New Clocks, Optimal Line Formation and Efficient Replication Population Protocols (Making Population Protocols Alive)

Leszek Gąsieniec

Department of Computer Science, University of Liverpool, Ashton Street, L69 38X, U.K.

Paul Spirakis

Department of Computer Science, University of Liverpool, Ashton Street, L69 38X, U.K.

Grzegorz Stachowiak

Institute of Computer Science, University of Wrocław Joliot Curie 15, Wrocław, Poland

November 23, 2021

Abstract

We consider the model of population protocols permitting presence of dynamically changing edges connecting agents. Our main contribution is a new constant space phase clock allowing to count parallel time $O(n \log n)$ whp in the adopted model. This clock admits confirmation of slow leader election and in turn construction of a line (and a ring) comprising every agent in the optimal parallel time $\Theta(n \log n)$ and constant space. This improves on the currently best known upper bound $O(n^2)$.

We also discuss a variant of the new clock in which utilisation of edges is replaced by interaction of agents with a unique leader. This variant provides a universal (for models with and without edges) synchronisation mechanism and is adopted in some of our efficient line replication protocols.

1 Introduction

The model of *population protocols* originates from the seminal paper of Angluin et al. [4]. This model provides tools for the formal analysis of *pairwise interactions* between simple indistinguishable entities referred to as *agents*. The agents are equipped with limited storage, communication and computation capabilities. When two agents engage in a direct interaction their memory content is assessed and their states are modified according to the predefined transition function that forms an integral part of the population protocol. In the *probabilistic variant* of population protocols adopted in this paper, in each step the

random scheduler selects a pair of agents uniformly at random. In this variant, in addition to *state utilisation* one is also interested in the *running time* of the proposed solutions. In more recent work on population protocols the focus is on *parallel time* defined as the total number of pairwise interactions leading to the solution divided by the size (in our case n) of the population. The parallel time provides also a good estimation on the number of interactions each agent was involved in throughout the computation process.

Unless stated otherwise we assume that any protocol starts in the *initial configuration* where all agents are in the same predefined *initial state*. A population protocol *terminates with success* if the whole population eventually stabilises, i.e., it arrives at and stays indefinitely in the *final configuration* of states reflecting the desired property of the solution.

1.1   Constructors

While in the standard population protocol model the population of agents remains unstructured, in the model introduced in [18] and adopted in this paper any two agents may become connected by establishing a physical (edge) connections between them. The two connected agents may later choose to drop their connection when they meet again. In this way the agents can self organize themselves into a desired temporary or more definite structures. Such distributed and dynamically structured systems based on population protocols are called *network constructors* or simply *constructors* [18].

Note that the expected number of interactions to fix (establish or modify) a particular connection is $\Theta(n^2)$ as each pair of agents is selected uniformly at random. And this cost often dominates the time complexity of the constructors protocols. But there are some exceptions where the focus is not on very specific but more arbitrary connections randomly selected from a large group of available connections. A good example is the construction of almost balanced trees from [11] requiring $o(n^2)$ interactions.

The model When possible we will use capital letters to denote states of the agents. In order to accommodate management of edge connections we need to extend the transition function in which every rule is of the following type

$$P + Q + S \ \longrightarrow \ P' + Q' + S'$$

The first two terms on both sides of the rule refer to the states P and Q of interacting agents before and P' and Q' after the interaction. The third term S before and S' after the interaction indicates the status of connection between the two where 1 means the interacting agents are connected or not denoted by 0. Note that the state of an agent can be more complex, e.g., represented as a tuple with several components, where only some components may change during an interaction. In such compound cases but also when the states are denoted by longer terms or numbers we will use vector representation with $<, >$ brackets.

One of the central problems in network constructors is formation of a stable line (or ring) comprising all agents. When no leader is initially assumed, the

fastest known protocol for a stable line construction requires $\Theta(n^3)$ interactions in expectation, i.e. $\Theta(n^2)$ parallel time, see [18]. Until now, no bounds were known for efficient ring construction. In this paper we propose fastest possible protocols for the line and ring construction, while preserving constant space utilisation in agents.

Our clock, line and ring constructions are always correct and they stabilise *with high probability (whp)* which we define as follows. Let η be a universal constant referring to the reliability of our protocols. We say that an event occurs with *negligible* probability if it occurs with probability at most $n^{-\eta}$, and an event occurs with high probability if it occurs with probability at least $1 - n^{-\eta}$. This estimate is of an asymptotic nature, i.e., we assume n is large enough to validate the results. Similarly, we say that an algorithm succeeds with high probability if it succeeds with probability at least $1 - n^{-\eta}$. When we refer to the probability of failure p different to $n^{-\eta}$, we say directly *with probability at least* $1 - p$. Our protocols make heavy use of Chernoff bounds and the new tail bounds for sums of geometric random variables derived in [17]. We refer to this new bound as to Chernoff-Janson bound.

1.2 Related work

One of the main tools used in this paper refers to the central problem of *leader election*. In this problem the final configuration comprises a single agent in the *leader* state and all other agents in the *follower* state. The leader election problem received in recent years greater attention in the context of population protocols. In particular, the results from [10, 12] laid down the foundation for the proof that leader election cannot be solved in a sublinear time with agents utilising a fixed number of states [14]. In further work [3], Alistarh and Gelashvili studied the relevant upper bound, where they proposed a new leader election protocol stabilising in time $O(\log^3 n)$ assuming $O(\log^3 n)$ states per agent.

In a very recent work Alistarh *et al.* [1] consider a more general trade-off between the number of states used by agents and the time complexity of stabilisation. In particular, the authors provide a separation argument distinguishing between *slowly stabilising* population protocols which utilise $o(\log \log n)$ states and *rapidly stabilising* protocols with $O(\log n)$ states per agent. This result nicely coincides with another fundamental observation by Chatzigiannakis *et al.* [9] which states that population protocols utilizing $o(\log \log n)$ states are limited to semilinear predicates, while the availability of $O(\log n)$ states admits computation of symmetric predicates. Further developments include also a protocol which elects the leader in time $O(\log^2 n)$ w.h.p. and in expectation utilizing $O(\log^2 n)$ states [8]. The number of states was later reduced to $O(\log n)$ by Alistarh et al. in [2] and by Berenbrink et al. in [7] through the application of two types of synthetic coins.

