

Stone Age Distributed Computing

Yuval Emek*

Roger Wattenhofer†

Abstract

A new model that depicts a network of randomized finite state machines operating in an asynchronous environment is introduced. This model, that can be viewed as a hybrid of the message passing model and cellular automata is suitable for applying the distributed computing lens to the study of networks of sub-microprocessor devices, e.g., biological cellular networks and man-made nano-networks. Although the computation and communication capabilities of each individual device in the new model are, by design, much weaker than those of an abstract computer, we show that some of the most important and extensively studied distributed computing problems can still be solved efficiently.

Keywords: finite state machines, efficient algorithms, message passing, cellular automata.

1 Introduction

Due to the major role that the Internet plays today, models targeted at understanding the fundamental properties of networks focus mainly on “Internet-capable” devices. Indeed, the standard network model in distributed computing is the so called *message passing* model, where nodes may exchange large messages with their neighbors, and perform arbitrary local computations. Recently, there is a trend to study distributed computing aspects in networks of sub-microprocessor devices, e.g., networks of biological cells [3, 18] or nano-scale mechanical devices [4]. However, the suitability of the message passing model to these types of networks is far from being certain: Do tiny bio/nano nodes “compute” and/or “communicate” essentially the same as a computer? Since such nodes will be fundamentally more limited than silicon-based devices, we believe that there is a need for a network model, where nodes are by design below the computation and communication capabilities of Turing machines.

*ETH Zurich, Switzerland. E-mail: yemek@ethz.ch.

†Microsoft Research, Redmond, WA and ETH Zurich, Switzerland. E-mail: wattenhofer@ethz.ch.

Networked finite state machines. In this paper, we introduce a new model, referred to as *networked finite state machines (nFSM)*, that depicts a network of randomized finite state machines (a.k.a. automata) progressing in asynchronous steps (refer to Section 2 for a formal description). Under the nFSM model, nodes communicate by transmitting messages belonging to some finite communication alphabet Σ such that a message $\sigma \in \Sigma$ transmitted by node u is delivered to its neighbors (the same σ to all neighbors) in an asynchronous fashion. Each neighbor v of u has a port corresponding to u in which the last message delivered from u is stored.

The access of node v to its ports is limited: In each step of v 's execution, the next state and the message transmitted by v at this step are determined by v 's current state and by the current number $\sharp(\sigma)$ of appearances of each letter $\sigma \in \Sigma$ in v 's ports. The crux of the model is that $\sharp(\sigma)$ is calculated according to the *one-two-many*¹ principle: the node can only count up to some predetermined *bounding parameter* $b \in \mathbb{Z}_{>0}$ and any value of $\sharp(\sigma)$ larger than b cannot be distinguished from b .

The nFSM model satisfies the following *model requirements*, that we believe, make it more applicable to the study of general networks consisting of weaker devices such as those mentioned above:

- (M1) The model is applicable to arbitrary network topologies.
- (M2) All nodes run the same protocol executed by a (randomized) FSM.
- (M3) The network operates in a fully asynchronous environment with adversarial node activation and message delivery delays.
- (M4) All features of the FSM (specifically, the state set Q , message alphabet Σ , and bounding parameter b) are of constant size independent of any parameter of the network (including the degree of the node executing the FSM).

The last requirement is perhaps the most interesting one as it implies that a node cannot perform any calculation that involves variables beyond some predetermined constant. This comes in contrast to many distributed algorithms operating under the message passing model that strongly rely on the ability of a node to perform such calculations (e.g., count up to some parameter of the network or a function thereof).

Results. Our investigation of the new nFSM model begins by observing that the computational power of a network operating under this model is not stronger than (and in some sense equivalent to) that of a randomized Turing machine with linear space bound (cf. linear bounded automaton). Since the computational power of a network operating under the message passing model is trivially equivalent to that of a (general) Turing machine, there exist distributed problems that can be solved by a message passing algorithm in constant time but cannot be solved by an nFSM algorithm

¹ The one-two-many theory states that some small isolated cultures (e.g., the Piraha tribe of the Amazon [29]) did not develop a counting system that goes beyond 2. This is reflected in their languages that include words for “1”, “2”, and “many” that stands for any number larger than 2.

at all. Nevertheless, as the main technical contribution of this paper, we show that some of the most important problems in distributed computing admit efficient — namely, with run-time polylogarithmic in the number of nodes — nFSM algorithms. Specifically, problems such as maximal independent set, node coloring, and maximal matching that have been extensively studied since the early 1980s, always assuming a network of some sort of abstract computers, can in fact be solved, and fast, when each device is nothing but a FSM.

Applicability to biological cellular networks. Regardless of the theoretical interest in implementing efficient algorithms using weaker assumptions, we believe that our new model and results should be appealing to anyone interested in understanding the computational aspects of biological cellular networks. A basic dogma in biology (see, e.g., [45]) states that all cells communicate and that they do so by emitting special kinds of ligand molecules (e.g. the DSL family of proteins) that bind in a reversible, non-covalent, fashion to designated receptors (e.g. the NOTCH family of transmembrane receptors) in neighboring cells. These, in turn, release some intracellular signaling proteins that penetrate the nucleus to modify gene expression (determining the cell’s actions).

Translated to the language of the nFSM model, the different types of ligand molecules correspond to the different letters in the communication alphabet, where an emission of a ligand corresponds to transmitting a letter. The effect that the ligands-receptors binding has on the concentration level of the signaling proteins in the nucleus corresponds to the manner in which a node in our model interprets the content of its ports. (The one-two-many counting is the discrete analogue of the ability to distinguish between different concentration levels, considering the fact that once the concentration exceeds some threshold, a further increase cannot be detected.) Using an FSM as the underlying computational model of the individual node seems to be the right choice especially in the biological setting as demonstrated by Benenson et al. [19] who showed that essentially all FSMs can be implemented by enzymes found in cells’ nuclei. One may wonder if the specific problems studied in the current paper have any relevance to biology. Indeed, Afek et al. [3] discovered that a biological process that occurs during the development of the nervous system of the *Drosophila melanogaster* is in fact equivalent to solving the MIS problem.

Related work and comparison to other models. As mentioned above, the message passing model is the gold standard when it comes to understanding distributed network algorithms. Several variants exist for this model, differing mainly in the bounds imposed on the message size and the level of synchronization. Perhaps the most popular message passing variants are the fully synchronous *local* and *congest* models [36, 44, 48], assuming that in each round, a node can send messages to its neighbors (different messages to different neighbors), receive and interpret the messages sent to it from its neighbors, and perform an arbitrary local computation determining, in particular, the messages sent in the next round. The difference between the two variants is cast in the size of the communicated messages: the local model does not impose any restrictions on the

message size, hence it can be used for the purpose of establishing general lower bounds, whereas the congest model is more information-theoretic, with a (typically logarithmic) bound on the message size. Indeed, most theoretical literature dealing with distributed network algorithms relies on one of these two models. Our nFSM model adopts the concept of an asynchronous message-based communication scheme from the message passing literature.

As the traditional message passing model allows for sending different messages to different neighbors in each round of the execution, it was too powerful for many settings. In particular, with the proliferation of wireless networks, more restrictive message passing models appeared such as the *radio network* model [22] that was the first model to take interference considerations into account. Over the years, several variants of the radio network model were introduced, the most extreme one in terms of its weak communication capabilities is the *beeping* model [27, 25], where in each round a node can either beep or stay silent, and can only distinguish between the case in which no node in its neighborhood beeps and the case in which at least one node beeps. Efficient algorithms and lower bounds for the MIS problem under the beeping model were developed by Afek et al. [3, 2]. Note that the beeping model resembles our nFSM model in the sense that the “beeping rule” can be viewed as counting under the one-two-many principle with bounding parameter $b = 1$. However, it is much stronger in other perspectives: (i) the beeping model assumes synchronous communication and does not seem to have a natural asynchronous variant; and (ii) the local computation is performed by a Turing machine whose memory is allowed to grow with time and with the network size (this is crucial for the algorithms of Afek et al. [3, 2]). In that regard, the beeping model is still too strong to capture the behavior of biological cellular networks.

Our nFSM model is also closely related to (and inspired by) the extensively studied *cellular automaton* model [51, 28, 52] that captures a network of FSMs, arranged in a grid topology (some other highly regular topologies were also considered), where the transition of each node depends on its current state and the states of its neighbors. Still, the nFSM model differs from the cellular automaton model in many aspects. In particular, the latter model is not applicable to non-regular network topologies and although a small fraction of the cellular automata literature is dedicated to cellular automata with asynchronous node activation [42, 43] (see also [1]), these do not support asynchronous message delivery. As such, cellular automata do not seem to provide a good abstraction for the sub-microprocessor networks we would like to focus on. Moreover, the main goal of our work is to study to what extent can such networks compute solutions quickly, a goal that lies outside the typical interest of the cellular automata community.

Another model that resembles the nFSM model is that of *communicating automata* [20]. This model also assumes that each node in the network operates an FSM in an asynchronous manner, however the steps of the FSMs are message driven: for each state q of node v and for each message m that node v may receive from an adjacent node u while residing in state q , the transition function of v should have an entry characterized by the 3-tuple (q, u, m) that determines its next move. Consequently, different nodes would typically operate different FSMs and the size of the FSM

operated by node v inherently depends on the degree of v and thus, the communicating automata model is still too strong to faithfully represent biological cellular networks.

The *population protocols* model, introduced by Angluin et al. [8] (see also [10, 41]), depicts a network of finite state machines communicating through pairwise rendezvous controlled by a fair adversarial scheduler. While in Section 4.3 we reveal some interesting connections between the nFSM model and population protocols, there are two conceptual differences between these models: First, population protocols are only required to *eventually converge* to a correct output and are allowed to return arbitrary (wrong) outputs beforehand. This provides population protocols with the power to solve, e.g., the consensus problem in arbitrary networks, in contrast to nFSM protocols that should irrevocably return a correct output (see Section 2) and therefore, cannot solve this problem (cf. [26]). Second, while the focus of the current paper is mainly on run-time complexity, there is an inherent problem with establishing run-time bounds for population protocols due to the nature of the adversarial scheduler that can delay the execution for arbitrarily long periods. Indeed, the population protocols literature is typically concerned with what can be computed, rather than how fast. An exception is the probabilistic variant of the model [9, 21], where interactions are selected randomly, rather than adversarially, but this variant is no longer fully asynchronous (in the adversarial sense). On top of that, the rendezvous based communication does not seem to fit the communication mechanisms in most biological cellular networks.

2 Model

Throughout, we assume a network represented by a finite undirected graph $G = (V, E)$. Under the *networked finite state machines (nFSM)* model, each node $v \in V$ runs a protocol depicted by the 8-tuple

$$\Pi = \langle Q, Q_I, Q_O, \Sigma, \sigma_0, b, \lambda, \delta \rangle,$$

where

- Q is a finite set of *states*;
- $Q_I \subseteq Q$ is the subset of *input states*;
- $Q_O \subseteq Q$ is the subset of *output states*;
- Σ is a finite *communication alphabet*;
- $\sigma_0 \in \Sigma$ is the *initial letter*;
- $b \in \mathbb{Z}_{>0}$ is a *bounding parameter*; let $B = \{0, 1, \dots, b-1, \geq b\}$ be a set of $b+1$ distinguishable symbols;
- $\lambda : Q \rightarrow \Sigma$ assigns a *query letter* $\sigma \in \Sigma$ to every state $q \in Q$; and
- $\delta : Q \times B \rightarrow 2^{Q \times (\Sigma \cup \{\varepsilon\})}$ is the *transition function*.

It is important to point out that protocol Π is oblivious to the graph G . In fact, the number of states in Q , the size of the alphabet Σ , and the bounding parameter b are all assumed to be universal constants, independent of any parameter of the graph G . In particular, the protocol executed by node $v \in V$ does not depend on the degree of v in G . We now turn to describe the semantics of the nFSM model.

Communication. Node v communicates with its adjacent nodes in G by *transmitting* messages. A transmitted message consists of a single letter $\sigma \in \Sigma$ and it is assumed that this letter is delivered to all neighbors u of v . Each neighbor u has a *port* $\psi_u(v)$ (a different port for every adjacent node v) in which the last message σ received from v is stored. At the beginning of the execution, all ports store the initial letter σ_0 . It will be convenient to consider the case in which v does not transmit any message (and hence does not affect the corresponding ports at the adjacent nodes) as a transmission of the special *empty symbol* ε .

Execution. The execution of node v progresses in discrete *steps* indexed by the positive integers. In each step $t \in \mathbb{Z}_{>0}$, node v resides in some state $q \in Q$. Let $\lambda(q) = \sigma \in \Sigma$ be the query letter that λ assigns to state q and let $\sharp(\sigma)$ be the number of appearances of σ in v 's ports in step t . Then, the pair (q', σ') of state $q' \in Q$ in which v resides in step $t+1$ and message $\sigma' \in \Sigma \cup \{\varepsilon\}$ transmitted by v in step t (recall that ε indicates that no message is transmitted) is chosen *uniformly at random*² (and independently of all other random choices) among the pairs in

$$\delta(q, \beta_b(\sharp(\sigma))) \subseteq Q \times (\Sigma \cup \{\varepsilon\}) ,$$

where $\beta_b : \mathbb{Z}_{\geq 0} \rightarrow B$ is defined as

$$\beta_b(x) = \begin{cases} x & \text{if } 0 \leq x \leq b-1 ; \\ \geq b & \text{otherwise .} \end{cases}$$

Informally, this can be thought of as if v queries its ports for appearances of σ and “observes” the exact value of $\sharp(\sigma)$ as long as it is smaller than the bounding parameter b ; otherwise, v merely “observes” that $\sharp(\sigma) \geq b$ which is indicated by the symbol $\geq b$.

