it might produce additional copies of \bar{d}_i if it appears as an output state of some transitions $\tau_1, \dots, \tau_{i-1}$ that were added. Let \mathbf{c}_1^Δ denote these counts. Since $\mathbf{c}_1^\Delta(\bar{d}_1) = 0$, we won't need to add any more τ_1 . Repeat the naïve approach a second time to consume \mathbf{c}_1^Δ , which will result in \mathbf{c}_2^Δ , where $\mathbf{c}_2^\Delta(\bar{d}_1) = \mathbf{c}_2^\Delta(\bar{d}_2) = 0$. Repeating this d times consumes all of Δ .

We now formally define matrices that will help to account for the exact changes in state counts that result from executing this path. First, we define a matrix $\mathbf{T} \in \mathbb{N}^{d \times d}$. Intuitively, if $\mathbf{c}^\Delta \in \mathbb{N}^\Delta$ represents counts of states in Δ , then the vector $\mathbf{t} \in \mathbb{N}^d$ defined by $\mathbf{t} = \mathbf{T} \cdot \mathbf{c}^\Delta$ represents counts of transitions τ_1, \dots, τ_d in the path p such that $\mathbf{c}^\Delta + \mathbf{e} \implies_p \mathbf{z}^\Gamma$. In particular, $\mathbf{t}(\bar{d}_i)$ will represent the total number of transitions τ_i that we add to path p , in order to consume all copies of \bar{d}_i , not only the $\mathbf{c}^\Delta(\bar{d}_i)$ present initially, but also any added because of transitions τ_j for $j < i$ appearing previously in p , if one of the outputs of τ_j is \bar{d}_i .

Define the $d \times d$ matrix \mathbf{T}_1 as follows. Intuitively, \mathbf{T}_1 is a matrix such that, if $\mathbf{t} \in \mathbb{N}^d$ represents “counts of transition executions”, i.e., $\mathbf{t}(j)$ means “execute transition τ_j $\mathbf{t}(j)$ times”, then $\mathbf{T}_1 \cdot \mathbf{t} \in \mathbb{N}^d$ (equivalently, \mathbb{N}^Δ) represents the total count of output states *in* Δ that would be *produced as outputs* of these transitions. It does not account for the number of input states consumed, nor the number of output states in Γ produced.

Formally, \mathbf{T}_1 is a strictly lower diagonal matrix (0's on and above the diagonal). Column j is 0's, other than potentially up to two positive entries, described below.

- If τ_j has output states $\bar{d}_k, \bar{d}_{k'}$ where $j < k \neq k'$, then $\mathbf{T}(k, j) = \mathbf{T}(k', j) = 1$.
- If τ_j has output states \bar{d}_k, \bar{d}_k where $j < k$, then $\mathbf{T}(k, j) = 2$.
- If τ_j has output states \mathbf{d}_k, \bar{g} , where $j < k$ and $\bar{g} \in \Gamma$, then $\mathbf{T}(k, j) = 1$.

By the fact that \mathcal{P} is Δ -ordered via τ_1, \dots, τ_d , there are no other forms the transitions can take. For example, if we have transitions

$$\begin{aligned}
\tau_1 &: \bar{d}_1, \bar{d}_3 &\rightarrow \bar{d}_2, \bar{d}_4 \\
\tau_2 &: \bar{d}_2, \bar{g}_1 &\rightarrow \bar{d}_3, \bar{d}_4 \\
\tau_3 &: \bar{d}_3, \bar{d}_6 &\rightarrow \bar{d}_5, \bar{g}_1 \\
\tau_4 &: \bar{d}_4, \bar{d}_6 &\rightarrow \bar{g}_1, \bar{g}_2 \\
\tau_5 &: \bar{d}_5, \bar{g}_2 &\rightarrow \bar{d}_6, \bar{d}_6 \\
\tau_6 &: \bar{d}_6, \bar{g}_1 &\rightarrow \bar{g}_1, \bar{g}_2
\end{aligned}$$

where $\bar{g}_1, \bar{g}_2 \in \Gamma$, then

$$\mathbf{T}_1 = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 2 & 0 \end{bmatrix}$$

We define \mathbf{T} based on \mathbf{T}_1 . Naïvely, to consume states in \mathbf{c}^Δ , for each $i \in \{1, \dots, d\}$ one would add $\mathbf{c}^\Delta(i)$ copies of τ_i . Since \mathcal{P} is Δ -ordered this would indeed result in count 0 of \bar{d}_1 . However, although for $i \in \{2, \dots, d\}$, this consumes $\mathbf{c}^\Delta(\bar{d}_i)$ copies of \bar{d}_i , it also produces $(\mathbf{T}_1 \cdot \mathbf{c}^\Delta)(i)$ copies of \bar{d}_i , which is positive if \bar{d}_i is an output of some transition in $\tau_1, \dots, \tau_{i-1}$.

Applying the naïve idea a second time, to consume the states that were produced on the first step, for each $i \in \{1, \dots, d\}$ we add $(\mathbf{T}_1 \cdot \mathbf{c}^\Delta)(\bar{d}_i)$ copies of τ_i . (Note that $(\mathbf{T}_1 \cdot \mathbf{c}^\Delta)(\bar{d}_1) = 0$ so this second step adds no additional copies of τ_1 .) Thus, this results in count 0 of \bar{d}_2 , and although it consumes the copies of $\bar{d}_3, \dots, \bar{d}_d$ that remained after the first step, it also produces $\mathbf{T}_1 \cdot (\mathbf{T}_1 \cdot \mathbf{c}^\Delta)(\bar{d}_i)$ additional copies of \bar{d}_i . The number of transitions after two steps is then described by summing steps 1 and 2: $\mathbf{c}^\Delta + \mathbf{T}_1 \cdot \mathbf{c}^\Delta$. We iterate this procedure a total of d steps, where the transitions added in step i are described by the vector $\mathbf{T}_1^{i-1} \cdot \mathbf{c}^\Delta$.

Since the i 'th step results in getting to count 0 of $\bar{d}_1, \dots, \bar{d}_i$, all $\bar{d}_1, \dots, \bar{d}_d$ will have count 0 after d steps. The total number of transitions applied over all steps is then described by the vector obtained by summing the d vectors indicating transition counts for each step 1 through d : $\mathbf{c}^\Delta + \mathbf{T}_1 \cdot \mathbf{c}^\Delta + \mathbf{T}_1^2 \cdot \mathbf{c}^\Delta + \mathbf{T}_1^3 \cdot \mathbf{c}^\Delta + \dots + \mathbf{T}_1^{d-1} \cdot \mathbf{c}^\Delta$. Thus, taking \mathbf{T}_1^0 to be the $d \times d$ identity matrix, we can define the matrix $\mathbf{T} = \sum_{i=0}^{d-1} \mathbf{T}_1^i$. Since each column of \mathbf{T}_1 is either all 0, has one or two 1's, or has a single 2, a simple induction shows that for all $i \in \{1, \dots, d-1\}$, $\max(\mathbf{T}_1^i) \leq 2^i$. Thus $\max(\mathbf{T}) \leq \sum_{i=0}^{d-1} \max(\mathbf{T}_1^i) \leq \sum_{i=0}^{d-1} 2^i < 2^d$. (This bound is nearly tight; e.g., transitions $\tau_i : \bar{d}_i, \bar{s} \rightarrow \bar{d}_{i+1}, \bar{d}_{i+1}$ result in $\mathbf{T}(d, 1) = 2^{d-1}$.)

Now that we have defined \mathbf{T} , which tells us that we will have $(\mathbf{T} \cdot \mathbf{c}^\Delta)(i)$ copies of τ_i in path p , it is easy to define \mathbf{C}_1 and \mathbf{C}_2 based on \mathbf{T} . For each copy of $\tau_i : \bar{d}_i, \bar{s} \rightarrow \bar{o}, \bar{o}'$, we add a copy of \bar{s} to \mathbf{e} . Thus, define the matrix $\mathbf{S} \in \mathbb{N}^{\Lambda \times \Delta}$ so that, for all $i \in \{1, \dots, d\}$, $\mathbf{S}(\bar{s}, \bar{d}_i) = 1$ if transition τ_i , which by definition has one input state \bar{d}_i , has \bar{s} as its other input state. All other entries of \mathbf{S} are 0. Then $\mathbf{C}_2 = \mathbf{S} \cdot \mathbf{T}$, and $\max(\mathbf{C}_2) = \max(\mathbf{T}) < 2^d$.

It remains to define \mathbf{C}_1 , so that $\mathbf{C}_1 \cdot \mathbf{c}^\Delta$ describes the vector \mathbf{z}^Γ of states in Γ produced by path p . First, define the matrix $\mathbf{G} \in \mathbb{N}^{\Gamma \times \Delta}$ as follows, which intuitively maps a count vector of transitions in τ_1, \dots, τ_d to a total count of states in Γ produced as output by the transitions. Let $j \in \{1, \dots, d\}$ and let $\tau_j : \bar{d}_j, \bar{s} \rightarrow \bar{o}, \bar{o}'$. If $\bar{o}, \bar{o}' \in \Delta$, then the j 'th column of \mathbf{G} is all 0. If exactly one (w.l.o.g.) $\bar{o} \in \Gamma$, then $\mathbf{G}(\bar{o}, \bar{d}_j) = 1$ and the remainder of the j 'th column of \mathbf{G} is all 0. If both $\bar{o}, \bar{o}' \in \Gamma$, and $\bar{o} = \bar{o}'$, then $\mathbf{G}(\bar{o}, \bar{d}_j) = 2$ and the remainder of the j 'th column of \mathbf{G} is all 0. If both $\bar{o}, \bar{o}' \in \Gamma$, and $\bar{o} \neq \bar{o}'$, then $\mathbf{G}(\bar{o}, \bar{d}_j) = \mathbf{G}(\bar{o}', \bar{d}_j) = 1$ and the remainder of the j 'th column of \mathbf{G} is all 0. Then $\mathbf{C}_1 = \mathbf{G} \cdot \mathbf{T}$, and $\max(\mathbf{C}_1) \leq 2 \cdot \max(\mathbf{T}) < 2^{d+1}$. \square

We demonstrate Lemma 3.7 with a concrete example. Consider a Population Protocol \mathcal{P} defined by transitions

$$\begin{aligned} \tau_1 : \bar{d}_1, \bar{d}_3 &\rightarrow \bar{g}_1, \bar{d}_2 \\ \tau_2 : \bar{g}_1, \bar{d}_2 &\rightarrow \bar{g}_1, \bar{d}_3 \\ \tau_3 : \bar{g}_1, \bar{d}_3 &\rightarrow \bar{g}_1, \bar{g}_1 \\ \tau_4 : \bar{g}_1, \bar{g}_2 &\rightarrow \bar{g}_1, \bar{g}_1 \\ \tau_5 : \bar{g}_1, \bar{g}_2 &\rightarrow \bar{g}_2, \bar{g}_2 \end{aligned}$$

Let $\Delta = \{\bar{d}_1, \bar{d}_2, \bar{d}_3\}$. Then \mathcal{P} is Δ -ordered via τ_1, τ_2, τ_3 , because τ_2 does not reference \bar{d}_1 , and τ_3 does not reference \bar{d}_2 or \bar{d}_3 , and for all $i \in \{1, 2, 3\}$, τ_i contains exactly one reference to \bar{d}_i as an input. By the terminology of Lemma 3.7, $\Gamma = \{\bar{g}_1, \bar{g}_2\}$.

For a given \mathbf{c}^Δ , we can design a configuration $\mathbf{e} \in \mathbb{N}^\Delta$ such that $\mathbf{c}^\Delta + \mathbf{e} \implies \mathbf{z}^\Gamma$. We will remove agents in Δ using the transitions given to us by \mathcal{P} being Δ -ordered. For example, to remove n agents of \bar{d}_1 we will need transition τ_1 to occur n times. Similarly, to remove m agents of \bar{d}_2 , we will need τ_2 to occur at least m times, and we may need more if τ_1 generated additional copies of \bar{d}_2 .

Let \mathbf{T} be defined as in the proof of Lemma 3.7 and let

$$\mathbf{c}^\Delta = \begin{bmatrix} 5 \\ 1 \\ 2 \end{bmatrix}$$

In order to remove 5 copies of \bar{d}_1 , we need enough copies of \bar{d}_3 for τ_1 to occur 5 times. As such, we add $\mathbf{T} \cdot \mathbf{c}^\Delta(\bar{d}_1) = 5 \bar{d}_3$ agents to \mathbf{e} allowing τ_1 to occur 5 times, as shown in Figure 1. This effectively removes all copies of \bar{d}_1 and produces 5 extra copies of both \bar{g}_1 and \bar{d}_2 . Since \mathcal{P} is Δ -ordered, we know that the only states created will either be in Γ or they will be in Δ but further in the ordering - allowing us to remove the extra agents at a later time.

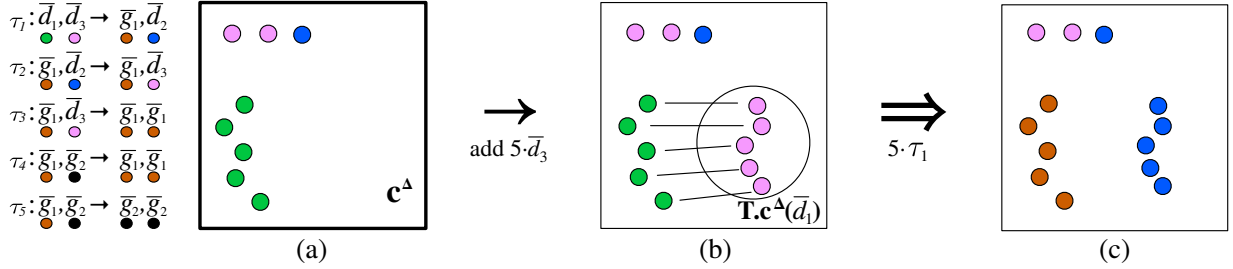


Figure 1: First step of the surgery detailed in Lemma 3.7. (a) Our initial configuration \mathbf{c}^Δ . (b) Add 5 \bar{d}_3 agents which will react with 5 \bar{d}_1 agents. (c) All 5 copies of \bar{d}_1 have been removed.

Now, we must remove 1 original copy of \bar{d}_2 plus 5 newly created copies created via τ_1 in the previous step. To do so, we can use τ_2 , which requires we add $\mathbf{T} \cdot \mathbf{c}^\Delta(\bar{d}_2) = 6$ additional copies of \bar{g}_1 to \mathbf{e} . Via τ_2 , an additional 6 copies of both \bar{d}_3 and \bar{g}_1 are created. This process is illustrated in Figure 2. Again, since \bar{d}_3 comes after \bar{d}_2 in the ordering, we can still remove all copies of \bar{d}_3 at a later step.

There are now 8 copies of \bar{d}_3 to remove due to additional agents being produced in the previous steps. We will add $\mathbf{T} \cdot \mathbf{c}^\Delta(\bar{d}_3) = 8$ copies of \bar{g}_1 to \mathbf{e} to allow 8 instances of τ_3 to take place. This will transition all instances of \bar{d}_3 into instances of \bar{g}_1 and leave us with a configuration \mathbf{z}^Γ of states only in Γ as shown in Figure 3.

While the process can be presented as several separate steps as above, \mathbf{e} will be the sum of all the agents added in the prior steps. In Figure 4, we add \mathbf{e} to \mathbf{c}^Δ , which will then transition to \mathbf{z}^Γ . We have removed all agents in Δ and arrived at a configuration \mathbf{z}^Γ as desired.

The next lemma works toward generating the vector of states \mathbf{e} needed to apply Lemma 3.7. The “cost” for Lemma 3.8 is that the path must be taken “in the context” of additional agents in states captured by \mathbf{p} . The intuitive reason \mathbf{p} is needed is this: In Lemma 3.7, we add new transitions and add enough new states (\mathbf{e}) to supply the inputs for all these transitions. Thus the resulting counts of states in Γ can only be larger than the the original path. However, the manipulation of

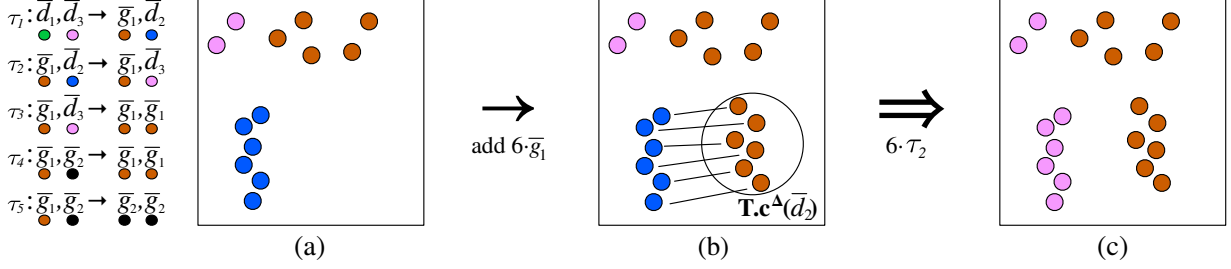


Figure 2: Second step of the surgery detailed in Lemma 3.7. (a) The same configuration as in panel (c) of Figure 1. (b) Add 6 \bar{g}_1 agents which will react with 6 \bar{d}_2 agents. (c) All 6 copies of \bar{d}_2 have been removed.

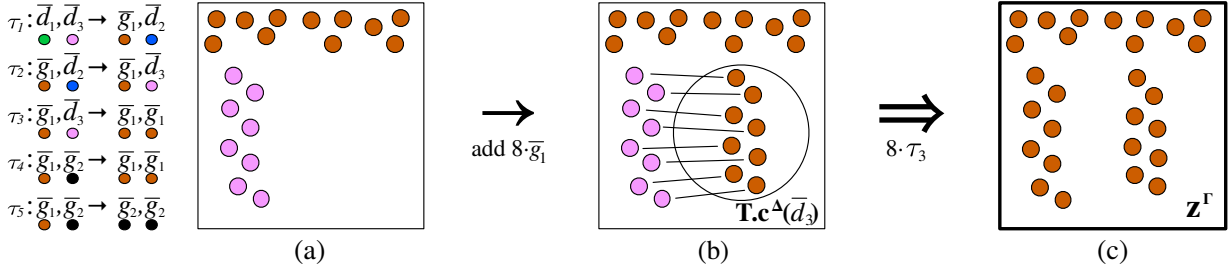


Figure 3: Third step of the surgery detailed in Lemma 3.7. (a) The same configuration as in panel (c) of Figure 2. (b) Add 8 \bar{g}_1 agents which will react with 8 \bar{d}_3 agents. (c) All 8 copies of \bar{d}_3 have been removed. The resulting configuration contains only states in Γ .

Lemma 3.8 does not add new states, so inputs of added transitions may have lower count. Also, unlike Lemma 3.8, the manipulation also involves *removing* transitions; thus the *outputs* of those transitions may have their counts lowered. The states in \mathbf{p} are used as a “buffer” to keep any state count from becoming negative, which could happen if the manipulation were applied to the path starting from just the original configuration \mathbf{x} . Because of this, we do not specify where precisely in the transition sequence certain transitions are added, nor which specific occurrences are removed. These could be chosen anywhere along the sequence, and the buffer \mathbf{p} ensures that the entire sequence remains applicable.

Importantly, the net effect of the path preserves \mathbf{p} , which will give a way to “interleave” Lemmas 3.7 and 3.8, in order to start from a configuration with large counts of all states and reach a configuration with count 0 of all states in Δ . Note that unlike matrices $\tilde{\mathbf{C}}_1$ and $\tilde{\mathbf{C}}_2$ of Lemma 3.7, the matrix $\tilde{\mathbf{C}}_3$ may have negative entries, since the resulting configuration is described as a difference from the configuration \mathbf{o} , and some states in Γ may have lower count in the new configuration than in \mathbf{o} .

Lemma 3.8 (adapted from [20]). *Let $b_1, b_2 \in \mathbb{N}$. Let $\mathbf{x}, \mathbf{o} \in \mathbb{N}^\Lambda$ such that $\mathbf{x} \implies_q \mathbf{o}$ via path q that does not contain a b_2 -bottleneck. Define $\Delta = \{\bar{d} \in \Lambda \mid \mathbf{o}(\bar{d}) \leq b_1\}$, $d = |\Delta|$, and $\Gamma = \Lambda \setminus \Delta$. Let $\mathbf{o}^\Delta = \mathbf{o} \upharpoonright \Delta$ and $\mathbf{o}^\Gamma = \mathbf{o} \upharpoonright \Gamma$. Then there is a matrix $\tilde{\mathbf{C}}_3 \in \mathbb{Z}^{\Gamma \times \Delta}$ with $\text{amax}(\tilde{\mathbf{C}}_3) < 2^{d+1}$, such that for all $\mathbf{e}^\Delta \in \mathbb{N}^\Delta$, if $b_2 \geq |\Lambda| \cdot b_1 + d2^d \cdot \max(b_1, \mathbf{e}^\Delta) \cdot |\Lambda|^2$, and if $\mathbf{x}(\bar{s}) \geq b_2$ for all $\bar{s} \in \Lambda$ and $\mathbf{o}(\bar{s}) \geq b_2$ for all $\bar{s} \in \Gamma$, letting $\mathbf{p} \in \mathbb{N}^\Lambda$ be defined $\mathbf{p}(\bar{s}) = d2^{d+1} \cdot \max(b_1, \mathbf{e}^\Delta)$ for all $\bar{s} \in \Lambda$, then $\mathbf{p} + \mathbf{x} \implies \mathbf{p} + \mathbf{o}^\Gamma + \tilde{\mathbf{C}}_3 \cdot \mathbf{o}^\Delta - \tilde{\mathbf{C}}_3 \cdot \mathbf{e}^\Delta + \mathbf{e}^\Delta$.*

Note that we consider configurations in which all counts in \mathbf{o}^Γ are arbitrarily large (see Lemma 3.9), whereas counts in \mathbf{o}^Δ and \mathbf{e}^Δ , as well as the entries of $\tilde{\mathbf{C}}_3$, are bounded. Thus, for sufficiently large

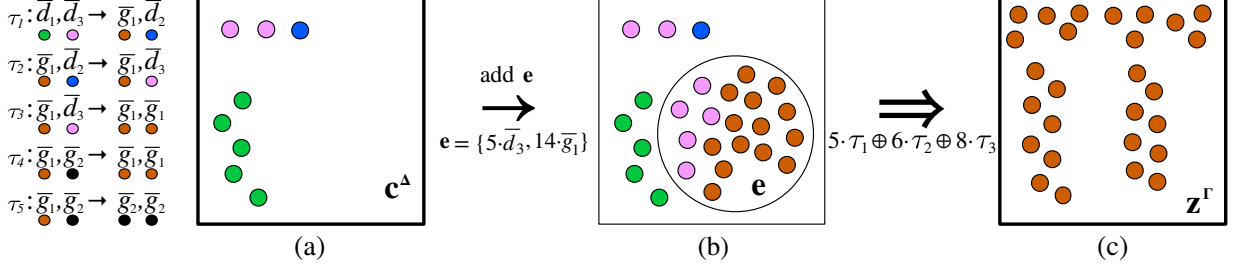


Figure 4: Illustration of the whole path surgery technique detailed in Lemma 3.7. We combine the steps shown in Figures 1, 2, and 3 to form a single transition sequence leading from \mathbf{c}^Δ to \mathbf{z}^Γ . (a) The same initial configuration \mathbf{c}^Δ as in panel (a) of Figure 1. (b) Add $\mathbf{e} = \{5\bar{d}_3\} + \{6\bar{g}_1\} + \{8\bar{g}_1\} = \{5\bar{d}_3, 14\bar{g}_1\}$ to \mathbf{c}^Δ . (c) All states in Δ are removed and we reach a configuration \mathbf{z}^Γ .

starting configurations, $\mathbf{o}^\Gamma \geq \tilde{\mathbf{C}}_3 \cdot \mathbf{o}^\Delta - \tilde{\mathbf{C}}_3 \cdot \mathbf{e}^\Delta + \mathbf{e}^\Delta$, justifying our earlier claim that in this lemma, we “get \mathbf{p} back at the end.”

Proof. By Lemma 3.4, \mathcal{P} is Δ -ordered via τ_1, \dots, τ_d , and each τ_i occurs at least $(b_2 - |\Lambda| \cdot b_1)/|\Lambda|^2$ times in q .

Intuitively, this proof is similar to that of Lemma 3.7, except that instead of targeting count 0 of all states in Δ , we target counts given by \mathbf{e}^Δ . Also, rather than constructing a new path consisting solely of transitions of type τ_1, \dots, τ_d , we alter the path q , which may contain other types of transitions (although we will only modify transitions of type τ_1, \dots, τ_d). Since $\mathbf{x} \implies_p \mathbf{o}$, we can think of our “starting value” for counts in Δ as being \mathbf{o}^Δ , and thus the total change in counts that we want to make is described by the vector $\mathbf{o}^\Delta - \mathbf{e}^\Delta$. In particular, since we may have $\mathbf{o}^\Delta(\bar{d}_i) < \mathbf{e}^\Delta(\bar{d}_i)$ for some $\bar{d} \in \Delta$, this may require *removing* transitions from q as well as adding them. Furthermore, since we have no \mathbf{e} at the start as in Lemma 3.7, when adding or removing transition $\bar{d}_i, \bar{s} \rightarrow \bar{o}, \bar{o}'$ to alter the count of \bar{d}_i , we must account not only for the effect this has on the output states \bar{o}, \bar{o}' , but also the effect on the other input state \bar{s} . This may result in a path that is not valid, in the sense that some counts may be negative after the modification. The extra states in \mathbf{p} have the purpose of keeping the entire path valid. The bound on \mathbf{p} will then be derived from the bound on the size of the changes to q that we make.

First, we define a matrix $\tilde{\mathbf{T}} \in \mathbb{Z}^{d \times d}$. Intuitively, if $\tilde{\mathbf{c}}^\Delta \in \mathbb{Z}^\Delta$ represents changes in counts of states in Δ that we wish to achieve through addition and removal of transitions τ_1, \dots, τ_d from the path q , then the vector $\tilde{\mathbf{t}} \in \mathbb{Z}^d$ defined by $\tilde{\mathbf{t}} = \tilde{\mathbf{T}} \cdot \tilde{\mathbf{c}}^\Delta$ represents changes in counts of transitions τ_1, \dots, τ_d in the path q to achieve this. More precisely, the counts in \mathbf{o} are given by \mathbf{o}^Δ , but we wish them to be \mathbf{e}^Δ instead. Letting $\tilde{\mathbf{c}}^\Delta = \mathbf{o}^\Delta - \mathbf{e}^\Delta$, then $\tilde{\mathbf{t}} = \tilde{\mathbf{T}} \cdot \tilde{\mathbf{c}}^\Delta$ describes how many transitions of each type to add or remove from q .

Define the $d \times d$ matrix $\tilde{\mathbf{T}}_1$ as follows. Intuitively, $\tilde{\mathbf{T}}_1$ is a matrix such that, if $\tilde{\mathbf{t}} \in \mathbb{Z}^d$ represents (possibly negative) “counts of transition executions”, i.e., $\tilde{\mathbf{t}}(j)$ means “execute transition τ_j an additional $\tilde{\mathbf{t}}(j)$ times” (where executing a transition an additional negative number of times means removing it from q), then $\tilde{\mathbf{T}}_1 \cdot \tilde{\mathbf{t}} \in \mathbb{Z}^\Delta$ (equivalently, \mathbb{Z}^Δ) represents the total count of states in Δ that would be *produced as outputs* of these transitions or *consumed as the second input*. Here, the “second” input means, for transition $\tau_i : \bar{d}_i, \bar{s} \rightarrow \bar{o}, \bar{o}'$, the input \bar{s} that is *not* \bar{d}_i . It does not account for the number of input states in Γ consumed, nor the number of output states in Γ produced.

Formally, $\tilde{\mathbf{T}}_1$ is a strictly lower diagonal matrix (0’s on and above the diagonal). Column j is 0’s, other than potentially up to two nonzero entries, described below.

- If τ_j has output states $\bar{d}_k, \bar{d}_{k'}$ where $j < k \neq k'$, then $\tilde{\mathbf{T}}(k, j) = \tilde{\mathbf{T}}(k', j) = 1$.

- If τ_j has output states \bar{d}_k, \bar{d}_k where $j < k$, then $\tilde{\mathbf{T}}(k, j) = 2$.
- If τ_j has output states \bar{d}_k, \bar{g} , where $j < k$ and $\bar{g} \in \Gamma$, then $\tilde{\mathbf{T}}(k, j) = 1$.
- If τ_j has second input state \bar{d}_k where $j < k$, then $\tilde{\mathbf{T}}(k, j) = -1$.