In more recent work of Gąsieniec and Stachowiak reduce memory utilisation to $O(\log \log n)$ while preserving the time complexity $O(\log^2 n)$ whp [15]. We also know that the high probability can be traded for faster leader election in the expected parallel time $O(\log n \log \log n)$, see [16]. This upper bound was

reduced to the optimal expected time $O(\log n)$ by Berenbrink *et al.* in [6]. In fact the main open problem is to establish whether one can elect a single leader in time $o(log^2 n)$ whp while preserving the optimal number of states $O(\log \log n)$.

The protocols and methods discussed in this paper are closely related to the concept of phase clocks. The term and the first analysis of a leader based $O(1)$ space phase clock was given by Angluin *et al.* in [5]. Further extensions including a junta based clock and nested clocks counting any time $O(\log^c n)$ whp, for any constant c, can be found in [15]. In very recent work [13] Doty *et al.* propose and analyse constant resolution clocks utilising $O(\log n)$ states as the main engine in the optimal *majority* computation protocols.

1.3 Our results

Our main contribution is a new constant space clock allowing (in the model with edge connections) to count parallel time $O(n \log n)$ whp. This clock is used to confirm the conclusion of the slow leader election protocol. The selected a leader is used to construct a line and a ring of agents in the optimal parallel time $O(n \log n)$ whp. We also propose and analyse the second clock based on the selected leader which also operates in time $O(n \log n)$. Please note that this clock is universal, i.e., it can be used in population protocols with and without edges. Thanks to periodic application of the second clock one can monitor efficient construction of lines (and rings due to lack of space discussed in the Appendix). All our new protocols use the optimal constant space and operate whp. With the exception of replication our protocols are also optimal wrt the time complexity $\Theta(n \log n)$. They are also always correct meaning that with a negligible probability they may operate longer, however, they never terminate with the wrong answer.

2 Two clocks and leader election

In order to compute a single leader in the population we are executing two protocols simultaneously. We run the slow (naive) leader election protocol to identify a unique leader in time $O(n \log n)$ whp, simultaneously with the new *matching based clock* (discussed below) which counts time $\Theta(n \log n)$ whp. When this clock concludes the remaining leader progresses to further stages where we compute the line, the ring and discuss the line replication protocol.

The transition rules of the considered protocols follow.

Slow leader election
$$L + L \rightarrow L + F$$
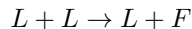
where L represents a (remaining) leader candidate, and F stands for a *follower* or a *free* agent. It is known that such naive leader election protocol operates in time $O(n \log n)$ whp.

Matching based clock The proposed matching based clock assumes the constructors model in which the transition function recognises whether to agents

4

are connected by an edge, indicated by 1, or not, indicated by 0. The agents begin in the predefined state $< start >$. When two agents in state $< start >$ interact they get connected and they enter the *counting stage* in which their counters are initially set to 0 and eventually reach the maximum value max. Note that these counters can either go up or down depending on the rule (of the transition function) is used during interaction. Note also that the number of agents present in the counting stage is always even. The counting stage protocol guaranties that the counters of all agents which enter this stage reach value max in time $\Theta(n \log n)$, see Theorem 1. And in the next interaction between the two connected agents in state $< max >$ the connection is removed and the states are updated to $< end >$ indicating the end of the counting stage.

The rules of the transition function used in the counting stage are as follows:

Initialisation

$$< start > + < start > \; + 0 \; \longrightarrow < 0 > + < 0 > + 1$$

Timid counting

- For all connected $i \leq j$ and $i < max$

$$< i > + < j > + 1 \; \longrightarrow < i + 1 > + < i + 1 > + 1$$

- For all disconnected $i < j$

$$< i > + < j > + 0 \; \longrightarrow < i > + < i + 1 > + 0$$

Maximum level epidemic

$$< max > + < i > + 0 \; \longrightarrow < max > + < max > + 0$$

Conclude and disconnect

$$< max > + < max > + 1 \; \longrightarrow < end > + < end > + 0$$

$$< start > + < end > + 0 \; \longrightarrow < end > + < end > + 0$$

Leader based clock We allocate separate constant memory to host the states of the leader based clock. This allows to run the actions of the two clocks simultaneously and independently. The followers in the leader based clock start with the counter set to 0 which is denoted by $< 0 >$, while L refers to the leader state. Note that state $< 0 >$ is initiated for the leader based clock as soon as the agent reaches state $< max >$ or $< end >$ in the matching based clock. The timid counting rules now refer to the interactions with the leader L.

Timid counting

- Leader interactions, where $i < max$

$$< i > + L \; \longrightarrow < i + 1 > + L$$

- Non-leader interactions, where $i < j$

$$< i > + < j > \longrightarrow < i > + < i+1 >$$

One can show that the two clocks have the same asymptotic time performance, see Section 3 for the relevant detail. Note that the leader based clock can be used independently from presence of edges in the population. In particular, this clock can be used to count time required to remove all edges used in the matching based clock as well as to count time needed to form the line and the ring of all nodes.

Periodic leader based clock One can extend the functionality of the leader based clock such that it paces multiple rounds, each operating in time $\Theta(n \log n)$, of some more complex process, e.g., line replication.

The extension uses three consecutive stages 0, 1 and 2, where each stage corresponds to one full execution of the leader based clock, and all three stages form a single round. Thus any round starts in stage 0 on which conclusion maximum level epidemic signal is accompanied by the message that stage 0 is concluded. Note that such message is distributed rapidly by one-way epidemic in time $O(\log n)$ whp. And when this happens all agents proceed to stage 1. Note that the signal to start stage 1 remains in the system throughout the whole stage but it will be wiped out by the signal announcing the beginning of stage 2. Analogously, on the conclusion of stage 2 the signal announcing stage 0 (and the new round) will wipe out the previous message. This way after at most $O(\log n)$ time delay (caused by the time of the epidemic) all agents will always run the clock in the same stage whp.