Input and output. Initially (in step 1), each node resides in one of the input states of Q_I . The choice of the initial state of node $v \in V$ reflects the input passed to v at the beginning of the execution. This allows our model to cope with distributed problems in which different nodes get different input symbols. When dealing with problems in which the nodes do not get any initial input (such as the graph theoretic problems addressed in this paper), we shall assume that Q_I contains a single state referred to as the *initial* state.

² The protocol is deterministic if the images under δ are always singleton subsets of $Q \times (\Sigma \cup \{\varepsilon\})$.

The output states Q_O are mapped to the possible output values of the problem. For each possible output value o , it is required that the subset $P_o \subseteq Q_O$ of output states mapped to o form a *sink* with respect to the transition function δ in the sense that a node v that moves to a P_o -state will remain in P_o indefinitely, in which case the output of v is determined (irrevocably) to be o . We say that the (global) execution of the protocol is in an *output configuration* if all nodes reside in output states of Q_O .

Asynchrony. The nodes are assumed to operate in an *asynchronous* environment. This asynchrony has two facets: First, for the sake of convenience, we assume that the actual application of the transition function in each step $t \in \mathbb{Z}_{>0}$ of node $v \in V$ is instantaneous (namely, lasts zero time) and occurs at the end of the step;³ the length of step t of node v , denoted $L_{v,t}$, is defined as the time difference between the application of the transition function in step $t - 1$ and that of step t . It is assumed that $L_{v,t}$ is finite, but apart from that, we do not make any other assumptions on this length, that is, the step length $L_{v,t}$ is determined by the adversary independently of all other step lengths $L_{v',t'}$. In particular, we do not assume any synchronization between the steps of different nodes whatsoever.

The second facet of the asynchronous environment is that a message transmitted by node v in step t (if such a message is transmitted) is assumed to reach the port $\psi_u(v)$ of an adjacent node u after a finite time delay, denoted $D_{v,t,u}$. We assume that if v transmits message $\sigma_1 \in \Sigma$ in step t_1 and message $\sigma_2 \in \Sigma$ in step $t_2 > t_1$, then σ_1 reaches u before σ_2 does. Apart from this “FIFO” assumption, we do not make any other assumptions on the delays $D_{v,t,u}$. In particular, this means that under certain circumstances, the adversary may overwrite message σ_1 with message σ_2 in port $\psi_u(v)$ of u so that u will never “know” that message σ_1 was transmitted.⁴

Consequently, a *policy* of the adversary is captured by: (1) the length $L_{v,t}$ of step t of node v for every $v \in V$ and $t \in \mathbb{Z}_{>0}$; and (2) the delay $D_{v,t,u}$ of the delivery of the transmission of node v in step t to an adjacent node u for every $v \in V$, $t \in \mathbb{Z}_{>0}$, and $u \in N(v)$.⁵ Assuming that the adversary is oblivious to the random coin tosses of the nodes, an adversarial policy is depicted by infinite sequences of $L_{v,t}$ and $D_{v,t,u}$ parameters.

For further information on asynchronous environments, we refer the reader to one of the standard textbooks [38, 44].

³ This assumption can be lifted at the cost of a more complicated definition of the adversarial policy described soon.

⁴ Often, much stronger assumptions are made in the literature. For example, a common assumption for asynchronous environments is that the port of node u corresponding to the adjacent node v is implemented by a buffer so that messages cannot be “lost”. We do not make any such assumption for our nFSM model.

⁵ We use the standard notation $N(v)$ for the *neighborhood* of node v in G , namely, the subset of nodes adjacent to v .

Correctness and run-time measures. A protocol Π for problem P is said to be *correct* under the nFSM model if for every instance of P and for every adversarial policy, Π reaches an output configuration w.p. 1, and for every output configuration reached by Π w.p. > 0 , the output of the nodes is a valid solution to P .⁶ Given a correct protocol Π , the complexity measure that interests us in the current paper is the *run-time* of Π for which we use the following standard definition (cf. [11, 44]).

Consider some instance \mathcal{I} of problem P . Given an adversarial policy \mathcal{A} and a sequence (actually an n -tuple of sequences) \mathcal{R} of random coin tosses that lead to an output configuration within finite time, the run-time $T_{\Pi}(\mathcal{I}, \mathcal{A}, \mathcal{R})$ of Π on \mathcal{I} with respect to \mathcal{A} and \mathcal{R} is defined as the (possibly fractional) number of *time units* that pass from the beginning of the execution until an output configuration is reached, where a time unit is defined⁷ to be the maximum among all step length parameters $L_{v,t}$ and delivery delay parameters $D_{v,t,u}$ appearing in \mathcal{A} before the output configuration is reached. Let $T_{\Pi}(\mathcal{I}, \mathcal{A})$ denote the random variable that depicts the run-time of Π on \mathcal{I} with respect to \mathcal{A} . Following the standard procedure in this regard, we say that the run-time of a correct protocol Π for problem P is $f(n)$ if for every n -node instance \mathcal{I} of P and for every adversarial policy \mathcal{A} , it holds that $T_{\Pi}(\mathcal{I}, \mathcal{A})$ is at most $f(n)$ in expectation and w.h.p. The protocol is said to be *efficient* if its run-time is polylogarithmic in the size of the network (cf. [36]).

3 Convenient transformations

In this section, we show that the nFSM protocol designer may, in fact, assume a slightly more “user-friendly” environment than the one described in Section 2. This is based on the design of black-box *compilers* transforming a protocol that makes strong assumptions on the environment into one that does not make any such assumptions. Specifically, the assumptions that can be lifted in that way are synchrony (Section 3.1), and multiple-letter queries (Section 3.2).

3.1 Implementing a synchronizer

As described in Section 2, the nFSM model assumes an asynchronous environment. Nevertheless, it will be convenient to extend the nFSM model to *synchronous* environments. One natural such extension augments the model described in Section 2 with the following two *synchronization properties* that should hold for every two adjacent nodes $u, v \in V$ and for every $t \in \mathbb{Z}_{>0}$:

(S1) when node u is in step t , node v is in step $t - 1$, t , or $t + 1$; and

(S2) at the end of step $t + 1$ of u , port $\psi_u(v)$ stores the message transmitted by v in step t of v ’s execution (or the last message transmitted by v prior to step t if v does not transmit any message

⁶ Throughout, w.p. and w.h.p. abbreviate “with probability” and “with high probability”, respectively.

⁷ Note that time units are defined solely for the purpose of the analysis. Under an asynchronous environment, the nodes have no notion of time and in particular, they cannot measure a single time unit.

in step t).

An environment in which properties (S1) and (S2) are guaranteed to hold is called a *locally synchronous* environment. Local-only communication can never achieve global synchrony, however, research in the message passing model has shown that local synchrony is often sufficient to provide efficient algorithms [11, 13, 12]. To distinguish a protocol assumed to operate in a locally synchronous environment from those making no such assumptions, we shall often refer to the execution steps of the former as *rounds* (cf. fully synchronized protocols). Our goal in this section is to establish the following theorem.

Theorem 3.1. *Every nFSM protocol $\Pi = \langle Q, Q_I, Q_O, \Sigma, \sigma_0, b, \lambda, \delta \rangle$ designed to operate in a locally synchronous environment can be simulated in an asynchronous environment by a protocol $\widehat{\Pi}$ with the same bounding parameter b at the cost of a constant multiplicative run-time overhead.*

The procedure in charge of the simulation promised in Theorem 3.1 is referred to as a *synchronizer* [11]. The remainder of Section 3.1 is dedicated to the design (and analysis) of a synchronizer for the nFSM model.

Overview. Round $t \in \mathbb{Z}_{>0}$ of node $v \in V$ under Π is simulated by $O(1)$ contiguous steps under $\widehat{\Pi}$; the collection of these steps is referred to as v 's *simulation phase* of round t . Protocol $\widehat{\Pi}$ is designed so that v maintains the value of $t \bmod 3$, referred to as the *trit* (trinary digit) of round t ; this trit is also encoded in the message transmitted by v in the last step of the simulation phase, which is the only message v transmits throughout the simulation phase. The main principle behind our synchronizer is that node v will not move to the simulation phase of round $t + 1$ while its ports still contain messages sent in a round whose trit is $t - 2 \bmod 3$.

Under Π , the decisions made by node v in round t should be based on the messages transmitted by all neighbors u of v in round $t - 1$. However, during v 's simulation phase of round t , port $\psi_v(u)$ may contain messages transmitted in round $t - 1$ or in round t under Π . The latter case is problematic since the message transmitted by u in the simulation phase of round $t - 1$ is overwritten by that transmitted in the simulation phase of round t . To avoid this obstacle, a message transmitted by node u under $\widehat{\Pi}$ at the end of the simulation phase of round t also encodes the message that u transmitted under Π in round $t - 1$.

So, if v resides in a state whose query letter is $\sigma \in \Sigma$ in round t under Π , then under $\widehat{\Pi}$, v should query for all $\widehat{\Sigma}$ -letters encoding a transmission of σ in round $t - 1$. Since there are several such letters, a carefully designed feature should be used so that $\widehat{\Pi}$ accounts for their combined number.

Protocol $\widehat{\Pi}$. Let

$$\widehat{\Pi} = \langle \widehat{Q}, \widehat{Q}_I, \widehat{Q}_O, \widehat{\Sigma}, \widehat{\sigma}_0, b, \widehat{\lambda}, \widehat{\delta} \rangle .$$

Consider node $v \in V$ and round $t \in \mathbb{Z}_{>0}$. As the name implies, node v 's *simulation phase* of round t under $\widehat{\Pi}$, denoted $\phi_v(t)$, simulates round t of Π . Protocol $\widehat{\Pi}$ is designed so that at every step in

$\phi_v(t)$ other than the last one, v does not transmit any message (indicated by transmitting ε), and at the last step of the simulation phase, v always transmits some message $\hat{\sigma} \in \widehat{\Sigma}$, denoted $M_v(t)$.

The alphabet $\widehat{\Sigma}$ is defined to be

$$\Sigma' = (\Sigma \cup \{\varepsilon\}) \times (\Sigma \cup \{\varepsilon\}) \times \{0, 1, 2\} .$$

The semantics of the message $M_v(t) = (\sigma, \sigma', j)$ sent by node v at the last step of the simulation phase $\phi_v(t)$ is that: (a) v transmits $\sigma \in \Sigma \cup \{\varepsilon\}$ at round $t-1$ under Π ; (b) v transmits $\sigma' \in \Sigma \cup \{\varepsilon\}$ at round t under Π ; and (c) $j = t \bmod 3$. Following that logic, we set $\hat{\sigma}_0 = (\varepsilon, \sigma_0, 0)$.

The state set \widehat{Q} of $\widehat{\Pi}$ is defined to be

$$\widehat{Q} = \left(\bigcup_{q \in Q} (P_q \cup S_q) \right) \times \{0, 1, 2\} ,$$

where $P_q \times \{j\}$ and $S_q \times \{j\}$, $q \in Q$, $j \in \{0, 1, 2\}$, are referred to as the *pausing* and *simulating* features, respectively, whose roles will be clarified soon. Suppose that v resides in state $q \in Q$ in step t under Π and that $j = t \bmod 3$. Then, throughout $\phi_v(t)$, node v resides in some state in $(P_q \cup S_q) \times \{j\}$. In particular, in the first steps of the simulation phase, v resides in states of the pausing feature $P_q \times \{j\}$, and then at some stage it switches to the simulating feature $S_q \times \{j\}$ and remains in its states until the end of the current simulation phase.

The input states of $\widehat{\Pi}$ are defined as

$$\widehat{Q}_I = \{(\hat{p}_q, 1) \mid q \in Q_I\} ,$$

where $(\hat{p}_q, 1)$ is the first state in the pausing feature $P_q \times \{1\}$; the output states of $\widehat{\Pi}$ are defined as

$$\widehat{Q}_O = \left(\bigcup_{q \in Q_O} (P_q \cup S_q) \right) \times \{0, 1, 2\} .$$

The association of \widehat{Q}_I with the input values of the problem and the mapping from \widehat{Q}_O to the output values of the problem carry out in the natural manner from those of Q_I and Q_O .