By the fact that \mathcal{P} is Δ -ordered via τ_1, \dots, τ_d , there are no other forms the transitions can take. For example, if we have transitions

$$\begin{aligned}
\tau_1 &: \bar{d}_1, \bar{d}_3 \rightarrow \bar{d}_2, \bar{d}_4 \\
\tau_2 &: \bar{d}_2, \bar{g}_1 \rightarrow \bar{d}_3, \bar{d}_4 \\
\tau_3 &: \bar{d}_3, \bar{d}_6 \rightarrow \bar{d}_5, \bar{g}_1 \\
\tau_4 &: \bar{d}_4, \bar{d}_6 \rightarrow \bar{g}_1, \bar{g}_2 \\
\tau_5 &: \bar{d}_5, \bar{g}_2 \rightarrow \bar{d}_6, \bar{d}_6 \\
\tau_6 &: \bar{d}_6, \bar{g}_1 \rightarrow \bar{g}_1, \bar{g}_2
\end{aligned}$$

where $\bar{g}_1, \bar{g}_2 \in \Gamma$, then

$$\tilde{\mathbf{T}}_1 = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ -1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & -1 & 2 & 0 \end{bmatrix}$$

We define $\tilde{\mathbf{T}}$ based on $\tilde{\mathbf{T}}_1$. Naïvely, to consume (respectively, produce) counts of states as described in $\tilde{\mathbf{c}}^\Delta$, for each $i \in \{1, \dots, d\}$ one would add $\tilde{\mathbf{c}}^\Delta(i)$ copies of τ_i (where adding a negative amount means removing from path q). Since \mathcal{P} is Δ -ordered this would indeed result in altering the count of \bar{d}_1 by $\tilde{\mathbf{c}}^\Delta(\bar{d}_1)$. However, although for $i \in \{2, \dots, d\}$, this consumes (resp. produces) $\mathbf{c}^\Delta(\bar{d}_i)$ copies of \bar{d}_i , it also produces $(\tilde{\mathbf{T}}_1 \cdot \tilde{\mathbf{c}}^\Delta)(i)$ copies of \bar{d}_i (where ‘‘producing’’ a negative number corresponds to consuming copies of \bar{d}_i). Note that $(\tilde{\mathbf{T}}_1 \cdot \tilde{\mathbf{c}}^\Delta)(1) = 0$ since \bar{d}_1 does not appear in any transition other than τ_1 .

We take the same approach as in the proof of Lemma 3.7, in which we take the vector $\tilde{\mathbf{T}}_1 \cdot \tilde{\mathbf{c}}^\Delta$, which represents the difference between the count of states in Δ , compared to our target *after* doing step 1 above. Step 2 consists of adding transitions as described by the vector $\tilde{\mathbf{T}}_1 \cdot \tilde{\mathbf{c}}^\Delta$, resulting in $\tilde{\mathbf{T}}_1^2 \cdot \tilde{\mathbf{c}}^\Delta$, in which $\tilde{\mathbf{T}}_1^2 \cdot \tilde{\mathbf{c}}^\Delta(1) = \tilde{\mathbf{T}}_1^2 \cdot \tilde{\mathbf{c}}^\Delta(2) = 0$. The i 'th step involves adding transitions according to the vector $\tilde{\mathbf{T}}_1^i \cdot \tilde{\mathbf{c}}^\Delta$. We define $\tilde{\mathbf{T}} = \sum_{i=0}^{d-1} \tilde{\mathbf{T}}_1^i$.

Thus, if transition τ_i appears a_i times in q , then in the altered path q' , it appears $a_i + (\tilde{\mathbf{T}} \cdot \tilde{\mathbf{c}}^\Delta)(i)$ times. Thus, so long as for each i , $a_i \geq -(\tilde{\mathbf{T}} \cdot \tilde{\mathbf{c}}^\Delta)(i)$, there are sufficiently many transitions of each type in q to potentially remove. Since each column of $\tilde{\mathbf{T}}_1$ has either a single 2, or two 1's, and at most one -1 , a simple induction shows that for each i , $\text{amax}(\tilde{\mathbf{T}}_1^i) \leq 2^i$. Thus $\text{amax}(\tilde{\mathbf{T}}) \leq \sum_{i=0}^{d-1} \tilde{\mathbf{T}}_1^i \leq \sum_{i=0}^{d-1} 2^i < 2^d$, which implies that $\max(-(\tilde{\mathbf{T}} \cdot \tilde{\mathbf{c}}^\Delta)) < d2^d \cdot \text{amax}(\tilde{\mathbf{c}}^\Delta)$. Thus, it suffices if $a_i \geq d2^d \cdot \text{amax}(\tilde{\mathbf{c}}^\Delta)$. By Lemma 3.4, $a_i \geq (b_2 - |\Lambda| \cdot b_1)/|\Lambda|^2$. Thus it suffices to show that $d2^d \cdot \text{amax}(\tilde{\mathbf{c}}^\Delta) \leq (b_2 - |\Lambda| \cdot b_1)/|\Lambda|^2$. Note that $b_1 \geq \max(\mathbf{o}^\Delta)$ by the definition of Δ in the statement of the lemma. Then we have $d2^d \cdot \text{amax}(\tilde{\mathbf{c}}^\Delta) = d2^d \cdot \text{amax}(\mathbf{o}^\Delta - \mathbf{e}^\Delta) \leq d2^d \cdot \max(\mathbf{o}^\Delta, \mathbf{e}^\Delta) \leq d2^d \cdot \max(b_1, \mathbf{e}^\Delta) \leq (b_2 - |\Lambda| \cdot b_1)/|\Lambda|^2$ where the last inequality follows from the assumption that $b_2 \geq |\Lambda| \cdot b_1 + d2^d \cdot \max(b_1, \mathbf{e}^\Delta) \cdot |\Lambda|^2$.

It remains to define $\tilde{\mathbf{C}}_3$, so that the vector $\tilde{\mathbf{C}}_3 \cdot \tilde{\mathbf{c}}^\Delta = \tilde{\mathbf{C}}_3 \cdot \mathbf{o}^\Delta - \tilde{\mathbf{C}}_3 \cdot \mathbf{e}^\Delta$ describes relative counts of states in Γ compared to \mathbf{o}^Γ . First, define the matrix $\mathbf{G} \in \mathbb{N}^{\Gamma \times \Delta}$ as follows, which intuitively maps a count vector of transitions in τ_1, \dots, τ_d to a total count of states in Γ either produced as output or consumed as a second input by the transitions. Let $j \in \{1, \dots, d\}$ and let $\tau_j : \bar{d}_j, \bar{s} \rightarrow \bar{o}, \bar{o}'$. Let

$i \in \{1, \dots, g\}$ (recall $g = |\Gamma|$). Let $\mathbf{G}(i, j) \in \{-1, 0, 1, 2\}$ denote the net number of \bar{g}_i produced by τ_j ; e.g., -1 if $\bar{s} = \bar{g}_i \neq \bar{o}, \bar{o}'$, 0 if $\bar{g}_i \neq \bar{s}, \bar{o}, \bar{o}'$ or if $\bar{g}_i = \bar{s} = \bar{o} \neq \bar{o}'$, etc. Then $\tilde{\mathbf{C}}_3 = \mathbf{G} \cdot \mathbf{T}$, and $\text{amax}(\tilde{\mathbf{C}}_3) \leq 2 \cdot \text{amax}(\mathbf{T}) < 2^{d+1}$.

Finally, recall that the new path may not be valid since, although the final configuration $\mathbf{o}^\Gamma + \tilde{\mathbf{C}}_3 \cdot \mathbf{o}^\Delta - \tilde{\mathbf{C}}_3 \cdot \mathbf{e}^\Delta + \mathbf{e}^\Delta \in \mathbb{N}^\Lambda$ is nonnegative, some intermediate configurations between \mathbf{x} and the final configuration could be negative. Define $\mathbf{G}' \in \mathbb{N}^{\Lambda \times \Delta}$ similarly to \mathbf{G} above, but reflecting the effect of a transition τ_1, \dots, τ_d on *every* state $\bar{s} \in \Lambda$ (not just those in Γ). Then letting $\tilde{\mathbf{C}}'_3 = \mathbf{G}' \cdot \mathbf{T}$, $\text{amax}(\tilde{\mathbf{C}}'_3 \cdot \tilde{\mathbf{c}}^\Delta) < d2^{d+1} \text{amax}(\tilde{\mathbf{c}}^\Delta)$ is an upper bound on how much any individual state count can change. Thus, for all $\bar{s} \in \Lambda$, letting $\mathbf{p}(\bar{s}) = d2^{d+1} \cdot \text{amax}(\tilde{\mathbf{c}}^\Delta) = d2^{d+1} \cdot \text{amax}(\mathbf{o}^\Delta - \mathbf{e}^\Delta) \leq d2^{d+1} \cdot \max(\mathbf{o}^\Delta, \mathbf{e}^\Delta) \leq d2^{d+1} \cdot \max(b_1, \mathbf{e}^\Delta)$ suffices to ensure that the whole path $\mathbf{p} + \mathbf{x} \implies \mathbf{p} + \mathbf{o}^\Gamma + \tilde{\mathbf{C}}_3 \cdot \mathbf{o}^\Delta - \tilde{\mathbf{C}}_3 \cdot \mathbf{e}^\Delta + \mathbf{e}^\Delta$ is valid. \square

The following example provides intuition for Lemma 3.8. Let $\Delta, \Gamma, \mathbf{x}, \mathbf{o}, \mathbf{p}$ be defined as in Lemma 3.8. For all \mathbf{e}^Δ , we can alter a transition sequence p — either by adding or removing transitions — to ensure we finish with only states in Γ plus exactly \mathbf{e}^Δ from Δ . In other words, given that we have additional agents in \mathbf{p} to use, we can manipulate a population protocol to have the exact amount of agents in Δ that we desire. Depending on the desired \mathbf{e}^Δ , doing so will effect the counts of states in Γ in a predictable way.

Let us continue with the above example using \mathcal{P} . In order to transition \mathbf{c}^Δ to \mathbf{z}^Γ , we needed to add \mathbf{e} , which contained 14 copies of \bar{g}_1 and 5 copies of \bar{d}_3 . So, for this example, we will show how to produce \mathbf{e}^Δ such that

$$\mathbf{e}^\Delta = \begin{bmatrix} 0 \\ 0 \\ 5 \end{bmatrix}$$

Let $\mathbf{x} \implies_p \mathbf{o}$ by transition sequence p and let

$$\mathbf{o}^\Delta = \begin{bmatrix} 3 \\ 1 \\ 2 \end{bmatrix}$$

So,

$$\tilde{\mathbf{c}}^\Delta = \mathbf{o}^\Delta - \mathbf{e}^\Delta = \begin{bmatrix} 3 \\ 1 \\ -3 \end{bmatrix}$$

representing the number of agents of each state in Δ that must be removed. The proof of Lemma 3.8 provides further detail on how to create matrices to determine exactly how many instances of each transition are added or removed based on $\tilde{\mathbf{c}}^\Delta$. We will need to remove two instances of τ_3 and add 3 instances of τ_1 and 4 instances of τ_2 .

Additional transition executions can be added to the end of the transition sequence without otherwise affecting the original transition sequence; however, you can only remove transition executions where they take place. Thus, to remove two instance of τ_3 , we need to alter p in the middle of the sequence.

Removing transitions has an affect on the overall counts of other states. We have extra states in \mathbf{p} to account for this. If removing a transition in the middle of the sequence causes the count of a state to become zero later in the sequence, and if a transition requiring that state was meant to take place after that point, the extra agents in \mathbf{p} can be used instead.

Our example population protocol \mathcal{P} includes τ_4 and τ_5 to demonstrate how \mathbf{p} may be used. While these transition rules may not seem particularly useful in practice, they allow us to show a simple situation in which the count of one state dips to zero.

In \mathcal{P} , removing two instances of τ_3 will increase the count of \bar{d}_3 by 2, at the same time reducing the count of \bar{g}_1 by 2. At some point in p , the count of \bar{g}_1 dips to 2. With 2 fewer \bar{g}_1 agents in our modified transition sequence, we instead arrive at a configuration where there are zero copies of \bar{g}_1 . In order for the next τ_4 transition to take place, we must use the extra agents provided by \mathbf{p} . Thus, the presence of additional agents allows the rest of the transition sequence to remain unchanged after removing transitions from the transition sequence. This is demonstrated in Figure 5.

We can then add the transitions necessary to remove \bar{d}_1 and \bar{d}_2 from the resulting configuration, since our final configuration \mathbf{e}^Δ contains no agents from these states. This process is similar to the one detailed above to describe Lemma 3.7. In Figure 6 we add 3 instances of τ_1 and in Figure 7 we add 4 instances of τ_3 , which completes the adjusted transition sequence.

In the final configuration of Figure 7, counts of states in Γ are altered and counts of states in Δ are exactly equal to those in \mathbf{e}^Δ as desired. This demonstrates how we can create \mathbf{e} as required by Lemma 3.7. This concludes the description of the example showing Lemma 3.8.

The following is a general lemma that uses Lemmas 3.7 and 3.8 to “steer” certain configurations (each of which is expressed as “twice a configuration \mathbf{x}_n where all counts are large ($\geq b_2$), plus a few more states described by \mathbf{d}^Δ ”) to a configuration with a “target” \mathbf{t}^Δ amount of states in Δ .

Intuitively, the proof goes like this: Letting \mathbf{c}^Δ in Lemma 3.7 equal $\mathbf{o}_n \upharpoonright \Delta + \mathbf{d}^\Delta$, the states in Δ in the final configuration \mathbf{o}_n , plus the states \mathbf{d}^Δ that we have added at the start, which we want to eliminate, use Lemma 3.7 to determine what states \mathbf{e} need to be added to apply the extra transitions of Lemma 3.7 and eliminate all of $\mathbf{d}^\Delta + \mathbf{o}_n \upharpoonright \Delta$. Then, apply Lemma 3.8 to produce these states \mathbf{e} from \mathbf{x}_n . However, Lemma 3.8 requires the presence of an extra “buffer” \mathbf{p} of states to enable \mathbf{e} to be produced. A second copy of \mathbf{x}_n serves as this extra buffer (for large enough n $\mathbf{x}_n \geq \mathbf{p}$, since $\mathbf{x}_n(\bar{s}) \geq \beta n$ for all state counts in \mathbf{x}_n , due to \mathbf{x}_n , when we employ Lemma 3.9, being the sequence of configurations from Lemma 3.6 in which all state counts are at least βn for a fixed $\beta > 0$). Finally, we can generalize to not only eliminate all of $\mathbf{d}^\Delta + \mathbf{o}_n \upharpoonright \Delta$ (i.e., set final counts of states in Δ all to 0), but to target other *positive* counts of states in Δ , represented by the vector \mathbf{t}^Δ . The states in \mathbf{t}^Δ can simply be generated by Lemma 3.8 alongside the states in \mathbf{e} . In this paper we use only Corollary 3.11, which sets $\mathbf{t}^\Delta = \mathbf{0}$, but it may be useful in other applications to be able to choose a nonzero \mathbf{t}^Δ .

Lemma 3.9. *Let $b_2 \in \mathbb{N}$. Let $N \subseteq \mathbb{N}$ be infinite. Let $(\mathbf{x}_n)_{n \in N}$ and $(\mathbf{o}_n)_{n \in N}$ be nondecreasing sequences of configurations. Let $(p_n)_{n \in N}$ be a sequence of paths such that, for all $n \in N$,*

1. $\|\mathbf{x}_n\| = \|\mathbf{o}_n\| = n$,
2. $\mathbf{x}_n(\bar{s}) \geq b_2$ for all $\bar{s} \in \Lambda$,
3. $\mathbf{x}_n \implies_{p_n} \mathbf{o}_n$, and
4. p_n has no b_2 -bottleneck transition.

Let $\Delta = \text{bdd}((\mathbf{o}_n)_{n \in N})$, $d = |\Delta|$, and $\Gamma = \Lambda \setminus \Delta$. Then there are matrices $\tilde{\mathbf{D}}_1 \in \mathbb{Z}^{\Gamma \times \Delta}$ and $\tilde{\mathbf{D}}_2 \in \mathbb{Z}^{\Gamma \times \Delta}$ with $\text{amax}(\tilde{\mathbf{D}}_1), \text{amax}(\tilde{\mathbf{D}}_2) \leq d2^{2d+2}$ and the following holds. For each $n \in N$, let $\mathbf{o}_n^\Delta = \mathbf{o}_n \upharpoonright \Delta$ and $\mathbf{o}_n^\Gamma = \mathbf{o}_n \upharpoonright \Gamma$. Let $b_1 = \max_{n \in N, \bar{d} \in \Delta} \mathbf{o}_n^\Delta(\bar{d})$. For all $b \in \mathbb{N}$, there is $n_b \in N$ such that for all $n \in N$ with $n \geq n_b$ and all $\mathbf{d}^\Delta, \mathbf{t}^\Delta \in \mathbb{N}^\Delta$ with $\mathbf{d}^\Delta \leq b$, $\mathbf{t}^\Delta \leq d2^d(b_1 + b)$, and $b_2 \geq \max(|\Lambda| \cdot b_1 + d2^d \cdot \max(b_1, d2^{d+1}(b_1 + b)) \cdot |\Lambda|^2, d^2 2^{2d+2}(b_1 + b))$, we have $2\mathbf{x}_n + \mathbf{d}^\Delta \implies 2\mathbf{o}_n^\Gamma + \tilde{\mathbf{D}}_1 \cdot \mathbf{o}_n^\Delta + \tilde{\mathbf{D}}_2 \cdot \mathbf{d}^\Delta + \mathbf{t}^\Delta$.

Proof. Choose $n \in N$ sufficiently large to satisfy the conditions below as needed. The bound on b_2 in the hypothesis is a max of two different bounds, each needed for its own purpose:

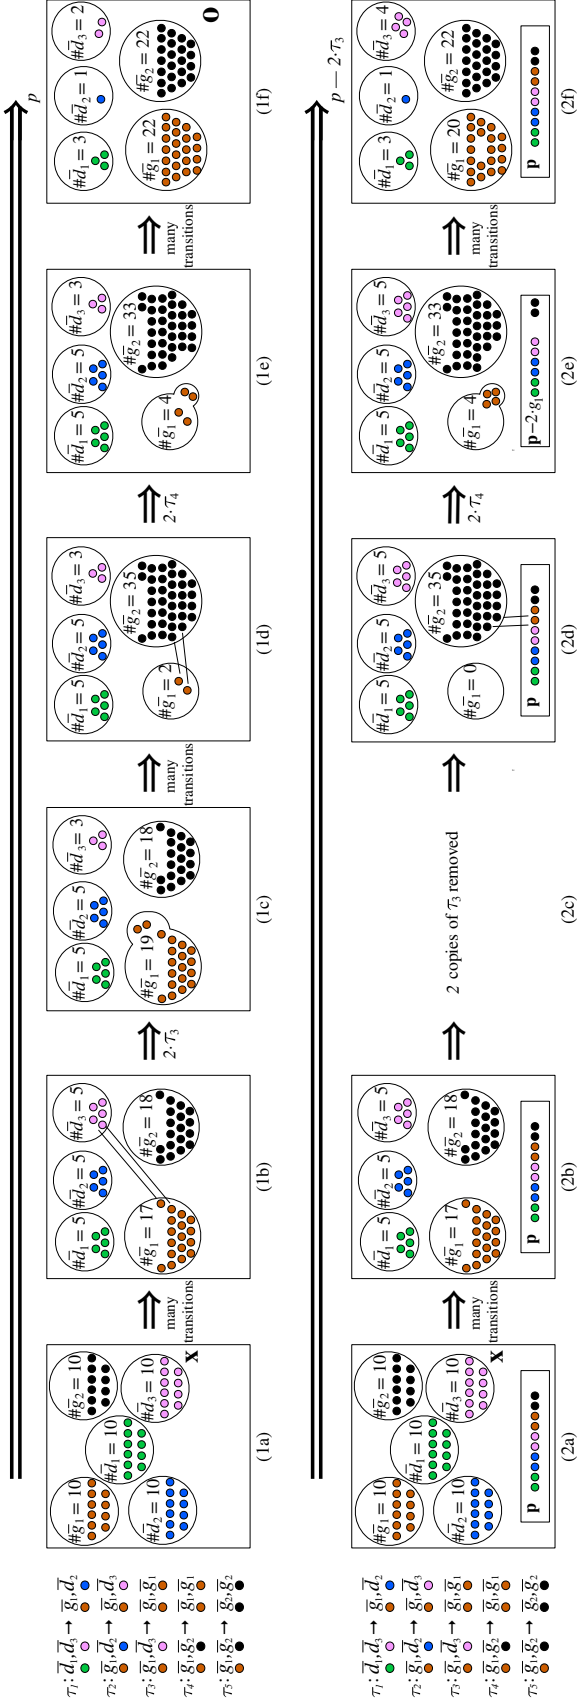


Figure 5: Surgery as described in Lemma 3.8 to produce two extra copies of \bar{d}_3 . This figure shows the difference between the original transition sequence and the sequence after surgery. The top diagram shows $\mathbf{x} \implies_p \mathbf{o}$ without any interference. **(1a)** The initial configuration \mathbf{x} . **(1b)** An intermediate configuration. **(1c)** 2 instances of τ_3 take place. **(1d)** A configuration after many instances of τ_5 have caused the counts of \bar{g}_1 to drop to 2. **(1e)** 2 instances of τ_4 take place bringing the count of \bar{g}_1 to 4. **(1f)** Finally, the transition sequence ends in \mathbf{o} . Consistent with the definition of Δ and Γ , the counts of states in Δ are low, and counts of both \bar{g}_1 and \bar{g}_2 are “large”. **(2a)** The same initial configuration \mathbf{x} from **(1a)**. **(2b)** the same intermediate configuration from **(1b)**. **(2c)** Remove 2 instances of τ_3 . Due to the removal of 2 instances of τ_3 , there are 2 less copies of \bar{g}_1 and 2 extra copies of \bar{d}_3 compared to **(1d)**. The transition sequence reaches a point where there are not enough agents outside of \mathbf{p} for τ_4 to execute. **(2e)** We must use the additional agents in \mathbf{p} to continue with the remainder of the transition sequence $p - 2\tau_3$. **(2f)** We are left with 2 additional copies of \bar{d}_3 and 2 less copies of \bar{g}_1 due to removing 2 instances of τ_3 . We have created some of the copies of \bar{d}_3 required for our desired \mathbf{e}^Δ . Note that \mathbf{p} can be “restored” by taking copies of \bar{g}_1 from \mathbf{o}^Γ since it has “large” counts.

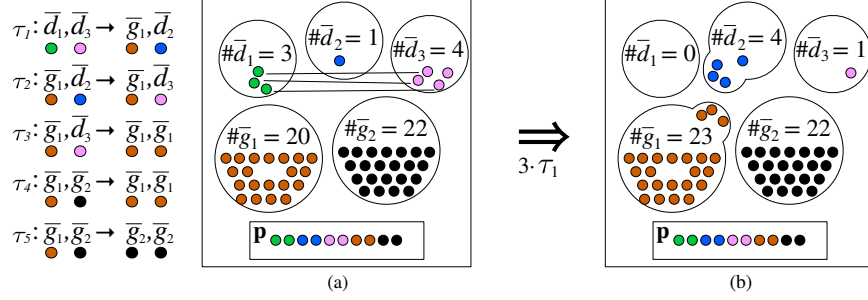


Figure 6: Removing all instances of \bar{d}_1 using techniques detailed in Lemma 3.8. Since we are adding transitions (not removing them as in Figure 5), this is similar to the surgery shown in Figures 1, 2, and 3. (a) The same configuration as seen in panel (2f) of Figure 5. (b) Add 3 instances of τ_1 to remove all copies of \bar{d}_1 as \mathbf{e}^Δ contains no copies of \bar{d}_1

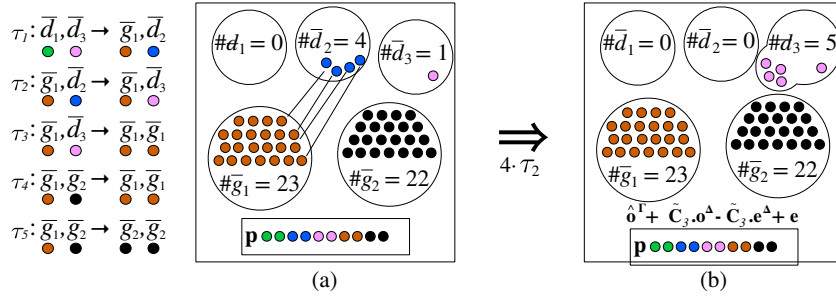


Figure 7: Removing all instances of \bar{d}_2 using techniques detailed in Lemma 3.8. This is similar to the surgery shown in Figures 1, 2, and 3. (a) The same configuration as seen in panel (b) of Figure 6. (b) Add 4 instances of τ_2 to remove all copies of \bar{d}_2 as \mathbf{e}^Δ contains no copies of \bar{d}_2 (c) The only agents from Δ that remain are exactly those needed in $\mathbf{e}^\Delta = \{5\bar{d}_3\}$.

1. $b_2 \geq |\Lambda| \cdot b_1 + d2^d \cdot \max(b_1, d2^{d+1}(b_1 + b)) \cdot |\Lambda|^2$ is necessary to apply Lemma 3.8.
2. $b_2 \geq d^2 2^{2d+2}(b_1 + b)$ is necessary to ensure that \mathbf{x}_n , which has $\mathbf{x}_n(\bar{s}) \geq b_2$ for all $\bar{s} \in \Lambda$ by hypothesis, obeys $\mathbf{x}_n \geq \mathbf{p}$ as defined below, where \mathbf{p} is used as in Lemma 3.8.

Choose $\tilde{\mathbf{C}}_3$ for Δ , Γ , \mathbf{x}_n , \mathbf{o}_n , and p_n as in Lemma 3.8. By Lemma 3.7, letting $\mathbf{c}^\Delta = \mathbf{d}^\Delta + \mathbf{o}_n^\Delta$, there are $\mathbf{C}_1 \in \mathbb{N}^{\Gamma \times \Delta}$ and $\mathbf{C}_2 \in \mathbb{N}^{\Lambda \times \Delta}$, so that, setting $\mathbf{e} = \mathbf{C}_2 \cdot \mathbf{c}^\Delta$, there is a transition sequence q_1 such that

$$\begin{aligned} \mathbf{d}^\Delta + \mathbf{o}_n^\Delta + \mathbf{e} &\Longrightarrow_{q_1} \mathbf{C}_1 \cdot (\mathbf{d}^\Delta + \mathbf{o}_n^\Delta) \\ &= \mathbf{C}_1 \cdot \mathbf{d}^\Delta + \mathbf{C}_1 \cdot \mathbf{o}_n^\Delta. \end{aligned}$$

Since $\max(C_1) < 2^{d+1}$, $\max(C_2) < 2^d$, $\max(\mathbf{o}_n^\Delta) \leq b_1$, and $\max(\mathbf{d}^\Delta) \leq b$, we have $\max(\mathbf{C}_1 \cdot (\mathbf{d}^\Delta + \mathbf{o}_n^\Delta)) < d2^{d+1}(b_1 + b)$ and $\max(\mathbf{e}) = \max(\mathbf{C}_2 \cdot (\mathbf{d}^\Delta + \mathbf{o}_n^\Delta)) < d2^d(b_1 + b)$.

Let $\mathbf{e}^\Delta = \mathbf{e} \upharpoonright \Delta$ and $\mathbf{e}^\Gamma = \mathbf{e} \upharpoonright \Gamma$. Let $\mathbf{e}_2^\Delta = \mathbf{e}^\Delta + \mathbf{t}^\Delta$. Since $\max(\mathbf{e}) < d2^d(b_1 + b)$ and $\max \mathbf{t}^\Delta \leq d2^d(b_1 + b)$, we have that $\max(\mathbf{e}_2^\Delta) \leq d2^{d+1}(b_1 + b)$.