3 The analysis

In this section we analyse the time complexity the two clocks from Section 2. We will prove the following theorem towards the end of this section, but we need first a collection of lemmas.

Theorem 1. *In either of the clocks the time in which state $<max>$ appears for the first time is $\Theta(n \log n)$ whp.*

We first focus on the matching based leaderless clock and later extend the reasoning to the leader based clock.

Let us define the *edge collector problem* in which one is asked to collect all edges of a given matching M of cardinality $n' > \frac{n}{4}$. This process concludes when the random scheduler generates all edges of a given matching via pairwise interactions of agents in the population.

Lemma 1. *For any cardinality $n' \in [n/4, n/2]$, the time complexity of the edge collector problem is $O(n \log n)$. In addition, the time needed to collect the last $0.05 \cdot n$ edges (of the matching) is at least $0.4 \cdot n \ln n$ whp.*

Proof. The probability of collecting an edge when i edges are still to be collected is $\frac{2i}{n(n-1)}$, and in turn the expected time needed to collect an edge is $\frac{n-1}{2i}$. Thus the expected time to collect $k - l$ edges, when k edges remain is

$$(n-1)\sum_{i=l}^{k}\frac{1}{2i} = \frac{n-1}{2}(H_k - H_l) \sim \frac{n-1}{2}\ln\frac{k}{l}.$$

By Chernoff-Janson bound when $k = n', l = 1$ this time is $O(n\ln n)$. Using the same bound for $k = 0.05n, l = n^{0.1}$ results in time exceeding $0.4n\ln n$ whp. $\square$

Lemma 2. *All matching edges are formed in the expected time $\Theta(n)$ and whp $O(n\log n)$.*

Proof. The probability of an interaction forming edge $i + 1$ when i edges are already present is $\frac{(n-2i)(n-2i-1)}{n(n-1)}$. So the number of interactions forming edge $i + 1$ has geometric distribution with the expected value $\frac{n(n-1)}{(n-2i)(n-2i-1)}$. Thus the expected time of forming all edges is

$$\frac{1}{n}\sum_{i=1}^{n}\frac{n(n-1)}{(n-2i)(n-2i-1)} \leq n.$$

By Chernoff-Janson bound this time is $O(n\log n)$ whp. $\square$

The following lemma refers to early interactions of the matching based clock.

Lemma 3. *After time 0.51 at least $\frac{n}{2}$ agents are paired in edges whp.*

Proof. Assume that so far we have formed i edges. The probability that in an interaction edge $i + 1$ is formed is $\frac{(n-2i)(n-2i-1)}{n(n-1)}$. So the expected number of interactions T_i of forming edge $i + 1$ is $\frac{n(n-1)}{(n-2i)(n-2i-1)}$. Thus the expected number of interactions T of forming first $n/4$ edges satisfies

$$T = T_0 + \cdots + T_{n/4} = \sum_{i=0}^{n/4}\frac{n(n-1)}{(n-2i)(n-2i-1)} \sim n\int_0^{1/4}\frac{dx}{(1-2x)^2} = \frac{n}{2}.$$

By Chernoff-Janson bound (for large enough n) this process requires at most $0.51n$ interactions whp. This is equivalent to parallel time at most 0.51. $\square$

There exists a positive constant d for which the following lemma holds.

Lemma 4. *In a time window of size n^a, where $0 < a < 1$, any edge in the matching is used in at most d interactions whp.*

Proof. By Union bound the probability that an edge is a subject to at least d interactions in time n^a does not exceed

$$\binom{n^{1+a}}{d}\left(\frac{2}{n(n-1)}\right)^d \leq \left(\frac{2n^a}{n-1}\right)^d.$$

and this value is smaller than $n^{-\eta}$ is for d big enough. $\square$

Lemma 5. *In time window of size n^a, where $0.1 \le a < 1$, there are at most $2.1n^a$ edge interactions whp.*

Proof. The probability that a given interaction is an edge interaction is $\frac{2n'}{n(n-1)}$. Thus in the time window of size n^a, there are expected $2n^a \frac{n'}{n-1}$ edge interactions. By Chernoff Bound the number of edge interactions is at most $2.1n^a$ whp. $\square$

Depending on the context and for the clarity of the presentation in what follows we will use the notions of a counter and a *level* interchangeably.

Lemma 6. *Assume that integer $k < max - d - 2$. Since time $t_0 = 0.51$ and for as long as at least one level $i < k$ is present in the clock, there is a subpopulation of at least $0.1n$ agents residing on levels $j \le k$ whp. Also during this time no agent reaches level max whp.*

Proof. As we proved in Lemma 3, during the initial time 0.51 at least $n/2$ agents enter the clock with state (0) whp. Some of these agents could also relocate to the higher levels. By Lemma 5 applied to the initial time period $n^{0.2}$ there are at most $2.1n^{0.2}$ of the latter whp. Thus in the time interval $[0.51, n^{0.2}]$ level 0 is the host of at least $0.4n < 0.5n - 2.1n^{0.2}$ agents constantly residing at this level whp. Also by Lemma 4 no agent gets to level max whp.

The proof is done by induction on t. Assume before time t there are at least $0.1n$ agents on levels $j \le k$. We will prove that the thesis of the Lemma also holds before time $t' = t + n^{0.1}$ whp.

We notice first that during period $[t - n^{0.1}, t]$ all agents which entered the clock are at least once on level $l \le k+1$ whp. And indeed during this period an agent avoids interactions with agents on levels $j \le k$ with probability at most

$$(1 - 0.1/n)^{n^{1.1}} < e^{0.1n^{0.1}}$$

Thus, in this period, any agent which entered the clock goes to level at most $k + 1$ whp. And Lemma 4 guarantees that no agent reaches level max during period $[t, t + n^{0.1}]$.