The pausing feature. For the simulation phase of round t , we denote the letters in $(\Sigma \cup \{\varepsilon\}) \times (\Sigma \cup \{\varepsilon\}) \times \{j-2\}$ as *dirty* and the letters in $(\Sigma \cup \{\varepsilon\}) \times (\Sigma \cup \{\varepsilon\}) \times \{j-1, j\}$ as *clean*, recalling that $j = t \bmod 3$.⁸ The purpose of the pausing feature $P_q \times \{j\}$ is to pause the execution of v until its ports do not contain any dirty letter. This is carried out by including in $P_q \times \{j\}$ a state $p_{\sigma, \sigma'}$ for every $\sigma, \sigma' \in \Sigma \cup \{\varepsilon\}$; the query letter of $p_{\sigma, \sigma'}$ is (the dirty letter) $\hat{\lambda}(p_{\sigma, \sigma'}) = (\sigma, \sigma', j-2)$ and the transition function $\hat{\delta}$ is designed so that v moves to the next (according to some fixed order) state in the feature $P_q \times \{j\}$ if and only if there are no ports storing the query letter.

⁸ Throughout this section, arithmetic involving the parameter j is done modulo 3.

We argue that the pausing feature guarantees synchronization property (S1). For the sake of the analysis, it is convenient to assume the existence of a fully synchronous simulation phase of a virtual round 0; upon completion of this simulation phase (at the beginning of the execution), every node $v \in V$ transmits the message $M_v(0) = \widehat{\sigma}_0$. We are now ready to establish the following lemma.

Lemma 3.2. *For every $t \in \mathbb{Z}_{>0}$, $v \in V$, and $u \in N(v)$, when v completes the pausing feature of $\phi_v(t)$, port $\psi_v(u)$ stores either $M_u(t-1)$ or $M_u(t)$.*

Proof. By induction on t . The base case of round $t = 0$ holds by our assumption that $\phi_v(0)$ and $\phi_u(0)$ are fully synchronous. Assume by induction that the assertion holds for round $t - 1$. Applying the inductive hypothesis to both u and v , we conclude that (1) when v completes the pausing feature of $\phi_v(t-1)$, port $\psi_v(u)$ stores either $M_u(t-2)$ or $M_u(t-1)$; and (2) when u completes the pausing feature of $\phi_u(t-1)$, port $\psi_u(v)$ stores either $M_v(t-2)$ or $M_v(t-1)$.

Let τ_u and τ_v denote the times at which u and v complete the pausing feature of $\phi_u(t)$ and $\phi_v(t)$, respectively. Since v cannot complete the pausing feature of $\phi_v(t)$ while $M_u(t-2)$ is still stored in $\psi_v(u)$, it follows that at time τ_v , port $\psi_v(u)$ stores the message $M_u(t')$ for some $t' \geq t-1$. Our goal in the remainder of this proof is to show that $t' \leq t$. If $\tau_v < \tau_u$, then t' must be exactly $t-1$, which concludes the inductive step for that case.

So, assume that $\tau_v > \tau_u$ and suppose by contradiction that $t' \geq t+1$. Using the same line of arguments as in the previous paragraph, we conclude that at time τ_u , port $\psi_u(v)$ stores the message $M_v(t-1)$. Node u cannot complete the pausing feature of $\phi_u(t+1)$ while $M_v(t-1)$ is still stored in $\psi_u(v)$, hence v must have transmitted $M_v(t)$ before u completed the pausing feature of $\phi_u(t+1)$. But this means that v completed the pausing feature of $\phi_v(t)$ before u could have transmitted $M_u(t+1)$, in contradiction to the assumption that $\psi_v(u)$ stores $M_u(t')$ for some $t' \geq t+1$ at time τ_v . The assertion follows. \square

Consider two adjacent nodes $u, v \in V$. If node u is at round $t-1$ when an adjacent node v is at round $t+1$, then v completed the pausing feature of $\phi_v(t)$ before u transmitted $M_u(t-1)$, in contradiction to Lemma 3.2. Therefore, our synchronizer satisfies synchronization property (S1). Furthermore, a similar argument shows that between the time v completed the pausing feature of $\phi_v(t)$ and the time v completed the simulation phase $\phi_v(t)$ itself, the content of $\psi_v(u)$ may change from $M_u(t-1)$ to $M_u(t)$ (if it was not already $M_u(t)$), but it will not store $M_u(t')$ for any $t' > t$. This fact is crucial for the implementation of the simulation feature.

The simulation feature. Upon completion of the pausing feature $P_q \times \{j\}$, v moves on to the simulation feature $S_q \times \{j\}$. The purpose of this feature is to perform the actual simulation of round t in v , namely, to determine the state (of Q) dominating the simulation phase of the next round and the message transmitted when moving from the simulation phase of the current round

to that of the next round.

To see how this works out, suppose that $\lambda(q) = \sigma \in \Sigma$. We would have wanted node v to count (up to the bounding parameter b) the number of appearances of $\widehat{\Sigma}$ -letters in its ports that correspond to the transmission of σ at round $t - 1$ under Π , that is, the number of appearances of letters in $\Gamma_{t-1} \cup \Gamma_t$, where

$$\Gamma_{t-1} = \{(\sigma', \sigma, j - 1) \mid \sigma' \in \Sigma \cup \{\varepsilon\}\} \quad \text{and} \quad \Gamma_t = \{(\sigma, \sigma', j) \mid \sigma' \in \Sigma \cup \{\varepsilon\}\} .$$

More formally, the application of the transition function $\widehat{\delta}$ at the end of the simulation phase $\phi_v(t)$ should be based on $\beta_b(\sum_{\gamma \in \Gamma_{t-1} \cup \Gamma_t} \#(\gamma))$, where $\#(\gamma)$ stands for the number of appearances of the letter γ in the ports of v at the end of $\phi_v(t)$.

Identifying the integer b with the symbol $\geq b$, we observe that the function $\beta_b : \mathbb{Z}_{\geq 0} \rightarrow B$ satisfies

$$\beta_b \left(\sum_{\gamma \in \Gamma_{t-1} \cup \Gamma_t} \#(\gamma) \right) = \min \left\{ \sum_{\gamma \in \Gamma_{t-1} \cup \Gamma_t} \beta_b(\gamma), b \right\} .$$

A naive attempt to compute $\beta_b(\sum_{\gamma \in \Gamma_{t-1} \cup \Gamma_t} \#(\gamma))$ would simply traverse a sequence of $|\Gamma_{t-1} \cup \Gamma_t|$ states, each dedicated to collecting the value of $\beta_b(\gamma)$ for some $\gamma \in \Gamma_{t-1} \cup \Gamma_t$, and adding it to the sum maintained so far. More formally, the feature $S_q \times \{j\}$ would include a state $s_{\gamma,i}$ for every letter $\gamma \in \Gamma_{t-1} \cup \Gamma_t$ and integer $i \in \{0, \dots, b\}$; the query letter of $s_{\gamma,i}$ would be $\widehat{\lambda}(s_{\gamma,i}) = \gamma$ and the transition function $\widehat{\delta}$ would be designed so that v moves from $s_{\gamma,i}$ to $s_{\gamma',i'}$, where γ' follows γ in some arbitrary (fixed) order of the letters in $\Gamma_{t-1} \cup \Gamma_t$ and $i' = \min\{i + \beta_b(\#(\gamma)), b\}$.

However, care must be taken with this approach since $\#(\gamma)$ may decrease (respectively, increase) during (the simulating feature of) phase $\phi_v(t)$ for $\gamma \in \Gamma_{t-1}$ (resp., for $\gamma \in \Gamma_t$) due to new incoming messages transmitted by neighbors of v when they completed their simulation of round $t - 1$ under Π . To avoid this obstacle, we design the feature $S_q \times \{j\}$ so that first, it computes $\varphi_1 \leftarrow \beta_b(\sum_{\gamma \in \Gamma_{t-1}} \#(\gamma))$; next, it computes $\varphi_2 \leftarrow \beta_b(\sum_{\gamma \in \Gamma_t} \#(\gamma))$; and finally, it computes “again” $\varphi_3 \leftarrow \beta_b(\sum_{\gamma \in \Gamma_{t-1}} \#(\gamma))$. If $\varphi_1 = \varphi_3$, then the current simulation phase is over and $\widehat{\delta}$ is applied, simulating $\delta(q, \beta_b(\varphi_1 + \varphi_2))$; otherwise, the feature $S_q \times \{j\}$ is invoked from scratch. Since the value of $\beta_b(\sum_{\gamma \in \Gamma_{t-1}} \#(\gamma))$ cannot increase during the simulation phase, and since $\varphi_1 \leq b$, the feature $S_q \times \{j\}$ is invoked at most b times throughout the execution of the simulation phase. By induction on t , we conclude that our synchronizer satisfies synchronization property (S2), which concludes the correctness proof of the simulation.

Accounting. It remains to show that all ingredients of protocol $\widehat{\Pi}$ are of constant size and that the run-time of protocol $\widehat{\Pi}$ incurs at most a constant multiplicative overhead on top of that of protocol Π . The former claim is established by following our synchronizer construction, observing that $|\widehat{\Sigma}| = O(|\Sigma|^2)$ and $|\widehat{Q}| = O(|Q| \cdot (|\Sigma|^2 + |\Sigma| \cdot b))$ (recall that the bounding parameter b remains unchanged). For the latter claim, we need the following definition: given some node subset $U \subseteq V$

and round $t \in \mathbb{Z}_{>0}$, let $\tau(U, t)$ denote the first time at which u completed simulation phase $\phi_u(t)$ for all nodes $u \in U$. The following proposition can now be established.

Proposition 3.3. *For every node $v \in V$ and round $t \in \mathbb{Z}_{>0}$, the time $\tau(\cdot, \cdot)$ satisfies*

$$\tau(\{v\}, t+1) \leq \tau(N(v) \cup \{v\}, t) + O(1).$$

Proof. Since each transmitted message has a delay of at most 1 unit of time, it follows that by time $\tau(N(v) \cup \{v\}, t) + 1$, message $M_u(t)$ must reach $\psi_v(u)$ for all $u \in N(v)$. The pausing and simulation features of $\phi_v(t+1)$ are then completed within $O(|\Sigma|^2)$ and $O(|\Sigma| \cdot b)$ steps, respectively. The assertion follows as each step lasts for at most 1 unit of time. \square

Employing Proposition 3.3, we conclude that $\tau(V, t+1) \leq \tau(V, t) + O(1)$ and hence, by induction on t , that $\tau(V, t) = O(t)$. Therefore, if the execution of protocol Π requires T rounds, then the execution of protocol $\widehat{\Pi}$ is completed within $O(T)$ time units. Theorem 3.1 follows.

3.2 Multiple-letter queries

Recall that according to the model presented in Section 2, each state $q \in Q$ is associated with a query letter $\lambda(q)$ and the application of the transition function when node v resides in state q is determined by $\beta_b(\sharp(\sigma))$, where $\sharp(\sigma)$ is the number of appearances of the letter σ in the ports of v . From the perspective of the protocol designer, it is often more convenient to assume that the node queries on all letters simultaneously, namely, that the application of the transition function is determined by the vector $(\beta_b(\sharp(\sigma)))_{\sigma \in \Sigma}$.

Now that we may assume a synchronous environment, this stronger multiple-letter queries assumption can easily be supported. Indeed, at the cost of increasing the number of states and the run-time by constant factors, one can subdivide each round into $|\Sigma|$ subrounds, dedicating each subround to a different letter in Σ , so that at the end of the round, the state of v reflects $\beta_b(\sharp(\sigma))$ for every $\sigma \in \Sigma$.

Theorem 3.4. *Every nFSM protocol with multiple-letter queries can be simulated by an nFSM protocol with single-letter queries and the same bounding parameter b at the cost of a constant multiplicative run-time overhead.*

4 Efficient algorithms

As stated earlier, the main technical contribution of this paper is cast in the development of efficient nFSM algorithms for some of the most important and extensively studied problems in distributed computing. These problems include maximal independent set, maximal 2-hop independent set, node coloring of bounded degree graphs with $\Delta + 1$ colors, node 2-hop coloring of bounded degree

graphs with $\Delta^2 + 1$ colors, node coloring of (undirected) trees with 3 colors, and maximal matching (where we use a small unavoidable modification of the model).

4.1 Maximal independent set

Given a graph $G = (V, E)$, the *maximal independent set (MIS)* problem asks for a node subset $U \subseteq V$ which is independent in the sense that $(U \times U) \cap E = \emptyset$, and maximal in the sense that $U' \subseteq V$ is not independent for every $U' \supset U$. The challenge of designing a fast distributed MIS algorithm was first posed by Valiant in the early 1980s [50]. Distributed MIS algorithms with logarithmic run-time operating in the message passing model were subsequently presented by Luby [37] and independently, by Alon et al. [5].⁹ Luby’s algorithm has since become a specimen of distributed algorithms; in the last 25 years, researchers have tried to improve it, if only e.g., with an improved bit complexity [40], on special graph classes [46, 35, 16], or in a weaker communication model [2]. An $\Omega(\sqrt{\log n})$ lower bound on the run-time of any distributed MIS algorithm operating in the message passing model was established by Kuhn et al. [33]. Our goal in this section is to establish the following theorem.