Apply Lemma 3.8 on \mathbf{e}_2^Δ , which says that, letting $\mathbf{p} \in \mathbb{N}^\Delta$ be defined

$$\begin{aligned}\mathbf{p}(\bar{s}) &= d2^{d+1} \cdot \max(b_1, \mathbf{e}_2^\Delta) \leq d2^{d+1} \cdot \max(b_1, d2^{d+1}(b_1 + b)) \\ &= d2^{d+1} \cdot d2^{d+1}(b_1 + b) \\ &= d^2 2^{2d+2}(b_1 + b)\end{aligned}$$

for all $\bar{s} \in \Lambda$, there is a transition sequence q_2 such that

$$\begin{aligned}\mathbf{p} + \mathbf{x}_n &\Longrightarrow_{q_2} \mathbf{p} + \mathbf{o}_n^\Gamma + \tilde{\mathbf{C}}_3 \cdot (\mathbf{o}_n^\Delta) - \tilde{\mathbf{C}}_3 \cdot \mathbf{e}_2^\Delta + \mathbf{e}_2^\Delta \\ &= \mathbf{p} + \mathbf{o}_n^\Gamma + \tilde{\mathbf{C}}_3 \cdot (\mathbf{o}_n^\Delta) - \tilde{\mathbf{C}}_3 \cdot \mathbf{e}_2^\Delta + \mathbf{e}^\Delta + \mathbf{t}^\Delta.\end{aligned}$$

Since $b_2 \geq \max(|\Lambda| \cdot b_1 + d2^d \cdot \max(b_1, \mathbf{e}_2^\Delta) \cdot |\Lambda|^2, (d^2 + 1)2^{2d+1}(b_1 + b))$, we have $b_2 \geq (d^2 + 1)2^{2d+1}(b_1 + b)$. Since $\mathbf{x}_n(\bar{s}) \geq b_2$ for all $\bar{s} \in \Lambda$, $\mathbf{x}_n \geq \mathbf{p}$, so by additivity

$$2\mathbf{x}_n \Longrightarrow_{q_2} \mathbf{x}_n + \mathbf{o}_n^\Gamma + \tilde{\mathbf{C}}_3 \cdot (\mathbf{o}_n^\Delta) - \tilde{\mathbf{C}}_3 \cdot \mathbf{e}^\Delta + \mathbf{e}^\Delta + \mathbf{t}^\Delta.$$

For large enough n , for all $\bar{s} \in \Gamma$, $\mathbf{o}_n^\Gamma(\bar{s}) \geq \mathbf{e}^\Gamma(\bar{s})$. Define $\hat{\mathbf{o}}_n^\Gamma = \mathbf{o}_n^\Gamma - \mathbf{e}^\Gamma$. Then $\mathbf{o}_n^\Gamma + \mathbf{e}^\Delta = \hat{\mathbf{o}}_n^\Gamma + \mathbf{e}$. Therefore $2\mathbf{x}_n \Longrightarrow_{q_2} \mathbf{x}_n + \hat{\mathbf{o}}_n^\Gamma + \tilde{\mathbf{C}}_3 \cdot \mathbf{o}_n^\Delta - \tilde{\mathbf{C}}_3 \cdot \mathbf{e}^\Delta + \mathbf{e}^\Delta + \mathbf{t}^\Delta$. Recall that the transition sequence p_n takes $\mathbf{x}_n \Longrightarrow_{p_n} \mathbf{o}_n$. By additivity we have (the relevant parts of the configuration needed for the subsequence transitions are underlined)

$$\begin{aligned}\underline{2\mathbf{x}_n} + \mathbf{d}^\Delta &\Longrightarrow_{q_2} \underline{\mathbf{x}_n} + \hat{\mathbf{o}}_n^\Gamma + \tilde{\mathbf{C}}_3 \cdot \mathbf{o}_n^\Delta - \tilde{\mathbf{C}}_3 \cdot \mathbf{e}^\Delta + \mathbf{e} + \mathbf{t}^\Delta + \mathbf{d}^\Delta \\ &\Longrightarrow_{p_n} \underline{\mathbf{o}_n} + \hat{\mathbf{o}}_n^\Gamma + \tilde{\mathbf{C}}_3 \cdot \mathbf{o}_n^\Delta - \tilde{\mathbf{C}}_3 \cdot \mathbf{e}^\Delta + \mathbf{d}^\Delta + \mathbf{e} + \mathbf{t}^\Delta \\ &= \mathbf{o}_n^\Gamma + \hat{\mathbf{o}}_n^\Gamma + \tilde{\mathbf{C}}_3 \cdot \mathbf{o}_n^\Delta - \tilde{\mathbf{C}}_3 \cdot \mathbf{e}^\Delta + \underline{\mathbf{o}_n^\Delta + \mathbf{d}^\Delta + \mathbf{e} + \mathbf{t}^\Delta} \\ &\Longrightarrow_{q_1} \mathbf{o}_n^\Gamma + \underline{\hat{\mathbf{o}}_n^\Gamma} + \tilde{\mathbf{C}}_3 \cdot \mathbf{o}_n^\Delta - \tilde{\mathbf{C}}_3 \cdot \mathbf{e}^\Delta + \mathbf{C}_1 \cdot \mathbf{d}^\Delta + \mathbf{C}_1 \cdot \mathbf{o}_n^\Delta + \mathbf{t}^\Delta \\ &= 2\mathbf{o}_n^\Gamma + \tilde{\mathbf{C}}_3 \cdot \mathbf{o}_n^\Delta - \tilde{\mathbf{C}}_3 \cdot \mathbf{e}^\Delta + \mathbf{C}_1 \cdot \mathbf{d}^\Delta + \mathbf{C}_1 \cdot \mathbf{o}_n^\Delta - \mathbf{e}^\Gamma + \mathbf{t}^\Delta\end{aligned}$$

Recall that $\mathbf{e} = \mathbf{C}_2 \cdot (\mathbf{d}^\Delta + \mathbf{o}_n^\Delta)$, so $\mathbf{e}^\Gamma = \mathbf{C}_2 \cdot (\mathbf{d}^\Delta + \mathbf{o}_n^\Delta) \upharpoonright \Gamma$. Thus, the last configuration above is

$$2\mathbf{o}_n^\Gamma + \tilde{\mathbf{C}}_3 \cdot \mathbf{o}_n^\Delta - \tilde{\mathbf{C}}_3 \cdot (\mathbf{C}_2 \cdot (\mathbf{d}^\Delta + \mathbf{o}_n^\Delta) \upharpoonright \Delta) + \mathbf{C}_1 \cdot \mathbf{d}^\Delta + \mathbf{C}_1 \cdot \mathbf{o}_n^\Delta - \mathbf{C}_2 \cdot (\mathbf{d}^\Delta + \mathbf{o}_n^\Delta) \upharpoonright \Gamma + \mathbf{t}^\Delta$$

Let \mathbf{C}_2^Γ be \mathbf{C}_2 with the rows corresponding to Δ removed (so that for any $\mathbf{v} \in \mathbb{N}^\Delta$, $\mathbf{C}_2^\Gamma \cdot \mathbf{v} \in \mathbb{Z}^\Gamma$; this has the same effect as restricting the output vector to Γ , as above with $\mathbf{C}_2 \cdot (\mathbf{d}^\Delta + \mathbf{o}_n^\Delta) \upharpoonright \Gamma$). Similarly, let \mathbf{C}_2^Δ be \mathbf{C}_2 with the rows corresponding to Γ removed, and let $\tilde{\mathbf{C}}_4 \in \mathbb{Z}^{\Gamma \times \Delta}$ be $\tilde{\mathbf{C}}_3 \cdot \mathbf{C}_2^\Delta$. Then the above configuration is

$$\begin{aligned}2\mathbf{o}_n^\Gamma + \tilde{\mathbf{C}}_3 \cdot \mathbf{o}_n^\Delta - \tilde{\mathbf{C}}_4 \cdot (\mathbf{d}^\Delta + \mathbf{o}_n^\Delta) + \mathbf{C}_1 \cdot \mathbf{d}^\Delta + \mathbf{C}_1 \cdot \mathbf{o}_n^\Delta - \mathbf{C}_2^\Gamma \cdot (\mathbf{d}^\Delta + \mathbf{o}_n^\Delta) + \mathbf{t}^\Delta \\ = 2\mathbf{o}_n^\Gamma + \tilde{\mathbf{C}}_3 \cdot \mathbf{o}_n^\Delta - \tilde{\mathbf{C}}_4 \cdot \mathbf{o}_n^\Delta + \mathbf{C}_1 \cdot \mathbf{o}_n^\Delta - \mathbf{C}_2^\Gamma \cdot \mathbf{o}_n^\Delta - \tilde{\mathbf{C}}_4 \cdot \mathbf{d}^\Delta + \mathbf{C}_1 \cdot \mathbf{d}^\Delta - \mathbf{C}_2^\Gamma \cdot \mathbf{d}^\Delta + \mathbf{t}^\Delta.\end{aligned}$$

Letting $\tilde{\mathbf{D}}_1 = \tilde{\mathbf{C}}_3 - \tilde{\mathbf{C}}_4 + \mathbf{C}_1 - \mathbf{C}_2^\Gamma$ and $\tilde{\mathbf{D}}_2 = -\tilde{\mathbf{C}}_4 + \mathbf{C}_1 - \mathbf{C}_2^\Gamma$, the above is $2\mathbf{o}_n^\Gamma + \tilde{\mathbf{D}}_1 \cdot \mathbf{o}_n^\Delta + \tilde{\mathbf{D}}_2 \cdot \mathbf{d}^\Delta + \mathbf{t}^\Delta$. Since $\max(\mathbf{C}_2) < 2^d$, and $\text{amax}(\tilde{\mathbf{C}}_3) < 2^{d+1}$, we can conclude that $\text{amax}(\tilde{\mathbf{C}}_4) < d2^{2d+1}$. Since $\max(\mathbf{C}_1) < 2^{d+1}$ and \mathbf{C}_1 and \mathbf{C}_2 are both nonnegative, we know that $\text{amax}(\mathbf{C}_1 - \mathbf{C}_2^\Gamma) < 2^{d+1}$. Thus, $\text{amax}(\tilde{\mathbf{D}}_1) < 2^{d+1} + d2^{2d+1} + 2^{d+1} < d2^{2d+2}$, and similarly for $\text{amax}(\tilde{\mathbf{D}}_2)$. \square

We will use the same population protocol used in the previous two examples, \mathcal{P} , to demonstrate how Lemma 3.7 and Lemma 3.8 can be used in conjunction with each other. We choose the special case of target $\mathbf{t}^\Delta = \mathbf{0}$. We will show that if a transition sequence arrives at a configuration without any b_2 -bottleneck transitions, then we can also bring all counts of states in Δ to zero. In the

statement of Lemma 3.9, we define Γ and Δ over an infinite sequence. Here, will we be looking at a single population protocol with configuration \mathbf{x}_n that satisfies the constraints of Lemma 3.9. Let \mathbf{C}_1 be defined as in the proof of Lemma 3.9

We will begin with two copies of \mathbf{x}_n plus \mathbf{d}^Δ which contains additional agents from Δ that need to be removed.

$$\mathbf{x}_n = \begin{bmatrix} 10 \\ 10 \\ 10 \\ 10 \\ 10 \end{bmatrix}$$

$$\mathbf{d}^\Delta = \begin{bmatrix} 2 \\ 0 \\ 0 \end{bmatrix}$$

The second copy of \mathbf{x}_n will serve the same purpose as \mathbf{p} in Lemma 3.8.

Using Lemma 3.8, one copy of \mathbf{x}_n will transition via transition sequence q_2 to $\hat{\mathbf{o}}_n^\Gamma + \tilde{\mathbf{C}}_3 \cdot \mathbf{o}_n^\Delta - \tilde{\mathbf{C}}_3 \cdot \mathbf{e}^\Delta + \mathbf{e}$. We use the same techniques here to build the desired \mathbf{e}^Δ as in Lemma 3.8. Then, the second copy of \mathbf{x}_n will transition “normally” via p_n to \mathbf{o}_n . The process is shown in Figure 8.

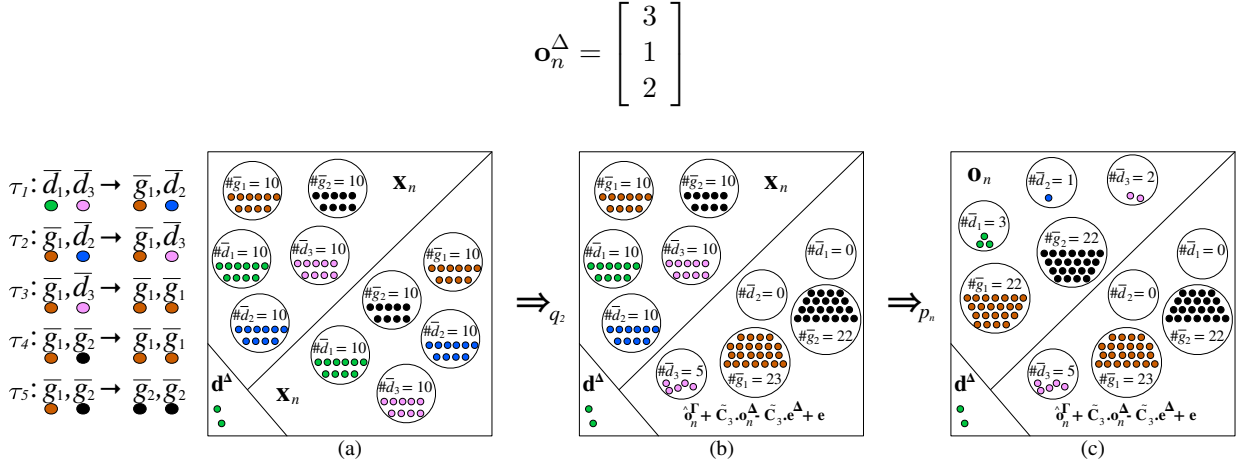


Figure 8: The first steps of Lemma 3.9 use techniques from Lemma 3.8 to push one copy of \mathbf{x}_n to the desired configuration and produce \mathbf{e} . (a) We start with two copies of \mathbf{x}_n plus \mathbf{d}^Δ . (b) Via transition sequence q_2 based on the techniques of Lemma 3.8, we produce the $\mathbf{e} = \{5\bar{d}_3, 14\bar{g}_1\}$ needed for Lemma 3.7. (c) Via transition sequence p_n , the second copy of \mathbf{x}_n transitions without interference to \mathbf{o}_n . We end in the configuration $\mathbf{d}^\Delta + \mathbf{o}_n + \hat{\mathbf{o}}_n^\Gamma + \tilde{\mathbf{C}}_3 \cdot \mathbf{o}_n^\Delta - \tilde{\mathbf{C}}_3 \cdot \mathbf{e}^\Delta + \mathbf{e}$. Figure 9 shows how to use Lemma 3.7 to use the \mathbf{e} produced in (b) to remove states in Δ in \mathbf{o}_n in (c).

At this point, our goal is to remove all agents from \mathbf{d}^Δ and \mathbf{o}_n^Δ . Using the techniques from Lemma 3.7 along with \mathbf{e} created previously using the techniques of Lemma 3.8, \mathbf{e} has exactly the counts of agents from Δ needed to remove all of $\mathbf{d}^\Delta + \mathbf{o}_n^\Delta$. And so, $\mathbf{d}^\Delta + \mathbf{o}_n^\Delta + \mathbf{e} \implies_{q_1} \mathbf{C}_1 \cdot (\mathbf{d}^\Delta + \mathbf{o}_n^\Delta) = \mathbf{C}_1 \cdot \mathbf{d}^\Delta + \mathbf{C}_1 \cdot \mathbf{o}_n^\Delta$. The resulting final configuration contains only agents in Γ . As shown in Figure 9, all agents in Δ have been removed, as desired.

Finally, we combine Lemmas 3.6 and 3.9 into a single lemma, which is the main technical result of this subsection, used (via Corollary 3.11, which sets $\mathbf{t}^\Delta = \mathbf{0}$) for proving time lower bounds in Theorems 4.4, 6.3, 6.4, 8.4, and 8.5.

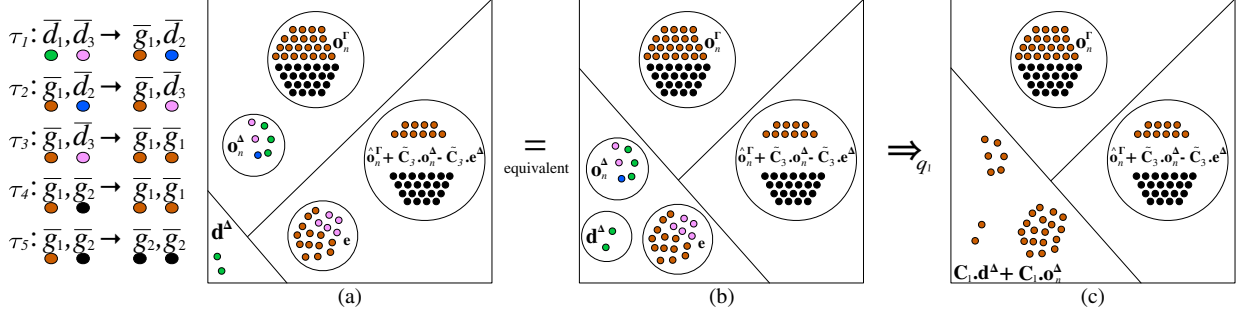


Figure 9: The second steps of Lemma 3.9 use techniques from Lemma 3.7 to remove all remaining states from Δ . (a) This is the same configuration as in panel (c) in Figure 8. (b) We visually separate \mathbf{e} from $\hat{\mathbf{o}}_n^\Gamma + \tilde{\mathbf{C}}_3 \cdot \mathbf{o}_n^\Delta - \tilde{\mathbf{C}}_3 \cdot \mathbf{e}^\Delta + \mathbf{e}$ and \mathbf{o}_n^Δ from \mathbf{o}_n^Γ . (c) Via transition sequence q_1 based on the techniques of Lemma 3.7, $\mathbf{d}^\Delta + \mathbf{o}_n^\Delta + \mathbf{e} \implies_{q_1} \mathbf{C}_1 \cdot \mathbf{d}^\Delta + \mathbf{C}_1 \cdot \mathbf{o}_n^\Delta$. The agents in \mathbf{e} react with agents from \mathbf{o}_n and \mathbf{d}^Δ to remove all remaining states in Δ . The final configuration has only states from Γ , as desired.

Lemma 3.10. *Let $\alpha > 0$. Suppose that for some set $S \subseteq \mathbb{N}^\Lambda$ and infinite set I of α -dense configurations, for all $\mathbf{i} \in I$, letting $n = \|\mathbf{i}\|$, $\Upsilon[\mathbf{i} \implies S] = o(n)$.*

Then there are matrices $\tilde{\mathbf{D}}_1 \in \mathbb{Z}^{\Gamma \times \Delta}$ and $\tilde{\mathbf{D}}_2 \in \mathbb{Z}^{\Gamma \times \Delta}$, an infinite set $N \subseteq \mathbb{N}$, and infinite nondecreasing sequences of configurations $(\mathbf{i}_n)_{n \in N}$ and $(\mathbf{o}_n)_{n \in N}$ such that the following holds. Let $\Delta = \text{bdd}((\mathbf{o}_n)_{n \in N})$, $d = |\Delta|$, and $\Gamma = \Lambda \setminus \Delta$. For each $n \in N$, let $\mathbf{o}_n^\Delta = \mathbf{o}_n \upharpoonright \Delta$ and $\mathbf{o}_n^\Gamma = \mathbf{o}_n \upharpoonright \Gamma$. Then $\text{amax}(\tilde{\mathbf{D}}_1), \text{amax}(\tilde{\mathbf{D}}_2) \leq d2^{2d+2}$ and

1. *For all $n \in N$, $\mathbf{i}_n \in I$, $\mathbf{o}_n \in S$, $\|\mathbf{i}_n\| = \|\mathbf{o}_n\| = n$, and $\mathbf{i}_n \implies \mathbf{o}_n$.*
2. *Let $b_1 = \max_{n \in N, \bar{d} \in \Delta} \mathbf{o}_n^\Delta(\bar{d})$. For all $b \in \mathbb{N}$, there is $n_b \in N$ such that, for all $n \in N$ such that $n \geq n_b$ and all $\mathbf{d}^\Delta, \mathbf{t}^\Delta \in \mathbb{N}^\Delta$ such that $\mathbf{d}^\Delta \leq b$ and $\mathbf{t}^\Delta \leq d2^d(b_1 + b)$, we have that $2\mathbf{i}_n + \mathbf{d}^\Delta \implies 2\mathbf{o}_n^\Gamma + \tilde{\mathbf{D}}_1 \cdot \mathbf{o}_n^\Delta + \tilde{\mathbf{D}}_2 \cdot \mathbf{d}^\Delta + \mathbf{t}^\Delta$.*

Proof. Let $t : \mathbb{N} \rightarrow \mathbb{N}$ be such that $\Upsilon[\mathbf{i}_n \implies S] \leq t(n)$. Define $b(n) = \frac{1}{4|\Lambda|} \sqrt{\frac{n}{t(n)}}$; note that $b(n) = \omega(1)$ since $t(n) = o(n)$ by hypothesis. By Lemma 3.6 there is $\beta > 0$ such that the following holds. There is an infinite set $N \subseteq \mathbb{N}$ and infinite sequences of configurations $(\mathbf{i}_n \in I)_{n \in N}$, $(\mathbf{x}_n \in \mathbb{N}^\Lambda)_{n \in N}$, $(\mathbf{o}_n \in \mathbb{N}^\Lambda)_{n \in N}$, where $(\mathbf{x}_n)_{n \in N}$ and $(\mathbf{o}_n)_{n \in N}$ are nondecreasing, and an infinite sequence of paths $(p_n)_{n \in N}$ such that, for all $n \in N$,

1. $\|\mathbf{i}_n\| = \|\mathbf{x}_n\| = \|\mathbf{o}_n\| = n$,
2. $\mathbf{i}_n \implies \mathbf{x}_n$,
3. $\mathbf{x}_n(\bar{s}) \geq \beta n$ for all $\bar{s} \in \Lambda$,
4. $\mathbf{x}_n \implies_{p_n} \mathbf{o}_n$, where $\mathbf{o}_n \in S$, and
5. p_n has no $b(n)$ -bottleneck transition.

Conditions (1) and (3) of Lemma 3.9 are two of the above. Let $b_1 = \max_{n \in N, \bar{d} \in \Delta} \mathbf{o}_n(\bar{d})$. Let $b_2 = \max(|\Lambda| \cdot b_1 + d2^d \cdot \max(b_1, d2^{d+1}(b_1 + b)) \cdot |\Lambda|^2, d^2 2^{2d+2}(b_1 + b))$. For sufficiently large n , $\beta n > b_2$, satisfying condition (2) of Lemma 3.9, and $b(n) \geq b_2$, satisfying condition (4) of Lemma 3.9.

Thus Lemma 3.9 tells us that there are $\tilde{\mathbf{D}}_1, \tilde{\mathbf{D}}_2 \in \mathbb{Z}^{\Gamma \times \Delta}$ such that for all $b \in \mathbb{N}$, there exists $n_b \in N$ such that for all $n \geq n_b$ such that $n \in N$ and all $\mathbf{d}^\Delta, \mathbf{t}^\Delta \in \mathbb{N}^\Delta$ such that $\mathbf{d}^\Delta \leq b$ and $\mathbf{t}^\Delta \leq d2^d(b_1 + b)$, letting $\mathbf{o}_n^\Delta = \mathbf{o}_n \upharpoonright \Delta$ and $\mathbf{o}_n^\Gamma = \mathbf{o}_n \upharpoonright \Gamma$, $2\mathbf{x}_n + \mathbf{d}^\Delta \implies 2\mathbf{o}_n^\Gamma + \tilde{\mathbf{D}}_1 \cdot \mathbf{o}_n^\Delta + \tilde{\mathbf{D}}_2 \cdot \mathbf{d}^\Delta + \mathbf{t}^\Delta$.

Since $\mathbf{i}_n \implies \mathbf{x}_n$, by additivity $2\mathbf{i}_n + \mathbf{d}^\Delta \implies 2\mathbf{o}_n^\Gamma + \tilde{\mathbf{D}}_1 \cdot \mathbf{o}_n^\Delta + \tilde{\mathbf{D}}_2 \cdot \mathbf{d}^\Delta + \mathbf{t}^\Delta$. Lemma 3.9 also gives that $\text{amax}(\tilde{\mathbf{D}}_1), \text{amax}(\tilde{\mathbf{D}}_2) \leq d2^{2d+2}$. \square

The following corollary, which sets $\mathbf{t}^\Delta = \mathbf{0}$, is the only result of this subsection employed for our time lower bounds.

Corollary 3.11. *Let $\alpha > 0$. Suppose that for some set $S \subseteq \mathbb{N}^\Delta$ and infinite set I of α -dense configurations, for all $\mathbf{i} \in I$, letting $n = \|\mathbf{i}\|$, $\top[\mathbf{i} \implies S] = o(n)$.*

Then there are matrices $\tilde{\mathbf{D}}_1 \in \mathbb{Z}^{\Gamma \times \Delta}$ and $\tilde{\mathbf{D}}_2 \in \mathbb{Z}^{\Gamma \times \Delta}$, an infinite set $N \subseteq \mathbb{N}$, and infinite nondecreasing sequences of configurations $(\mathbf{i}_n)_{n \in N}$ and $(\mathbf{o}_n)_{n \in N}$ such that the following holds. Let $\Delta = \text{bdd}((\mathbf{o}_n)_{n \in N})$, $d = |\Delta|$, and $\Gamma = \Lambda \setminus \Delta$. For each $n \in N$, let $\mathbf{o}_n^\Delta = \mathbf{o}_n \upharpoonright \Delta$ and $\mathbf{o}_n^\Gamma = \mathbf{o}_n \upharpoonright \Gamma$. Then $\text{amax}(\tilde{\mathbf{D}}_1), \text{amax}(\tilde{\mathbf{D}}_2) \leq d2^{2d+2}$ and

1. *For all $n \in N$, $\mathbf{i}_n \in I$, $\mathbf{o}_n \in S$, $\|\mathbf{i}_n\| = \|\mathbf{o}_n\| = n$, and $\mathbf{i}_n \implies \mathbf{o}_n$.*
2. *For all $b \in \mathbb{N}$, there is $n_b \in N$ such that, for all $n \in N$ such that $n \geq n_b$ and all $\mathbf{d}^\Delta \in \mathbb{N}^\Delta$ such that $\max(\mathbf{d}^\Delta) \leq b$, we have $2\mathbf{i}_n + \mathbf{d}^\Delta \implies 2\mathbf{o}_n^\Gamma + \tilde{\mathbf{D}}_1 \cdot \mathbf{o}_n^\Delta + \tilde{\mathbf{D}}_2 \cdot \mathbf{d}^\Delta$.*

For all $n \in N$, define $\mathbf{v}_n^\Gamma = 2\mathbf{o}_n^\Gamma + \tilde{\mathbf{D}}_1 \cdot \mathbf{o}_n^\Delta + \tilde{\mathbf{D}}_2 \cdot \mathbf{d}^\Delta$ to be the configuration reached in part (2) of the conclusion. We note two important properties of \mathbf{v}_n^Γ :

1. $\mathbf{v}_n^\Gamma \in \mathbb{N}^\Gamma$ (i.e., it has no states in Δ).⁸ This is crucial for the final statement of Observation 3.12 below arguing that \mathbf{v}_n^Γ is stable.
2. There is a constant $c \in \mathbb{N}$ such that for all $\bar{s} \in \Gamma$ and $n \in N$, $|\mathbf{v}_n^\Gamma(\bar{s}) - 2\mathbf{o}_n(\bar{s})| \leq c$, i.e., the counts of Γ states in \mathbf{v}_n^Γ are within a constant c of those in $2\mathbf{o}_n$. The constant c depends on the sequence (\mathbf{o}_n) , which defines $\max(\mathbf{o}_n^\Delta)$, and c also depends on the bound b on $\max(\mathbf{d}^\Delta)$, as well as the values of entries of $\tilde{\mathbf{D}}_1$ and $\tilde{\mathbf{D}}_2$, but crucially c does *not* depend on n .

3.5 Stable configurations and unbounded states

In this subsection, let \mathcal{C} be a function-computing or function-approximating population protocol with input states Σ and output state \bar{y} , taking “stable” in the next two observations to mean stable with respect to \bar{y} . The set Δ (and its complement Γ) referenced frequently in previous sections will play a key role the main proof as the set of states with bounded counts in some infinite sequence of reachable stable configurations.

The next observation states that if we have a stable configuration \mathbf{o} , and we modify it by reducing the counts of states that are already “small” (contained in Δ) and changing in either direction the counts of states that are “large” (contained in Γ), then the resulting configuration \mathbf{v} is also stable.

Observation 3.12. *If there is an infinite nondecreasing sequence (\mathbf{o}_n) of stable configurations such that $\Gamma = \text{unbdd}((\mathbf{o}_n))$ and $\Delta = \text{bdd}((\mathbf{o}_n))$, for every $\mathbf{v} \in \mathbb{N}^\Delta$ such that $\mathbf{v} \upharpoonright \Delta \leq \mathbf{o}_n \upharpoonright \Delta$ for some $n \in \mathbb{N}$, \mathbf{v} is stable. In particular, any $\mathbf{v} \in \mathbb{N}^\Gamma$ is stable.*

⁸ Compare this to Lemma 3.10, which allows one to “target” a particular nonzero count \mathbf{t}^Δ of Δ states to be present in the final configuration.

This follows since for sufficiently large n , $\mathbf{v} \leq \mathbf{o}_n$, and stability is closed downward. The following corollary is useful, which states that adding any amount \mathbf{u} of states in Γ to a stable configuration, as well as removing any amount \mathbf{w} of states (whether in Γ or not), keeps it stable.