In order to prove the first thesis of the lemma we consider two cases.

In the first case in time t there are at least $0.11n$ agents on levels not exceeding k. Since by Lemma 5 in time period $[t, t']$ at most $2.1n^{0.1}$ such agents can increase their level whp. And in turn, in time t' there are at least $0.1n > 0.11n - 2.1n^{0.1}$ agents on levels $j \le k$.

In the second case in time t the number of agents on levels at most k is between $0.1n$ and $0.11n$.

Let Y be the set of agents belonging to the levels above k in time t. If in time t the number of agents on levels smaller than k is bigger than $3n^{0.1}$, then by Lemma 5 the probability that in time window $[t, t']$ this number drops below $0.9n^{0.1}(= 3n^{0.1} - 2.1n^{0.1})$ is negligible. Consider any set X with $0.9n^{0.1}$ agents residing at levels smaller than k and estimate how many agents from set Y interact with them. For as long as $0.38n$ agents from Y do not interact with X, the probability of interaction between an unused (not in contact with X agents)

8

agent in Y and some agent in X is at least $0.68n^{-0.9}$. Any such interaction increases the number of agents on levels not exceeding k. Consider a sequence of $n^{1.1}$ zeroes and ones in which position ι is one (1) if and only if either

- interaction ι is between an unused agent in Y with some agent in X if there are more than $0.38n$ unused agents in Y,

- if this number is smaller than $0.38n$ value 1 is drawn with a fixed probability $0.68n^{-0.9}$.

By Chernoff bound the probability that this sequence has less than $0.6n^{0.2}$ ones is negligible. Since $0.12n < 0.11n + 0.6n^{0.2}$ this sequence has less than $0.6n^{0.2}$ ones only when the number of agents moved to levels not exceeding k is smaller than $0.6n^{0.2}$. Also by Lemma 5 during period $[t, t']$ at most $2.1n^{0.1}$ other agents may increase their level beyond k whp. So in this sub-case the number of agents on levels not exceeding k increases in period $[t, t']$ by at least $0.6n^{0.2} - 2.1n^{0.1}$.

If in time t the number of agents on levels below k is smaller than $3n^{0.1}$, then the probability of an interaction between such an agent and an agent in Y is at most $6n^{-0.9}$. Any such an interaction increases the number of agents on levels not exceeding k. By Chernoff bound the probability that this number of interactions exceeds $7n^{0.2}$ in $[t, t']$ is negligible. Thus in this sub-case the probability that the number of agents on levels at most k exceeds $0.12n > 0.11m + 7n^{0.2}$ is negligible.

We need the following two claims.

Claim 1: During period $[t, t']$ there are at most $0.26n^{0.1}$ agents located at levels $j \leq k$ which increment their level whp.

Indeed, for as long as there are at most $0.12n$ agents on levels not greater than k, the probability that such agent interacts as the initiator with a clock agent is at most $0.12/n$. Such an interaction increments the level of this clock agent with probability at most $0.12/n$. We prove that the probability of at least $0.13n^{0.1}$ such incrementations is negligible. Consider a sequence of $n^{1.1}$ zeroes and ones in which position ι is one if and only if either

- interaction ι increments initiator's level and there are at most $0.12n$ agents on levels not greater than k

- if this number is greater than $0.38n$ value 1 is drawn with a fixed probability $0.12/n$.

By Chernoff bound this sequence has less than $0.13n^{0.1}$ ones (1s) whp. On the other hand we have at most $0.12n$ agents on levels at most k whp. Thus whp at most $0.13n^{0.1}$ agents on levels not exceeding k can increment their levels in $[t, t']$ acting as initiators. Analogously we can prove that whp at most $0.13n^{0.1}$ agents on levels not exceeding k can increment their levels in $[t, t']$ acting as responders. So altogether at most $0.26n^{0.1}$ agents on levels $j \leq k$ increment their levels during period $[t, t']$ whp.

Claim 2: During period $[t, t']$ there are at least $0.75n^{0.1}$ interactions between agents on level $i < k$ and those residing on levels higher than k whp.

For as long as there are at most $0.12n$ agents on levels at most k, at least $0.38n = n/2 - 0.12n$ agents are on levels higher than k. The probability of an interaction of such agents with an agent on level $i < k$ is at least $0.76/n = 2 \cdot 0.38/n$. Any such an interaction increases the number of agents on levels not exceeding k. Consider a sequence of $n^{1.1}$ zeroes and ones in which at position ι is one (1) if and only if either

- there are at most $0.12n$ agents on levels not greater than k and interaction ι increases the number of such agents

- the number of agents on levels not exceeding k is greater than $0.12n$ and value 1 is drawn with a fixed probability $0.76/n$.

By Chernoff bound this sequence has more than $0.75n^{0.1}$ ones (1s) whp. On the other hand we have at most $0.12n$ agents on levels at most k whp. Thus whp at least $0.75n^{0.1}$ agents on levels exceeding k can reduce their levels to at most k during period $[t, t']$ while acting as initiators.