Theorem 4.1. *There exists an nFSM protocol with bounding parameter $b = 1$ that computes an MIS for any n -node graph in time $O(\log^2 n)$.*

Outline of the key technical ideas. The protocol promised by Theorem 4.1 is inspired by the existing message passing MIS algorithms. Common to all these algorithms is that they are based on the concept of grouping consecutive rounds into *phases*, where in each phase, nodes compete against their neighbors over the right to join the MIS. Existing implementations of such competitions require at least one of the following three capabilities: (1) performing calculations that involve super-constant variables; (2) communicating with each neighbor independently; or (3) sending messages of a logarithmic size. The first two capabilities are clearly out of the question for an nFSM protocol. The third one is also not supported by the nFSM model, but perhaps one can divide a message with a logarithmic number of bits over logarithmically many rounds, sending $O(1)$ bits per round (cf. Algorithm B in [40])?

This naive attempt results in super-constant long phases, while no FSM can count the rounds in such phases — a task essential for deciding if the current phase is over and the next one should begin. Furthermore, to guarantee fair competition, the phases must be aligned across the network, thus ruling out the possibility to start node v ’s phase i before phase $i - 1$ of some node $u \in N(v)$ is finished. In fact, an efficient algorithm that requires $\omega(1)$ long aligned phases cannot be implemented under the nFSM model. So, how can we decide if node v joins the MIS using constant size messages without the ability to maintain long aligned phases?

⁹ The focus of [37] and [5] was actually on the PRAM model, but their algorithms can be adapted to the message passing model.

This issue is resolved by relaxing the requirements that the phases are aligned and of a pre-determined length, introducing a feature referred to as a *tournament*. Our tournaments are only “softly” aligned and their lengths are determined probabilistically, in a manner that can be maintained under the nFSM model. Nevertheless, they enable a fair competition between neighboring nodes, as desired.

The protocol. Employing Theorems 3.1 and 3.4, we assume a locally synchronous environment and use multiple-letter queries. The state set of the protocol is $Q = \{\text{WIN}, \text{LOSE}, \text{DOWN}_1, \text{DOWN}_2, \text{UP}_0, \text{UP}_1, \text{UP}_2\}$, with $Q_I = \{\text{DOWN}_1\}$ (the initial state of all nodes) and $Q_O = \{\text{WIN}, \text{LOSE}\}$, where WIN (respectively, LOSE) indicates membership (resp., non-membership) in the MIS output by the protocol. The states in $Q_A = Q - Q_O$ are called the *active* states and a node in an active state is referred to as an *active* node. We take the communication alphabet Σ to be identical to the state set Q , where the letter transmissions are designed so that node v transmits letter q whenever it moves to state q from some state $q' \neq q$; no letter is transmitted in a round at which v remains in the same state. Letter DOWN_1 is the initial letter stored in all ports at the beginning of the execution. The bounding parameter is set to $b = 1$.

A schematic description of the transition function is provided in Figure 1; its logic is as follows. Each state $q \in Q_A$ has a subset $D(q) \subseteq Q_A - \{q\}$ of *delaying states*: node v remains in the current state q if and only if (at least) one of its neighbors is in some state in $D(q)$. This is implemented by querying on the letters (corresponding to the states) in $D(q)$, staying in state q as long as at least one of these letters is found in the ports. Specifically, state DOWN_1 is delayed by state DOWN_2 , which is delayed by all three UP states. State UP_j , $j = 0, 1, 2$, is delayed by state $\text{UP}_{j-1 \bmod 3}$, where state UP_0 is also delayed by state DOWN_1 .

States WIN and LOSE are sinks. Assuming that an active node v does not find any delaying letter in its ports, the logic of the UP and DOWN states is as follows. From state DOWN_1 , v moves to state UP_0 . From state DOWN_2 , v moves to state DOWN_1 if $\sharp(\text{WIN}) = 0$, that is, if it does not find any WIN letter in its ports; otherwise, it moves to state LOSE. When in state UP_j , v tosses a fair coin and proceeds as follows: if the coin turns head, then v moves to state $\text{UP}_{j+1 \bmod 3}$; if the coin turns tail, then v moves to state WIN if $\sharp(\text{UP}_j) = \sharp(\text{UP}_{j+1 \bmod 3}) = 0$ (note that $\sharp(\text{UP}_{j-1 \bmod 3})$ must be 0 as $\text{UP}_{j-1 \bmod 3} \in D(\text{UP}_j)$); and to state DOWN_2 otherwise. This completes the description of our nFSM protocol for the MIS problem.

Turns and tournaments. Our protocol is designed so that an active node v traverses the DOWN and UP states in a (double-)circular fashion: an inner loop of the UP states (moving from state UP_j to state $\text{UP}_{j+1 \bmod 3}$) nested within an outer loop consisting of the DOWN states and the inner loop. Of course, v may spend more than one round at each state $q \in Q_A$ (delayed by adjacent nodes in states $D(q)$); we refer to a maximal contiguous sequence of rounds that v spends in the same state $q \in Q_A$ as a q -*turn*, or simply as a *turn* if the actual state q is irrelevant. A maximal contiguous




Figure 1: The transition function of the MIS protocol with state names abbreviated by their first (capital) letter. The node stays in state q (a.k.a. *delayed*) as long as letter q' appears (at least once) in its ports for any state q' such that a $q' \rightarrow q$ transition is defined (for clarity, this is omitted from the figure). Assuming that the node is not delayed, each transition specified in the figure is associated with a condition on the number of appearances of the query letters in the ports (depicted by the corresponding lower-case letter) so that the transition is followed only if the condition is satisfied (an empty condition is satisfied by all port configurations); if some port configuration satisfies several transition conditions, then one of them is chosen uniformly at random.

sequence of turns that starts at a DOWN_1 -turn and does not include any other DOWN_1 -turn (i.e., a single iteration of the outer loop) is referred to as a *tournament*. We index the tournaments and the turns within a tournament by the positive integers. Note that by definition, every tournament i of v starts with a DOWN_1 -turn, followed by a non-empty sequence of UP -turns; tournament i can end with an UP -turn from which v moves to state WIN , with a DOWN_2 -turn from which v moves to state LOSE , or with a DOWN_2 -turn from which v moves to state DOWN_1 starting tournament $i + 1$. The following observation is established by induction on the rounds.

Observation 4.2. *Consider some active node $v \in V$ in turn $j \in \mathbb{Z}_{>0}$ of tournament $i \in \mathbb{Z}_{>0}$ and some active node $u \in N(v)$.*

- *If this is a DOWN_1 -turn of v ($j = 1$), then u is in either (a) the last (DOWN_2 -)turn of tournament $i - 1$; (b) turn 1 of tournament i ; or (c) turn 2 of tournament i .*
- *If this is an UP -turn of v ($j \geq 2$), then u is in either (a) turn $j - 1$ of tournament i ; (b) turn j of tournament i ; (c) turn $j + 1$ of tournament i ; or (d) the last (DOWN_2 -)turn $j' \leq j + 1$ of tournament i .*
- *If this is a DOWN_2 -turn of v (the last turn of this tournament), then u is in either (a) an*

UP-turn $j' \geq j - 1$ of tournament i ; (b) the last (DOWN₂-)turn of tournament i ; or (c) turn 1 of tournament $i + 1$.

Given some $U \subseteq V$ and $i, j \in \mathbb{Z}_{>0}$, let $T^U(i, j)$ denote the first time at which every node $v \in U$ satisfies either

- (1) v is inactive;
- (2) v is in tournament $i' > i$;
- (3) v is in the last (DOWN₂-)turn of tournament i ; or
- (4) v is in turn $j' \geq j$ of tournament i .

Note that $T^U(i, j)$ is well defined even if some node $v \in U$ does not reach turn j of tournament i . Employing Observation 4.2, the delaying states feature guarantees that

$$T^{\{v\}}(i, j + 1) \leq T^{N(v) \cup \{v\}}(i, j) + 1 \quad (1)$$

for every $v \in V$ and $i, j \in \mathbb{Z}_{>0}$. Since $T^U(i, j) \leq T^V(i, j)$ for every $U \subseteq V$, we can apply inequality (1) to each node $v \in V$, concluding that

$$T^V(i, j + 1) \leq T^V(i, j) + 1,$$

which immediately implies that

$$T^V(i, k + 1) \leq T^V(i, 1) + k. \quad (2)$$

Moreover, if no node in V goes beyond turn j of tournament i , then

$$T^V(i + 1, 1) = T^V(i, j + 1) \leq T^V(i, j) + 1. \quad (3)$$

The virtual graph G_i . Let V_i be the set of nodes for which tournament i exists and let $G_i = (V_i, E_i)$ be the subgraph induced on G by V_i , where $E_i = E \cap (V_i \times V_i)$. Given some node $v \in V_i$, let $N_i(v) = \{u \in V_i \mid (u, v) \in E_i\}$ be the neighborhood of node v in G_i and let $d_i(v) = |N_i(v)|$ be its degree. Note that the graph G_i is virtual in the sense that it is defined solely for the sake of the analysis: we do not assume that there exists some time at which the graph induced by any meaningful subset of the nodes (say, the nodes in tournament i) agrees with G_i .

Given some node $v \in V_i$, let $X^v(i)$ denote the number of UP-turns in tournament i of v and recall that the total number of turns in this tournament is at most $X^v(i) + 2$, accounting for the DOWN₁ turn in the beginning of the tournament and the DOWN₂-turn in its end. The logic of the UP states implies that $X^v(i)$ is a Geom(1/2)-random variable, namely, it obeys the geometric distribution with parameter 1/2. The key observation now is that conditioned on G_i , the random variables $X^v(i)$, $v \in V_i$, are independent. Moreover, the graph G_{i+1} is fully determined by the random variables $X^v(i)$, $v \in V_i$. Since the maximum of (at most) n independent Geom(1/2)-random variables is $O(\log n)$ w.h.p., inequalities (2) and (3) yield the following observation.

Observation 4.3. For every $i \in \mathbb{Z}_{>0}$, $T^V(i, 1)$ is finite w.p. 1 and

$$T^V(i + 1, 1) \leq T^V(i, 1) + O(\log n)$$

w.h.p.

Our protocol is designed so that node v moves to an output state (WIN or LOSE) in the end of each tournament w.p. > 0 . Moreover, the logic of state DOWN₂ guarantees that if node v moves to state WIN in the end of tournament i , then all its active neighbors move to state LOSE in the end of their respective tournaments i . The correctness of our protocol now follows from Observation 4.3: the protocol reaches an output configuration w.p. 1 and every output configuration reflects an MIS. It remains to bound the run-time of our protocol; the following lemma plays a major role in this task.

Lemma 4.4. *Our MIS protocol guarantees the existence of two constants $0 < p, c < 1$ such that $|E_{i+1}| \leq c|E_i|$ w.p. $\geq p$.*

We will soon prove Lemma 4.4, but first, let us explain why it suffices for the completion of our analysis. Define the random variable $Y = \min\{i \in \mathbb{Z}_{>0} : |E_i| = 0\}$. Lemma 4.4 implies that Y is stochastically dominated by a random variable that obeys distribution

$$O(\log n) + \text{NB}(O(\log n), 1 - p),$$

namely, a fixed term of $O(\log n)$ plus the negative binomial distribution with parameters $O(\log n)$ and $1 - p$, hence $Y = O(\log n)$ in expectation and w.h.p. Since the nodes in $V - V_i$ are all in an output state (and will remain in that state), and since the logic of the UP states implies that a degree-0 node in G_i will move to state WIN in the end of tournament i (w.p. 1) and thus, will not be included in V_{i+1} , we can employ Observation 4.3 to conclude that the run-time of our protocol is indeed $O(\log^2 n)$.

The remainder of this section is dedicated to establishing Lemma 4.4. The proof technique we use for that purpose resembles (a hybrid of) the techniques used in [5] and [40] for the analysis of their MIS algorithms. We say that node $v \in V_i$ is *good* in G_i if

$$|\{u \in N_i(v) \mid d_i(u) \leq d_i(v)\}| \geq d_i(v)/3,$$

i.e., if at least third of v 's neighbors in G_i have degrees smaller or equal to that of v . The following lemma is established in [5].

Lemma 4.5 ([5]). *More than half of the edges in E_i are incident on good nodes in G_i .*

Disjoint winning events. Consider some good node v in G_i with $d = d_i(v) > 0$ and let $\widehat{N}_i(v) = \{u \in N_i(v) \mid d_i(u) \leq d\}$. Recall that the definition of a good node implies that $|\widehat{N}_i(v)| \geq d/3$. We say that node $u \in \widehat{N}_i(v)$ *wins* v in tournament i if

$$X^u(i) > \max \left\{ X^w(i) \mid w \in N_i(u) \cup \widehat{N}_i(v) - \{u\} \right\}$$

and denote this event by $A_i(u, v)$. The main observation now is that if u wins v in tournament i , then in the end of their respective tournaments i , u moves to state WIN and v moves to state LOSE. Moreover, the events $A_i(u, v)$ and $A_i(w, v)$ are disjoint for every $u, w \in \widehat{N}_i(v)$, $u \neq w$.