Corollary 3.13. *If there is an infinite nondecreasing sequence (\mathbf{o}_n) of stable configurations such that $\Gamma = \text{unbdd}((\mathbf{o}_n))$, for every \mathbf{o}_n and every $\mathbf{u} \in \mathbb{N}^\Gamma$ and $\mathbf{w} \in \mathbb{N}^\Lambda$, $\mathbf{o}_n + \mathbf{u} - \mathbf{w}$ is stable.*

4 Predicate computation

In this section we show that a wide class of Boolean predicates cannot be stably computed in sublinear time by population protocols (without a leader). This is the class of predicates $\phi : \mathbb{N}^k \rightarrow \{0, 1\}$ that are not *eventually constant* (see definition below): for all $m \in \mathbb{N}$, there are two inputs $\mathbf{m}_0, \mathbf{m}_1 \in \mathbb{N}_{\geq m}^k$ such that $\phi(\mathbf{m}_0) \neq \phi(\mathbf{m}_1)$.

4.1 Definition of predicate computation

Computation of Boolean predicates $\phi : \mathbb{N}^k \rightarrow \{0, 1\}$ was the first type of computational problem studied in population protocols [4–7]. Compared to function computation (defined formally in Section 5), it is a bit more complex to define output, since we require a convention for converting several integer counts to a single Boolean value. However, the definition is also simpler because there is no need for initial configurations to contain quiescent states (see Section 5): whatever predicates are computable by population protocols, are computable from initial configurations containing *only* the input states [7]. Thus we have a 1-1 correspondence between inputs to ϕ and valid initial configurations.

It is worth mentioning that, using the output convention from the foundational work on predicate computation with population protocols [4–7], we cannot merely consider predicates a special case of functions with integer outputs in $\{0, 1\}$. If this were the case, then the results of this section would follow trivially from Theorem 8.5. The reason this does not work is that the output convention requires not merely to produce a single \bar{y} if and only if the answer is yes; instead it requires *all* agents to vote unanimously on a “yes” or “no” output.⁹

Formally, a *predicate-deciding leaderless population protocol* is a tuple $\mathcal{D} = (\Lambda, \delta, \Sigma, \Upsilon_1)$, where (Λ, δ) is a population protocol, $\Sigma \subseteq \Lambda$ is the set of *input states*, and $\Upsilon_1 \subseteq \Lambda$ is the set of *1-voters*. By convention, we define $\Upsilon_0 = \Lambda \setminus \Upsilon_1$ to be the set of *0-voters*. The *output* $\Phi(\mathbf{c})$ of a configuration $\mathbf{c} \in \mathbb{N}^\Lambda$ is $b \in \{0, 1\}$ if $\mathbf{c}(\bar{s}) = 0$ for all $\bar{s} \in \Upsilon_{1-b}$ (i.e., if the vote is unanimously b); the output is undefined if voters of both types are present. We say $\mathbf{o} \in \mathbb{N}^\Lambda$ is *stable* if $\Phi(\mathbf{o})$ is defined and for all $\mathbf{o}' \in \text{post}(\mathbf{o})$, $\Phi(\mathbf{o}') = \Phi(\mathbf{o})$. For all $\mathbf{m} \in \mathbb{N}^k$, define initial configuration $\mathbf{i}_\mathbf{m} \in \mathbb{N}^\Lambda$ by $\mathbf{i}_\mathbf{m} \upharpoonright \Sigma = \mathbf{m}$ and $\mathbf{i}_\mathbf{m} \upharpoonright (\Lambda \setminus \Sigma) = \mathbf{0}$. Call such an initial configuration *valid*. For any valid initial configuration $\mathbf{i}_\mathbf{m} \in \mathbb{N}^\Lambda$ and predicate $\phi : \mathbb{N}^k \rightarrow \{0, 1\}$, let $S_{\mathbf{i}_\mathbf{m}, \phi} = \{\mathbf{o} \in \mathbb{N}^\Lambda \mid \mathbf{i}_\mathbf{m} \Longrightarrow \mathbf{o}, \mathbf{o} \text{ is stable, and } \Phi(\mathbf{o}) = \phi(\mathbf{m})\}$. A population protocol *stably decides*¹⁰ a predicate $\phi : \mathbb{N}^k \rightarrow \{0, 1\}$ if, for any

⁹ The reason this issue is not trivially resolved by converting the “0/1-function” output convention to the “whole population votes” convention is two-fold: 1) If \bar{y} is absent, there is no straightforward way to detect this in order to ensure that 0-voters are produced. It turns out that this output convention is equivalent if time complexity is not an issue, although this is not straightforward to prove [11]. But this leads to the second issue: 2) Even with a symmetric convention where a single \bar{n} state is present to represent output 0, and a single \bar{y} state to represent output 1, it takes at least linear time to convert all other voters by a standard scheme where the single agent representing output directly interacts with all other agents.

¹⁰ The original definition [4] used the term *stably compute*, which we reserve for integer-valued function computation.

valid initial configuration $\mathbf{i}_m \in \mathbb{N}^\Lambda$, $\Pr[\mathbf{i}_m \implies S_{\mathbf{i}_m, \phi}] = 1$. This is equivalent to requiring that for all $\mathbf{c} \in \text{post}(\mathbf{i}_m)$, there is $\mathbf{o} \in \text{post}(\mathbf{c})$ such that \mathbf{o} is stable and $\Phi(\mathbf{o}) = \phi(\mathbf{m})$.

For example, the protocol defined by transitions

$$\begin{aligned} \bar{x}_1, \bar{x}_2 &\rightarrow \bar{q}_1, \bar{q}_2 \\ \bar{x}_1, \bar{q}_2 &\rightarrow \bar{x}_1, \bar{q}_1 \\ \bar{x}_2, \bar{q}_1 &\rightarrow \bar{x}_2, \bar{q}_2 \\ \bar{q}_1, \bar{q}_2 &\rightarrow \bar{q}_1, \bar{q}_1 \end{aligned}$$

if $\Upsilon_1 = \{\bar{x}_1, \bar{q}_1\}$ and $\Upsilon_0 = \{\bar{x}_2, \bar{q}_2\}$, decides whether $m_1 = \mathbf{i}(\bar{x}_1) \geq m_2 = \mathbf{i}(\bar{x}_2)$. The first transition stops once the less numerous input state is gone. If \bar{x}_1 (resp. \bar{x}_2) is left over, then the second (resp. third) transition converts \bar{q}_i states to its vote. If neither is left over (i.e., if $m_1 = m_2$, requiring output 1), the fourth transition converts all \bar{q}_2 states to \bar{q}_1 .

Relation to prior work. Alistarh, Aspnes, Eisenstat, Gelashvili, and Rivest [1] showed a linear-time lower bound on any leaderless population protocol deciding the majority predicate. Their technique is based on showing that after adding enough of the input in the minority to change it to the majority, the effect of this addition can be effectively nullified by surgery of the transition sequence, yielding a stable configuration with the original (now incorrect) answer. The technique can be extended easily to show various other specific predicates, such as equality and parity, also require linear time. We use the same technique of finding pairs of inputs with opposite correct answers and apply a similar transition sequence surgery. The main difficulty in showing Theorem 4.4, which covers the class of *all* predicates that are semilinear but not eventually constant (see below), is to identify a common characteristic, derived from the semilinear structure of the predicate computed, that can be exploited to find an infinite sequence of pairs of inputs that are all α -dense for some fixed $\alpha > 0$. Subsection 4.2 shows how this structure can be used.

4.2 Eventually constant predicates

Let $\phi : \mathbb{N}^k \rightarrow \{0, 1\}$, and for $b \in \{0, 1\}$, define $\phi^{-1}(b) = \{\mathbf{m} \in \mathbb{N}^k \mid \phi(\mathbf{m}) = b\}$ to be the set of inputs on which ϕ outputs b . We say ϕ is *eventually constant* if there is $m_0 \in \mathbb{N}$ such that ϕ is constant on $\mathbb{N}_{\geq m_0}^k = \{\mathbf{m} \in \mathbb{N}^k \mid (\forall i \in \{1, \dots, k\}) \mathbf{m}(i) \geq m_0\}$, i.e., either $\phi^{-1}(0) \cap \mathbb{N}_{\geq m_0}^k = \emptyset$ or $\phi^{-1}(1) \cap \mathbb{N}_{\geq m_0}^k = \emptyset$. In other words, although ϕ may have an infinite number of each output, “sufficiently far from the boundary” (where all coordinates exceed m_0), only one output appears.

The main result of this section, Theorem 4.4, concerns eventually constant predicates as defined above. However, our proof technique requires reasoning about infinitely many inputs $\mathbf{m} \in \mathbb{N}^k$ that are α -dense for some $\alpha > 0$. A predicate $\phi : \mathbb{N}^k \rightarrow \{0, 1\}$ can be not eventually constant, yet for any fixed $\alpha > 0$, have all but finitely many α -dense inputs map to a single output. For example, the predicate $\phi(\mathbf{m}) = 1 \iff \mathbf{m}(1) = \mathbf{m}(2)^2$ is not eventually constant, yet for any fixed $\alpha > 0$, all but finitely many α -dense inputs \mathbf{m} have $\phi(\mathbf{m}) = 0$. The rest of this section shows that for *semilinear* predicates ϕ , if ϕ is not eventually constant, then we *can* find infinitely many α -dense inputs mapping to each output. The actual requirements we need to prove Theorem 4.4 are a bit more technical and are captured in Corollary 4.3.

Given $D \subseteq \mathbb{N}^k$, we say ϕ is *almost constant on D* if either $|\phi^{-1}(0) \cap D| < \infty$ or $|\phi^{-1}(1) \cap D| < \infty$. In other words, ϕ is constant on D except for a finite number of counterexamples in D .¹¹ For $\alpha > 0$

¹¹ Note that almost constant is a stricter requirement than eventually constant, since the latter allows infinitely many counterexamples so long as at least one component is “small”.

and $k \in \mathbb{N}$, let $D_\alpha^k = \{\mathbf{m} \in \mathbb{N}^k \mid \mathbf{m} \text{ is } \alpha\text{-dense}\}$. Say that a predicate $\phi : \mathbb{N}^k \rightarrow \{0, 1\}$ is α -densely almost constant if ϕ is almost constant on D_α^k . Say that ϕ is densely almost constant if for all $\alpha > 0$, ϕ is α -densely almost constant.

The following proof uses the definition of semilinear given in Section 8 in terms of finite unions of periodic cosets.

Lemma 4.1. *If $\phi : \mathbb{N}^k \rightarrow \{0, 1\}$ is semilinear and densely almost constant, then ϕ is eventually constant.*

Proof. We prove this by contrapositive. Assume ϕ is semilinear and not eventually constant. We will show ϕ is not densely almost constant.

For $b \in \{0, 1\}$, let $I_b = \phi^{-1}(b)$ be the set of inputs mapping to output b . Since ϕ is not eventually constant, for all $m_0 \in \mathbb{N}$, for both $b \in \{0, 1\}$, $|I_b \cap \mathbb{N}_{\geq m_0}^k| \neq \emptyset$. If for some m_0 , $|I_b \cap \mathbb{N}_{\geq m_0}^k| < \infty$, then for sufficiently large m'_0 , we would have $|I_b \cap \mathbb{N}_{\geq m'_0}^k| = \emptyset$, a contradiction. So in fact, for all $m_0 \in \mathbb{N}$, for both $b \in \{0, 1\}$, $|I_b \cap \mathbb{N}_{\geq m_0}^k| = \infty$.

Since ϕ is semilinear, so is I_b , so expressible as a finite union $I_b = \bigcup_{i=1}^p P_i$, of p periodic cosets P_1, \dots, P_p . Since $|I_b \cap \mathbb{N}_{\geq m_0}^k| = \infty$ for all $m_0 \in \mathbb{N}$, without loss of generality, $|P_1 \cap \mathbb{N}_{\geq m_0}^k| = \infty$ for all $m_0 \in \mathbb{N}$ as well. Let $\mathbf{b}, \mathbf{p}_1, \dots, \mathbf{p}_l \in \mathbb{N}^k$ be such that $P_1 = \{\mathbf{b} + n_1 \mathbf{p}_1 + \dots + n_l \mathbf{p}_l \mid n_1, \dots, n_l \in \mathbb{N}\}$.

Let $\mathbf{p} = \mathbf{p}_1 + \dots + \mathbf{p}_l$. Then $\mathbf{p}(i) > 0$ for all $i \in \{1, \dots, k\}$. (Otherwise, we would have $(\exists i)(\forall j) \mathbf{p}_j(i) = 0$ and P_1 could not be arbitrarily large on component i , so it could not intersect $\mathbb{N}_{\geq m_0}^k$ for all m_0 .) Letting $\alpha' = \min(\mathbf{p})/\|\mathbf{p}\|$, we have that for all $r \in \mathbb{N}$, $r\mathbf{p}$ is α' -dense. Let $\alpha = \alpha'/2$. Then for sufficiently large r , $\mathbf{b} + r\mathbf{p}$ is α -dense. Since $(\mathbf{b} + r\mathbf{p}) \in P_1 \subseteq I_b$, this shows that I_b has infinitely many α -dense points. Since $b \in \{0, 1\}$ was arbitrary, ϕ is not α -densely almost constant, hence not densely almost constant. \square

For $i \in \{1, \dots, k\}$, let $\mathbf{u}_i \in \mathbb{N}^k$ be the unit vector such that $\mathbf{u}_i(i) = 1$ and $\mathbf{u}_i(j) = 0$ for all $j \neq i$.

Lemma 4.2. *Let $\phi : \mathbb{N}^k \rightarrow \{0, 1\}$. If ϕ is not densely almost constant, then there is $\alpha > 0$ and an infinite subset $D \subseteq D_\alpha^k$ so that one of the following two conditions holds.*

1. For all $\mathbf{m} \in D$, $\phi(\mathbf{m}) \neq \phi(2\mathbf{m})$.
2. There is $i \in \{1, \dots, k\}$ such that for all $\mathbf{m} \in D$, $\phi(\mathbf{m}) \neq \phi(\mathbf{m} + \mathbf{u}_i)$ and $\phi(\mathbf{m}) \neq \phi(2(\mathbf{m} + \mathbf{u}_i))$.

Proof. Since ϕ is not densely almost constant, for some $\alpha > 0$, $|\phi^{-1}(0) \cap D_\alpha^k| = |\phi^{-1}(1) \cap D_\alpha^k| = \infty$. So for infinitely many $\mathbf{m} \in D_\alpha^k$, there is $i_{\mathbf{m}} \in \{1, \dots, k\}$ such that $\phi(\mathbf{m}) \neq \phi(\mathbf{m} + \mathbf{u}_{i_{\mathbf{m}}})$. In other words, since each output $b \in \{0, 1\}$ is supported on infinitely many points in D_α^k , there must be infinitely many pairs of adjacent points with opposite output (“adjacent” meaning that the points differ on exactly one coordinate, and that difference is 1). By the pigeonhole principle there is $i \in \{1, \dots, k\}$ so that some infinite subset of these \mathbf{m} use the same coordinate $i_{\mathbf{m}} = i$, so consider only this infinite subset $D' \subseteq D_\alpha^k$.

If $\phi(\mathbf{m}) \neq \phi(2\mathbf{m})$ for infinitely many $\mathbf{m} \in D'$, take D to be this infinite subset, and we are done.

Otherwise, $\phi(\mathbf{m}) = \phi(2\mathbf{m})$ for all but finitely many $\mathbf{m} \in D'$. If $\phi(2\mathbf{m}) \neq \phi(2(\mathbf{m} + \mathbf{u}_i))$ for infinitely many of these \mathbf{m} , then $\phi(\mathbf{m}) \neq \phi(2(\mathbf{m} + \mathbf{u}_i))$ and we are done. In the remaining case, we have that $\phi(2(\mathbf{m} + \mathbf{u}_i)) = \phi(2\mathbf{m}) \neq \phi(\mathbf{m} + \mathbf{u}_i)$ for infinitely many $\mathbf{m} \in D'$. In this case, replace each such \mathbf{m} in D' with $\mathbf{m}' = \mathbf{m} + \mathbf{u}_i$, calling the resulting set D , noting that \mathbf{m}' satisfies condition (1) of the lemma since $\phi(\mathbf{m}') \neq \phi(2\mathbf{m}')$. \square

Combining Lemma 4.2 and the contrapositive of Lemma 4.1 gives the following corollary.

Corollary 4.3. *Let $\phi : \mathbb{N}^k \rightarrow \{0, 1\}$ be semilinear and not eventually constant. Then there is $\alpha > 0$ and an infinite subset $D \subseteq D_\alpha^k$ so that one of the following two conditions holds.*

1. *For all $\mathbf{m} \in D$, $\phi(\mathbf{m}) \neq \phi(2\mathbf{m})$.*
2. *There is $i \in \{1, \dots, k\}$ such that for all $\mathbf{m} \in D$, $\phi(\mathbf{m}) \neq \phi(\mathbf{m} + \mathbf{u}_i)$ and $\phi(\mathbf{m}) \neq \phi(2(\mathbf{m} + \mathbf{u}_i))$.*

4.3 Time lower bound for non-eventually-constant predicates

The following theorem shows that unless a predicate is eventually constant, it cannot be stably decided in sublinear time by a leaderless population protocol.

Theorem 4.4. *Let $\phi : \mathbb{N}^k \rightarrow \{0, 1\}$ and \mathcal{D} be a predicate-deciding leaderless population protocol that stably decides ϕ . If ϕ is not eventually constant, then \mathcal{D} takes expected time $\Omega(n)$.*

The high level intuition behind our proof technique is as follows. Sublinear time computation requires avoiding “bottlenecks”—having to go through a transition in which both states are present in small count (constant independent of the number of agents n). Traversing even a single such transition requires linear time. Corollary 3.11 shows that bottleneck-free execution sequences from α -dense initial configurations (i.e., initial configurations where every state that is present is present in at least αn count) are amenable to predictable “surgery”. Using Corollary 3.11, we show how to consume additional input states but still drive the system to the same output stable answer, and thus fool the population protocol into giving the wrong answer. Using Corollary 3.11 is rather technical and requires finding infinitely many candidate execution sequences and respective α -dense initial configurations that are “close” to other initial configurations on which the computed predicate is supposed to evaluate to the opposite answer. The reason that Theorem 4.4 holds only for not eventually constant predicates is that the initial configurations susceptible to surgery need to be α -dense, and thus we can only fool the population protocol if the predicate evaluates to both 0 and 1 “far away” from the boundaries of \mathbb{N}^k .

A bit of care is needed in picking the pairs of inputs that give a different answer. In particular, we need to ensure that any input states \bar{x}_i that we add in \mathbf{d}^Δ when applying Corollary 3.11 are actually contained in Δ , the set of states with bounded counts in all the output stable configurations in the sequence. Not all input states are contained in Δ in some cases. For instance, consider the majority-computing protocol

$$\begin{aligned} \bar{x}_1, \bar{x}_2 &\rightarrow \bar{y}, \bar{y} \\ \bar{x}_1, \bar{n} &\rightarrow \bar{x}_1, \bar{y} \\ \bar{x}_2, \bar{y} &\rightarrow \bar{x}_2, \bar{n} \\ \bar{y}, \bar{n} &\rightarrow \bar{y}, \bar{y}, \end{aligned}$$

which decides whether $\#\bar{x}_1 \geq \#\bar{x}_2$ initially, if $\phi(\bar{x}_1) = \phi(y) = 1$ and $\phi(\bar{x}_2) = \phi(n) = 0$. If the sequence of inputs picked were such that $\#\bar{x}_1 = 2\#\bar{x}_2$, then $\#\bar{x}_1$ would grow unboundedly in stable configurations, hence $\bar{x}_1 \in \Gamma$. Note, however, that $\bar{x}_2 \in \Delta$ since it would have count 0 in all such stable configurations. This helps to see how Δ depends on the choice of infinite sequences of inputs; if instead $\#\bar{x}_1 = \#\bar{x}_2/2$ in all such inputs, then $\bar{x}_1 \in \Delta$ but $\bar{x}_2 \in \Gamma$. If $\#\bar{x}_1 = \#\bar{x}_2$, then $\bar{x}_1, \bar{x}_2 \in \Delta$.

We choose such inputs \mathbf{m} as in Lemma 4.2 to be such that either $\phi(\mathbf{m}) \neq \phi(2\mathbf{m})$, or $\phi(\mathbf{m}) \neq \phi(\mathbf{m} + \mathbf{u}_i)$ and $\phi(\mathbf{m}) \neq \phi(2(\mathbf{m} + \mathbf{u}_i))$, where $\mathbf{u}_i \in \{0, 1\}^k$. In the first case, we don’t need any inputs to be in Δ ; we get a contradiction since Corollary 3.11 with $\mathbf{d}^\Delta = \mathbf{0}$ gives us a way to drive from $2\mathbf{m}$ to a stable configuration very close to (therefore having the same output as) twice the stable

configuration reached from \mathbf{m} , which gives a contradiction. In the second case, the contradiction arises between inputs \mathbf{m} and $2(\mathbf{m} + \mathbf{u}_i)$, but unlike the first case, we need to apply Corollary 3.11 with $\mathbf{d}^\Delta \neq \mathbf{0}$. The \mathbf{d}^Δ we choose corresponds to the positive entries of \mathbf{u}_i . The assumption that $\phi(\mathbf{m}) \neq \phi(\mathbf{m} + \mathbf{u}_i)$ is used to justify that those input states that are positive in \mathbf{u}_i are in fact contained in Δ , so that we are able to send their counts to 0 via Corollary 3.11 and conclude via Corollary 3.13 that the resulting configuration is stable.

Proof. Suppose for the sake of contradiction that \mathcal{D} takes expected time $o(n)$. Since \mathcal{D} stably decides ϕ , ϕ is semilinear [5], so by Corollary 4.3 there is $\alpha > 0$ and an infinite subset $D \subseteq D_\alpha^k$ so that one of the following two conditions holds.

1. For all $\mathbf{m} \in D$, $\phi(\mathbf{m}) \neq \phi(2\mathbf{m})$.
2. There is $i \in \{1, \dots, k\}$ such that for all $\mathbf{m} \in D$, $\phi(\mathbf{m}) \neq \phi(\mathbf{m} + \mathbf{u}_i)$ and $\phi(\mathbf{m}) \neq \phi(2(\mathbf{m} + \mathbf{u}_i))$.

Let $I = \{\mathbf{i}_m \in \mathbb{N}^\Lambda \mid \mathbf{m} \in D\}$ denote the set of initial configurations corresponding to D . Let $S = \{\mathbf{o} \mid (\exists \mathbf{i} \in I) \mathbf{i} \Longrightarrow \mathbf{o} \text{ and } \mathbf{o} \text{ is stable}\}$ be the set of stable configurations reachable from some initial configuration in I . By assumption we have that for each $\mathbf{i} \in I$, $\top[\mathbf{i} \Longrightarrow S] = o(n)$.

Apply Corollary 3.11. Then there are matrices $\tilde{\mathbf{D}}_1 \in \mathbb{Z}^{\Gamma \times \Delta}$ and $\tilde{\mathbf{D}}_2 \in \mathbb{Z}^{\Gamma \times \Delta}$, an infinite set $N \subseteq \mathbb{N}$, and infinite nondecreasing sequences of configurations $(\mathbf{i}_n)_{n \in N}$ and $(\mathbf{o}_n)_{n \in N}$ such that the following holds. Let $\Delta = \text{bdd}((\mathbf{o}_n)_{n \in N})$, $d = |\Delta|$, and $\Gamma = \Lambda \setminus \Delta$. For each $n \in N$, let $\mathbf{o}_n^\Delta = \mathbf{o}_n \upharpoonright \Delta$ and $\mathbf{o}_n^\Gamma = \mathbf{o}_n \upharpoonright \Gamma$. Then $\text{amax}(\tilde{\mathbf{D}}_1), \text{amax}(\tilde{\mathbf{D}}_2) \leq d2^{2d+2}$ and

1. For all $n \in N$, $\mathbf{i}_n \in I$, $\mathbf{o}_n \in S$, $\|\mathbf{i}_n\| = \|\mathbf{o}_n\| = n$, and $\mathbf{i}_n \Longrightarrow \mathbf{o}_n$.
2. For all $b \in \mathbb{N}$, there is $n_b \in N$ such that, for all $n \in N$ such that $n \geq n_b$ and all $\mathbf{d}^\Delta \in \mathbb{N}^\Delta$ such that $\max(\mathbf{d}^\Delta) \leq b$, we have $2\mathbf{i}_n + \mathbf{d}^\Delta \Longrightarrow 2\mathbf{o}_n^\Gamma + \tilde{\mathbf{D}}_1 \cdot \mathbf{o}_n^\Delta + \tilde{\mathbf{D}}_2 \cdot \mathbf{d}^\Delta$.

Let $b = 2$; below we ensure that $\max(\mathbf{d}^\Delta) \leq b$. Let $\mathbf{m} = \mathbf{i}_{n_b} \upharpoonright \Sigma$ and $\mathbf{o} = \mathbf{o}_{n_b}$.

Suppose we are in case (1) of Lemma 4.2; then $\phi(\mathbf{m}) \neq \phi(2\mathbf{m})$. Let $\mathbf{v}^\Gamma = 2\mathbf{o}_{n_b}^\Gamma + \tilde{\mathbf{D}}_1 \cdot \mathbf{o}_{n_b}^\Delta$. Then applying the above reachability argument with $\mathbf{d}^\Delta = \mathbf{0}$ gives that $2\mathbf{i}_{n_b} \Longrightarrow \mathbf{v}^\Gamma$, which is stable and has the same output as is correct for \mathbf{m} , a contradiction since it is reachable from $2\mathbf{i}_{n_b}$ representing input $2\mathbf{m}$, which should have the opposite output.

Suppose we are in case (2) of Lemma 4.2; then there is $i \in \{1, \dots, k\}$ such that for all $\mathbf{m} \in D$, $\phi(\mathbf{m}) \neq \phi(\mathbf{m} + \mathbf{u}_i)$ and $\phi(\mathbf{m}) \neq \phi(2(\mathbf{m} + \mathbf{u}_i))$. Let \bar{x}_i be the input corresponding to unit vector \mathbf{u}_i . We claim that $\bar{x}_i \in \Delta$. To see why, recall that $\phi(\mathbf{m}) \neq \phi(\mathbf{m} + \mathbf{u}_i)$. If $\bar{x}_i \notin \Delta$, then Corollary 3.13 tells us that $\mathbf{o} + \{\bar{x}_i\}$ is stable (hence the same output $\phi(\mathbf{m})$ as \mathbf{o}). However, $\mathbf{i}_{n_b} + \{\bar{x}_i\} \Longrightarrow \mathbf{o} + \{\bar{x}_i\}$, a contradiction since $\phi(\mathbf{m}) \neq \phi(\mathbf{m} + \mathbf{u}_i)$. This shows that $\bar{x}_i \in \Delta$.

Let $\mathbf{d}^\Delta = \{2\bar{x}_i\} \in \mathbb{N}^\Delta$. By the above, letting $\mathbf{v}^\Gamma = 2\mathbf{o}_{n_b}^\Gamma + \tilde{\mathbf{D}}_1 \cdot \mathbf{o}_{n_b}^\Delta + \tilde{\mathbf{D}}_2 \cdot \mathbf{d}^\Delta$, we have that $2\mathbf{i}_{n_b} + \mathbf{d}^\Delta \Longrightarrow \mathbf{v}^\Gamma$. By Corollary 3.13, \mathbf{v}^Γ is stable and reachable from the initial configuration $2\mathbf{i}_{n_b} + \mathbf{d}^\Delta$ corresponding to input $2\mathbf{m} + \mathbf{u}_i$, which should have the opposite output as \mathbf{m} . This contradicts the correctness of \mathcal{D} . \square

5 Definition of function computation with population protocols

In this section we formally define error-free (a.k.a. *stable*) function computation by population protocols. This mode of computation was discussed briefly in the first population protocols paper [4, Section 3.4], which focused more on Boolean predicate computation, and more extensively in the more general model of chemical reaction networks [15, 19].

5.1 Issues with definition of approximation and computation

There are subtle issues in formalizing these definitions for population protocols, as well as extending from exact to approximate computation, which we discuss before stating our definitions.

We focus on single-output functions $f : \mathbb{N}^k \rightarrow \mathbb{N}$. It is simple to apply all of our results to multi-output functions $f : \mathbb{N}^k \rightarrow \mathbb{N}^l$, simply by considering f as the result of combining l different single-output functions all computed in parallel by independent protocols.

The model of chemical reaction networks allows the creation and destruction of molecules via reactions such as $X \rightarrow 2Y$ and $2X \rightarrow Y$. However, all population protocol transitions occur between two agents, so to compute a function such as $f(m) = 2m$ via the transition $\bar{x}, \bar{q} \rightarrow \bar{y}, \bar{y}$ requires starting with m agents in an input state \bar{x} and at least m additional agents in a “quiescent” state \bar{q} , to ensure that there are enough agents to represent the output of the function.