Because of both Claims 1 and 2 after time $n^{0.1}$ there are at least $0.51n^{0.1}(= 0.75n^{0.1} - 0.24n^{0.1})$ more agents with state $<j>: j \leq k$ than in time t. This proves that in time t' there are at least $0.1n$ agents with state $<j>: j \leq k$. $\quad\square$

Lemma 7. *The time in which the first agent achieves level max is larger than $(max - d - 2) \cdot 0.4n \ln n$ whp.*

Proof. Let t_i be the time when for the first time there are no agents available at levels lower than i. By Lemma 6 during period $[0.51, t_i]$, there are at least $0.1n$ agents on level i or lower. Let $n_i \in [n/4, n/2]$ be the number of edges in time t_i. Thus between time t_i and t_{i+1} at least $0.1n$ agents have to increment their levels to $i + 1$. This is done by collecting (interacting via) edges adjacent to them. By Lemma 1 this takes time at least $0.4n \ln n$. This process has to be repeated for $max - d - 2$ levels when no agents reach state (max) whp. $\quad\square$

Lemma 8. *The first agent moves to level max in time $O(n \log n)$ whp.*

Proof. The total time to initiate $\lfloor n/2 \rfloor$ edges is $O(n \log n)$ whp by Lemma 2. If the first agent achieves level max earlier the lemma remains true. If this is not the case, the time $O(n \log n)$ is determined by collection of all $\lfloor n/2 \rfloor$ edges which needs to be repeated max times resulting in the total time $O(n \log n)$. $\quad\square$

Now we are ready to prove the thesis of Theorem 1. The thesis for matching based clock follows directly from Lemmas 7 and 8. The thesis for the leader based clock can be proved by a sequence of lemmas almost identical to Lemmas 6, 7 and 8. In the analog of Lemma 6 we can take $n - 2$ followers instead of n' edges. This is because Lemma 2 assures that the time counted by the matching based clock is long enough to form all edges whp. Note that $n - 2$ agents are initiated at level 0 of the leader based clock in time $O(\log n)$ whp by the epidemic resulting in dismantling of the matching based clock. And in turn we can use the initial time $O(n \log n)$ instead of 0.51 in the analog of Lemma 6.

4   Fast formation of lines

Line formation We define and analyse a new optimal line formation protocol which operates in time $\Theta(n \log n)$ whp. while utilising a constant number of extra states (not mixed with other protocols including clocks). The protocol is preceded by leader election confirmed by the matching based clock. And when this happens, the periodic leader based clock starts running together with the following line formation protocol based on two main rules defined below.

Form head and tail

$$L + F + 0 \rightarrow H + T + 1$$

This rule creates the initial head in state H and the tail in state T of the newly formed line. Note that since the line formation process uses separate memory the leader in the leader based clock remains in the leadership state, i.e., it is the head state H is used solely in the line formation protocol.

Extend the line

$$H + F + 0 \rightarrow R + H + 1$$

This rule extends the current line by addition of an extra agent from the head end of the line.

Theorem 2. *The line formation protocol operates in time $O(n \log n)$ whp.*

Proof. The probability of an interaction adding agent $i + 1$ to the line when i agents are already present is $\frac{2(n-i)}{n(n-1)}$. So the number of interactions to add agent $i+1$ has geometric distribution with the expected value $\frac{n(n-1)}{2(n-i)}$. Thus the expected time of forming the line is

$$\frac{1}{n} \sum_{i=1}^{n} \frac{n(n-1)}{2(n-i)} \sim \frac{n}{2} \sum_{i=1}^{n} \frac{1}{i} \sim \frac{n \ln n}{2}.$$

By Chernoff-Janson bound this time is $O(n \log n)$ whp. $\square$

In order to make the line formation protocol always correct we need some backup rules for the unlikely case of desynchronisation when two or more leaders survive to the line formation stage. In such case we need to continue leader elimination.

$$L + L + 0 \; \longrightarrow \; L + F + 0$$

Also when a leader meets already formed head.

$$L + H + 0 \rightarrow F + H + 0$$

Finally we have to dismantle excessive lines if two or more lines are formed. This is done using extra state D which dismantles the line edge by edge starting from the head.

$$H + H + 0 \rightarrow H + D + 0$$
$$D + R + 1 \rightarrow F + D + 0$$
$$D + T + 1 \rightarrow F + F + 0$$

11

5 Fast formation of rings

Ring Formation After the line formation is completed we need one extra round controlled by the periodic leader based clock to close the ring.

Close the ring

$$H + T + 0 \rightarrow H_R + T_R + 1$$

Theorem 3. *The ring formation protocol operates in time $O(n \log n)$ whp.*

Proof. In the first round (clocked by the periodic leader based clock) the ring formation protocol executes a line is formed and in the second it the head of the line connects with the tail, where both rounds operate in time $\Theta(n \log n)$. $\square$

In order to make the ring formation protocol always correct we need extra (on the top of line formation) backup rules for the unlikely case of desynchronisation when two or more leaders survive, and in turn two or more lines (possibly already closed in rings or cross-connected) are formed. And indeed when two heads closed in two different rings meet one of accepts dismantling role D_R.

$$H_R + H_R + 0 \rightarrow H_R + D_R + 0$$

And when it meets the tail in its ring forms a line to be dismantled.

$$D_R + T_R + 1 \rightarrow D + T + 0$$

When the head on a ring meets the head of a line, the ring head goes to state H_D meaning that the ring has to be disconnected, and the line head starts dismantling the line.

$$H_R + H + 0 \rightarrow H_D + D + 0$$

When the head on a ring meets a leader or a free agent, the ring head goes to state H_D as before, and the other agent becomes or stays free.

$$H_R + L|F + 0 \rightarrow H_D + F + 0$$

Finally, When the head on a ring in state H_D meets the tail on the same ring the ring is replaced by the line still open to accept new nodes until the end of the round clocked by the periodic leader based clock.

$$H_D + T_R + 1 \rightarrow H + T + 0$$

6 Line replication

In this section we consider a *raw line replication mechanism* allowing multiple replication of one or more lines of agents. When the replication process starts a small number of agents belong to some relatively short lines (sequences) of agents. Each such line has *the head* (the first agent on the line) and *the tail* (the last agent on the line). We refer to all agents between the head and the tail of

a line as regular agents. We also assume that each agent on a line carries one bit (0 or 1) of information, so in fact one can interpret each line as a binary sequence which carries one or more messages. In our presentation we focus on replication of a single line, but our protocol can be simultaneously applied to many lines of the same or different content.

The agents utilise a constant number of states organised in triplets

$$< Role, B_i, Buffer >,$$

where

Role is either *the head* of the line H, *the tail* of the line T, or *a regular* internal agent R.