Fix some node $u \in \widehat{N}_i(v)$. Let u_1, \dots, u_k be the nodes in $N_i(u) \cup \widehat{N}_i(v)$, where $0 < k \leq 2d$ by the definition of a good node. Let $B_i(u, v)$ denote the event that the maximum of $\{X^{u_\ell}(i) \mid 1 \leq \ell \leq k\}$ is attained at a single $1 \leq \ell \leq k$. Since $X^{u_1}(i), \dots, X^{u_k}(i)$ are independent random variables that obey distribution $\text{Geom}(1/2)$, it follows that $\mathbb{P}(B_i(u, v)) \geq 2/3$ and therefore,

$$\mathbb{P}(A_i(u, v)) = \mathbb{P}(A_i(u, v) \mid B_i(u, v)) \cdot \mathbb{P}(B_i(u, v)) \geq \frac{1}{k} \cdot \frac{2}{3}.$$

Given that v is good in G_i and recalling the disjointness of the $A_i(u, v)$ events, the last inequality implies that

$$\mathbb{P}(v \notin V_{i+1}) \geq \mathbb{P}\left(\bigvee_{u \in \widehat{N}_i(v)} A_i(u, v)\right) = \sum_{u \in \widehat{N}_i(v)} \mathbb{P}(A_i(u, v)) \geq \frac{d}{3} \cdot \frac{1}{2d} \cdot \frac{2}{3} = \frac{1}{9}.$$

Combined with Lemma 4.5, we conclude that $\mathbb{E}[|E_{i+1}|] < \frac{17}{18} |E_i|$. Lemma 4.4 now follows by Markov's inequality, thus establishing Theorem 4.1.

4.2 Node coloring

Given a graph $G = (V, E)$, the *coloring* problem asks for an assignment of colors to the nodes such that no two neighboring nodes have the same color. A coloring using at most k colors is called a *k-coloring*. The smallest number of colors needed to color graph G is referred to as the *chromatic number* of G , denoted by $\chi(G)$. In general, $\chi(G)$ is difficult to compute even in a centralized model [17]. As such, the distributed computing community is generally satisfied already with a $(\Delta + 1)$ -, $O(\Delta)$ - or even $\Delta^{O(1)}$ -coloring, where $\Delta = \Delta(G)$ is the largest degree in the graph G , with possibly $\Delta(G) \gg \chi(G)$ [24, 36, 49, 14, 32, 39, 15, 47]. As the output of each node under the nFSM model is taken from a predetermined constant size set, we cannot hope to solve these problems for general graphs. Instead, we observe that the following simple nFSM protocol colors any *bounded degree* graph in logarithmic time.

Consider some predetermined constant upper bound d on the node degrees. Node $v \in V$ maintains a list $A(v)$ of available colors, where initially, $A(v) \leftarrow \{1, \dots, d + 1\}$. In round $2t - 1$ (employing Theorem 3.1, we assume a locally synchronous environment), an uncolored node v picks some color $c(v) \in A(v)$ uniformly at random and transmits a “proposing color c ” message. If v does not receive any “proposing color c ” message from its neighbors in round $2t$, then it irrevocably colors itself c and transmits a “my color is c ” message; following that, each uncolored neighbor u of v updates its list $A(u) \leftarrow A(u) - \{c\}$. Since w.p. $\geq 1/e$, no neighbor of v picks color c in round $2t - 1$, a constant fraction of the nodes are expected to be (irrevocably) colored in round $2t$, which establishes the following theorem.

Theorem 4.6. *Given some constant d , there exists an nFSM protocol with bounding parameter $b = 1$ that $(d + 1)$ -colors any n -node graph satisfying $\Delta \leq d$ in time $O(\log n)$.*

As Δ may grow quickly with n even for relatively simple graph classes, the remainder of this section is dedicated to a natural graph class that features a small chromatic number regardless of Δ : trees. Any tree T has a chromatic number $\chi(T) = 2$. Unfortunately, it is easy to show that in general, the task of 2-coloring trees requires run-time proportional to the diameter of the tree even under the message passing model, and hence cannot be achieved by an efficient distributed algorithm. The situation improves dramatically once 3 colors are allowed; indeed, Cole and Vishkin [24] presented a distributed algorithm that 3-colors directed paths, and in fact, any directed tree (directed in the sense that each node knows the port leading to its unique parent), in time $O(\log^* n)$. Linial [36] showed that this is asymptotically optimal.

Since it is not clear how to represent directed trees in the nFSM model, we focus on undirected trees. A lower bound result of Kothapalli et al. [31] shows that under the *anonymous* (namely, the nodes are not assumed to have unique identifiers) message passing model, 3-coloring undirected trees requires $\Omega(\log n)$ time as long as the size of each message is $O(1)$. Here we show that this lower bound is tight under the nFSM model.

Theorem 4.7. *There exists an nFSM protocol with bounding parameter $b = 3$ that 3-colors any n -node (undirected) tree in time $O(\log n)$.*

The description of most of the nFSM protocols in the remainder of this paper (including the one promised by Theorem 4.7) will not dwell into the level of defining the states and transition function as we did in Section 4.1 for the MIS protocol, but the reader will be easily convinced that these protocols can indeed be implemented under the nFSM model.

The modes. Employing Theorems 3.1 and 3.4, we assume a locally synchronous environment and use multiple-letter queries. At all times, each node $v \in V$ is in one of the following three *modes*.

- (1) Mode **COLORED**: the color of v is determined (v is in an output state) and it no longer takes an active part in the protocol.
- (2) Mode **ACTIVE**: the color of v has not been determined yet and v takes an active part in the protocol.
- (3) Mode **WAITING**: the color of v has not been determined yet and v is waiting for one of its neighbors to be colored before it resumes taking an active part in the protocol (going back to mode **ACTIVE**).

Initially, all nodes are in mode **ACTIVE**. When an **ACTIVE** node moves to mode **COLORED**, assigned with color $c \in \{1, 2, 3\}$, it transmits a ‘my color is c ’ message and it does not transmit any more messages; when an **ACTIVE** node moves to mode **WAITING**, it transmits an ‘I am **WAITING**’ message and it does not transmit any more messages until it returns to mode **ACTIVE**, in which

case it transmits an ‘I am ACTIVE’ message. Therefore, the message stored in the port of node v corresponding to neighbor u of v always indicates (perhaps among other things) the current mode of u .

The phases. The execution of the protocol is divided into *phases* indexed by the positive integers, where each phase consists of 4 rounds. Consider some phase $i \in \mathbb{Z}_{>0}$. Let V^i be the set of ACTIVE nodes at the beginning of phase i and let F^i be the forest induced on T by V^i (F^i may contain one or more trees), referred to as the ACTIVE forest. Given some node $v \in V^i$, let $N^i(v) = \{u \in V^i \mid (u, v) \in E\}$ be the neighborhood of v in F^i and let $d^i(v) = |N^i(v)|$ be its degree.

The structure of the phases is as follows. Consider some node $v \in V^i$. In round 1 of the phase, v transmits an ‘I am ACTIVE’ message. Setting the bounding parameter of the protocol to $b = 3$, we conclude that in round 2, v can distinguish between the cases $d^i(v) = 0$, $d^i(v) = 1$, $d^i(v) = 2$, and $d^i(v) \geq 3$ simply by querying its ports for ‘I am ACTIVE’ messages; in other words, v “knows” $\beta_3(d^i(v))$, i.e., its degree calculated with respect to the one-two-many principle with bounding parameter $b = 3$. Employing this “knowledge”, v transmits $\beta_3(d^i(v))$ in round 2 of phase i , so in round 3, the port of v corresponding to u stores a message indicating $\beta_3(d^i(u))$ for every node $u \in N^i(v)$.

Rounds 3 and 4 of phase i are dedicated to Procedure **RandColor** that we will describe soon. Whether or not v runs Procedure **RandColor** depends on the degree of v and on the degrees of its ACTIVE neighbors. Specifically, v runs Procedure **RandColor** if: (1) $d^i(v) = 0$; (2) $d^i(v) = 1$ with $N^i(v) = \{u\}$ and $d^i(u) = 1$; or (3) $d^i(v) = 2$ with $N^i(v) = \{u_1, u_2\}$ and $d^i(u_1), d^i(u_2) \leq 2$. In contrast, if $d^i(v) = 1$ with $N^i(v) = \{u\}$ and $d^i(u) \geq 2$, then v moves to mode **WAITING** without running Procedure **RandColor**, in which case we say (just for the sake of the analysis) that v *waits on u* . Otherwise ($d^i(v) \geq 3$ or $d^i(v) = 2$ with some neighbor $u \in N^i(v)$ such that $d^i(u) \geq 3$), v remains in mode **ACTIVE** without running Procedure **RandColor**.

As stated beforehand, the **COLORED** nodes do not take an active part in the protocol. A **WAITING** node v moves to mode **ACTIVE** in the end of phase i if some neighbor u of v , $u \in V^i$, moves to mode **COLORED** during phase i (v spots this event by querying on ‘my color is c ’ messages).

Procedure RandColor. Responsible for the actual color assignments, Procedure **RandColor** takes 2 rounds (rounds 3 and 4 of some phase). Only an ACTIVE node may run the procedure, and when the procedure is over, the node either stays in mode **ACTIVE** or moves to mode **COLORED**. Consider some node v running the procedure and let $C(v) \subseteq \{1, 2, 3\}$ be the subset of colors which are not yet assigned to the neighbors of v in T . (Our analysis shows that if v is ACTIVE, then $C(v) \neq \emptyset$.) As every **COLORED** node transmits a message indicating its color, v can determine $C(v)$ by querying its ports.

In the first round of Procedure **RandColor**, v picks some color $c \in C(v)$ uniformly at random

and transmits a ‘proposing color c ’ message. In the second round of the procedure, if v finds a ‘proposing color c ’ (with the same c) in its ports, then it remains in mode ACTIVE. Otherwise (no neighbor of v competes with v over color c), it moves to mode COLORED and transmits a ‘my color is c ’ message. This completes the description of our protocol.

The waiting hierarchy. The ‘waits on’ relation induces a hierarchy referred to as the *waiting hierarchy* which is represented by a (collection of) directed tree(s) defined over a subset of the edges of the tree T . Our protocol is designed so that if v waits on u , moving to mode WAITING in phase i , then in phases $1, \dots, i$, u was ACTIVE, and in phase $i + 1$, u is either ACTIVE or COLORED. Moreover, if u is ACTIVE and $v \in N(u)$ is WAITING, then v must be waiting on u . Note also that if v waits on u and u moves to mode COLORED in phase j , then v moves back to mode ACTIVE in (the beginning of) phase $j + 1$ and $d^{j+1}(v) = 0$.

Observation 4.8. *In the beginning of phase i , $|C(v)| \geq \min\{d^i(v) + 1, 3\}$ for every $i \in \mathbb{Z}_{>0}$ and node $v \in V^i$.*

Proof. As long as $d^i(v) \geq 3$, no neighbor of v can run Procedure RandColor, and hence no neighbor of v can move to mode COLORED. Therefore, $C(v) = \{1, 2, 3\}$ in the beginning of the first phase $i \in \mathbb{Z}_{>0}$ such that $d^i(v) \leq 2$. From that moment on, every ACTIVE neighbor of v that moves to mode COLORED decreases both $|C(v)|$ and $d^i(v)$ by 1. The assertion is completed by recalling that non-ACTIVE neighbors of v must be waiting on v and hence, cannot move to mode COLORED before v does. \square

Corollary 4.9. *Consider some node $v \in V^i$ that runs Procedure RandColor. If $d^i(v) = 0$, then v moves to mode COLORED w.p. 1. Otherwise ($d^i(v)$ is either 1 or 2), v moves to mode COLORED with a positive constant probability.*

Let \tilde{V}^i be the restriction of V^i to nodes v that were ACTIVE in all phases $1, \dots, i$; this is, \tilde{V}^i does not include WAITING nodes that became ACTIVE again (recall that these will move to mode COLORED in the next phase w.p. 1). Let \tilde{F}^i be the forest induced on T by \tilde{V}^i . Given some node $v \in \tilde{V}^i$, let $\tilde{N}^i(v) = \{u \in \tilde{V}^i \mid (u, v) \in E\}$ be the neighborhood of v in \tilde{F}^i and let $\tilde{d}^i(v) = |\tilde{N}^i(v)|$ be its degree. Observe that if $v \in \tilde{V}^i$, then $v \in V^i$ and $\tilde{d}^i(v) = d^i(v)$. Therefore, if $v \in V^i - \tilde{V}^i$, then $d^i(v) = 0$, in which case v runs Procedure RandColor in phase i and Corollary 4.9 guarantees that $v \notin V^{i+1}$.

The correctness of the protocol can now be established: The logic of Procedure RandColor implies that every output configuration is a legal coloring. Since ACTIVE leaves are removed from \tilde{F}^i w.p. 1 and since every tree has at least two leaves, it follows that $\hat{V}^{1+k} = \emptyset$ for $k = \lceil n/2 \rceil$. Combining the properties of the waiting hierarchy with Corollary 4.9, we conclude that the execution reaches an output configuration within at most k additional phases. It remains to analyze the run-time of our protocol.

Good nodes. Consider some tree T' . We say that node v of T' is *good* if v is a leaf or if the degree of v is 2 and both neighbors of v are of degree at most 2.