Next, the time complexity models are slightly different between chemical reaction networks and population protocols. In addition to the distinction between continuous time in chemical reaction networks and discrete time in population protocols, the former model has an extra parameter called the *volume*, which dictates the rate of bimolecular reactions. To make the models have the same expected time, the “volume” of a population protocol is implicitly the total number of agents n . However, as we noted above, the total number of agents n in some cases necessarily differs from the number of agents representing *input*. This leads to a definitional quandary: is it more “fair” to measure the time as a function of the input size m (number of agents representing input), or as a function of the total system size $n \geq m$ (number of total agents, including non-input agents)? We will handle this by requiring the protocol to work no matter how large n is compared to m , but will measure expected time as a function of n , and our main theorem applies to show an expected time *lower* bound only when $n = O(m)$. This will lead to the same *asymptotic* time measurement within a multiplicative constant, whether time is measured as a function of m or n .

Finally, given the specific nature of the problem we study, that of *approximating* functions with population protocols, we must take care in what are defined as valid initial configurations. We observed that the protocol $\bar{a}, \bar{x} \rightarrow \bar{b}, \bar{y}$ and $\bar{b}, \bar{x} \rightarrow \bar{a}, \bar{q}$ can approximate the function $f(m) = \lfloor m/2 \rfloor$ within multiplicative factor γ from the initial configuration $\mathbf{i}_{m,\gamma}$ defined by $\mathbf{i}_{m,\gamma}(\bar{x}) = m$ and $\mathbf{i}_{m,\gamma}(\bar{a}) = \lfloor \gamma m \rfloor$. The initial count of state \bar{a} serves to control the approximation factor; setting it lower makes the approximation better but the stabilization time longer. What is the appropriate way to generalize this to a complete definition of function approximation? Note that the following does *not* work: declaring the set the valid initial configurations to be $\{\mathbf{i}_{m,\gamma} \mid m \in \mathbb{N}\}$, where γ is a fixed constant. For then, we could simply set $\gamma = \frac{1}{2}$ and let $\bar{y} = \bar{a}$, and have no transitions whatsoever: the initial configuration would already contain the correct amount of \bar{y} , with the computation having been done not by the protocol itself, but merely by the specification of the initial configuration. To avoid this sort of cheating by “sneaking computation into the initial configuration,” we allow the designer of the protocol to set a constant lower bound (not dependent on the input value), but *not* an upper bound, on how many agents have initial state \bar{a} .

Informally, a population protocol exactly computes a function $f : \mathbb{N}^k \rightarrow \mathbb{N}$ if, starting in a configuration with counts of agents in “input” states $\bar{x}_1, \dots, \bar{x}_k$ described by a vector $\mathbf{m} \in \mathbb{N}^k$, and sufficiently large counts of other agents in a “quiescent” state \bar{q} , the protocol is guaranteed to stabilize to exactly $f(\mathbf{m})$ agents in “output” state \bar{y} . Unlike predicate computation with a Boolean output (as studied in the foundational population protocol papers [4–7]), it is necessary to allow initial agents in the quiescent state \bar{q} , beyond those representing the input, if $f(\mathbf{m}) > \mathbf{m}(1) + \dots + \mathbf{m}(k)$, to ensure there are enough agents to represent the output. For example, the function $f(m) = 2m$ is computable via the transition $\bar{x}, \bar{q} \rightarrow \bar{y}, \bar{y}$ if $\#\bar{q} \geq \#\bar{x}$ initially. However, the protocol must work in any sufficiently large population, i.e., for any sufficiently large initial count

of \bar{q} .

Informally, a protocol approximates a function if it is guaranteed to get $\#\bar{y}$ “close” to the correct output value, where the closeness is controlled by the initial count of an “approximation” state \bar{a} . (Actually, there is no requirement that the approximation error *must* depend on $\#\bar{a}$, but the definition is motivated by the protocols of Theorem 7.2, which do use the approximation state in this way.) In particular, the protocol may stabilize to different values of $\#\bar{y}$ on different runs starting from the same initial configuration (so the protocol may not stably compute any function).

However, the protocol is required to stabilize its output count to *some* value on each run. An alternative formulation would relax this requirement, and merely require that $\#\bar{y}$ eventually enters, and never again leaves, an interval surrounding the correct value, while continuing to change the value of $\#\bar{y}$ indefinitely. However, our proof technique is based on the idea that certain (carefully chosen) input states must be absent, or very low count, in stable configurations. Under this more relaxed definition, this would not be the case. For example, for any $\gamma > 0$, there is a $t \in \mathbb{R}_{\geq 1}$ such that, for any m , starting with $\#\bar{x} = m$, the transition $\bar{x}, \bar{x} \rightarrow \bar{y}, \bar{q}$ in expected time t gets $\#\bar{y}$ to the interval $[m/2 - \gamma m/2, m/2]$, so in *constant* time the protocol “stabilizes to that interval”, even though the value of $\#\bar{y}$ may still change within that interval, so it is not stable by the stricter definition. However, at the time that $\#\bar{y}$ enters the interval, $\#\bar{x} = \gamma m$, whereas our proof technique requires $\#\bar{x} = O(1)$ in any stable configuration. In this example, once $\#\bar{y}$ stabilizes (i.e., stops changing), then $\#\bar{x}$ is either 0 or 1.

Sections 5.2 and 5.3 give formal definitions of these concepts.

5.2 Definition of exact function computation

A *function-computing leaderless population protocol* is a tuple $\mathcal{C} = (\Lambda, \delta, \Sigma, \bar{y}, \bar{q})$, where (Λ, δ) is a population protocol, $\Sigma = \{\bar{x}_1, \dots, \bar{x}_k\} \subset \Lambda$ is the set of *input states*, $\bar{y} \in \Lambda$ is the *output state*, and $\bar{q} \in \Lambda \setminus \Sigma$ is the *quiescent state*. We say that a configuration $\mathbf{o} \in \mathbb{N}^\Lambda$ is *stable* if, for all $\mathbf{o}' \in \text{post}(\mathbf{o})$, $\mathbf{o}(\bar{y}) = \mathbf{o}'(\bar{y})$, i.e., the count of \bar{y} cannot change once \mathbf{o} is reached.

Let $f : \mathbb{N}^k \rightarrow \mathbb{N}$, $\mathbf{i} \in \mathbb{N}^\Lambda$, and let $\mathbf{m} = \mathbf{i} \upharpoonright \Sigma$. We say that \mathcal{C} *stably computes* f from \mathbf{i} if, for all $\mathbf{c} \in \text{post}(\mathbf{i})$, there exists a stable $\mathbf{o} \in \text{post}(\mathbf{c})$ such that $\mathbf{o}(\bar{y}) = f(\mathbf{m})$, i.e., \mathcal{C} stabilizes to the correct output from the initial configuration \mathbf{i} . However, for any input $\mathbf{m} \in \mathbb{N}^k$, there are many initial configurations $\mathbf{i} \in \mathbb{N}^\Lambda$ representing it (i.e., such that $\mathbf{i} \upharpoonright \Sigma = \mathbf{m}$). We now formalize what sort of initial configurations \mathcal{C} is required to handle.

We say a function $q_0 : \mathbb{N}^k \rightarrow \mathbb{N}$ is *linearly bounded* if there is a constant $c \in \mathbb{N}$ such that, for all $\mathbf{m} \in \mathbb{N}^k$, $q_0(\mathbf{m}) \leq c \|\mathbf{m}\|$. We say that \mathcal{C} *stably computes* f if there is a linearly bounded function $q_0 : \mathbb{N}^k \rightarrow \mathbb{N}$ such that, for any $\mathbf{i} \in \mathbb{N}^\Lambda$, defining $\mathbf{m} = \mathbf{i} \upharpoonright \Sigma$, if $\mathbf{i}(\bar{q}) \geq q_0(\mathbf{m})$ and $\mathbf{i}(\bar{s}) = 0$ for all $\bar{s} \in \Lambda \setminus (\Sigma \cup \{\bar{q}\})$, then \mathcal{C} stably computes f from \mathbf{i} . We say that an initial configuration \mathbf{i} so defined is *valid*. Since all semilinear functions are linearly bounded [15], a linearly bounded q_0 suffices to ensure there are enough agents to represent the output of a semilinear function, even if we choose $\mathbf{i}(\bar{q}) = q_0(\mathbf{i} \upharpoonright \Sigma)$. If q_0 were *not* linearly bounded, and thus a super-linear count of state \bar{q} is required, we would essentially need to do non-semilinear computation just to initialize the population protocol.

It is well-known [7] that this is equivalent to requiring, under the randomized model in which the next interaction is between a pair of agents picked uniformly at random, that the protocol stabilizes on the correct output with probability 1. More formally, given $f : \mathbb{N}^k \rightarrow \mathbb{N}$ and $\mathbf{m} \in \mathbb{N}^k$, defining $S_{f, \mathbf{m}}^{\mathcal{C}} = \{\mathbf{o} \in \mathbb{N}^\Lambda \mid \mathbf{o} \text{ is stable and } \mathbf{o}(\bar{y}) = f(\mathbf{m})\}$, \mathcal{C} stably computes f if and only if, for all \mathbf{m} , defining \mathbf{i} with $\mathbf{i} \upharpoonright \Sigma = \mathbf{m}$ as above with $\mathbf{i}(\bar{q})$ sufficiently large, $\Pr \left[\mathbf{i} \Longrightarrow S_{f, \mathbf{m}}^{\mathcal{C}} \right] = 1$. It is also equivalent to requiring that every fair infinite execution leads to a correct stable configuration,

where an execution is *fair* if every configuration infinitely often reachable appears infinitely often in the execution.

Let $f : \mathbb{N}^k \rightarrow \mathbb{N}$ and $t : \mathbb{N} \rightarrow \mathbb{N}$. Given a function-computing leaderless population protocol \mathcal{C} that stably computes f , we say \mathcal{C} *stably computes f in expected time t* if, for all valid initial configurations \mathbf{i} of \mathcal{C} , letting $\mathbf{m} = \mathbf{i} \upharpoonright \Sigma$, $\mathbb{T} \left[\mathbf{i} \Longrightarrow S_{f,\mathbf{m}}^{\mathcal{C}} \right] \leq t(n)$.

Note that unstability is closed upwards in \mathbb{N}^Λ . In other words, for any $\mathbf{c} \in \mathbb{N}^\Lambda$, if \mathbf{c} is not stable, then no $\mathbf{c}' \geq \mathbf{c}$ is stable either. This is because \mathbf{c} is unstable if and only if there is a path p such that $\mathbf{c} \Longrightarrow_p \mathbf{d}$ and $\mathbf{c}(\bar{y}) \neq \mathbf{d}(\bar{y})$. Thus p is applicable to any $\mathbf{c}' \geq \mathbf{c}$, also changing $\# \bar{y}$ by the amount $\mathbf{d}(\bar{y}) - \mathbf{c}(\bar{y})$, so \mathbf{c}' is also unstable. By contrapositive, the set of stable configurations is then closed *downward*: for any stable \mathbf{o} , if $\mathbf{o}' \leq \mathbf{o}$, then \mathbf{o}' is also stable.

5.3 Definition of function approximation

A *function-approximating leaderless population protocol* is a tuple $\mathcal{A} = (\Lambda, \delta, \Sigma, \bar{y}, \bar{q}, \bar{a})$, where $(\Lambda, \delta, \Sigma, \bar{y}, \bar{q})$ is a function-computing population protocol and $\bar{a} \in \Lambda \setminus (\Sigma \cup \{\bar{y}, \bar{q}\})$ is the *approximation state*. Let $\epsilon, \tau \in \mathbb{N}$; intuitively τ represents the “target” (or “true”) function output, and ϵ represents the allowed approximation error. We say that a configuration $\mathbf{o} \in \mathbb{N}^\Lambda$ is ϵ - τ -*correct* if $|\mathbf{o}(\bar{y}) - \tau| \leq \epsilon$.

Let $f : \mathbb{N}^k \rightarrow \mathbb{N}$, $\epsilon \in \mathbb{N}$, $\mathbf{i} \in \mathbb{N}^\Lambda$, and let $\mathbf{m} = \mathbf{i} \upharpoonright \Sigma$. We say that \mathcal{A} *stably ϵ -approximates f from \mathbf{i}* if, for all $\mathbf{c} \in \text{post}(\mathbf{i})$, there exists a $\mathbf{o} \in \text{post}(\mathbf{c})$ that is stable and ϵ - $f(\mathbf{m})$ -correct, i.e., from the initial configuration \mathbf{i} , \mathcal{A} gets the output to stabilize to a value at most ϵ from the correct output. Let $S_{f,\mathbf{m},\epsilon}^{\mathcal{A}} = \{ \mathbf{o} \in \mathbb{N}^\Lambda \mid \mathbf{o} \text{ is stable and } \epsilon$ - $f(\mathbf{m})$ -correct $\}$. Note that \mathcal{A} stably ϵ -approximates f from \mathbf{i} if and only if $\Pr \left[\mathbf{i} \Longrightarrow S_{f,\mathbf{m},\epsilon}^{\mathcal{A}} \right] = 1$.

Let $\mathcal{E} : \mathbb{N} \rightarrow \mathbb{N}$; the choice of \mathcal{E} as a function instead of a constant reflects the idea that the approximation error is allowed to depend on the initial count $\mathbf{i}(\bar{a})$ of the approximation state \bar{a} , i.e., $\mathcal{E}(\mathbf{i}(\bar{a}))$ is the desired approximation error. We say that \mathcal{A} *stably \mathcal{E} -approximates f* if there are $a_0 \in \mathbb{N}$ and linearly bounded $q_0 : \mathbb{N}^{k+1} \rightarrow \mathbb{N}$ such that, for any $\mathbf{i} \in \mathbb{N}^\Lambda$, defining $\mathbf{m} = \mathbf{i} \upharpoonright \Sigma$, if $\mathbf{i}(\bar{a}) \geq a_0$, $\mathbf{i}(\bar{q}) \geq q_0(\mathbf{m}, \mathbf{i}(\bar{a}))$, and $\mathbf{i}(\bar{s}) = 0$ for all $\bar{s} \in \Lambda \setminus (\Sigma \cup \{\bar{q}, \bar{a}\})$, then \mathcal{A} stably $\mathcal{E}(\mathbf{i}(\bar{a}))$ -approximates f from \mathbf{i} .¹² An initial configuration \mathbf{i} so defined is *valid*.

Note that we allow the protocol to require at least a certain initial amount of \bar{a} in order to function correctly. For example, the protocol with transitions $\bar{a}, \bar{x} \rightarrow \bar{b}, \bar{y}$ and $\bar{b}, \bar{x} \rightarrow \bar{a}, \bar{q}$ stably \mathcal{E} -approximates $f(m) = \lfloor m/2 \rfloor$, where $\mathcal{E} : \mathbb{N} \rightarrow \mathbb{N}$ is the identity function, *if* at least one \bar{a} is present initially. However, setting a_0 does not imply the protocol is able to use \bar{a} as a leader: since a_0 is just a lower bound. The protocol must work correctly for *any* initial value of $\mathbf{i}(\bar{a}) \geq a_0$, where $\mathcal{E}(\mathbf{i}(\bar{a}))$ defines how close the output must be to $f(\mathbf{m})$ to be “correct”.

Note also that in this example, the approximation error increases with $\mathbf{i}(\bar{a})$ (i.e., \mathcal{E} is monotonically increasing), while the expected time to stabilization decreases with $\mathbf{i}(\bar{a})$. It is conceivable for the approximation error to *decrease* with $\mathbf{i}(\bar{a})$, or even not to depend on $\mathbf{i}(\bar{a})$, although we do not know of any examples of such protocols. Our main theorem lower bounds the approximation error as a linear function of the count of the lowest-count state present in the initial configuration, whether that is \bar{a} or not, so our proof works regardless of the precise form of \mathcal{E} .

¹²I.e., the initial count $\mathbf{i}(\bar{a})$ can influence the initial required count $\mathbf{i}(\bar{q})$, since adding more initial \bar{a} may imply that more quiescent agents are required as “fuel”. However, a_0 is constant, not a function of \mathbf{m} .

6 Sublinear-time, sublinear-error approximation of linear functions with negative or non-integer coefficients is impossible

We say a function $f : \mathbb{N}^k \rightarrow \mathbb{N}$ is \mathbb{N} -linear if there are $c_1, \dots, c_k \in \mathbb{N}$ such that for all $\mathbf{m} \in \mathbb{N}^k$, $f(\mathbf{m}) = \sum_{i=1}^k c_i \mathbf{m}(i)$.

As we consider leaderless population protocols, we need to make sure that \bar{a} does not act as a small count “leader”. Consistent with the rest of this paper, we reason about initial configurations with $\mathbf{i}(\bar{a}) \geq \alpha n$ for some $\alpha > 0$ to ensure α -density.

Let $f : \mathbb{N}^k \rightarrow \mathbb{N}$. In defining running time for function-approximating population protocols, we express the expected time as a function of *both* the total number of agents $n = \|\mathbf{i}\|$ and the initial count $\mathbf{i}(\bar{a})$ of approximation states. Let $\mathcal{E} : \mathbb{N} \rightarrow \mathbb{N}$ and $t : \mathbb{N}^2 \rightarrow \mathbb{N}$. Given a function-approximating population protocol \mathcal{A} that \mathcal{E} -approximates f , we say \mathcal{A} \mathcal{E} -approximates f in expected time t if, for all valid initial configurations \mathbf{i} of \mathcal{A} , letting $\mathbf{m} = \mathbf{i} \upharpoonright \Sigma$, $\top \left[\mathbf{i} \Longrightarrow S_{f, \mathbf{m}, \mathcal{E}(\mathbf{i}(\bar{a}))}^{\mathcal{A}} \right] \leq t(n, \mathbf{i}(\bar{a}))$.

The following theorem states that given any linear function f and any population protocol \mathcal{P} , if f has a non-integer or negative coefficient, then \mathcal{P} requires at least linear time to approximate f with sublinear error. It states this by contrapositive: if the protocol takes sublinear time, then the error $\mathcal{E} : \mathbb{N} \rightarrow \mathbb{N}$ must grow at least linearly with the initial count of approximation state \bar{a} . In particular, the initial configurations \mathbf{i} (letting $n = \|\mathbf{i}\|$) on which our argument maximizes the error have $\mathbf{i}(\bar{a}) = \Omega(n)$. Thus, the fact that $\mathcal{E}(a) \geq \gamma a$ implies that on these \mathbf{i} , the error is $\Omega(n)$.

Theorem 6.1. *Let $f : \mathbb{N}^k \rightarrow \mathbb{N}$ be a linear function that is not \mathbb{N} -linear. Let $\mathcal{E} : \mathbb{N} \rightarrow \mathbb{N}$. Let \mathcal{A} be a function-approximating leaderless population protocol that stably \mathcal{E} -approximates f in expected time t , where for some $\alpha > 0$, $t(n, \alpha n) = o(n)$. Then there is a constant $\gamma > 0$ such that, for infinitely many $a \in \mathbb{N}$, $\mathcal{E}(a) \geq \gamma a$.*

A protocol stably computing f also stably \mathcal{E} -approximates f for $\mathcal{E}(a) = 0$, so we have:

Corollary 6.2. *Let $f : \mathbb{N}^k \rightarrow \mathbb{N}$ be a linear function $f(\mathbf{m}) = \sum_{i=1}^k \lfloor c_i \mathbf{m}(i) \rfloor$, where $c_i \notin \mathbb{N}$ for some $i \in \{1, \dots, k\}$. Let \mathcal{C} be a function-computing leaderless population protocol that stably computes f . Then \mathcal{C} takes expected time $\Omega(n)$.*

This gives a complete classification of the asymptotic efficiency of computing linear functions $f(\mathbf{m}) = \sum_{i=1}^k \lfloor c_i \mathbf{m}(i) \rfloor$ with population protocols. If $c_i \in \mathbb{N}$ for all $i \in \{1, \dots, k\}$, then f is stably computable in logarithmic time by Observation 7.1. Otherwise, f requires linear time to stably compute by Corollary 6.2, or even to approximate within sublinear error by Theorem 6.1. The remainder of Section 6 is devoted to proving Theorem 6.1.

Theorem 6.1 follows directly from the two theorems in this section.

First, we show that subtraction takes linear time to approximate with sublinear error.

Theorem 6.3. *Let $f : \mathbb{N}^k \rightarrow \mathbb{N}$ be a linear function $f(\mathbf{m}) = \sum_{i=1}^k \lfloor c_i \mathbf{m}(i) \rfloor$, where $c_i < 0$ for some $i \in \{1, \dots, k\}$. Let $\mathcal{E} : \mathbb{N} \rightarrow \mathbb{N}$ and $t : \mathbb{N}^2 \rightarrow \mathbb{N}$. Let \mathcal{A} be a function-approximating leaderless population protocol that stably \mathcal{E} -approximates f in expected time t . Suppose there is $\alpha > 0$ such that $t(n, \alpha n) = o(n)$. Then there is a constant $\gamma > 0$ such that, for infinitely many $a \in \mathbb{N}$, $\mathcal{E}(a) \geq \gamma a$.*

Proof. Assume without loss of generality that $f(\mathbf{m}) = c_1 \mathbf{m}(1) - c_2 \mathbf{m}(2)$ for $c_1, c_2 > 0$; a function with more inputs can have those inputs set to 0, and the remaining two re-ordered to be the first two, to result in this f . Assume for notational ease that $c_1 = c_2 = 1$. The extension to other coefficients in $\mathbb{Q}_{>0}$ is routine, requiring us only to modify the set I_α^- below to contain configurations \mathbf{i} with $c_1 \mathbf{i}(x_1) = c_2 \mathbf{i}(x_2)$ instead of merely $\mathbf{i}(x_1) = \mathbf{i}(x_2)$.

Let $\mathcal{E} : \mathbb{N} \rightarrow \mathbb{N}$ and $t : \mathbb{N}^2 \rightarrow \mathbb{N}$. Let \mathcal{C} be a function-approximating population protocol that stably \mathcal{E} -approximates f in expected time t . Let $\alpha > 0$ be such that $t(n, \alpha n) = o(n)$ and, for all $m \in \mathbb{N}$, there is a valid α -dense initial configuration $\mathbf{i} \in \mathbb{N}^\Lambda$ such that $\mathbf{i}(\bar{x}_1) = \mathbf{i}(\bar{x}_2) = m$. Such an α exists since the quiescent threshold $q_0 : \mathbb{N}^k \rightarrow \mathbb{N}$ is linearly bounded.

Let I_α^- be the set of all such α -dense valid \mathbf{i} , where $\mathbf{i}(x_1) = \mathbf{i}(x_2)$ (i.e., the set of α -dense initial configurations representing an input \mathbf{m} to f such that $f(\mathbf{m}) = 0$). Let $S = \{\mathbf{o} \mid (\exists \mathbf{i} \in I_\alpha^-) \mathbf{i} \Longrightarrow \mathbf{o} \text{ and } \mathbf{o} \text{ is stable}\}$ be the set of stable configurations reachable from some initial configuration in I_α^- . By assumption we have that for each $\mathbf{i} \in I_\alpha^-$, $\top[\mathbf{i} \Longrightarrow S] = o(n)$.

Apply Corollary 3.11. Then there are matrices $\tilde{\mathbf{D}}_1 \in \mathbb{Z}^{\Gamma \times \Delta}$ and $\tilde{\mathbf{D}}_2 \in \mathbb{Z}^{\Gamma \times \Delta}$, an infinite set $N \subseteq \mathbb{N}$, and infinite nondecreasing sequences of configurations $(\mathbf{i}_n)_{n \in N}$ and $(\mathbf{o}_n)_{n \in N}$ such that the following holds. Let $\Delta = \text{bdd}((\mathbf{o}_n)_{n \in N})$, $d = |\Delta|$, and $\Gamma = \Lambda \setminus \Delta$. For each $n \in N$, let $\mathbf{o}_n^\Delta = \mathbf{o}_n \upharpoonright \Delta$ and $\mathbf{o}_n^\Gamma = \mathbf{o}_n \upharpoonright \Gamma$. Then $\text{amax}(\tilde{\mathbf{D}}_1), \text{amax}(\tilde{\mathbf{D}}_2) \leq d2^{2d+2}$ and

1. For all $n \in N$, $\mathbf{i}_n \in I$, $\mathbf{o}_n \in S$, $\|\mathbf{i}_n\| = \|\mathbf{o}_n\| = n$, and $\mathbf{i}_n \Longrightarrow \mathbf{o}_n$.
2. For all $b \in \mathbb{N}$, there is $n_b \in N$ such that, for all $n \in N$ such that $n \geq n_b$ and all $\mathbf{d}^\Delta \in \mathbb{N}^\Delta$ such that $\max(\mathbf{d}^\Delta) \leq b$, we have $2\mathbf{i}_n + \mathbf{d}^\Delta \Longrightarrow 2\mathbf{o}_n^\Gamma + \tilde{\mathbf{D}}_1 \cdot \mathbf{o}_n^\Delta + \tilde{\mathbf{D}}_2 \cdot \mathbf{d}^\Delta$.

Let $b = 1$; below we ensure that $\max(\mathbf{d}^\Delta) \leq b$.

There are two cases:

$\bar{y} \in \Delta$: We claim that $\bar{x}_1 \in \Delta$. To see why, observe that increasing the initial amount of \bar{x}_1 by 1 should increase the output by 1. If we had $\bar{x} \in \Gamma$, then by Corollary 3.13, for any stable configuration \mathbf{o} , the configuration $\mathbf{o} + \{\bar{x}_1\}$ would also be stable, but would have the same output. Let \mathbf{i} be any initial configuration representing an input with $x_1 = x_2$, such that $\mathbf{i} \Longrightarrow \mathbf{o}$. Then we should have $\mathbf{o}(\bar{y}) = 0$. Then $\mathbf{i} + \{\bar{x}_1\} \Longrightarrow \mathbf{o} + \{\bar{x}_1\}$, a contradiction since the correct output to compute from $\mathbf{i} + \{\bar{x}_1\}$ is 1, but the output is also 0 in stable configuration $\mathbf{o} + \{\bar{x}_1\}$. This shows the claim that $\bar{x}_1 \in \Delta$.

Let $\mathbf{d}^\Delta = \{\bar{x}_1\}$. Let $\mathbf{w} = 2\mathbf{o}_{n_b}^\Gamma + \tilde{\mathbf{D}}_1 \cdot \mathbf{o}_{n_b}^\Delta$ and let $\mathbf{c}_1 = \tilde{\mathbf{D}}_2 \cdot \mathbf{d}^\Delta$. The above implies that $2\mathbf{i}_{n_b} + \{\bar{x}_1\} \Longrightarrow \mathbf{w} + \mathbf{c}_1$. For all $l \in \mathbb{N}$, let $\mathbf{v}_l = l(\mathbf{w} + \mathbf{c}_1)$ and $\mathbf{i}'_l = l(2\mathbf{i}_{n_b} + \{\bar{x}_1\})$. Thus $\mathbf{i}'_l \Longrightarrow \mathbf{v}_l$.

Let $\mathbf{m} = (2lm + l, 2lm) = \mathbf{i}'_l \upharpoonright \Sigma$, with correct output $f(2lm + l, 2lm) = l$, but $\mathbf{v}_l(\bar{y}) = 0$.

$\bar{y} \in \Gamma$: Then for all sufficiently large $n \geq n_b$, $2\mathbf{o}_n(\bar{y}) + \left(\min_{n' \in N} \left(\tilde{\mathbf{D}}_1 \cdot \mathbf{o}_{n'}^\Delta \right) (\bar{y}) \right) > 0$. Fix such an n .

Let $\mathbf{w} = 2\mathbf{o}_n^\Gamma + \tilde{\mathbf{D}}_1 \cdot \mathbf{o}_n^\Delta$; note that $\mathbf{w}(\bar{y}) > 0$. Letting $\mathbf{d}^\Delta = \mathbf{0}$, the above implies that $2\mathbf{i}_n \Longrightarrow \mathbf{w}$. For all $l \in \mathbb{N}$, let $\mathbf{v}_l = l\mathbf{w}$ and $\mathbf{i}'_l = 2l\mathbf{i}_n$. Thus $\mathbf{i}'_l \Longrightarrow \mathbf{v}_l$.

Let $\mathbf{m} = (2lm, 2lm) = \mathbf{i}'_l \upharpoonright \Sigma$, with correct output $f(2lm, 2lm) = 0$, but $\mathbf{v}_l(\bar{y}) \geq l$.

In each case \mathbf{v}_l is stable by Observation 3.12 and reachable from an initial configuration \mathbf{i}'_l representing input \mathbf{m} , with $|\mathbf{v}_l(\bar{y}) - f(\mathbf{m})| \geq l$, so $\mathcal{E}(2l\mathbf{i}_{n_b}(\bar{a})) \geq l$. Since $2\mathbf{i}_{n_b}(\bar{a})$ is constant with respect to l , choosing $\gamma = \frac{1}{2\mathbf{i}_{n_b}(\bar{a})}$ proves the theorem. \square

Now we show that division by a constant takes linear time to approximate with sublinear error.