B_i refers to combined notation B_i for the bit of information stored at position i on the line. Please note that this location is computed modulo 3 counting from the head's position 0. The second agent has position 1, the third 2, the forth 0 etc. Note that while this enumeration is limited it allows an agent carrying B_i distinguish between the two neighbours on the line carrying B_{i-1} (on the way to H) and B_{l+1} (towards T). Finally, by $|B_i|$ we denote the sole value of the bit without its location, i.e., $B_i =< |B_i|, i >$.

Buffer is used to carry either bit or control messages. When an agent is not currently involved in any action its buffer is in the neutral state ϕ.

In the *replicated (old) line*, when the buffer is empty but the agent support transfer towards H the buffer is in state ϕ^H. Similarly, when the buffer is occupied by a bit $|B_x|$ moving towards H the relevant value is $|B_x|^H$.

In the *newly formed (new) line*, when the buffer is empty but the agent support transfer towards T the buffer is in state ψ^T. Similarly, when the buffer is occupied by a bit $|B_x|$ moving towards T the relevant value is $|B_x|^T$. We also distinguish another value ψ^N which indicates that a new node is expected at the current tail of the line.

For example when a regular (internal) agent at position i is *neutral*, i.e., it is not currently involved in the replication process, it's state is set to $< R, B_i, \phi >$.

(R1) Start of the line replication The process begins when the head H of a line in the neutral state meets a free agent in state F. This interaction is governed by the following rule:

$$< H, B_0, \phi > + F + 0 \longrightarrow$$
$$< H, B_0, \phi^H > + < H, B_0, \psi > + 1$$

When this rule is applied, in the old line a signal ϕ^H (pipeline all bits towards the head) is created, and in the new line signal ψ means await further instructions, i.e., either to add a new agent or conclude the replication process.

In what follows, we first explain how the information (the sequence of bits) is

transferred from the old line onto the new one. We later discuss how the new line is being built simultaneously.

(R2) Create $|B_i|^H$ or $|B_T|^H$ message When signal ϕ^H arrives at the $(i-1)^{th}$ agent and the i^{th} agent is neutral, message $|B_i|^H$ is placed in the buffer of the latter.

$$< R, B_i, \phi > + < R|H, B_{i-1}, \phi^H > + 1 \longrightarrow$$
$$< R, B_i, |B_i|^H > + < R|H, B_{i-1} + \phi^H > + 1$$

A similar action is taken at the tail agent in neutral state $< T, B_T, \phi >$

$$< T, B_T, \phi > + < R|H, B_{i-1}, \phi^H > + 1 \longrightarrow$$
$$< T, B_T, |B_T|^H > + < R|H, B_{i-1}, \phi^H > + 1$$

The use of these two rules allows to propagate the request to pipeline all bit messages towards the head H. The next two rules explain how the bit messages are moved towards the head H.

(R3) Move a non-tail message $|B_x|^H$ towards H

$$< R, B_i, |B_x|^H > + < R|H, B_{i-1}, \phi^H > + 1 \longrightarrow$$
$$< R, B_i, \phi^H > + < R|H, B_{i-1}, |B_x|^H > +1$$

Note that when the bit message $|B_x|$ is moved the request ϕ^H for further bit messages remains in the i^{th} agent.

(R4) Move the tail message $|B_T|^H$ towards H

$$< T|R, B_i, |B_T|^H > + < R|H, B_{i-1}, \phi^H > + 1 \longrightarrow$$
$$< T|R, B_i, \phi > + < R|H, B_{i-1}, |B_T|^H > + 1$$

Note that when the tail message $|B_T|^H$ is moved the neutrality of the tail agent is restored. Eventually, thanks to the final transfer of the tail message all buffers in the old line are reset to ϕ. And the role of the old line in the replication process concludes when this message is moved to the new line.

The following two rules govern transfer of messages between the old and the new line.

(R5) Transfer a non-tail message $|B_x|^H$ to the head of the new line

$$< H, B_0, |B_x|^H > + < H, B_0, \psi^T > + 1 \longrightarrow$$
$$< H, B_0, \phi^H > + < H, B_0, |B_x|^T > + 1$$

Note that during such transfer the direction of the message is changed towards the tail of the new line.

(R6) Transfer the tail message $|B_T|^H$ to the head of the new line

$$< H, B_0, |B_T|^H > + < H, B_0, \psi^T > + 1 \ \longrightarrow$$
$$< H, B_0, \phi > + < H, B_0, |B_T|^T > + 0$$

As indicated earlier, due to the transfer of the tail message the neutrality of the new line is restored. In addition the two lines get disconnected and the old line is now ready to replicate again.

Finally, we show how the new line is constructed with the help of bit messages arriving from the old line. Recall that the buffer message ψ at the current end of the new line indicates that the line can be still extended.

(R7) Move a non-tail message $|B_x|^T$ towards (non-existent) tail T

$$< H|R, B_i, |B_x|^T > + < R, B_{i+1}, \psi^T > + 1 \ \longrightarrow$$
$$< H|R, B_i, \psi^T > + < R, B_{i+1}, |B_x|^T > + 1$$

After this move the buffer in the i^{th} agent expects further messages.

(R8) Move the tail message $|B_T|^T$ towards (non-existent) tail T

$$< H|R, B_i, |B_T|^T > + < R, B_{i+1}, \psi^T > + 1 \ \longrightarrow$$
$$< H|R, B_i, \phi > + < R, B_{i+1}, |B_x|^T > + 1 >$$

In this case the neutrality of the i^{th} agent is restored.

When there is no room for a message coming from the head of the new line yet another agent has to be added. This is done in two stages. We first request addition of a new agent using message ψ^N.

(R9) Request extension via message ψ^N (non-tail message $|B_x|^T$)

$$< H|R, B_i, |B_x|^T > + < R, B_{i+1}, \psi > + 1 \ \longrightarrow$$
$$< H|R, B_i, |B_x|^T > + < R, B_{i+1}, \psi^N > + 1$$

and a specific rule requesting extension beyond the head of the new line.

$$< H, B_0, |B_1|^H > + < H|R, B_0, \psi > + 1 \ \longrightarrow$$
$$< H, B_i, |B_1|^H > + < H, B_0, \psi^N > + 1$$

And when this message is already present the new agent is added from the pool of free agents.