Observation 4.10. *In every tree, at least a $(1/5)$ -fraction of the nodes are good.*

Consider some $i \in \mathbb{Z}_{>0}$ and some node $v \in \tilde{V}^i$. Let T' be the tree to which v belongs in \tilde{F}^i . We argue that if v is good in T' , then $v \notin \tilde{V}^i$ with a positive constant probability. Indeed, if v is a leaf in T' , which means that $\tilde{d}^i(v) = d^i(v) = 1$, then it either moves to mode **WAITING** w.p. 1 (if the neighbor of v has a higher degree) or it runs Procedure **RandColor**, in which case Corollary 4.9 guarantees that v moves to mode **COLORED** with a positive constant probability; if $\tilde{d}^i(v) = d^i(v) = 2$ and both neighbors of v in F^i (and in T') are of degree at most 2, then v runs Procedure **RandColor**, in which case Corollary 4.9 again guarantees that v moves to mode **COLORED** with a positive constant probability. Since Corollary 4.9 also guarantees that nodes of degree 0 in \tilde{F}^i move to mode **COLORED** w.p. 1, we can employ Observation 4.10 and Markov's inequality to establish the following observation.

Observation 4.11. *There exists two constants $0 < p, c < 1$ such that $|\tilde{V}^{i+1}| \leq c|\tilde{V}^i|$ w.p. $\geq p$.*

Similarly to the analysis in Section 4.1, define the random variable $Y = \min\{i \in \mathbb{Z}_{>0} : |\tilde{V}^i| = 0\}$. Observation 4.11 implies that Y is stochastically dominated by a random variable that obeys distribution $\text{NB}(O(\log n), 1 - p) + O(\log n)$, namely, a fixed term of $O(\log n)$ plus the negative binomial distribution with parameters $O(\log n)$ and $1 - p$, hence $Y = O(\log n)$ in expectation and w.h.p. Since Y bounds from above the depth of the waiting hierarchy, it follows that the execution reaches an output configuration within $2Y$ phases, thus establishing Theorem 4.7.

4.3 The square graph

Consider some graph $G = (V, E)$, node $v \in V$, and positive integer k . We define the *k -hop neighborhood* of v in G , denoted by $N^k(v)$, as the set of all nodes $u \in V$, $u \neq v$, at distance at most k from v . The *k^{th} power* of G , denoted by G^k , is the graph obtained from G by extending its edge set so that v is adjacent to all nodes in $N^k(v)$ for every $v \in V$. The second power G^2 of G is called the *square* of G .

Lemma 4.12. *For every n FSM protocol Π with bounding parameter $b = 1$, there exists an n FSM protocol Π^2 with bounding parameter $b = 2$ such that for every graph G , the execution of Π^2 on G simulates the execution of Π on G^2 with a constant multiplicative run-time overhead.*

Proof. Employing Theorems 3.1 and 3.4, we assume a locally synchronous environment and use multiple-letter queries. Moreover, denoting the communication alphabet of protocol Π by Σ , we assume that Π transmits some Σ -letter in every round (i.e., the empty symbol ε is never used); this is a valid assumption since a node can always replace ε with the last transmitted Σ -letter without affecting the execution.

In the remainder of this proof, when we mention protocols Π and Π^2 , we refer to their executions

on G^2 and G , respectively. Round $t \in \mathbb{Z}_{>0}$ of node $v \in V$ under Π , is simulated by 2 rounds under Π^2 , specifically, rounds $2t - 1$ and $2t$. Assume by induction on t that in round $2t - 1$ under Π^2 , node v is provided with all the information it has in round t under Π . Then, v transmits in round $2t - 1$ under Π^2 the same Σ -letter it transmits in round t under Π . Round $2t$ under Π^2 is then dedicated to gathering information from the 2-hop neighborhoods, as follows.

We augment the communication alphabet of Π^2 with *auxiliary letters*, identified with the functions in $\{0, 1, \geq 2\}^\Sigma$; these auxiliary letters are used in round $2t$ under Π^2 . Let $\#(\sigma)$ be the number of appearances of the letter $\sigma \in \Sigma$ in the ports of node v in round $2t$ under Π^2 and let σ_v be the Σ -letter transmitted by v in round $2t - 1$ under Π^2 (and in round t under Π). Then, the letter $f \in \{0, 1, \geq 2\}^\Sigma$ transmitted by v in round $2t$ is defined so that

$$f(\sigma) = \begin{cases} \beta_2(\#(\sigma) + 1) & \text{if } \sigma = \sigma_v \\ \beta_2(\#(\sigma)) & \text{otherwise,} \end{cases}$$

where we recall the definition of β_2 from Section 2.

We argue that at least one of v 's ports contains the letter $\sigma \in \Sigma - \{\sigma_v\}$ (respectively, the letter σ_v) in round $t+1$ under Π if and only if at least one of v 's neighbors in G transmitted some function $f \in \{0, 1, \geq 2\}^\Sigma$ satisfying $f(\sigma) \in \{1, \geq 2\}$ (resp., satisfying $f(\sigma_v) = \geq 2$) in round $2t$ under Π^2 . To see why this argument holds, consider some letter $\sigma \in \Sigma$ and recall that it appears in v 's ports in round $t + 1$ under Π if and only if some node $u \in N^2(v)$ transmitted it in round t under Π , which means that u transmitted σ in round $2t - 1$ under Π^2 . This, in turns, holds if and only if there exists some node $u' \in N(v)$ (that can be u itself or a neighbor of both u and v in G) such that the function $f \in \{0, 1, \geq 2\}^\Sigma$ transmitted by u' in round $2t$ under Π^2 satisfies $f(\sigma) \in \{1, \geq 2\}$, which establishes our argument for the case $\sigma \neq \sigma_v$. If, on the other hand, $\sigma = \sigma_v$, then every neighbor of v has at least one appearance of σ in its ports in round $2t$ under Π^2 (contributed by the port corresponding to v), and therefore some node $u \in N^2(v)$ transmitted σ in round $2t - 1$ under Π^2 if and only if there exists some node $u' \in N(v)$ (that, again, can be u itself or a neighbor of both u and v in G) such that the function $f \in \{0, 1, \geq 2\}^\Sigma$ transmitted by u' in round $2t$ under Π^2 satisfies $f(\sigma) = \geq 2$; this establishes our argument for the case $\sigma = \sigma_v$. The inductive step follows since v queries for each one of the functions in $\{0, 1, \geq 2\}^\Sigma$, which completes the proof. \square

A node subset $U \subseteq V$ is a *k-hop independent set* of the graph $G = (V, E)$ if $v \in U$ implies that $u \notin U$ for every $u \in N^k(v) - \{v\}$; the *maximal k-hop independent set (MkIS)* problem asks for a *k-hop independent set* $U \subseteq V$ which is maximal in the sense that $U' \subseteq V$ is not a *k-hop independent set* for any $U' \supset U$. Likewise, the *k-hop coloring* problem asks for an assignment of colors to the nodes such that the color of v differs from the color of u for every $u \in N^k(v) - \{v\}$. Cast in this terminology, the MIS and coloring problems are special cases of the MkIS set and *k-hop coloring* problems, respectively, for $k = 1$. It is shown in [26] that the MkIS and *k-hop coloring* problems cannot be solved for $k \geq 3$ by a distributed algorithm even under the much stronger anonymous

message passing model. In contrast, the $k = 2$ case is resolved positively by plugging Theorems 4.1 and 4.6 into Lemma 4.12.

Corollary 4.13. *Given some constant d , there exists an nFSM protocol with bounding parameter $b = 2$ that computes a 2-hop coloring with $(d^2 + 1)$ colors for any n -node graph satisfying $\Delta \leq d$ in time $O(\log n)$.*

Corollary 4.14. *There exists an nFSM protocol with bounding parameter $b = 2$ that computes an M2IS for any n -node graph in time $O(\log^2 n)$.*

Corollary 4.13 essentially provides us with the power to implement independent communication along each edge in bounded degree graphs: Given a 2-hop coloring c , we can append the pair $(c(u), c(v))$ to a message originated at node u whose destination is node $v \in N(u)$. This way, node v can detect that the message was sent by u , whereas any node $w \in N(u) - \{v\}$ can ignore it.

Simulating population protocols. Corollary 4.14 also has an interesting implication that takes us back to the comparison between the nFSM model and the population protocols model (see Section 1) as it allows us to simulate the rendezvous based communication of any population protocol (in arbitrary interaction graphs). To explain how this is done, we need to introduce the notion of an *oriented matching* consisting of a set $M = \{(x_1, y_1), \dots, (x_k, y_k)\}$ of ordered node pairs satisfying (1) $(x_i, y_i) \in E$ for every $1 \leq i \leq k$; and (2) $(x_i, y_j) \notin E$ for every $1 \leq i, j \leq k, i \neq j$. We refer to the nodes x_1, \dots, x_k as *leaders* and to the nodes y_1, \dots, y_k as *subordinates*.

Based on the M2IS protocol promised in Corollary 4.14, we will soon present an nFSM protocol that computes an oriented matching M ; specifically, upon termination of this protocol, each node knows if it is a leader, a subordinate, or neither. Moreover, for every edge $e \in E$ in each one of its two possible orientations, it holds that M is a singleton $M = \{e\}$ w.p. > 0 . By the definition of an oriented matching, each leader x_i (respectively, subordinate y_i) can now safely communicate with its subordinate y_i (resp., leader x_i) without risking interference from other leaders (resp., subordinates). Therefore, we can apply the rule of the simulated population protocol to each (x_i, y_i) pair and update x_i and y_i 's states under this simulated protocol accordingly. With an appropriate node delaying mechanism (see, e.g., Section 4.1), this process can be repeated indefinitely, thus yielding a schedule that must be fair (cf. [10]) due to the probabilistic guarantees of M .

Lemma 4.15. *The model obtained from nFSM by relaxing the correctness requirement to an eventually converging correctness is at least as strong, in terms of its computational power, as the population protocols model (with the same interaction graph).*

The nFSM protocol for the computation of an oriented matching M works as follows. We first employ the M2IS protocol promised in Corollary 4.14 and take the nodes in the M2IS to be the *potential* leaders of M . Each potential leader $x \in V$ then attempts to appoint her matching subordinate y chosen uniformly at random among her neighbors $N(x)$ through a *subordinate selection process* (SSP). If the SSP succeeds and a y is appointed, then the oriented edge (x, y) joins M ;

otherwise, x does not become a leader in M .

The SSP works as follows. Initially, all nodes $z \in N$ are classified as *candidates*. In each round of the SSP, the potential leader x chooses a bit $b_x \in \{0, 1\}$ uniformly at random (and independently of all other random choices) and transmits it. Each node $z \in N(x)$ that is still a candidate also chooses some bit $b_z \in \{0, 1\}$ uniformly at random (and independently); if $b_z \neq b_x$, then z turns into a non-candidate and transmits a designated non-candidate message which is delivered, in particular, to the corresponding port $\psi_x(z)$ of x . The SSP reaches a successful end if x has exactly one candidate neighbor z , in which case z is appointed as x 's subordinate. If a configuration with zero candidates is reached, then the SSP fails and x does not become a leader in M .

Note that this implementation of the SSP under the nFSM model is made possible only because the leaders form an M2IS which guarantees that disjoint SSPs do not interfere. The last thing to observe now is that since each node is included in the MIS constructed by the nFSM protocol designed in Section 4.1 w.p. > 0 , the same also holds for the M2IS constructed in the current section. Therefore, each node $x \in V$ is a potential leader that initiates the SSP w.p. > 0 and each edge $(x, y) \in E$ is included in M w.p. > 0 . Moreover, w.p. > 0 , all other SSPs fail in which case $M = \{e\}$ is a singleton.

4.4 Maximal matching

Given a graph $G = (V, E)$, an edge subset $M \subseteq E$ is called a *matching* if every node $v \in V$ is incident on at most one edge in M . The *maximal matching (MM)* problem asks for a matching which is maximal in the sense that $M' \subseteq E$ is not a matching for every $M' \supset M$. A message passing MM algorithm with logarithmic run-time was developed by Israeli and Itai [30], whereas the $\Omega(\sqrt{\log n})$ lower bound of [33] applies to the MM problem as well.

We would like to design an nFSM protocol for the MM problem. However, the nFSM model as defined in Section 2 is not expressive enough for this problem: in a complete bipartite graph with n nodes in each side for example, the number of maximal matchings is $n!$, while an nFSM protocol with c output states can only specify $c^{2n} \ll n!$ different (global) outputs. In other words, the nFSM model is geared towards problems whose output is specified by labeling the graph's nodes, whereas the output in the MM problem requires assigning (binary) labels to the graph's edges.

This obstacle is tackled by slightly extending the nFSM model in a manner that enhances it with the power to cope with edge labeling problems such as MM without violating model requirements (M1)–(M4) (see Section 1). To that end, we augment the 8-tuple $\Pi = \langle Q, Q_I, Q_O, \Sigma, \sigma_0, b, \lambda, \delta \rangle$ introduced in Section 2 with a finite *internal alphabet* Γ and with a *port transition function*

$$\eta : Q \times \Gamma \times \Sigma \rightarrow 2^\Gamma .$$

We also change the function $\lambda : Q \rightarrow \Sigma$ to $\lambda : Q \rightarrow \Gamma$ and the initial letter $\sigma_0 \in \Sigma$ to $\gamma_0 \in \Gamma$.