Theorem 6.4. *Let $f : \mathbb{N}^k \rightarrow \mathbb{N}$ be a linear function $f(\mathbf{m}) = \sum_{i=1}^k [c_i \mathbf{m}(i)]$, where $c_i \notin \mathbb{Z}$ for some $i \in \{1, \dots, k\}$. Let $\mathcal{E} : \mathbb{N} \rightarrow \mathbb{N}$ and $t : \mathbb{N}^2 \rightarrow \mathbb{N}$. Let \mathcal{A} be a function-approximating leaderless population protocol that stably \mathcal{E} -approximates f in expected time t . Suppose there is $\alpha > 0$ such that $t(n, \alpha n) = o(n)$. Then there is a constant $\gamma > 0$ such that, for infinitely many $a \in \mathbb{N}$, $\mathcal{E}(a) \geq \gamma a$.*

Proof. If $c_i < 0$ for some i , then Theorem 6.3 applies and we are done, so assume all $c_i \geq 0$. Assume without loss of generality that $f(m) = \lfloor cm \rfloor$, where $c \in \mathbb{Q}_{>0} \setminus \mathbb{N}$; a function with more inputs can have those inputs set to 0 to result in this f . Write c in lowest terms as $\frac{p}{r}$ for $p, r \in \mathbb{Z}_{\geq 1}$, and note that $r \geq 2$ since $c \notin \mathbb{Z}$. Let the input state \bar{x}_1 be denoted simply \bar{x} .

Let $\mathcal{E} : \mathbb{N} \rightarrow \mathbb{N}$. Let $\delta > 0$ and let \mathcal{A} be a function-approximating population protocol that stably \mathcal{E} -approximates f in expected time $o(n)$. Let $\alpha > 0$ be such that $t(n, \alpha n) = o(n)$ and, for all $m \in \mathbb{N}$, there is a valid α -dense initial configuration $\mathbf{i} \in \mathbb{N}^\Delta$ with $\mathbf{i}(\bar{x}) = m$. Such an α exists since the quiescent threshold $q_0 : \mathbb{N}^k \rightarrow \mathbb{N}$ is linearly bounded.

Let I_α be the set of all such α -dense valid \mathbf{i} . Let $S = \{ \mathbf{o} \mid (\exists \mathbf{i} \in I_\alpha) \mathbf{i} \Longrightarrow \mathbf{o} \text{ and } \mathbf{o} \text{ is stable and } \mathcal{E}(\mathbf{i}(\bar{x})) - f(\mathbf{i}(\bar{x}))\text{-correct} \}$ be the set of stable “ ϵ - t ”-correct configurations reachable from some configuration $\mathbf{i} \in I_\alpha$, and where the choice of ϵ and t to define ϵ - t -correct depends on \mathbf{i} . By hypothesis we have that for each $\mathbf{i} \in I_\alpha$, $\top[\mathbf{i} \Longrightarrow S] = o(n)$.

Apply Corollary 3.11. Then there are matrices $\tilde{\mathbf{D}}_1 \in \mathbb{Z}^{\Gamma \times \Delta}$ and $\tilde{\mathbf{D}}_2 \in \mathbb{Z}^{\Gamma \times \Delta}$, an infinite set $N \subseteq \mathbb{N}$, and infinite nondecreasing sequences of configurations $(\mathbf{i}_n)_{n \in N}$ and $(\mathbf{o}_n)_{n \in N}$ such that the following holds. Let $\Delta = \text{bdd}((\mathbf{o}_n)_{n \in N})$, $d = |\Delta|$, and $\Gamma = \Lambda \setminus \Delta$. For each $n \in N$, let $\mathbf{o}_n^\Delta = \mathbf{o}_n \upharpoonright \Delta$ and $\mathbf{o}_n^\Gamma = \mathbf{o}_n \upharpoonright \Gamma$. Then $\text{amax}(\tilde{\mathbf{D}}_1), \text{amax}(\tilde{\mathbf{D}}_2) \leq d2^{2d+2}$ and

1. For all $n \in N$, $\mathbf{i}_n \in I$, $\mathbf{o}_n \in S$, $\|\mathbf{i}_n\| = \|\mathbf{o}_n\| = n$, and $\mathbf{i}_n \Longrightarrow \mathbf{o}_n$.
2. For all $b \in \mathbb{N}$, there is $n_b \in N$ such that, for all $n \in N$ such that $n \geq n_b$ and all $\mathbf{d}^\Delta \in \mathbb{N}^\Delta$ such that $\max(\mathbf{d}^\Delta) \leq b$, we have $2\mathbf{i}_n + \mathbf{d}^\Delta \Longrightarrow 2\mathbf{o}_n^\Gamma + \tilde{\mathbf{D}}_1 \cdot \mathbf{o}_n^\Delta + \tilde{\mathbf{D}}_2 \cdot \mathbf{d}^\Delta$.

Note that $\bar{y} \in \Gamma$ since $f(m)$ grows unboundedly with m and $m = \mathbf{i}_n(\bar{x})$ grows unboundedly with n . We claim that $\bar{x} \in \Delta$. To see why, observe that increasing the initial amount of \bar{x} by r should increase the output by p . If we had $\bar{x} \in \Gamma$, then by Corollary 3.13, for any stable configuration \mathbf{o} , the configuration $\mathbf{o} + \{r\bar{x}\}$ would also be stable, but would have the same output. Let \mathbf{i} be any initial configuration such that $\mathbf{i} \Longrightarrow \mathbf{o}$. Then $\mathbf{i} + \{r\bar{x}\} \Longrightarrow \mathbf{o} + \{r\bar{x}\}$, a contradiction since the correct output to compute from $\mathbf{i} + \{r\bar{x}\}$ is p larger than the correct output for \mathbf{i} . This shows the claim that $\bar{x} \in \Delta$.

Let $b = 1$; below we ensure that $\max(\mathbf{d}^\Delta) \leq b$. Let $\mathbf{w} = r \left(2\mathbf{o}_{n_b}^\Gamma + \tilde{\mathbf{D}}_1 \cdot \mathbf{o}_{n_b}^\Delta \right)$ and $\mathbf{c}_1 = \tilde{\mathbf{D}}_2 \cdot \mathbf{d}^\Delta$. The above argument shows that $2r\mathbf{i}_{n_b} \Longrightarrow \mathbf{w}$ (letting $\mathbf{d}^\Delta = \mathbf{0}$) and $2r\mathbf{i}_{n_b} + r\{\bar{x}\} \Longrightarrow \mathbf{w} + r\mathbf{c}_1$ (letting $\mathbf{d}^\Delta = \{\bar{x}\}$). Since $\mathbf{w} \in \mathbb{N}^\Gamma$ and $\mathbf{c}_1 \in \mathbb{N}^\Gamma$, both \mathbf{w} and $\mathbf{w} + r\mathbf{c}_1$ are stable by Observation 3.12. Recall $f(2rm) = \lfloor 2rmp/r \rfloor = 2mp$. We have two cases:

$\mathbf{w}(\bar{y}) \neq 2mp$: Thus $|\mathbf{w}(\bar{y}) - 2mp| \geq 1$. For all $l \in \mathbb{N}$, let $\mathbf{v}_l = l\mathbf{w}$ and $m_l = 2lrm$. Then $2lri_{n_b} \Longrightarrow \mathbf{v}_l$, and

$$|\mathbf{v}_l(\bar{y}) - f(m_l)| = |l\mathbf{w}(\bar{y}) - 2lmp| = l|\mathbf{w}(\bar{y}) - 2mp| \geq l.$$

$\mathbf{w}(\bar{y}) = 2mp$: For all $l \in \mathbb{N}$, let $\mathbf{v}_l = l(\mathbf{w} + r\mathbf{c}_1)$ and $m_l = l(2rm + r)$. Thus, $l(2r\mathbf{i}_{n_b} + r\{\bar{x}\}) \Longrightarrow \mathbf{v}_l$. Also, $f(m_l) = \lfloor lp(2rm + r)/r \rfloor = 2lmp + lu$. Note that $\mathbf{c}_1(\bar{y}) \in \mathbb{Z}$. There are two subcases:

$\mathbf{c}_1(\bar{y}) \geq 1$: Then $\mathbf{v}_l(\bar{y}) = l(\mathbf{w}(\bar{y}) + r\mathbf{c}_1(\bar{y})) \geq l(2mp + r) = 2lmp + lr$, so

$$\mathbf{v}_l(\bar{y}) - f(m_l) \geq 2lmp + lr - (2lmp + l) = lr - l \geq l.$$

$\mathbf{c}_1(\bar{y}) \leq 0$: Then $\mathbf{v}_l(\bar{y}) = l(\mathbf{w}(\bar{y}) + r\mathbf{c}_1(\bar{y})) \leq 2lmp$, so

$$f(m_l) - \mathbf{v}_l(\bar{y}) \geq 2lmp + l - 2lmp = l.$$

In each case \mathbf{v}_l is stable by Observation 3.12 and reachable from an initial configuration representing input m_l , with count $2lri_{n_b}(\bar{a})$ of \bar{a} , but $|\mathbf{v}_l(\bar{y}) - f(m_l)| \geq l$, so $\mathcal{E}(2lri_{n_b}(\bar{a})) \geq l$. Since $2lri_{n_b}(\bar{a})$ is constant with respect to l , choosing $\gamma = \frac{1}{2ri_{n_b}(\bar{a})}$ proves the theorem. \square

7 Logarithmic-time, linear-error approximation of linear functions with nonnegative rational coefficients is possible

Recall the definition of \mathbb{N} -linear function from Section 6. It is easy to see that any \mathbb{N} -linear function f can be stably computed in logarithmic time. Recall that $\bar{x}, \bar{q} \rightarrow \bar{y}, \bar{y}$ stably computes $f(m) = 2m$ in expected time $O(\log n)$. The extension to larger coefficients, e.g., $f(m) = 4m$, uses a series of transitions:

$$\begin{aligned}\bar{x}, \bar{q} &\rightarrow \bar{y}, \bar{x}' \\ \bar{x}', \bar{q} &\rightarrow \bar{y}, \bar{x}'' \\ \bar{x}'', \bar{q} &\rightarrow \bar{y}, \bar{y}\end{aligned}$$

The extension to multiple inputs, e.g., $f(m_1, m_2, m_3) = 4m_1 + m_2 + 2m_3$, uses similar transitions for each input:

$$\begin{aligned}\bar{x}_1, \bar{q} &\rightarrow \bar{y}, \bar{x}'_1 \\ \bar{x}'_1, \bar{q} &\rightarrow \bar{y}, \bar{x}''_1 \\ \bar{x}''_1, \bar{q} &\rightarrow \bar{y}, \bar{y} \\ \bar{x}_2, \bar{q} &\rightarrow \bar{y}, \bar{q} \\ \bar{x}_3, \bar{q} &\rightarrow \bar{y}, \bar{y}\end{aligned}$$

We summarize this in the following observation.

Observation 7.1. *Let $f : \mathbb{N}^k \rightarrow \mathbb{N}$ be an \mathbb{N} -linear function. There is a function-computing leaderless population protocol that stably computes f in expected time $O(\log n)$.*

A function is $\mathbb{Q}_{\geq 0}$ -linear if there are $c_1, \dots, c_k \in \mathbb{Q}_{\geq 0}$ such that for all $\mathbf{m} \in \mathbb{N}^k$, $f(\mathbf{m}) = \sum_{i=1}^k \lfloor c_i \mathbf{m}(i) \rfloor$. We now describe how to stably approximate $\mathbb{Q}_{\geq 0}$ -linear functions with a linear approximation error, in logarithmic time. (It is open to do this for negative coefficients, e.g., $f(m_1, m_2) = m_1 - m_2$). Recall the following simple example of a population protocol that approximately divides by 2 (that is, with probability 1 it outputs a value guaranteed to be a certain distance to the correct output), with a linear approximation error, and is fast ($O(\log n)$ time) with initial counts $\#\bar{x} = m$, $\#\bar{a} = \gamma m$, and $\#\bar{q} = \#\bar{y} = 0$:

$$\begin{aligned}\bar{a}, \bar{x} &\rightarrow \bar{b}, \bar{y} \\ \bar{b}, \bar{x} &\rightarrow \bar{a}, \bar{q}\end{aligned}$$

which stabilizes $\#\bar{y}$ to somewhere in the interval $\{m/2, m/2 + 1, \dots, m/2 + \gamma m\}$.

To see that the protocol is correct, note that the transition sequence can make $\#\bar{y}$ closer to one endpoint of the interval or the other depending on which transitions are chosen to consume the *last* γm of \bar{x} , but no matter what, the first transition executes at least as many times as the second, but not more than γm times more.

If $\#\bar{a} = 1$ initially, the above protocol stably computes $\lfloor m/2 \rfloor$ (taking linear time just for the last transition; and in total takes $\Theta(n \log n)$ time, by a coupon collector argument).

To see that the protocol takes $O(\log n)$ time if $\#\bar{a} = \gamma m$ initially, note $n = m + \gamma m \leq 2m$. Observe that $\#\bar{a} + \#\bar{b} = \gamma m$ in any reachable configuration. Thus the probability any given interaction is one of the above two transitions is $\approx \frac{\gamma m \#\bar{x}}{n^2}$, so the expected number of interactions

until such a transition occurs is $\frac{n^2}{\gamma m \#\bar{x}}$. After m such transitions occur, all the input \bar{x} is gone and the protocol stabilizes, which by linearity of expectation takes expected number of interactions

$$\sum_{\#\bar{x}=1}^m \frac{n^2}{\gamma m \#\bar{x}} = \frac{n^2}{\gamma m} \sum_{\#\bar{x}=1}^m \frac{1}{\#\bar{x}} \approx \frac{n^2}{\gamma m} \ln m \leq \frac{n^2}{\gamma n/2} \ln n = \frac{2n}{\gamma} \ln n,$$

i.e., expected parallel time $\frac{2}{\gamma} \ln n$. Thus this shows a tradeoff between accuracy and speed in a single protocol, adjustable by the initial count of \bar{a} . In this case, the approximation error increases, and the expected time to stabilization decreases, with increasing initial $\#\bar{a}$.

More generally, we can prove the following. In particular, if $a = \Omega(n)$, then $t(n, a) = O(\log n)$. Also, if $a = o(n)$, then the approximation error is $o(n)$, and if $a = \omega(\log n)$, then the expected time is $o(n)$ also. This does not contradict Theorem 6.1 since setting $a = o(n)$ implies the initial configurations are not all α -dense for a fixed $\alpha > 0$.

Theorem 7.2. *Let $f : \mathbb{N}^k \rightarrow \mathbb{N}$ be a $\mathbb{Q}_{\geq 0}$ -linear function. Let $\mathcal{E} : \mathbb{N} \rightarrow \mathbb{N}$ be the identity function. Define $t : \mathbb{N}^2 \rightarrow \mathbb{N}$ by $t(n, a) = \frac{n}{a} \log n$. Then there is a function-approximating leaderless population protocol \mathcal{A} that \mathcal{E} -approximates f in expected time $O(t)$.*

Proof. Write each rational coefficient $c_i = \frac{p_i}{r_i}$, where $p_i, r_i \in \mathbb{Z}_{\geq 1}$. (Since f is linear, the case $p_i = 0$ is easy to handle by simply ignoring \bar{x}_i as an input, so we assume all coefficients are strictly positive.) The initial configuration \mathbf{i} has $\mathbf{i}(\bar{x}_i) = \mathbf{m}(i)$ for each $i \in \{1, \dots, k\}$, $\mathbf{i}(\bar{a}) > 1$ and $\mathbf{i}(\bar{q}) \geq 2k\mathbf{i}(\bar{a}) + \sum_{i=1}^k p_i \mathbf{i}(\bar{x}_i)$.

First, for each $i \in \{1, \dots, k\}$, we have transitions that multiply $\#\bar{x}_i = \mathbf{m}(i)$ by p_i . If $p_i = 1$, we have the transition $\bar{x}_i, \bar{q} \rightarrow \bar{y}_i, \bar{q}$. If $p_i = 2$, we have the transition $\bar{x}_i, \bar{q} \rightarrow \bar{y}_i, \bar{y}_i$. If $p_i \geq 3$, we have a sequence of transitions that multiply $\#\bar{x}_i$ by p_i :

$$\begin{aligned} \bar{x}_i, \bar{q} &\rightarrow \bar{y}_i, \bar{x}_i^{(p_i-1)} \\ \bar{x}_i^{(p_i-1)}, \bar{q} &\rightarrow \bar{y}_i, \bar{x}_i^{(p_i-2)} \\ \bar{x}_i^{(p_i-2)}, \bar{q} &\rightarrow \bar{y}_i, \bar{x}_i^{(p_i-3)} \\ &\dots \\ \bar{x}_i^{(2)}, \bar{q} &\rightarrow \bar{y}_i, \bar{y}_i \end{aligned}$$

This ensures that for each $i \in \{1, \dots, k\}$, eventually $\#\bar{y}_i = p_i \mathbf{i}(\bar{x}_i)$. Note that each agent in state \bar{x}_i eventually causes p_i agents in state \bar{q} to be consumed as inputs to the above transitions, hence at least $\sum_{i=1}^k p_i \mathbf{m}(i)$ agents in state \bar{q} are required initially for the above to complete.

Next, we split up the approximation states to have a specific one for each input \bar{x}_i . In other words, the initial configuration has $\mathbf{i}(\bar{a})$ agents in state \bar{a} , and we want to reach a configuration with $\mathbf{i}(\bar{a})$ agents in state \bar{a}_1 , $\mathbf{i}(\bar{a})$ other agents in state $\bar{a}_2, \dots, \mathbf{i}(\bar{a})$ other agents in state \bar{a}_k . If $k = 1$ then we simply use \bar{a} directly in place of the state $\bar{a}_1^{(1)}$ below. Otherwise, let $l \in \mathbb{N}$ be such that $2^{l-1} < k \leq 2^l$ (i.e., 2^l is the next power of 2 at least k).

We have the transitions that make copies of \bar{a} for each input:

$$\begin{aligned}
\bar{a}, \bar{q} &\rightarrow \bar{a}_0, \bar{a}_1 \\
\bar{a}_0, \bar{q} &\rightarrow \bar{a}_{00}, \bar{a}_{01} \\
\bar{a}_1, \bar{q} &\rightarrow \bar{a}_{10}, \bar{a}_{11} \\
\bar{a}_{00}, \bar{q} &\rightarrow \bar{a}_{000}, \bar{a}_{001} \\
\bar{a}_{01}, \bar{q} &\rightarrow \bar{a}_{010}, \bar{a}_{011} \\
\bar{a}_{10}, \bar{q} &\rightarrow \bar{a}_{100}, \bar{a}_{101} \\
\bar{a}_{11}, \bar{q} &\rightarrow \bar{a}_{110}, \bar{a}_{111} \\
\bar{a}_{000}, \bar{q} &\rightarrow \bar{a}_{0000}, \bar{a}_{0001} \\
&\dots
\end{aligned}$$

and so on, up to transitions whose output (right-hand side) states are binary strings of length l . Also, for each $i \in \{1, \dots, k\}$, we have the transition $\bar{a}_{b(i)}, \bar{q} \rightarrow \bar{a}_i^{(1)}, \bar{q}$, where $b(i)$ is the i 'th string in the lexicographical ordering of $\{0, 1\}^l$. If k is not a power of 2, then we will simply let $\bar{a}_{b(i)}$ go unused for $i \in \{k+1, \dots, 2^l\}$.

Note that each agent in state \bar{a} eventually causes $\sum_{j=1}^l 2^j = 2^{l+1} - 1 < 2k$ agents in state \bar{q} to be consumed as inputs to the above transitions, hence at least $\mathbf{i}(\bar{a})(2^{l+1} - 1)$ agents in state \bar{q} are required initially for the above to complete. Combined with the lower bound on $\mathbf{i}(\bar{q})$ associated to the first set of transitions, at least $\mathbf{i}(\bar{a})(2^{l+1} - 1) + \sum_{i=1}^k p_i \mathbf{i}(\bar{x}_i) \leq 2k\mathbf{i}(\bar{a}) + \sum_{i=1}^k p_i \mathbf{i}(\bar{x}_i)$ agents in state \bar{q} are required initially. Note that this is $O(\|\mathbf{i} \upharpoonright (\Sigma \cup \{\bar{a}\})\|)$.

Next, we approximately divide $\#\bar{y}_i$ by r_i , using the state $\bar{a}_i^{(1)}$ produced by the second sequence of transitions. If $r_i = 1$, this is easily handled (exactly) by the transition $\bar{y}_i, \bar{q} \rightarrow \bar{y}, \bar{q}$. Otherwise, for each i such that $r_i \geq 2$, we have the transitions that divide $\#\bar{y}_i$ by r_i :

$$\begin{aligned}
\bar{a}_i^{(1)}, \bar{y}_i &\rightarrow \bar{a}_i^{(2)}, \bar{y} \\
\bar{a}_i^{(2)}, \bar{y}_i &\rightarrow \bar{a}_i^{(3)}, \bar{q} \\
\bar{a}_i^{(3)}, \bar{y}_i &\rightarrow \bar{a}_i^{(4)}, \bar{q} \\
&\dots \\
\bar{a}_i^{(r_i)}, \bar{y}_i &\rightarrow \bar{a}_i^{(1)}, \bar{q}
\end{aligned}$$

which consumes r_i copies of \bar{y}_i (of which there are $p_i \mathbf{m}(i)$) for each \bar{y} produced. Since for all i , the first transition above outputs the same state \bar{y} , the final count of \bar{y} will be the sum $\sum_{i=1}^k \left\lfloor \frac{p_i}{r_i} \mathbf{m}(i) \right\rfloor$, i.e., the value $f(\mathbf{m})$.

It remains to prove the stated bound on expected time. Above we stated a lower bound on $\mathbf{i}(\bar{q})$ required for the protocol to be correct. To obtain the stated expected time, we double this bound and add $\|\mathbf{i} \upharpoonright (\Sigma \cup \{\bar{a}\})\|$, which ensures that $\#\bar{q}$ is always at least $n/2$ in any reachable configuration. Note that $n \geq \sum_{i=1}^k p_i \mathbf{m}(i)$. Let $\hat{p} = \max_i p_i$.

We analyze a simpler process that stochastically dominates the actual Markov process. We assume that the first set of transitions (multiplying $\#\bar{x}_i$ by p_i) completes before the start of the second (making copies of \bar{a} for each input), and that these then complete before the start of the third group (dividing $\#\bar{y}_i$ by r_i). Also, transitions that can go in parallel within these groups (such as the first set of transitions for different values of i) can have their expected times analyzed independently and summed to obtain a loose bound on the time for them all to terminate.

Let $i \in \{1, \dots, k\}$. For the first group, we assume that for each $\bar{x}_i^{(p_i-j)}, \bar{q} \rightarrow \bar{y}_i, \bar{x}_i^{(p_i-j-1)}$ completes (consuming all available $\bar{x}_i^{(p_i-1)}$ states) before the next $\bar{x}_i^{(p_i-j-1)}, \bar{q} \rightarrow \bar{y}_i, \bar{x}_i^{(p_i-j-2)}$ starts, since this process stochastically dominates the actual process that allows the transitions to go in parallel. For notational brevity we just analyze the first transition $\bar{x}_i, \bar{q} \rightarrow \bar{y}_i, \bar{x}_i^{(p_i-1)}$. By our choice of $\mathbf{i}(\bar{q})$, each interaction has probability at least $\frac{1}{2}$ to have a \bar{q} state, and conditioned on this, probability at least $\frac{\#\bar{x}_i}{n}$ for the other state to be \bar{x}_i , so probability at least $\frac{\#\bar{x}_i}{2n}$ to be of the form $\bar{x}_i, \bar{q} \rightarrow \bar{y}_i, \bar{x}_i^{(p_i-1)}$. Thus the expected number of interactions until this transition happens is at most $\frac{2n}{\#\bar{x}_i}$. By linearity of expectation, the expected number of interactions until all \bar{x}_i are consumed is at most $\sum_{\#\bar{x}_i=1}^{\mathbf{m}(i)} \frac{2n}{\#\bar{x}_i} \approx 2n \ln \mathbf{m}(i)$. Thus the expected parallel time to complete these transitions is $2 \ln \mathbf{m}(i)$. Thus for each i , the expected time for all p_i of these transitions to complete is at most $2p_i \ln \mathbf{m}(i)$, so the expected time for all to complete is at most $2 \sum_{i=1}^k p_i \ln \mathbf{m}(i) \leq 2k\hat{p} \ln n$.

The second group of transitions can be analyzed similarly, by assuming they each complete (consume all non- \bar{q} input states) before the next starts. Each then takes $\ln \mathbf{i}(\bar{a})$ expected time to complete. There are $2^{l+1} - 1 < 2k$ total transitions so $2k \ln \mathbf{i}(\bar{a}) \leq 2k \ln n$ expected time is required.

The third group of transitions is a bit different. Since we do not begin analysis until the second group of transitions completes, we start with $\#\bar{a}_i^{(1)} = \mathbf{i}(\bar{a})$ and $\#\bar{y}_i = p_i \mathbf{m}(i)$. Also note that $\mathbf{i}(\bar{a}) = \sum_{j=1}^{r_i} \#\bar{a}_i^{(j)}$ in any subsequently reachable configuration. Thus, in any such configuration the probability that the next transition involves an $\bar{a}_i^{(j)}$ state is $\frac{\mathbf{i}(\bar{a})}{n}$, and conditioned on this, the probability that the *other* input state is \bar{y}_i is $\frac{\#\bar{y}_i}{n}$, so probability $\frac{\mathbf{i}(\bar{a})\#\bar{y}_i}{n^2}$ that the next transition is *one* of the r_i transitions. Thus the expected number of transitions until such a transition happens is $\frac{n^2}{\mathbf{i}(\bar{a})\#\bar{y}_i}$. By linearity of expectation, the expected number of interactions until all \bar{y}_i are consumed is at most $\sum_{\#\bar{y}_i=1}^{p_i \mathbf{m}(i)} \frac{n^2}{\mathbf{i}(\bar{a})\#\bar{y}_i} \approx \frac{n^2}{\mathbf{i}(\bar{a})} \ln(p_i \mathbf{m}(i))$, so expected parallel time $\frac{n}{\mathbf{i}(\bar{a})} \ln(p_i \mathbf{m}(i))$. Thus the expected time for all k groups of transitions to complete is at most $\frac{n}{\mathbf{i}(\bar{a})} \sum_{i=1}^k \ln(p_i \mathbf{m}(i)) = \frac{n}{\mathbf{i}(\bar{a})} k \ln(\hat{p}n)$.

Thus, summing the above three bounds, the total expected time to stabilization is at most $2k\hat{p} \ln n + 2k \ln n + \frac{n}{\mathbf{i}(\bar{a})} k \ln(\hat{p}n)$. Since k and \hat{p} are constant with respect to n and $\mathbf{i}(\bar{a})$, the bound is dominated by the last term, which is $O\left(\frac{n}{\mathbf{i}(\bar{a})} \ln n\right)$. \square

8 Exact computation of nonlinear functions

In Section 6, we obtained a precise characterization of the linear functions stably computable in sublinear time by population protocols and furthermore show that those not exactly computable in sublinear time are not even approximable with sublinear error in sublinear time. However, the class of functions stably computable (in any amount of time) by population protocols is known to contain non-linear functions such as $f(m_1, m_2) = \max(m_1, m_2)$, or $f(m) = m$ if m is even and $f(m) = 2m$ if m is odd. In fact a function is stably computable by a population protocol if and only if its *graph* $\{(\mathbf{m}, f(\mathbf{m})) \mid \mathbf{m} \in \mathbb{N}^k\}$ is a semilinear set [5, 15]. A set $A \subseteq \mathbb{N}^k$ is *semilinear* if and only if [23] it is expressible as a finite number of unions, intersections, and complements of sets of one of the following two forms: *threshold sets* of the form $\{\mathbf{x} \mid \sum_{i=1}^k a_i \cdot \mathbf{x}(i) < b\}$ for some constants $a_1, \dots, a_k, b \in \mathbb{Z}$ or *mod sets* of the form $\{\mathbf{x} \mid \sum_{i=1}^k a_i \cdot \mathbf{x}(i) \equiv b \pmod{c}\}$ for some constants $a_1, \dots, a_k, b, c \in \mathbb{N}$.

Say that a set $P \subseteq \mathbb{N}^k$ is a *periodic coset* if there exist $\mathbf{b}, \mathbf{p}_1, \dots, \mathbf{p}_l \in \mathbb{N}^k$ such that $P = \{\mathbf{b} + n_1 \mathbf{p}_1 + \dots + n_l \mathbf{p}_l \mid n_1, \dots, n_l \in \mathbb{N}\}$. (These are typically called “linear” sets, but we wish to avoid confusion with linear functions.) Equivalently, a set is semilinear if and only if it is a finite

union of periodic cosets.