(R10) Extend the new line

$$< H|R, B_i, \psi^N > + F \ + \ 0 \ \longrightarrow$$
$$< H|R, B_i, \psi^T > + < R, *, \psi > + 1$$

Note that after this rule is applied the newly added agent still awaits its bit message which is denoted by $*$. The new bit message arrives with the help of the following two rules.

(R11) Arrival of a non-tail bit message

$$< R, B_i, |B_x|^T > + < R, *, \psi > + 1 \ \longrightarrow$$
$$< R, B_i, \psi^T > + < R, B_x, \psi > + 1$$

As a non-tail bit arrived the line will be still extended which is denoted by messages ψ^T (expect more bit messages) in the i^{th} agent and ψ (expect further extension). The situation is different when the tail bit message arrives.

(R12) Arrival of the tail bit message

$$< R, B_i, |B_T|^T > + < R, *, \psi > + 1 \ \longrightarrow$$
$$< R, B_i, \phi > + < T, B_T, \phi > + 1$$

After this rule is applied the neutrality at the tail end of the new line is restored. Note, however, that since the neutrality of the agents closer to the head of this line was restored earlier the front of the new line can be already involved in the next line replication process. But since we use different messages for the transfers in the old and the new lines, the two simultaneously run processes will not interrupt one another.

We conclude with the following theorem.

Theorem 4. *The population protocol based on rules* **R1-R12** *is a correct line replication protocol.*

Proof. We argue first about correction of the replication protocol for the old line, where

- The actions of the tail node are governed by rules **R2** and **R4**. The first rules creates the bit message $|B_T|^H$ and the second moves this message towards the head of the line restoring the neutrality of the tail agent T.

- The actions of a regular node require also rule **R3** which support movement of multiple non-tail bit messages towards H.

- The actions of the head H are more complex. The process begins with application of rule **R1** which encapsulates three different actions: adding the head of the new line, replication of its bit in the newly created head. Transfer of non-tail bit messages to the new line is managed by rules **R3** and **R5** and the tail message by rules **R4** and **R6**, when after the application of the latter the old line concludes the replication process.

For the full cycles of rules used by agents see Figure 1.

The new line formation requires different organisation of states and transitions. Note that all agents introduced to the line must originate from the state F, see Figure 2.
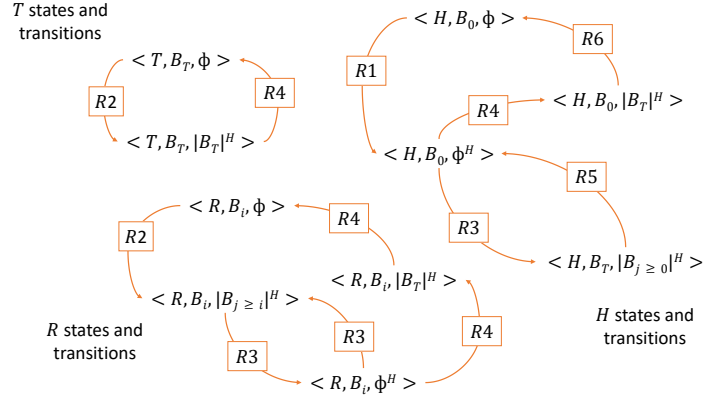
Figure 1: The old line states and transitions

- The formation of the tail agent requires only two rules: **R10** to add new agent and **R12** to equip this agents with the bit B_T.

- The situation with the regular nodes is more complex as they have to accept their own bit B_i (rule **R11**) add additional agent (rules **R9** and **R10**) and keep pipelining non-tail bit messages (rule **R7**) until the tail bit message arrives (rule **R8**) and finally establish the neutrality of the agent (rule **R8** or rule **R12** when neighbour of the tail agent).

- Rule **R1** creates the head of the new line, rules **R9** and **R10** add a new agent, rules **R5** and **R7** pipeline non-tail bit messages in the direction of the tail node until the tail bit message arrives (rule **R6**) when the neutrality of the tail node is reached (rule **R8**).

$\square$

With respect to the time complexity it is relatively easy to observe that the proposed replication protocol requires at most $O(k^2 \cdot n)$ parallel time as transfer of each bit message requires $O(n \cdot k)$ parallel time in expectation. As all transfers are done (pipelined) simultaneously we would like to claim that this process concludes faster in time $O(n \cdot k)$ in expectation. However we can only prove the bound $O(kn \log n)$ for which we need the following theorem.

Theorem 5. *The aforementioned raw line replication protocol can be amended to operate in parallel time $O(n \cdot k \log n)$ whp.*

Proof. We can amend the states of the agents such that two bit messages: the *main* and the *newly arrived*, can be stored in the buffer simultaneously. Note that this does not violate the assumed constant space limit. The bit messages pipelining process works in synchronised rounds where each round lasts parallel
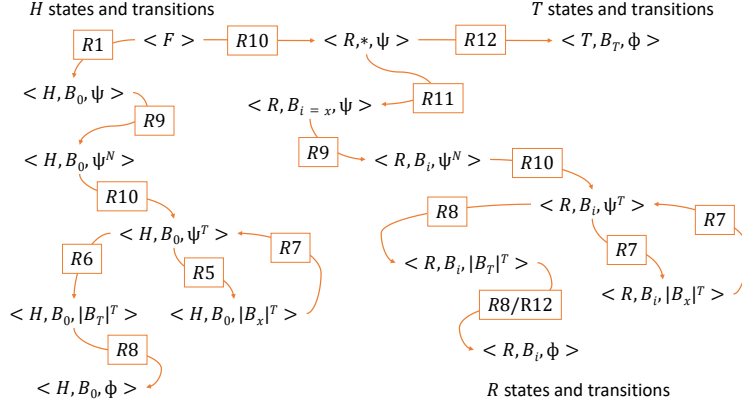
17

Figure 2: The new line states and transitions

time $\Theta(n \log n)$ paced by the periodic leader based clock, see Section 2. The pace of this clock is set to accommodate single addition of an agent to the new line and a constant number of interactions along each edge of the line.