The new semantics is as follows. While the communication alphabet Σ is still used for message transmission, each port now contains some letter of the internal alphabet Γ , hence the function λ now maps Q to a query letter in Γ . Given some node $v \in V$ and port $\psi_v(u)$ corresponding to neighbor u of v , the port transition function η takes the current state $q \in Q$ of v , the letter $\gamma \in \Gamma$ currently stored in $\psi_v(u)$, and the new letter $\sigma \in \Sigma$ delivered to v from u , and returns some letter $\gamma' \in \Gamma$ chosen uniformly at random from $\eta(q, \gamma, \sigma) \subseteq \Gamma$; the letter γ' is then stored in $\psi_v(u)$ (replacing γ). This can be viewed as a FSM that controls the letters stored in each one of v 's ports (the same FSM for all ports).

Using this *extended nFSM* model, we can now specify edge labels through the ports of the nodes residing in output states. In particular, an MM protocol can specify its output matching M as follows. The protocol designer designates some letter $\gamma_{\text{out}} \in \Gamma$ for the purpose of outputting M . Node $v \in V$ residing in an output state may have at most one port $\Psi_v(u)$ storing the letter γ_{out} in which case it is guaranteed that port $\Psi_u(v)$ also stores γ_{out} , thus marking that $(u, v) \in M$. Note that this cannot be achieved under the (non extended) nFSM model, where, by definition, if $\psi_v(u)$ stores the letter $\sigma \in \Sigma$ at the end of the execution, then $\psi_w(u)$ also stores σ for every $w \in N(u)$.

Theorem 4.16. *There exists an extended nFSM protocol with bounding parameter $b = 2$ that computes an MM for any n -node graph in time $O(\log^2 n)$.*

Our first step towards establishing Theorem 4.16 is to observe that the synchronizer designed in Section 3.1 can be applied to extended nFSM protocols too (following essentially the same construction). Moreover, the line of arguments that leads to Theorem 3.4 can easily be adjusted to fit the extended nFSM model as well. Therefore, we subsequently assume a locally synchronous environment and use multiple-letter queries.

Similar to the MIS protocol presented in Section 4.1, in the context of the MM protocol we also distinguish between those nodes which are already in output states (actually, we only need one output state), referred to as *inactive*, and the rest of the nodes, referred to as *active*. Given some active node $v \in V$, we refer to the ports $\psi_v(u)$ of v corresponding to its active (respectively, *inactive*) neighbors u as *active ports* (resp., *inactive ports*). The protocol is designed so that an active node can distinguish between its active and inactive ports by the letters they store.

The phases. The execution of the active nodes proceeds in *phases* indexed by the non-negative integers $j \in \mathbb{Z}_{\geq 0}$ and each active node maintains the value of $j \bmod 5$. The phases are classified into the following 5 types executed one after the other in a round-robin fashion:

- **Selecting gender:** phases $j = 0 \bmod 5$.

Active node v chooses to be either a *male* or a *female* uniformly at random and independently of all other random choices and announces its choice.

Postcondition: port $\psi_v(u)$ of v is classified as either a *male port* or a *female port* according to the gender of u .

- **Proposing:** phases $j = 1 \pmod 5$.

Male v chooses some female $p(v)$ among its female neighbors uniformly at random and independently of all other random choices; for convenience, if v has no female neighbors, then set $p(v) = \text{null}$.

Postcondition: port $\psi_v(p(v))$ is designated.

- **Collecting proposals:** phases $j = 2 \pmod 5$.

Female u learns the set $P(u)$ of males $v \in N(u)$ such that $p(v) = u$.

Postcondition: the ports of u corresponding to $P(u)$ are designated.

- **Accepting:** phases $j = 3 \pmod 5$.

If $P(u) \neq \emptyset$, then female u chooses some (arbitrary) $q(u) \in P(u)$.

Postcondition: port $\psi_u(q(u))$ is designated.

- **Matching:** phases $j = 4 \pmod 5$.

Male v learns whether $q(p(v)) = v$, in which case edge $(v, p(v))$ joins the MM.

Postcondition: male v and female $p(v)$ become inactive.

The protocol is designed so that upon completion of phase j , active node v transmits a designated *completed-phase- $(j \pmod 5)$* message and does not start phase $j+1$ before receiving completed-phase- $(j \pmod 5)$ messages from all its neighbors. Once the execution of phase $j+1$ begins, no further delays occur until the phase is completed. This is in contrast to the delaying mechanism introduced in Section 4.1 for the MIS protocol, where each individual step of the execution may be delayed due to the states of the neighboring nodes.

Implementation. We now turn to describe the implementation of the 5 phases under the extended nFSM model. The gender selection phase is trivial to implement as it only requires node v to locally toss a fair coin and announce its outcome. Implementing the proposing phase is slightly more challenging since the (extended) nFSM model does not provide male v with the power to explicitly choose one random port. Instead, the female ports will compete with each other over the right to be chosen: The female ports $\psi_v(u)$ of male v are classified either as *candidate* ports or as *non-candidate* ports, where at the beginning of the phase, all female ports are candidates. In every round of the proposing phase, a candidate port remains a candidate w.p. $1/2$; otherwise, it turns into a non-candidate. The proposal phase of v is completed once v has exactly 1 candidate port $\psi_v(u)$, in which case we set $p(v) = u$. To ensure that this indeed occurs, if some round during v 's proposal phase starts with at least 2 candidate ports and all of them turn into non-candidates, then the outcome of this round is revoked (this is implemented by introducing an intermediate temporary internal stage between candidate and non-candidate).

Observation 4.17. *Upon completion of the proposing phase, variable $p(v)$ of male v is chosen uniformly at random from the set of female neighbors of v . Moreover, the proposing phase lasts $O(\log n)$ rounds in expectation and w.h.p.*

The implementation of the proposals collecting phase is perhaps the most interesting part of

the MM protocol. Although this phase is executed by the females, the males also take an active part in it. The male ports of a female node u are classified either as *valid* or as *invalid*, where at the beginning of the phase, all male ports are valid. Likewise, the female ports of a male node v other than port $\psi_v(p(v))$ are classified either as *informed* or as *uninformed*, where at the beginning of the phase, all female ports other than $\psi_v(p(v))$ are uninformed. In every odd round $2t - 1$ of the proposals collecting phase, female u chooses some bit $b_u \in \{0, 1\}$ uniformly at random and transmits it. In the subsequent even round $2t$, every male v such that $p(v) = u$, transmits the same bit $b_v = b_u$.

Consider some male v and let $u = p(v)$ be the female chosen by v in the proposing phase. Any female $u' \in N(v)$ such that $b_{u'} \neq b_u$ realizes after round $2t$ that $p(v) \neq u'$ since $b_v \neq b_{u'}$, hence she marks port $\psi_{u'}(v)$ as invalid. Moreover, observing that $b_{u'} \neq b_u$, male v knows in round $2t$ that u' will soon reach this conclusion, thus he marks port $\psi_v(u')$ as informed.

Once all ports $\psi_v(u')$, $u' \neq p(v)$, are informed, male v transmits a designated *all-informed* message. When female u receives an all-informed message from male v such that port $\psi_u(v)$ is valid, she adds v to $P(u)$ (by storing some designated letter in $\psi_u(v)$). From the perspective of female u , the proposals collecting phase is completed once all her male ports either correspond to nodes in $P(u)$ or are invalid.

Observation 4.18. *Upon completion of the proposals collecting phase, the set $P(u)$ maintained by a female node u consists of every male $v \in N(u)$ such that $p(v) = u$. Moreover, the proposing phase lasts $O(\log n)$ rounds in expectation and w.h.p.*

The implementation of the accepting phase is identical to that of the proposing phase, only that this time, female u chooses some male in $v \in P(u)$. Actually, male v is chosen uniformly at random from $P(u)$, but this is an artifact of the implementation and will not be needed in the analysis.

Observation 4.19. *Upon completion of the accepting phase, if $P(u) \neq \emptyset$, then $q(u)$ is set to be some male in $P(u)$. Moreover, the accepting phase lasts $O(\log n)$ rounds in expectation and w.h.p.*

Finally, the matching phase is implemented by using the technique employed in the implementation of the proposals collecting phase with reversed roles: male v chooses a random bit b_v in the odd rounds and female u transmits back the same bit $b_u = b_v$ in the subsequent even round if $v = q(u)$.

Observation 4.20. *Upon completion of the matching phase, male v with $p(v) = u$ knows whether $q(u) = v$, in which case the edge (u, v) joins the matching. Moreover, the matching phase lasts $O(\log n)$ rounds in expectation and w.h.p.*

Analysis. Given some $j \in \mathbb{Z}_{\geq 0}$, let $T(j)$ denote the first time at which every active node is in some phase $j' \geq j$. Combining the run-time bounds promised in Observations 4.17, 4.18, 4.19, and 4.20, we conclude that

$$T(j + 1) \leq T(j) + O(\log n),$$

hence, by induction on j ,

$$T(j) \leq j \cdot O(\log n). \quad (4)$$

Fix some $i \in \mathbb{Z}_{\geq 0}$. Similarly to the analysis of the MIS protocol (see Section 4.1), let V_i be the set of nodes that participate in the gender selection phase $5i$, namely, V_i consists of all nodes that did not become inactive in some phase $j < 5i$. Let $G_i = (V_i, E_i)$ be the subgraph induced on G by V_i , where $E_i = E \cap (V_i \times V_i)$. Given some node $v \in V_i$, let $N_i(v) = \{u \in V_i \mid (u, v) \in E_i\}$ be the neighborhood of node v in G_i and let $d_i(v) = |N_i(v)|$ be its degree. The main lemma in our analysis is the following MM version of Lemma 4.4.

Lemma 4.21. *Our MM protocol guarantees the existence of two constants $0 < p, c < 1$ such that $|E_{i+1}| \leq c|E_i|$ w.p. $\geq p$.*

Proof. Similarly to the proof of Lemma 4.4, we say that node $u \in V_i$ is *good* in G_i if

$$|\{v \in N_i(u) \mid d_i(v) \leq d_i(u)\}| \geq d_i(u)/3.$$

Consider some good node u in G_i and let $I(u)$ denote the event that u becomes inactive in matching phase $5i + 4$, that is, u will not be included in V_{i+1} . Our goal is to bound $\mathbb{P}(I(u))$ from below. To that end, let $I'(u)$ denote the event that u is a female and one of its male neighbors v chooses $p(v) = u$ in the proposing phase $5i + 1$. Observe that this implies that u is bound to accept some proposal in the accepting phase $5i + 3$, thus $\mathbb{P}(I(u)) \geq \mathbb{P}(I'(u))$.

Let $F(u)$ denote the event that u chooses to be a female in the gender selection phase $5i$. Fix $d = d_i(u)$ and let $\widehat{N}_i(u) = \{v \in N_i(u) \mid d_i(v) \leq d\}$. Recall that the definition of a good node implies that $|\widehat{N}_i(u)| \geq d/3$. Given some $v \in \widehat{N}_i(u)$, let $C(v, u)$ denote the event that node v chooses to be a male in the gender selection phase $5i$ and chooses $p(v) = u$ in the proposing phase $5i + 1$, conditioned on $F(u)$.

The analysis would have been simplified if we could have assumed that the events $C(v, u)$, $v \in \widehat{N}_i(u)$, are independent. Unfortunately, this is not necessarily true. To overcome this obstacle, we consider a slightly modified protocol, where in the proposing phase, male v chooses $p(v)$ uniformly at random among all its active neighbors (males and females). Clearly, this can only decrease the probability of event $I'(u)$ as each neighboring male has a larger set of nodes to choose from. The key observation now is that under this modified protocol, events $C(v, u)$, $v \in \widehat{N}_i(u)$, are indeed independent. Moreover, $\mathbb{P}(C(v, u)) = \frac{1}{d_i(v)} \cdot \frac{1}{2} \geq \frac{1}{2d}$ for every $v \in \widehat{N}_i(u)$, hence

$$\mathbb{P}\left(\bigwedge_{v \in \widehat{N}_i(u)} \neg C(v, u)\right) \leq \left(1 - \frac{1}{2d}\right)^{d/3} \leq e^{-1/6},$$

which implies that

$$\mathbb{P}(I'(u)) = \mathbb{P}(I'(u) \mid F(u)) \cdot \mathbb{P}(F(u)) \geq \frac{1 - e^{-1/6}}{2}.$$

Employing Lemma 4.5, we conclude that $\mathbb{E}[|E_{i+1}|] < \left(1 - \frac{1-e^{-1/6}}{4}\right) |E_i|$, which establishes the assertion by Markov’s inequality. \square

Similarly to the analysis carried out in Section 4.1, we define the random variable $Y = \min\{i \in \mathbb{Z}_{\geq 0} : |E_i| = 0\}$. Lemma 4.21 implies that Y is stochastically dominated by a random variable that obeys distribution $\text{NB}(O(\log n), 1 - p) + O(\log n)$, hence $Y = O(\log n)$ in expectation and w.h.p. Theorem 4.16 follows due to inequality (4).

5 Computability

Up until now, we focused on the run-time complexity of distributed problems under the nFSM model. In this section, we ask ourselves what can be computed under this model regardless of run-time considerations. As nodes under the nFSM model do not possess unique identifiers, impossibility results that hold for the anonymous message passing model immediately imply impossibility under the nFSM model. These impossibility results include some famous distributed problems such as *leader election* [6] and *consensus* [26].