Although our technique fails to completely characterize the efficient computability of all semilinear functions, we show that a wide class of semilinear functions cannot be stably computed in sublinear time: functions that are not eventually \mathbb{N} -linear. The only exceptions, for which we cannot prove linear time is required, yet neither is there known a counterexample protocol stably computing the function in sublinear time, are functions whose “non-integral-linearities are near the boundary of \mathbb{N}^k ”. For example, the function $f(m) = 0$ if $m \leq 3$ and $f(m) = m$ otherwise is non-linear (although it is semilinear, so stably computable), but restricted to the domain of inputs > 3 , it is linear with positive integer coefficients. Thus it is an example of a function whose “population protocol time complexity” is unknown.

Corollary 6.2 and Observation 7.1 imply that a linear function is stably computable in sublinear time by a population protocol if and only if it is \mathbb{N} -linear. Theorem 8.5 generalizes the forward direction (restricted to nonlinear functions) to eventually \mathbb{N} -linear functions.

8.1 Eventually affine functions

Say that a function $f : \mathbb{N}^k \rightarrow \mathbb{N}$ is *eventually \mathbb{N} -affine* if there exist $n_0, c_1, \dots, c_k \in \mathbb{N}$ and $b \in \mathbb{Z}$ such that, for all $\mathbf{m} \in \mathbb{N}_{\geq n_0}^k$, $f(\mathbf{m}) = b + \sum_{i=1}^k c_i \mathbf{m}(i)$.

The function $f : \mathbb{N}^k \rightarrow \mathbb{Q}$ is affine if and only if all points on the graph of f lie on a k -dimensional hyperplane.¹³ This holds if and only if, for all $\mathbf{m}, \mathbf{v} \in \mathbb{N}^k$, $f(\mathbf{m} + \mathbf{v}) - f(\mathbf{m}) = f(\mathbf{m} + 2\mathbf{v}) - f(\mathbf{m} + \mathbf{v})$. In other words, the change in output resulting by moving by a vector \mathbf{v} , starting from \mathbf{m} , is the same if we move by \mathbf{v} a second time. The next lemma, due to Sungjin Im [25], shows that a function $f : \mathbb{N}^k \rightarrow \mathbb{N}$ is eventually \mathbb{N} -affine if the above holds for any sufficiently large input \mathbf{m} and any $0/1$ -valued vector $\mathbf{v} \in \{0, 1\}^k$. In other words, if f “looks \mathbb{N} -affine” when moving by small amounts (each component at most 1), then it is \mathbb{N} -affine. This appears obvious, but some care is needed: for example, it is false if we assume that \mathbf{v} are only unit vectors. For example, the non-affine function $f(m_1, m_2) = m_1 \cdot m_2$ obeys $f(\mathbf{m} + \mathbf{u}) - f(\mathbf{m}) = f(\mathbf{m} + 2\mathbf{u}) - f(\mathbf{m} + \mathbf{u})$ for each $\mathbf{m} \in \mathbb{N}^2$ and each $\mathbf{u} \in \{(0, 1), (1, 0)\}$, but fails if $\mathbf{u} = (1, 1)$.

Lemma 8.1. *Let $f : \mathbb{N}^k \rightarrow \mathbb{N}$. Suppose there exists $n_0 \in \mathbb{N}$ such that for all $\mathbf{m} \in \mathbb{N}_{\geq n_0}^k$ and $\mathbf{v} \in \{0, 1\}^k$, $f(\mathbf{m} + \mathbf{v}) - f(\mathbf{m}) = f(\mathbf{m} + 2\mathbf{v}) - f(\mathbf{m} + \mathbf{v})$. Then f is eventually \mathbb{N} -affine: there are $n_0, b, c_1, \dots, c_k \in \mathbb{N}$ such that, for all $\mathbf{m} \in \mathbb{N}_{\geq n_0}^k$, $f(\mathbf{m}) = b + \sum_{i=1}^k c_i \mathbf{m}(i)$.*

Proof. Assume the hypothesis and let $\mathbf{v} \in \{0, 1\}^k$ and $\mathbf{m} \in \mathbb{N}_{\geq n_0}^k$. A straightforward induction shows that, for all $n \geq 1$, $f(\mathbf{m} + n\mathbf{v}) - f(\mathbf{m} + (n-1)\mathbf{v}) = f(\mathbf{m} + \mathbf{v}) - f(\mathbf{m})$, and thus

$$f(\mathbf{m} + n\mathbf{v}) = f(\mathbf{m}) + n \cdot (f(\mathbf{m} + \mathbf{v}) - f(\mathbf{m})). \quad (8.1)$$

Let $i \in \{1, \dots, k\}$, fix values of $m'_{i+1}, \dots, m'_k \geq n_0$, and define $f_i : \mathbb{N}^i \rightarrow \mathbb{N}$ by $f_i(m_1, \dots, m_i) = f(m_1, \dots, m_{i-1}, m_i, m'_{i+1}, \dots, m'_k)$. We show by induction on i that f_i is \mathbb{N} -affine on inputs $m_1, \dots, m_i \geq n_0$, i.e., $f(m_1, \dots, m_i) = b + c_1 m_1 + \dots + c_i m_i$ for some $b, c_1, \dots, c_i \in \mathbb{N}$.

Base case: For any fixed $m'_2, \dots, m'_k \geq n_0$, let $f_1(m_1) = f(m_1, m'_2, \dots, m'_k)$. We must show f_1 is an \mathbb{N} -affine function of m_1 (assuming $m_1 \geq n_0$). Let $\mathbf{v} = (1, 0, \dots, 0)$ and let $\mathbf{m} = (n_0, m'_2, \dots, m'_k)$. Then using eq. (8.1), we obtain that for any $m_1 \geq n_0$,

$$\begin{aligned} f_1(m_1) &= f(n_0, m'_2, \dots, m'_k) + \\ &\quad (m_1 - n_0) \cdot [f(n_0 + 1, m'_2, \dots, m'_k) - f(n_0, m'_2, \dots, m'_k)], \end{aligned}$$

¹³ We let the range be \mathbb{Q} instead of \mathbb{N} here to get a bidirectional implication, but the direction we are interested in requires only range \mathbb{N} : if $f : \mathbb{N}^k \rightarrow \mathbb{N}$ is affine, then all points lie on a k -dimensional hyperplane.

which is an affine function $b + cm_1$ of m_1 , with offset $b = f(n_0, m'_2, \dots, m'_k) - n_0 \cdot [f(n_0 + 1, m'_2, \dots, m'_k) - f(n_0, m'_2, \dots, m'_k)]$ and coefficient $c = f(n_0 + 1, m'_2, \dots, m'_k) - f(n_0, m'_2, \dots, m'_k)$.

Inductive case: Fix values $m'_{i+1}, \dots, m'_k \geq n_0$. By the inductive hypothesis, $f_i(m_1, \dots, m_i) = f(m_1, \dots, m_i, m'_{i+1}, \dots, m'_k)$ is an \mathbb{N} -affine function of m_1, \dots, m_i (if all $m_j \geq n_0$). We want to show that $f_{i+1}(m_1, \dots, m_{i+1}) = f(m_1, \dots, m_{i+1}, m'_{i+2}, \dots, m'_k)$ is an \mathbb{N} -affine function of m_1, \dots, m_{i+1} (if all $m_j \geq n_0$).

Let $m_1, \dots, m_{i+1} \geq n_0$. By eq. (8.1), letting $\mathbf{m} = (m_1, \dots, m_i, n_0, m'_{i+2}, \dots, m'_k)$ and $\mathbf{v} = (\underbrace{0, \dots, 0}_i, 1, 0, \dots, 0)$, we obtain that:

$$\begin{aligned} & f_{i+1}(m_1, \dots, m_{i+1}) \\ &= f(m_1, \dots, m_i, m_{i+1}, m'_{i+2}, \dots, m'_k) \\ &= f(m_1, \dots, m_i, n_0, m'_{i+2}, \dots, m'_k) + \\ & \quad (m_{i+1} - n_0) \cdot [f(m_1, \dots, m_i, n_0 + 1, m'_{i+2}, \dots, m'_k) - f(m_1, \dots, m_i, n_0, m'_{i+2}, \dots, m'_k)] \end{aligned} \quad (8.2)$$

By the inductive hypothesis, defining $f_{i,0}, f_{i,1} : \mathbb{N}^i \rightarrow \mathbb{N}$ by $f_{i,0}(m_1, \dots, m_i) = f(m_1, \dots, m_i, n_0, m'_{i+2}, \dots, m'_k)$ and $f_{i,1}(m_1, \dots, m_i) = f(m_1, \dots, m_i, n_0 + 1, m'_{i+2}, \dots, m'_k)$, $f_{i,0}$ and $f_{i,1}$ are both \mathbb{N} -affine. Thus for some $a_1, \dots, a_i, b \in \mathbb{N}$ and $c_1, \dots, c_i, d \in \mathbb{N}$,

$$f_{i,0}(m_1, \dots, m_i) = b + a_1 m_1 + \dots + a_i m_i \quad (8.3)$$

$$f_{i,1}(m_1, \dots, m_i) = d + c_1 m_1 + \dots + c_i m_i \quad (8.4)$$

Substituting (8.3) and (8.4) into (8.2), defining $f_{i+1} : \mathbb{N}^{i+1} \rightarrow \mathbb{N}$ for all m_1, \dots, m_{i+1} by

$$\begin{aligned} f_{i+1}(m_1, \dots, m_{i+1}) &= f(m_1, \dots, m_i, m_{i+1}, m'_{i+2}, \dots, m'_k) \\ &= (b + a_1 m_1 + \dots + a_i m_i) + \\ & \quad (m_{i+1} - n_0) \cdot [(d + c_1 m_1 + \dots + c_i m_i) - (b + a_1 m_1 + \dots + a_i m_i)]. \end{aligned}$$

Let $a'_{i+1} = d - b$ and $b' = b - a'_{i+1} n_0$. For $j \in \{1, \dots, i\}$, let $e_j = c_j - a_j$ and $a'_j = a_j - e_j n_0$. Then

$$\begin{aligned} & f_{i+1}(m_1, \dots, m_{i+1}) \\ &= b + a_1 m_1 + \dots + a_i m_i + (m_{i+1} - n_0) \cdot [a_{i+1} + e_1 m_1 + \dots + e_i m_i] \\ &= b + a_1 m_1 + \dots + a_i m_i + a_{i+1} m_{i+1} + e_1 m_1 m_{i+1} + e_2 m_2 m_{i+1} + \dots + e_i m_i m_{i+1} \\ & \quad - a_{i+1} n_0 - e_1 m_1 n_0 - e_2 m_2 n_0 - \dots - e_i m_i n_0 \\ &= b - a_{i+1} n_0 + \sum_{j=1}^{i+1} a_j m_j + \sum_{j=1}^i e_j m_j m_{i+1} - \sum_{j=1}^i e_j m_j n_0 \\ &= b' + \sum_{j=1}^{i+1} a'_j m_j + \sum_{j=1}^i e_j m_j m_{i+1} \end{aligned} \quad (8.5)$$

To prove f_{i+1} is \mathbb{N} -affine, we must show that $e_1 = e_2 = \dots = e_i = 0$ and that $a'_{i+1} \geq 0$ (the inductive hypothesis implies that $b, a_1, \dots, a_i \geq 0$, so if each $e_j = 0$, then $a'_j = a_j - e_j n_0 = a_j \geq 0$). First, we claim that $e_1, \dots, e_i \geq 0$. To see why, suppose for the sake of contradiction that some $e_l < 0$, and consider fixing m_j for each $j \notin \{l, i+1\}$ and taking the limit of $m_l = m_{i+1} = n$ as $n \rightarrow \infty$. Equation (8.5) will be dominated by the quadratic term in n , i.e., $e_l m_l m_{i+1} = e_l n^2$, and thus f_{i+1} would be negative for large enough n , a contradiction.

We will eventually show that $e_1, \dots, e_i = 0$. Once this is shown, a similar argument shows that $a'_{i+1} \geq 0$, or else with increasing m_{i+1} , f_{i+1} would become negative. (The assumption $a'_{i+1} \geq 0$ is not used subsequently in the proof of the inductive case; it is only needed at the conclusion.)

To show that $e_1 = e_2 = \dots = e_i = 0$, define $\mathbf{v} = (\underbrace{1, \dots, 1}_{i+1}, \underbrace{0, \dots, 0}_{k-i-1}) \in \mathbb{N}^k$. Let $\mathbf{m} = (\underbrace{n_0, \dots, n_0}_{i+1}, m'_{i+2}, \dots, m'_k)$. Let $\mathbf{m}_{i+1} = (\underbrace{n_0, \dots, n_0}_{i+1}) \in \mathbb{N}^{i+1}$ be the first $i+1$ coordinates of \mathbf{m} . Similarly, let $\mathbf{v}_{i+1} = (\underbrace{1, \dots, 1}_{i+1}) \in \mathbb{N}^{i+1}$ be the first $i+1$ coordinates of \mathbf{v} . Then $f(\mathbf{m}) = f_{i+1}(\mathbf{m}_{i+1})$, $f(\mathbf{m} + \mathbf{v}) = f_{i+1}(\mathbf{m}_{i+1} + \mathbf{v}_{i+1})$, and $f(\mathbf{m} + 2\mathbf{v}) = f_{i+1}(\mathbf{m}_{i+1} + 2\mathbf{v}_{i+1})$. Applying the hypothesis of the lemma and these identities,

$$f_{i+1}(\mathbf{m}_{i+1} + \mathbf{v}_{i+1}) - f_{i+1}(\mathbf{m}_{i+1}) = f_{i+1}(\mathbf{m}_{i+1} + 2\mathbf{v}_{i+1}) - f_{i+1}(\mathbf{m}_{i+1} + \mathbf{v}_{i+1})$$

Substituting eq. (8.5),

$$\begin{aligned} f_{i+1}(\mathbf{m}_{i+1} + \mathbf{v}_{i+1}) - f_{i+1}(\mathbf{m}_{i+1}) &= \left[b' + \sum_{j=1}^{i+1} a'_j(n_0 + 1) + \sum_{j=1}^i e_j(n_0 + 1)^2 \right] \\ &\quad - \left[b' + \sum_{j=1}^{i+1} a'_j n_0 + \sum_{j=1}^i e_j n_0^2 \right] \\ &= \sum_{j=1}^{i+1} a'_j + \sum_{j=1}^i e_j(2n_0 + 1) \end{aligned}$$

Similarly, the difference $f(\mathbf{m} + 2\mathbf{v}) - f(\mathbf{m} + \mathbf{v})$ can be expressed as

$$\begin{aligned} f_{i+1}(\mathbf{m}_{i+1} + 2\mathbf{v}_{i+1}) - f_{i+1}(\mathbf{m}_{i+1} + \mathbf{v}_{i+1}) &= \left[b' + \sum_{j=1}^{i+1} a'_j(n_0 + 2) + \sum_{j=1}^i e_j(n_0 + 2)^2 \right] \\ &\quad - \left[b' + \sum_{j=1}^{i+1} a'_j(n_0 + 1) + \sum_{j=1}^i e_j(n_0 + 1)^2 \right] \\ &= \sum_{j=1}^{i+1} a'_j + \sum_{j=1}^i e_j(2n_0 + 3) \end{aligned}$$

Since the first difference equals the second, $\sum_{j=1}^{i+1} e_j(2n_0 + 1) = \sum_{j=1}^{i+1} e_j(2n_0 + 3)$. Since each $e_j \geq 0$, this implies each $e_j = 0$, so $a'_j = a_j$. Thus $f_{i+1}(m_1, \dots, m_{i+1}) = b' + \sum_{j=1}^{i+1} a_j m_j$, an \mathbb{N} -affine function, proving the inductive case. \square

The contrapositive of Lemma 8.1 states that if f is not eventually \mathbb{N} -affine, then for all $m_0 \in \mathbb{N}$, there is $\mathbf{v} \in \{0, 1\}^k$ and $\mathbf{m} \in \mathbb{N}_{\geq m_0}^k$ such that $f(\mathbf{m} + \mathbf{v}) - f(\mathbf{m}) \neq f(\mathbf{m} + 2\mathbf{v}) - f(\mathbf{m} + \mathbf{v})$. Since $\{0, 1\}^k$ is finite, by the pigeonhole principle, some infinite subset of such \mathbf{m} 's can be found that agree on the same \mathbf{v} . Thus we have the following corollary.

Corollary 8.2. *Let $f : \mathbb{N}^k \rightarrow \mathbb{N}$. If f is not eventually \mathbb{N} -affine, then there is $\mathbf{v} \in \{0, 1\}^k$ such that for all $m_0 \in \mathbb{N}$, there is $\mathbf{m} \in \mathbb{N}_{\geq m_0}^k$ such that $f(\mathbf{m} + \mathbf{v}) - f(\mathbf{m}) \neq f(\mathbf{m} + 2\mathbf{v}) - f(\mathbf{m} + \mathbf{v})$.*

Finally, we need not only that we can find an infinite set of “arbitrarily large counter-examples to affine” given by Corollary 8.2, but *furthermore* that this set is α -dense for some $\alpha > 0$. This is not true for arbitrary functions. For example, consider the function $f(m_1, m_2) = m_1 + m_2$ unless $m_1 = 2^{m_2}$, in which case $f(m_1, m_2) = 0$ instead. However, if f is a semilinear function, then we can find an infinite set of “counter-examples to affine” that *is* α -dense for some $\alpha > 0$.

Lemma 8.3. *Let $f : \mathbb{N}^k \rightarrow \mathbb{N}$ be semilinear but not eventually \mathbb{N} -affine. Then there are $\mathbf{v} \in \{0, 1\}^k$, $\alpha > 0$, and infinitely many α -dense $\mathbf{m} \in \mathbb{N}^k$ such that $f(\mathbf{m} + \mathbf{v}) - f(\mathbf{m}) \neq f(\mathbf{m} + 2\mathbf{v}) - f(\mathbf{m} + \mathbf{v})$.*

Proof. For all $\mathbf{v} \in \{0, 1\}^k$, let $C_{\mathbf{v}} = \{\mathbf{m} \in \mathbb{N}^k \mid f(\mathbf{m} + \mathbf{v}) - f(\mathbf{m}) \neq f(\mathbf{m} + 2\mathbf{v}) - f(\mathbf{m} + \mathbf{v})\}$. Corollary 8.2 tells us that there is $\mathbf{v} \in \{0, 1\}^k$ such that, for all $m \in \mathbb{N}$, we have $\mathbb{N}_{\geq m}^k \cap C_{\mathbf{v}} \neq \emptyset$.

This implies that for all $m \in \mathbb{N}$, $|C_{\mathbf{v}} \cap \mathbb{N}_{\geq m}^k| = \infty$. Otherwise, since $\mathbb{N}_{\geq m}^k \subset \mathbb{N}_{\geq m+1}^k$ for all m , if $|C_{\mathbf{v}} \cap \mathbb{N}_{\geq m}^k| < \infty$, then for sufficiently large m' , $C_{\mathbf{v}} \cap \mathbb{N}_{\geq m'}^k = \emptyset$, contradicting the fact that C intersects *all* $\mathbb{N}_{\geq m'}^k$.

We denote the graph of f by $G(f) = \{(\mathbf{m}, f(\mathbf{m})) \mid \mathbf{m} \in \mathbb{N}^k\} \subset \mathbb{N}^{k+1}$. Since f is semilinear, $G(f)$ is a semilinear set. Thus it is a finite union $G(f) = \bigcup_{j=0}^{p-1} P_j$ of p periodic cosets P_0, \dots, P_{p-1} . Intuitively, the lemma will be proven as follows. $C_{\mathbf{v}}$, the “counterexamples to \mathbb{N} -affineness”, occur in these periodic cosets. Adding period vectors to a point keeps it in the same periodic coset. However, if we add the same multiple of period vectors each time, and if their sum is positive on all coordinates, then this sets a minimum α -denseness that we cannot fall below. This is how we will show that $C_{\mathbf{v}}$ has infinitely many α -dense points.

For all $\mathbf{m} \in \mathbb{N}^k$, let $\mathbf{m}^f = (\mathbf{m}, f(\mathbf{m})) \in \mathbb{N}^{k+1}$ denote the point in $G(f)$ corresponding to input \mathbf{m} . By the pigeonhole principle, there are $x, y, z \in \{0, \dots, p-1\}$ and an infinite subset $C_P \subseteq C_{\mathbf{v}}$ such that, for all m_0 , there is $\mathbf{m} \in C_P \cap \mathbb{N}_{> m_0}^k$ such that $\mathbf{m}^f \in P_x$, $(\mathbf{m} + \mathbf{v})^f \in P_y$, $(\mathbf{m} + 2\mathbf{v})^f \in P_z$. For notational convenience we assume that $x = 0, y = 1, z = 2$.¹⁴ In other words, we can find an infinite subset C_P of $C_{\mathbf{v}}$ in which, in the finite union of periodic cosets defining $G(f)$, \mathbf{m}^f is always part of the same periodic coset P_0 , $(\mathbf{m} + \mathbf{v})^f$ is always part of the same periodic coset P_1 , and $(\mathbf{m} + 2\mathbf{v})^f$ is always part of the same periodic coset P_2 .

Let $\mathbf{m}_S^f, \mathbf{m}_L^f \in C_P$ such that $\min(\mathbf{m}_L - \mathbf{m}_S) \geq 1$; i.e., \mathbf{m}_L is strictly larger than \mathbf{m}_S on all coordinates. By the fact that C_P contains arbitrarily large points (i.e., points in $\mathbb{N}_{\geq m}^k$ for all m), such \mathbf{m}_S and \mathbf{m}_L must exist.

Let $\mathbf{b}_0, \mathbf{p}_0^{(1)}, \dots, \mathbf{p}_0^{(l_0)} \in \mathbb{N}^{k+1}$ be such that

$$P_0 = \left\{ \mathbf{b}_0 + n^{(1)}\mathbf{p}_0^{(1)} + \dots + n^{(l_0)}\mathbf{p}_0^{(l_0)} \mid n^{(1)}, \dots, n^{(l_0)} \in \mathbb{N} \right\},$$

and similarly for $\mathbf{b}_1, \mathbf{p}_1^{(1)}, \dots, \mathbf{p}_1^{(l_1)}, \mathbf{b}_2, \mathbf{p}_2^{(1)}, \dots, \mathbf{p}_2^{(l_2)} \in \mathbb{N}^{k+1}$ and P_1, P_2 , respectively.

For $i \in \{0, 1, 2\}$, let $\mathbf{d}_i^f = (\mathbf{m}_L + i\mathbf{v})^f - (\mathbf{m}_S + i\mathbf{v})^f$, and let \mathbf{d}_i denote \mathbf{d}_i^f restricted to the first k coordinates, so that $\mathbf{d}_i = (\mathbf{m}_L + i\mathbf{v}) - (\mathbf{m}_S + i\mathbf{v}) = \mathbf{m}_L - \mathbf{m}_S$. In other words, $\mathbf{d}_0^f, \mathbf{d}_1^f, \mathbf{d}_2^f$ differ only on their last coordinate, representing $f(\mathbf{m}_L + i\mathbf{v}) - f(\mathbf{m}_S + i\mathbf{v})$. Then $\mathbf{d}_0^f = \sum_{j=1}^{l_0} n_0^{(j)} \mathbf{p}_0^{(j)}$ for some $n_0^{(1)}, \dots, n_0^{(l_0)} \in \mathbb{N}$ since $\mathbf{m}_S^f, \mathbf{m}_L^f \in P_0$. Similarly, there must exist $n_1^{(1)}, \dots, n_1^{(l_1)} \in \mathbb{N}$ and $n_2^{(1)}, \dots, n_2^{(l_2)} \in \mathbb{N}$ such that $\mathbf{d}_1^f = \sum_{j=1}^{l_1} n_1^{(j)} \mathbf{p}_1^{(j)}$ and $\mathbf{d}_2^f = \sum_{j=1}^{l_2} n_2^{(j)} \mathbf{p}_2^{(j)}$.

For all $n \in \mathbb{N}$, define $\mathbf{m}_n^f = \mathbf{m}_S^f + n\mathbf{d}_0^f$ (note that $\mathbf{m}_L^f = \mathbf{m}_S^f$ and $\mathbf{m}_S^f = \mathbf{m}_0^f$), letting \mathbf{m}_n be \mathbf{m}_n^f restricted to the first k coordinates. By the definition of P_0 , for all $n \in \mathbb{N}$, $\mathbf{m}_n^f \in P_0$, since $\mathbf{m}_S^f \in P_0$ and \mathbf{d}_0^f is a nonnegative multiple of period vectors in P_0 . Our goal next is to show that

¹⁴ This can be assumed without loss of generality even if some of x, y, z are equal, since we can have “duplicate” periodic cosets that are equal despite having different subscripts.

all sufficiently large \mathbf{m}_n^f are in $C_P \subseteq C_{\mathbf{v}}$. After showing that, we argue that \mathbf{m}_n are all α -dense for some $\alpha > 0$, proving the lemma.

For $i \in \{0, 1, 2\}$, let $\Delta y_i = \mathbf{d}_i^f(k+1) = f(\mathbf{m}_L + i\mathbf{v}) - f(\mathbf{m}_S + i\mathbf{v})$. Note that for all $n \in \mathbb{N}$, $f(\mathbf{m}_S + i\mathbf{v} + n\mathbf{d}_i) = f(\mathbf{m}_S + i\mathbf{v}) + n\Delta y_i$.

We claim that for all $n \in \mathbb{N}$, $\mathbf{m}_n \in C_{\mathbf{v}}$. Let $c_0 = f(\mathbf{m}_S + \mathbf{v}) - f(\mathbf{m}_S)$ and $c_1 = f(\mathbf{m}_S + 2\mathbf{v}) - f(\mathbf{m}_S + \mathbf{v})$. Note that $c_0 \neq c_1$ since $\mathbf{m}_S \in C_{\mathbf{v}}$. For all $n \in \mathbb{N}$,

$$\begin{aligned} f(\mathbf{m}_n + \mathbf{v}) - f(\mathbf{m}_n) &= f(\mathbf{m}_S + \mathbf{v} + n\mathbf{d}_1) - f(\mathbf{m}_S + n\mathbf{d}_0) \\ &= f(\mathbf{m}_S + \mathbf{v}) + n\Delta y_1 - (f(\mathbf{m}_S) + n\Delta y_0) \\ &= c_0 + n(\Delta y_1 - \Delta y_0), \end{aligned}$$

and

$$\begin{aligned} f(\mathbf{m}_n + 2\mathbf{v}) - f(\mathbf{m}_n + \mathbf{v}) &= f(\mathbf{m}_S + 2\mathbf{v} + n\mathbf{d}_2) - f(\mathbf{m}_S + \mathbf{v} + n\mathbf{d}_1) \\ &= f(\mathbf{m}_S + 2\mathbf{v}) + n\Delta y_2 - (f(\mathbf{m}_S + \mathbf{v}) + n\Delta y_1) \\ &= c_1 + n(\Delta y_2 - \Delta y_1). \end{aligned}$$

Thus

$$\begin{aligned} [f(\mathbf{m}_n + \mathbf{v}) - f(\mathbf{m}_n)] - [f(\mathbf{m}_n + 2\mathbf{v}) - f(\mathbf{m}_n + \mathbf{v})] &= [c_0 + n(\Delta y_1 - \Delta y_0)] - [c_1 + n(\Delta y_2 - \Delta y_1)] \\ &= (c_0 - c_1) + n(2\Delta y_1 - \Delta y_2 - \Delta y_0). \end{aligned}$$

There are two cases to complete the claim that $\mathbf{m}_n \in C_{\mathbf{v}}$:

$2\Delta y_1 - \Delta y_2 - \Delta y_0 = 0$: Then for all n , the above difference is $c_0 - c_1 \neq 0$, so $f(\mathbf{m}_n + \mathbf{v}) - f(\mathbf{m}_n) \neq f(\mathbf{m}_n + 2\mathbf{v}) - f(\mathbf{m}_n + \mathbf{v})$, implying $\mathbf{m}_n \in C_{\mathbf{v}}$.

$2\Delta y_1 - \Delta y_2 - \Delta y_0 \neq 0$: Then for all sufficiently large n , $(c_0 - c_1) + n(2\Delta y_1 - \Delta y_2 - \Delta y_0) \neq 0$, so $f(\mathbf{m}_n + \mathbf{v}) - f(\mathbf{m}_n) \neq f(\mathbf{m}_n + 2\mathbf{v}) - f(\mathbf{m}_n + \mathbf{v})$, implying $\mathbf{m}_n \in C_{\mathbf{v}}$.