A round begins when the *first stage* of the clock is initiated. During this stage a new agent is added towards the end of the new line, if such a need arises. During the *second stage* the main bit message currently present in the buffer of an agent moves to the neighbour agent (where it resides as newly arrived bit message) on the line towards the destination (H in the old line or the current tail in the new one). In the last *third stage* the newly arrived bit message (if any) is rebranded as the main bit message to make it ready for further transfer in the next round. Note that since all bit messages had enough time $O(n \log n)$ (coupon collector problem argument) to move one position closer to their destination such rebranding is safe. Finally, after $O(k)$ rounds (the tail bit message has to move across $2k - 1$ edges) all bit messages originating in the old line reach their destination, which translates to the parallel time $O(kn \log n)$ with high probability. □

One can further amend this protocol to obtain *fully amended* replication protocol with buffers of size 1. In this protocol an agent restrains its actions until two rounds after its neighbour (towards the destination) released its bit message. And when this happens the message originating in this agent moves along the line (towards its destination) by one position in each round. Since all moving bit messages are always at distance two at the end of each round no buffer conflicts are observed. And since the last bit message starts moving after $O(k)$ rounds and it moves to destination in $O(k)$ rounds, the total time complexity of fully amended replication protocol is $O(kn \log n)$ whp.

Theorem 6. *The fully amended line replication protocol operates whp in time* $O(n \cdot k \log n)$.

We conclude with the following corollary.

Corollary 1. *The raw line replication protocol operates in time $O(kn \log n)$ whp.*

Proof. We say that configuration A *dominates* another configuration B (denoted $A \geq B$) iff configuration A can be obtained from B through a finite number of steps (rule applications) of the raw replication protocol.

We first recall that fully amended replication protocol requires $O(k)$ synchronised rounds, see the proof of Theorem 6. Note also that application of a single round preserves domination relationship between arbitrary configurations. I.e., if two configurations $A \geq B$ enter a round, the respective configurations on the conclusion of the round A' and B' also satisfy $A' \geq B'$. We also observed earlier that during each round of the (fully) amended protocol each message moves one position closer to the destination. Thus since the configurations in the raw replication protocol are always dominating the configurations of the fully amended one, and the time complexity of the latter is $O(kn \log n)$, we conclude that the time complexity of the raw protocol is also $O(kn \log n)$. $\square$

References

[1] D. Alistarh, J. Aspnes, D. Eisenstat, R. Gelashvili, and R.L. Rivest. Time-space trade-offs in population protocols. In *Proc. SODA 2017*, pages 2560–2579.

[2] D. Alistarh, J. Aspnes, and R. Gelashvili. Space-optimal majority in population protocols. In *Proc. SODA 2018*, pages 2221–2239.

[3] D. Alistarh and R. Gelashvili. Polylogarithmic-time leader election in population protocols. In *Proc. ICALP 2015*, pages 479–491.

[4] D. Angluin, J. Aspnes, Z. Diamadi, M.J. Fischer, and R. Peralta. Computation in networks of passively mobile finite-state sensors. In *Proc. PODC 2004*, pages 290–299.

[5] D. Angluin, J. Aspnes, and D. Eisenstat. Fast computation by population protocols with a leader. *Distributed Comput.*, 21(3):183–199, 2008.

[6] P. Berenbrink, G. Giakkoupis, and P. Kling. Optimal time and space leader election in population protocols. In *Proc. STOC 2020*, pages 119–129.

[7] P. Berenbrink, D. Kaaser, P. Kling, and L. Otterbach. Simple and efficient leader election. In *Proc. SOSA 2018*, volume 61 of *OASICS*, pages 9:1–9:11.

[8] A. Bilke, C. Cooper, R. Elsässer, and T. Radzik. Brief announcement: Population protocols for leader election and exact majority with $O(\log^2 n)$ states and $O(\log^2 n)$ convergence time. In *Proc. PODC 2017*, pages 451–453.

[9] I. Chatzigiannakis, O. Michail, S. Nikolaou, A. Pavlogiannis, and P.G. Spirakis. Passively mobile communicating machines that use restricted space. *Theor. Comput. Sci.*, 412(46):6469–6483, 2011.

[10] H-L Chen, R. Cummings, D. Doty, and D. Soloveichik. Speed faults in computation by chemical reaction networks. In *Proc. DISC 2014*, pages 16–30.

[11] M. Connor, O. Michail, and P. Spirakis. On the distributed construction of stable networks in polylogarithmic parallel time. *Information*, 12(6):254–266, 2021.

[12] D. Doty. Timing in chemical reaction networks. In *Proc. SODA 2014*, pages 772–784.

[13] D. Doty, M. Eftekhari, L. Gąsieniec, E.E. Severson, G. Stachowiak, and P. Uznański. A time and space optimal stable population protocol solving exact majority. *CoRR*, abs/2106.10201, 2021, to appear at FOCS 2021.

[14] D. Doty and D. Soloveichik. Stable leader election in population protocols requires linear time. In *Proc. DISC 2015*, pages 602–616.

[15] L. Gąsieniec and G. Stachowiak. Enhanced phase clocks, population protocols, and fast space optimal leader election. *J. ACM*, 68(1):2:1–2:21, 2021.

[16] L. Gąsieniec, G. Stachowiak, and P. Uznański. Almost logarithmic-time space optimal leader election in population protocols. In *Proc. SPAA 2019*, pages 93–102.

[17] Svante Janson. Tail bounds for sums of geometric and exponential variables. *Satistics and Probability Letters*, 135(1):1–6, 2018.

[18] O. Michail and P. Spirakis. Simple and efficient local codes for distributed stable network construction. *Distributed Computing*, 29(3):207–237, 2016.