Next, we turn to the following question: What is the computational power of a network under the nFSM model compared to a central (sequential) algorithm? Answering this question requires defining some computational models. A *deterministic linear bounded automaton (dLBA)* is a (deterministic) Turing machine whose working tape is restricted to the cells specifying the input (this is equivalent to a $\text{DSPACE}(O(n))$ Turing machine). A *non-deterministic linear bounded automaton*, a.k.a., *linear bounded automaton (LBA)*, is the non-deterministic version of a dLBA, and a *randomized linear bounded automaton (rLBA)* is the randomized version. Kuroda [34] proved that the class of languages that can be decided by an LBA is exactly the context-sensitive languages, corresponding to the Type-1 grammars in Chomsky’s hierarchy of formal languages [23]. Whether LBAs are equivalent to dLBAs (a.k.a. the first LBA problem) and where exactly do rLBAs lie between the two are major open questions in computational complexity. The following three observations compare the computational power of an nFSM protocol to that of an rLBA (cf. [8]).

Observation 5.1. *An nFSM protocol on a graph G of arbitrary topology can be simulated by an rLBA.*

Proof. The input for the Turing machine is the graph G , given as an adjacency list. In order to simulate the execution of the nFSM protocol, we store some additional information in the entries of the adjacency list as follows: For each node v , we store its current state and the next letter it transmits. For every node u in the list of neighbors $N(v)$ attached to v , we store the entry of u ’s port that corresponds to v . In each round of the nFSM protocol, the rLBA performs two sweeps of the list of nodes: The first sweep serves to calculate v ’s next state q and transmitted letter σ for all nodes v , based on v ’s current state and the messages in its ports, according to the

nFSM state machine, which is hard-wired in the rLBA. However, the calculated letter σ is not being “transmitted” yet, so the calculations for subsequent nodes in the list are not messed up, but rather stored in the corresponding place next to v . In the second sweep, for every node v , the letter σ is being “transmitted”, that is, the lists of neighbors are traversed, and at each appearance of v , the current letter is replaced by σ . This way, we simulate every round of the nFSM protocol. In total, our simulation requires additional $O(1)$ space per node and $O(1)$ space per edge, hence it can be implemented with an rLBA. The assertion follows. \square

Observation 5.2. *An rLBA can be simulated by an nFSM protocol on a path.*

Proof. Let n be the number of cells in the tape of the rLBA. Then, the path network has n nodes, each corresponding to one cell of the tape, i.e., we identify a node v of the path nFSM with a certain cell on the tape. Let Γ be the working alphabet and P be the state space of the rLBA. The nFSM protocol is designed so that the state of node v indicates: (1) which letter from Γ is written in v ; (2) if the head of the rLBA currently points to v ; (3) the current state of the rLBA, which is allowed to be incorrect if (2) is false; and (4) if the head is currently located to the left or to the right of v . Hence, we fix $Q = \Gamma \times \{0, 1\} \times P \times \{L, R\}$. The alphabet of the nFSM is $\Sigma = \{L, R\} \times P$.

Suppose that the input to the rLBA is $\gamma_1 \dots \gamma_n \in \Gamma^n$. Then, we assume that the initial state of the i th node in the path is (γ_i, h, p_0, L) , where p_0 is the initial state of the Turing machine and

$$h = \begin{cases} 1 & \text{if } i = 1 \\ 0 & \text{if } i > 1. \end{cases}$$

Note that the distinction between the initial state of the first node in the path and the initial states of all other nodes is without loss of generality. Indeed, as the first and last nodes have degree 1 and all interior nodes have degree 2, it is easy for a node to “decide” (under the nFSM model) if it is an interior node. Distinguishing between the first and last nodes is unavoidable if one wants to distinguish between the inputs $\gamma_1 \dots \gamma_n$ and $\gamma_n \dots \gamma_1$.

At all times, we maintain the invariant that exactly one node is in a state in $\Gamma \times \{1\} \times P \times \{L, R\}$ — denote this node as *active* — whereas all other nodes are in a state in $\Gamma \times \{0\} \times P \times \{L, R\}$. Only the active node can transmit messages; all other nodes remain silent and listen. If a non-active node v receives a message indicating that the head should move to the left (respectively, right), and v ’s state indicates that the head is currently to its right (resp., left), then v becomes the active node; otherwise, v does not react to this message. Now, the nodes simulate the behavior of the rLBA by calculating the next state of the rLBA based on the rLBA’s transition function (which is hard-wired in the FSM) and updating their own states accordingly. The assertion follows. \square

Observation 5.3. *Every property of bounded degree graphs that can be computed by an rLBA can also be computed by an nFSM protocol with eventually converging correctness.*

Proof. Follows immediately from Lemma 4.15 due to known results stating that every property of bounded degree graphs that can be computed by an rLBA can also be computed by a population protocol [7, 10]. \square

References

- [1] H. Abelson, D. Allen, D. Coore, C. Hanson, G. Homsy, T. F. Knight, Jr., R. Nagpal, E. Rauch, G. J. Sussman, and R. Weiss. Amorphous computing. *Commun. ACM*, 43(5):74–82, May 2000.
- [2] Y. Afek, N. Alon, Z. Bar-Joseph, A. Cornejo, B. Haeupler, and F. Kuhn. Beeping a maximal independent set. In *DISC*, pages 32–50, 2011.
- [3] Y. Afek, N. Alon, O. Barad, E. Hornstein, N. Barkai, and Z. Bar-Joseph. A Biological Solution to a Fundamental Distributed Computing Problem. *Science*, 331(6014):183–185, Jan. 2011.
- [4] I. F. Akyildiz, J. M. Jornet, and M. Pierobon. Nanonetworks: a new frontier in communications. *Commun. ACM*, 54(11):84–89, Nov. 2011.
- [5] N. Alon, L. Babai, and A. Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *J. Algorithms*, 7:567–583, December 1986.
- [6] D. Angluin. Local and global properties in networks of processors (extended abstract). In *STOC*, pages 82–93, New York, NY, USA, 1980. ACM.
- [7] D. Angluin, J. Aspnes, M. Chan, M. J. Fischer, H. Jiang, and R. Peralta. Stably computable properties of network graphs. In *DCOSS*, pages 63–74, 2005.
- [8] D. Angluin, J. Aspnes, Z. Diamadi, M. J. Fischer, and R. Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed Computing*, pages 235–253, mar 2006.
- [9] D. Angluin, J. Aspnes, and D. Eisenstat. Fast computation by population protocols with a leader. *Distributed Computing*, 21(3):183–199, sep 2008.
- [10] J. Aspnes and E. Ruppert. An introduction to population protocols. In B. Garbinato, H. Miranda, and L. Rodrigues, editors, *Middleware for Network Eccentric and Mobile Applications*, pages 97–120. Springer-Verlag, 2009.
- [11] B. Awerbuch. Complexity of network synchronization. *J. ACM*, 32(4):804–823, 1985.
- [12] B. Awerbuch, B. Patt-Shamir, D. Peleg, and M. E. Saks. Adapting to asynchronous dynamic networks (extended abstract). In *STOC*, pages 557–570, 1992.
- [13] B. Awerbuch and D. Peleg. Network synchronization with polylogarithmic overhead. In *FOCS*, pages 514–522, 1990.
- [14] L. Barenboim and M. Elkin. Distributed $(\delta+1)$ -coloring in linear (in δ) time. In *STOC*, pages 111–120, 2009.
- [15] L. Barenboim and M. Elkin. Deterministic distributed vertex coloring in polylogarithmic time. *J. ACM*, 58(5):23, 2011.

- [16] L. Barenboim, M. Elkin, S. Pettie, and J. Schneider. The locality of distributed symmetry breaking. *CoRR*, abs/1202.1983, 2012.
- [17] M. Bellare, O. Goldreich, and M. Sudan. Free bits, pcps, and nonapproximability-towards tight results. *SIAM J. Comput.*, 27(3):804–915, 1998.
- [18] Y. Benenson. Biomolecular computing systems: principles, progress and potential. *Nat Rev Genet*, 13(7):455–468, July 2012.
- [19] Y. Benenson, T. Paz-Elizur, R. Adar, E. Keinan, Z. Livneh, and E. Shapiro. Programmable and autonomous computing machine made of biomolecules. *Nature*, 414(6862):430–434, Nov. 2001.
- [20] D. Brand and P. Zafiropulo. On communicating finite-state machines. *J. ACM*, 30:323–342, April 1983.
- [21] I. Chatzigiannakis and P. G. Spirakis. The dynamics of probabilistic population protocols. In *DISC*, DISC '08, pages 498–499, Berlin, Heidelberg, 2008. Springer-Verlag.
- [22] I. Chlamtac and S. Kutten. On Broadcasting in Radio Networks—Problem Analysis and Protocol Design. *Communications, IEEE Transactions on [legacy, pre - 1988]*, 33(12):1240–1246, 1985.
- [23] N. Chomsky. Three models for the description of language. *IRE Transactions on Information Theory*, 2:113–124, 1956. <http://www.chomsky.info/articles/195609-.pdf>.
- [24] R. Cole and U. Vishkin. Deterministic coin tossing with applications to optimal parallel list ranking. *Inf. Control*, 70(1):32–53, July 1986.
- [25] A. Cornejo and F. Kuhn. Deploying wireless networks with beeps. In *DISC*, pages 148–162, 2010.
- [26] Y. Emek, J. Seidel, and R. Wattenhofer. Distributed computability: Anonymity, revocability, and randomization. A manuscript.
- [27] R. Flury and R. Wattenhofer. Slotted Programming for Sensor Networks. In *IPSN*, April 2010.
- [28] M. Gardner. The fantastic combinations of John Conway’s new solitaire game ‘life’. *Scientific American*, 223(4):120–123, 1970.
- [29] P. Gordon. Numerical Cognition Without Words: Evidence from Amazonia. *Science*, 306(5695):496–499, Oct. 2004.
- [30] A. Israeli and A. Itai. A fast and simple randomized parallel algorithm for maximal matching. *Inf. Process. Lett.*, 22(2):77–80, 1986.

- [31] K. Kothapalli, C. Scheideler, M. Onus, and C. Schindelhauer. Distributed Coloring in $\tilde{O}(\sqrt{\log n})$ Bit Rounds. In *IPDPS*, 2006.
- [32] F. Kuhn. Weak graph colorings: distributed algorithms and applications. In *SPAA*, pages 138–144, New York, NY, USA, 2009. ACM.
- [33] F. Kuhn, T. Moscibroda, and R. Wattenhofer. What cannot be computed locally! In *PODC*, pages 300–309, 2004.
- [34] S.-Y. Kuroda. Classes of languages and linear-bounded automata. *Information and Control*, 7(2):207–223, 1964.
- [35] C. Lenzen and R. Wattenhofer. MIS on trees. In *PODC*, pages 41–48, New York, NY, USA, 2011.
- [36] N. Linial. Locality in distributed graph algorithms. *SIAM J. Comput.*, 21:193–201, Feb. 1992.
- [37] M. Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM J. Comput.*, 15:1036–1055, November 1986.
- [38] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1st edition, 1996.
- [39] Y. Métivier, J. M. Robson, N. Saheb-Djahromi, and A. Zemmari. About randomised distributed graph colouring and graph partition algorithms. *Inf. Comput.*, 208(11):1296–1304, Nov. 2010.
- [40] Y. Métivier, J. M. Robson, N. Saheb-Djahromi, and A. Zemmari. An optimal bit complexity randomised distributed MIS algorithm. *Distributed Computing*, 23(5-6):331–340, Jan. 2011.
- [41] O. Michail, I. Chatzigiannakis, and P. G. Spirakis. *New Models for Population Protocols*. Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool Publishers, 2011.
- [42] K. Nakamura. Asynchronous cellular automata and their computational ability. *Syst Comput Controls*, 5(5):58–66, 1974.
- [43] C. L. Nehaniv. Asynchronous automata networks can emulate any synchronous automata network. *Journal of Algebra*, pages 1–21, Dec. 2003.
- [44] D. Peleg. *Distributed computing: a locality-sensitive approach*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000.
- [45] D. Sadava. *Life: The Science of Biology*. Sinauer Associates, 2011.
- [46] J. Schneider and R. Wattenhofer. An optimal maximal independent set algorithm for bounded-independence graphs. *Distributed Computing*, 22(5-6):349–361, 2010.

- [47] J. Schneider and R. Wattenhofer. Distributed Coloring Depending on the Chromatic Number or the Neighborhood Growth. In *SIROCCO*, June 2011.
- [48] J. Suomela. Survey of local algorithms. *To appear in ACM Computing Surveys*, 2012. <http://www.cs.helsinki.fi/u/josuomel/doc/local-survey.pdf>.
- [49] M. Szegedy and S. Vishwanathan. Locality based graph coloring. In *STOC*, pages 201–207, 1993.
- [50] L. G. Valiant. Parallel computation. In *7th IBM Symp. on Math. Foundations of Computer Science*, 1982.
- [51] J. von Neumann. *Theory of Self-Reproducing Automata*. University of Illinois Press, Champaign, IL, USA, 1966.
- [52] S. Wolfram. *A new kind of science*. Wolfram Media, Champaign, Illinois, 2002.