Letting $\alpha' = \min(\mathbf{d}_0)/\|\mathbf{d}_0\|$, we have that for all $n \in \mathbb{N}$, $n\mathbf{d}_0$ is α' -dense. Note that $\alpha' > 0$ since we chose $\mathbf{d}_0 = \mathbf{m}_L - \mathbf{m}_S$ to be positive on all coordinates. Let $\alpha = \alpha'/2$. Then for all sufficiently large n , $\mathbf{m}_n = \mathbf{m}_S + n\mathbf{d}_0$ is α -dense. Since $\mathbf{m}_n \in C_P \subseteq C_{\mathbf{v}}$, this proves the lemma. \square

8.2 Time lower bounds for stably computing nonlinear functions

Recall that definition of \mathbb{N} -linear functions from Section 6. We say a function $f : \mathbb{N}^k \rightarrow \mathbb{N}$ is *eventually \mathbb{N} -linear* if it is eventually \mathbb{N} -affine with offset $b = 0$, i.e., if there are $c_1, \dots, c_k \in \mathbb{N}$ and $m_0 \in \mathbb{N}$ such that for all $\mathbf{m} \in \mathbb{N}_{\geq m_0}^k$, $f(\mathbf{m}) = \sum_{i=1}^k c_i \mathbf{m}(i)$.

The following lemma is a special case of the main result of this Section, Theorem 8.5, showing that the non-linear functions in a sense “closest” to being eventually \mathbb{N} -linear, the eventually \mathbb{N} -affine functions, are not stably computable in sublinear time. This includes functions as simple as $f(m) = m - 1$ (computable slowly via transition $\bar{x}, \bar{x} \rightarrow \bar{x}, \bar{y}$).

Theorem 8.4. *Let $f : \mathbb{N}^k \rightarrow \mathbb{N}$ be eventually \mathbb{N} -affine but not eventually \mathbb{N} -linear. Then no function-computing leaderless population protocol stably computes f in sublinear time.*

Proof. Since f is eventually \mathbb{N} -affine, there $m_0 \in \mathbb{N}$ and $b, c_1, \dots, c_k \in \mathbb{N}$ such that, for all $\mathbf{m} \in \mathbb{N}_{\geq m_0}^k$, $f(\mathbf{m}) = b + \sum_{i=1}^k c_i \mathbf{m}(i)$. Since f is not eventually \mathbb{N} -linear, $b \neq 0$. Thus, for all $\mathbf{m} \in \mathbb{N}_{\geq m_0}^k$, $f(2\mathbf{m}) \neq 2f(\mathbf{m})$.

Let I be the set of all α -dense valid initial configurations \mathbf{i} representing input in $\mathbb{N}_{\geq m_0}^k$ such that $f(2\mathbf{m}) > 0$. Let $S = \{\mathbf{o} \mid (\exists \mathbf{i} \in I_\alpha) \mathbf{i} \Longrightarrow \mathbf{o} \text{ and } \mathbf{o} \text{ is stable}\}$. By hypothesis we have that for each $\mathbf{i} \in I$, $\top[\mathbf{i} \Longrightarrow S] = o(n)$.

Apply Corollary 3.11. Then there are matrices $\tilde{\mathbf{D}}_1 \in \mathbb{Z}^{\Gamma \times \Delta}$ and $\tilde{\mathbf{D}}_2 \in \mathbb{Z}^{\Gamma \times \Delta}$, an infinite set $N \subseteq \mathbb{N}$, and infinite nondecreasing sequences of configurations $(\mathbf{i}_n)_{n \in N}$ and $(\mathbf{o}_n)_{n \in N}$ such that the following holds. Let $\Delta = \text{bdd}((\mathbf{o}_n)_{n \in N})$, $d = |\Delta|$, and $\Gamma = \Lambda \setminus \Delta$. For each $n \in N$, let $\mathbf{o}_n^\Delta = \mathbf{o}_n \upharpoonright \Delta$ and $\mathbf{o}_n^\Gamma = \mathbf{o}_n \upharpoonright \Gamma$. Then $\text{amax}(\tilde{\mathbf{D}}_1), \text{amax}(\tilde{\mathbf{D}}_2) \leq d2^{2d+2}$ and

1. For all $n \in N$, $\mathbf{i}_n \in I$, $\mathbf{o}_n \in S$, $\|\mathbf{i}_n\| = \|\mathbf{o}_n\| = n$, and $\mathbf{i}_n \Longrightarrow \mathbf{o}_n$.
2. For all $b \in \mathbb{N}$, there is $n_b \in N$ such that, for all $n \in N$ such that $n \geq n_b$ and all $\mathbf{d}^\Delta \in \mathbb{N}^\Delta$ such that $\max(\mathbf{d}^\Delta) \leq b$, we have $2\mathbf{i}_n + \mathbf{d}^\Delta \Longrightarrow 2\mathbf{o}_n^\Gamma + \tilde{\mathbf{D}}_1 \cdot \mathbf{o}_n^\Delta + \tilde{\mathbf{D}}_2 \cdot \mathbf{d}^\Delta$.

Let $b = 0$; below we ensure that $\max(\mathbf{d}^\Delta) \leq b$. Let $\mathbf{m} = \mathbf{i}_{n_b} \upharpoonright \Sigma$. Let $\mathbf{s}_0 = 2\mathbf{o}_{n_b}^\Gamma + \tilde{\mathbf{D}}_1 \cdot \mathbf{o}_{n_b}^\Delta$. Letting $\mathbf{d}^\Delta = \mathbf{0}$, the above argument shows that $2\mathbf{i}_{n_b} \Longrightarrow \mathbf{s}_0$. Since $\mathbf{s}_0 \in \mathbb{N}^\Gamma$, it is stable by Observation 3.12. Since the protocol is correct, $\mathbf{s}_0(\bar{y}) = f(2\mathbf{m}) > 0$. Then $4\mathbf{i}_{n_b} \Longrightarrow 2\mathbf{s}_0$, also stable, with $2\mathbf{s}_0(\bar{y}) = 2f(2\mathbf{m}) \neq f(4\mathbf{m})$, a contradiction. \square

Finally, we can prove the main result of this section, that any protocol requires linear expected time to stably compute a non-eventually- \mathbb{N} -linear function.

Intuitively, the proof has a similar shape to that of Theorem 4.4. However, applying the surgery lemmas to fool the population protocol is more difficult, since with the predicate output convention, it is immediate that if two configurations have a well-defined output and share at least one state in common, they must have the same output (equal to the vote of that state). With the integer-count output convention, things are trickier. The surgery of Corollary 3.11 to consume additional input states affects the count of the output state—how do we know that the effect of the surgery on the output is not consistent with the desired output of the function? In order to arrive at a contradiction we develop two techniques. The first involves showing that the slope of the change in the count of the output state as a function of the input states is inconsistent. The second involves exposing the semilinear structure of the graph of the function being computed, and forcing it to enter the “wrong piece” (i.e., periodic coset).

Theorem 8.5. *Let $f : \mathbb{N}^k \rightarrow \mathbb{N}$, and let \mathcal{C} be a function-computing leaderless population protocol that stably computes f . If f is not eventually \mathbb{N} -linear then \mathcal{C} takes expected time $\Omega(n)$.*

Proof. If f is eventually \mathbb{N} -affine but not eventually \mathbb{N} -linear, then Theorem 8.4 implies that f requires expected time $\Omega(n)$ to stably compute. So assume that f is not eventually \mathbb{N} -affine.

Since f is semilinear, Lemma 8.3 tells us that there is $\mathbf{v} \in \{0, 1\}^k$ and $\alpha > 0$ such that there are infinitely many α -dense $\mathbf{m} \in \mathbb{N}^k$ such that $f(\mathbf{m} + \mathbf{v}) - f(\mathbf{m}) \neq f(\mathbf{m} + 2\mathbf{v}) - f(\mathbf{m} + \mathbf{v})$. Let C be the set of all such \mathbf{m} .

Note that for every point $\mathbf{m}' \in \mathbb{N}^k$, there exists $\mathbf{m} \in \mathbb{N}^k$ and $\mathbf{u} \in \{0, 1\}^k$ such that $\mathbf{m}' = 2\mathbf{m} + \mathbf{u}$; i.e., every point \mathbf{m}' is equal some point $2\mathbf{m}$ with all even coordinates, plus some binary vector \mathbf{u} (to reach the odd coordinates in \mathbf{m}').

Let I be a set of α -dense valid initial configurations \mathbf{i} representing an input \mathbf{m} such that there is some $\mathbf{m}' \in C$ and $\mathbf{u} \in \{0, 1\}^k$ such that $\mathbf{m}' = 2\mathbf{m} + \mathbf{u}$. Let $S = \{\mathbf{o} \mid (\exists \mathbf{i} \in I) \mathbf{i} \Longrightarrow \mathbf{o} \text{ and } \mathbf{o} \text{ is stable}\}$ be the set of stable configurations reachable from some initial configuration in I . Assume for the sake of contradiction that for each $\mathbf{i} \in I$, $\top[\mathbf{i} \Longrightarrow S] = o(n)$.

Apply Corollary 3.11. Then there are matrices $\tilde{\mathbf{D}}_1 \in \mathbb{Z}^{\Gamma \times \Delta}$ and $\tilde{\mathbf{D}}_2 \in \mathbb{Z}^{\Gamma \times \Delta}$, an infinite set $N \subseteq \mathbb{N}$, and infinite nondecreasing sequences of configurations $(\mathbf{i}_n)_{n \in N}$ and $(\mathbf{o}_n)_{n \in N}$ such that the following holds. Let $\Delta = \text{bdd}((\mathbf{o}_n)_{n \in N})$, $d = |\Delta|$, and $\Gamma = \Lambda \setminus \Delta$. For each $n \in N$, let $\mathbf{o}_n^\Delta = \mathbf{o}_n \upharpoonright \Delta$ and $\mathbf{o}_n^\Gamma = \mathbf{o}_n \upharpoonright \Gamma$. Then $\text{amax}(\tilde{\mathbf{D}}_1), \text{amax}(\tilde{\mathbf{D}}_2) \leq d2^{2d+2}$ and

1. For all $n \in N$, $\mathbf{i}_n \in I$, $\mathbf{o}_n \in S$, $\|\mathbf{i}_n\| = \|\mathbf{o}_n\| = n$, and $\mathbf{i}_n \implies \mathbf{o}_n$.
2. For all $b \in \mathbb{N}$, there is $n_b \in N$ such that, for all $n \in N$ such that $n \geq n_b$ and all $\mathbf{d}^\Delta \in \mathbb{N}^\Delta$ such that $\max(\mathbf{d}^\Delta) \leq b$, we have $2\mathbf{i}_n + \mathbf{d}^\Delta \implies 2\mathbf{o}_n^\Gamma + \tilde{\mathbf{D}}_1 \cdot \mathbf{o}_n^\Delta + \tilde{\mathbf{D}}_2 \cdot \mathbf{d}^\Delta$.

Some of the inputs $\bar{x}_1, \dots, \bar{x}_k$ are in Δ , and others are in Γ . Let $\mathbf{u}^\Delta = \mathbf{u} \upharpoonright \Delta$, $\mathbf{u}^\Gamma = \mathbf{u} \upharpoonright \Gamma$, $\mathbf{v}^\Delta = \mathbf{v} \upharpoonright \Delta$, and $\mathbf{v}^\Gamma = \mathbf{v} \upharpoonright \Gamma$. Let $b = 3$; below we ensure that $\max(\mathbf{d}^\Delta) \leq b$. Let $\mathbf{m} \in \mathbb{N}^k$ be the input represented by \mathbf{i}_{n_b} . Let $\mathbf{m}' \in C$ and $\mathbf{u} \in \{0, 1\}^k$ such that $\mathbf{m}' = 2\mathbf{m} + \mathbf{u}$. Let $\mathbf{s}_0 = 2\mathbf{o}_n^\Gamma + \tilde{\mathbf{D}}_1 \cdot \mathbf{o}_n^\Delta + \tilde{\mathbf{D}}_2 \cdot \mathbf{u}^\Delta$. Letting $\mathbf{d}^\Delta = \mathbf{u}^\Delta$ in Lemma 3.10, we have $2\mathbf{i}_{n_b} + \mathbf{u}^\Delta \implies \mathbf{s}_0$. By correctness, since initial configuration $2\mathbf{i}_{n_b} + \mathbf{u}$ represents input \mathbf{m}' , we have $\mathbf{s}_0(\bar{y}) = f(\mathbf{m}')$.

Let $\mathbf{s}_1 = \mathbf{s}_0 + \tilde{\mathbf{D}}_2 \cdot \mathbf{d}^\Delta$. Then letting $\mathbf{d}^\Delta = \mathbf{u}^\Delta + \mathbf{v}^\Delta$ in Lemma 3.10, we have that $2\mathbf{i}_{n_b} + \mathbf{u}^\Delta + \mathbf{d}^\Delta \implies \mathbf{s}_1$. By additivity this implies that $2\mathbf{i}_{n_b} + \mathbf{u}^\Delta + \mathbf{u}^\Gamma + \mathbf{v}^\Gamma \implies \mathbf{s}_1 + \mathbf{u}^\Gamma + \mathbf{v}^\Gamma$. By correctness, $\mathbf{s}_1(\bar{y}) = f(\mathbf{m}' + \mathbf{v})$. By Corollary 3.13, $\mathbf{s}_1 + \mathbf{u}^\Gamma + \mathbf{v}^\Gamma$ is stable.

Let $\mathbf{s}_2 = \mathbf{s}_0 + 2\tilde{\mathbf{D}}_2 \cdot \mathbf{d}^\Delta$. Then letting $\mathbf{d}^\Delta = \mathbf{u}^\Delta + 2\mathbf{v}^\Delta$ in Lemma 3.9, we have that $2\mathbf{i}_{n_b} + \mathbf{u}^\Delta + 2\mathbf{v}^\Delta \implies \mathbf{s}_2$. By additivity this implies that $2\mathbf{i}_{n_b} + \mathbf{u}^\Delta + 2\mathbf{v}^\Delta + \mathbf{u}^\Gamma + 2\mathbf{v}^\Gamma \implies \mathbf{s}_2 + \mathbf{u}^\Gamma + 2\mathbf{v}^\Gamma$. By Corollary 3.13, $\mathbf{s}_2 + \mathbf{u}^\Gamma + 2\mathbf{v}^\Gamma$ is stable.

By linearity, $\mathbf{s}_2(\bar{y}) = \mathbf{s}_0(\bar{y}) + 2\mathbf{q}(\bar{y})$. However, since $f(\mathbf{m}' + \mathbf{v}) - f(\mathbf{m}') \neq f(\mathbf{m}' + 2\mathbf{v}) - f(\mathbf{m}' + \mathbf{v})$, this implies $\mathbf{s}_2(\bar{y}) \neq f(\mathbf{m}' + 2\mathbf{v})$. However, since \mathbf{s}_2 is stable and reachable from $2\mathbf{i}_m + \mathbf{u} + 2\mathbf{v}$, representing input $\mathbf{m}' + 2\mathbf{v}$, this implies that the protocol stabilizes on the wrong output. \square

9 Conclusion

Some interesting questions remain open.

Time complexity of non-eventually- \mathbb{N} -linear functions and non-eventually-constant predicates. The most obvious open question is to determine the optimal stabilization time complexity of computing semilinear functions and predicates not satisfying the hypotheses of Theorems 4.4 and 8.5; namely the *eventually \mathbb{N} -linear* functions, (e.g., $f(m) = 0$ if $m < 3$ and $f(m) = m$ otherwise) and *eventually constant* predicates (e.g., $\phi(m) = 1$ iff $m \geq 2$). (See Sections 8 and 4 for formal definitions.) The only known examples of eventually \mathbb{N} -linear functions computable in sublinear time are the \mathbb{N} -linear functions (e.g., $f(m) = 2m$), computable in logarithmic time by Observation 7.1. The only known examples of eventually constant predicates that are computable in sublinear time are the *detection predicates* studied in [14] (and observed to be decidable in logarithmic time): predicates whose value depends only on the presence or absence of certain inputs, but not on their exact positive values (e.g., $\phi(m) = 1$ iff $m \geq 1$).

Allowing more than $O(1)$ states. Alistarh, Aspnes, Eisenstat, Gelashvili, and Rivest [1] showed time lower bounds for leader election and majority with superconstant states (i.e., the number of states is allowed to grow with the population size). In this paper we have shown time lower bounds for more general function and predicate computation with a *constant* set of states. It is natural to ask whether more general function and predicate computation time lower bounds can be proven for superconstant states using similar techniques. Notably, Chatzigiannakis, Michail, Nikolaou, Pavlogiannis, and Spirakis [13] have shown that if the number of states $\lambda_n = o(\log n)$, then the protocol's computational power remains limited to computing semilinear predicates, the same limitation that applies when $\lambda_n = O(1)$.¹⁵ However, with $\lambda_n = \Omega(\log n)$, the computational ability of population protocols moves beyond semilinear predicates [13]. Thus, there is a wider

¹⁵The bound on states is described in [13] in terms of space available to a Turing machine that computes the transition function, permitting $o(\log \log n)$ space. Since a Turing machine with space s has $2^{O(s)}$ configurations, this

class of functions and predicates under this relaxed constraint, which may require new techniques to analyze their population protocol time complexity.

Convergence time without a leader. Although we measure computation time with respect to stabilization—the ultimate goal of stable computation—some work uses a different goalpost for completion. Consider a protocol stably computing a function, and consider one particular transition sequence that describes its history. We can say the transition sequence *converged* at the point when the output count is the same in every subsequently reached configuration. In contrast, recall that the point of stabilization is when the output count is the same in every subsequently *reachable* configuration (whether actually reached in the transition sequence or not). In other words, after stabilization, even an adversarial scheduler cannot change the output. Measuring time to stabilization in the randomized model, as we do here, measures the expected time until the probability of changing the output becomes 0.

Our proof shows only that stabilization must take expected $\Omega(n)$ time for “most” predicates and functions if there is no initial leader. However, convergence could occur much earlier in a transition sequence than stabilization. Indeed, Kosowski and Uznański [27] show that all semilinear predicates can be computed without an initial leader, converging in $O(\text{polylog } n)$ time if a small probability of error is allowed, and converging in $O(n^\epsilon)$ time with probability 1, where ϵ can be made arbitrarily close to 0 by changing the protocol.

Stabilization time with a leader. It remains open to determine the optimal stabilization time for stably computing semilinear predicates and functions *with* an initial leader. The stably computing protocols converging in $O(\log^5 n)$ time [6, 15] provably require expected time $\Omega(n)$ to stabilize, and it is unknown whether faster stabilization is possible.

High-probability computation of non-semilinear functions/predicates without a leader. Going beyond stable computation, the open question of Angluin, Aspnes, and Eisenstat [6] asks whether their efficient high-probability simulation of a space-bounded Turing machine by a population protocol could remove the assumption of an initial leader. That simulation has some small probability $\epsilon > 0$ of failure, so if one could elect a leader with a small probability $\epsilon' > 0$ of error and subsequently use it to drive the simulation, by the union bound the total probability of error would be at most $\epsilon + \epsilon'$ (i.e., still close to 0). Recent work by Kosowski and Uznański [27] may be relevant, which shows a protocol for leader election converging in $O(\text{polylog } n)$ time with a small probability of error.

Non-dense initial configurations without a leader. Our general negative result applies to α -dense initial configurations. However, is sublinear time stable computation possible from other kinds of initial configurations that satisfy our intuition of not having preexisting leaders? In particular, our technique for proving time lower bounds fails when the input states themselves have counts sufficiently skewed to prevent α -denseness, even without any other state initially present. For example, given any function $f : \mathbb{N} \rightarrow \mathbb{N}$ computable quickly with a leader, imagine a function $f_2 : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that $f_2(m, 1) = f(m)$ for all $m \in \mathbb{N}$. It is difficult to rule out the possibility of a protocol using the unique agent in state x_2 as a leader to simulate the fast protocol computing f . However, it is not clear how to ensure that f_2 is computed correctly on other inputs where $m_2 \neq 1$, since it is unknown how to compute whether $m_2 = 1$ in sublinear time.

is equivalent to requiring that the number of states be at most $o(\log n)$ (and furthermore that the transition function be uniformly computed by a single Turing machine.)

Functions approximable in unbounded time. We showed in Theorem 7.2 that linear functions all coefficients in $\mathbb{Q}_{\geq 0}$ can be approximated in logarithmic time with linear error, but we also showed in Theorem 6.1 that linear functions with some coefficient not in \mathbb{N} cannot even be approximated in sublinear time with sublinear error. However, this is the first study of approximate function computation with population protocols, so it is not even clear what functions can be approximated at all with unbounded time. Clearly it is not limited to semilinear functions, or even Turing-computable functions. For example, the identity function is a close approximation of the uncomputable function $f(m) = m$ if the m 'th Turing machine halts and $f(m) = m + 1$ otherwise.

A sensible way to formulate the question might be this: require (as we do) for a protocol on input $\mathbf{m} \in \mathbb{N}^k$ to stabilize, but nondeterministically to one of a set of possible outputs $Y_{\mathbf{m}} \subseteq \mathbb{N}$. Define the nondeterministic function it computes to be $f : \mathbb{N}^k \rightarrow \mathcal{P}(\mathbb{N})$ defined by $f(\mathbf{m}) = Y_{\mathbf{m}}$. What range of such functions can be computed? This is applicable to approximation when the output set is a small interval.

It might also be interesting to consider multi-valued nondeterministic functions, where the protocol has ℓ different output states $\bar{y}_1, \dots, \bar{y}_\ell$ and computes a function $f : \mathbb{N}^k \rightarrow \mathcal{P}(\mathbb{N}^\ell)$. In fact, this question is interesting even when there is no input, in the generalized model of chemical reaction networks that permit some reactions (transitions) such as $\bar{x} \rightarrow \bar{y}, \bar{y}$ or $\bar{x}_1, \bar{x}_2 \rightarrow \bar{x}_1$ to alter the population size. The system starts with a single \bar{x} , and it nondeterministically stabilizes to some counts of $\bar{y}_1, \dots, \bar{y}_\ell$. It might appear as though the set of counts that can be reached is semilinear, but in fact non-semilinear sets are possible [24].

Alternate Boolean output conventions. The standard model by which a population protocol reports a Boolean output is by consensus: each state “votes” 0 or 1, and output is defined only if all states present in a configuration vote unanimously. Brijder, Doty, and Soloveichik [10] showed that under some reasonable alternative output conventions, the class of Boolean predicates stably computable remains the semilinear predicates. However, it may not be the case that the *efficiency* of predicates remains unchanged. For example, one alternative is the *democratic* output convention: the output is $b \in \{0, 1\}$ if the b voters outnumber the $1-b$ voters. In this case, the majority predicate becomes trivially computable in constant time, merely by having each input state \bar{x}_1 and \bar{x}_2 vote for itself, and no transitions are even required: the initial configuration already reports the correct answer.

Acknowledgements. The authors are grateful to Sungjin Im for the proof of Lemma 8.1. We also thank anonymous reviewers for their helpful comments.

References

- [1] Dan Alistarh, James Aspnes, David Eisenstat, Rati Gelashvili, and Ronald L. Rivest. Time-space trade-offs in molecular computation. In *SODA 2017: Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2017. to appear.
- [2] Dan Alistarh, James Aspnes, and Rati Gelashvili. Space-optimal majority in population protocols. In *SODA 2018: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2221–2239. SIAM, 2018.
- [3] Dan Alistarh and Rati Gelashvili. Polylogarithmic-time leader election in population protocols. In *ICALP 2015: Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming, Kyoto, Japan*, 2015.

- [4] Dana Angluin, James Aspnes, Zoë Diamadi, Michael Fischer, and René Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed Computing*, 18:235–253, 2006. Preliminary version appeared in PODC 2004.
- [5] Dana Angluin, James Aspnes, and David Eisenstat. Stably computable predicates are semi-linear. In *PODC 2006: Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*, pages 292–299, New York, NY, USA, 2006. ACM Press.
- [6] Dana Angluin, James Aspnes, and David Eisenstat. Fast computation by population protocols with a leader. *Distributed Computing*, 21(3):183–199, September 2008. Preliminary version appeared in DISC 2006.
- [7] Dana Angluin, James Aspnes, David Eisenstat, and Eric Ruppert. The computational power of population protocols. *Distributed Computing*, 20(4):279–304, 2007.
- [8] Alexandre Baccouche, Kevin Montagne, Adrien Padirac, Teruo Fujii, and Yannick Rondelez. Dynamic dna-toolbox reaction circuits: a walkthrough. *Methods*, 67(2):234–249, 2014.
- [9] James M Bower and Hamid Bolouri. *Computational modeling of genetic and biochemical networks*. MIT press, 2004.
- [10] Robert Brijder, David Doty, and David Soloveichik. Robustness of expressivity in chemical reaction networks. In *DNA 2016: Proceedings of the 22th International Meeting on DNA Computing and Molecular Programming*, pages 52–66, 2016.
- [11] Robert Brijder, David Doty, and David Soloveichik. Democratic, existential, and consensus-based output conventions in stable computation by chemical reaction networks. *Natural Computing*, 17(1):97–108, 2018. Special issue of invited papers from DNA 2016.
- [12] E. Cardoza, Richard J. Lipton, and Albert R. Meyer. Exponential space complete problems for Petri nets and commutative semigroups (preliminary report). In *STOC 1976: Proceedings of the 8th annual ACM Symposium on Theory of Computing*, pages 50–54. ACM, 1976.
- [13] Ioannis Chatzigiannakis, Othon Michail, Stavros Nikolaou, Andreas Pavlogiannis, and Paul G. Spirakis. Passively mobile communicating machines that use restricted space. *Theoretical Computer Science*, 412(46):6469–6483, October 2011.
- [14] Ho-Lin Chen, Rachel Cummings, David Doty, and David Soloveichik. Speed faults in computation by chemical reaction networks. *Distributed Computing*, 2015. to appear. Special issue of invited papers from DISC 2014.
- [15] Ho-Lin Chen, David Doty, and David Soloveichik. Deterministic function computation with chemical reaction networks. *Natural Computing*, 13(4):517–534, 2014. Preliminary version appeared in DNA 2012.
- [16] Yuan-Jyue Chen, Neil Dalchau, Niranjana Srinivas, Andrew Phillips, Luca Cardelli, David Soloveichik, and Georg Seelig. Programmable chemical controllers made from DNA. *Nature Nanotechnology*, 8(10):755–762, 2013.
- [17] Leonard E. Dickson. Finiteness of the odd perfect and primitive abundant numbers with n distinct prime factors. *American Journal of Mathematics*, 35(4):413–422, October 1913.

- [18] David Doty. Timing in chemical reaction networks. In *SODA 2014: Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 772–784, January 2014.
- [19] David Doty and Monir Hajiaghayi. Leaderless deterministic chemical reaction networks. *Natural Computing*, 14(2):213–223, 2015. Preliminary version appeared in DNA 2013.
- [20] David Doty and David Soloveichik. Stable leader election in population protocols requires linear time. *Distributed Computing*, 31(4):257–271, 2018. Special issue of invited papers from DISC 2015.
- [21] Leszek Gasieneć and Grzegorz Stachowiak. Fast space optimal leader election in population protocols. In *SODA 2018: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 2653–2667, 2018.
- [22] Leszek Gasieneć, Grzegorz Stachowiak, and Przemysław Uznanski. Almost logarithmic-time space optimal leader election in population protocols. *CoRR*, abs/1802.06867, 2018.
- [23] Seymour Ginsburg and Edwin H. Spanier. Semigroups, Presburger formulas, and languages. *Pacific Journal of Mathematics*, 16(2):285–296, 1966.
- [24] John E. Hopcroft and Jean-Jacques Pansiot. On the reachability problem for 5-dimensional vector addition systems. *Theoretical Computer Science*, 8:135–159, 1979.
- [25] Sungjin Im. Personal communication.
- [26] Richard M Karp and Raymond E Miller. Parallel program schemata. *Journal of Computer and System Sciences*, 3(2):147–195, 1969.
- [27] Adrian Kosowski and Przemysław Uznanski. Brief announcement: Population protocols are fast. In *PODC 2018: Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing*, pages 475–477. ACM, 2018.
- [28] 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.
- [29] Othon Michail and Paul G Spirakis. How many cooks spoil the soup? In *International Colloquium on Structural Information and Communication Complexity*, pages 3–18. Springer, 2016.
- [30] Carl A Petri. Communication with automata. Technical report, DTIC Document, 1966.
- [31] Niranjana Srinivas. *Programming chemical kinetics: Engineering dynamic reaction networks with DNA strand displacement*. PhD thesis, California Institute of Technology, 2015.
- [32] Chris Thachuk, Erik Winfree, and David Soloveichik. Leakless DNA strand displacement systems. In *DNA 2015: Proceedings of the 21st International Conference on DNA Computing and Molecular Programming*, pages 133–153. Springer, 2015.
- [33] Vito Volterra. Variazioni e fluttuazioni del numero dindividui in specie animali conviventi. *Mem. Acad. Lincei Roma*, 2:31–113, 1926.
- [34] David Yu Zhang and Georg Seelig. Dynamic DNA nanotechnology using strand-displacement reactions. *Nature Chemistry*, 3(2):103–113, 2011.