

# Constant Space and Non-Constant Time in Distributed Computing

Tuomo Lempiäinen and Jukka Suomela

Department of Computer Science, Aalto University, Finland

## Abstract

While the relationship of time and space is an established topic in traditional centralised complexity theory, this is not the case in distributed computing. We aim to remedy this by studying the time and space complexity of algorithms in a weak message-passing model of distributed computing. While a constant number of communication rounds implies a constant number of states visited during the execution, the other direction is not clear at all. We consider several graph families and show that indeed, there exist non-trivial graph problems that are solvable by constant-space algorithms but that require a non-constant running time. This provides us with a new complexity class for distributed computing and raises interesting questions about the existence of further combinations of time and space complexity.

## 1 Introduction

This work studies the relationship between time and space complexity in distributed computing. While the relationships between various time complexity classes have been studied a lot in many of the established works of the field [11, 13, 15, 16] and also very recently [3, 4, 5, 6, 10], space complexity has not received that much attention. We aim to remedy that by establishing the existence of a new complexity class: constant-space with non-constant running time.

### 1.1 Distributed computing

We consider a model of computation where each node of a graph is a computational unit. The same graph is a communication network and an input to the algorithm. Adjacent nodes communicate with each other in synchronous rounds, and eventually each node outputs its own part of the output. All the nodes run the same deterministic algorithm. In our model, the nodes are anonymous, that is, they do not have access to unique identifiers, and furthermore, they cannot distinguish between their neighbours—they broadcast the same message to everyone and they receive the messages in a set.

We define running time as the number of communication rounds until all the nodes have halted, while space usage is defined as the number of bits per node needed to represent all the states of the algorithm. The complexity measures are considered as a function of  $n$ , the number of nodes in the graph. The amount of local computation is not limited in any way, nor is the size of messages. More formally, distributed algorithms can be defined as state machines; see Section 2.1 for the definitions.

### 1.2 Contributions

We prove the existence of graph problems, and consequently algorithms, that exhibit a constant space usage but non-constant running time. Our first theorem, proved in Section 3 shows that when considering a graph family containing graphs of maximum degree 3, there exists a graph

problem solvable by an algorithm with a constant number of states, while the time complexity of the problem is in  $\Omega(\log n)$ .

In our main theorem presented in Section 4, we turn our attention to the case of the class of graphs of maximum degree 2. While more non-trivial, it turns out that also in this case we can establish the existence of a graph problem solvable in constant space but requiring a linear number of communication rounds. Our result makes use of the properties of the Thue–Morse sequence, which we introduce in Section 2.5.

We emphasise that while it is straightforward to come up with such problems if we have a promise that the input graph is for example a path, we do not need to make such assumptions. It is also essential that we require the algorithms to halt in all nodes of any input graph. Again, this kind of results would be easy to achieve, if nodes were allowed to continue running indefinitely on some input graphs.

### 1.3 Motivation and related work

Traditionally, the focus of distributed computing has been on the time complexity of algorithms, and occasionally on message complexity, but space complexity has not been studied much. We aim to change this by introducing space complexity, constant space in particular, as a new dimension in the classification of computational problems.

While space complexity is especially interesting from the theory perspective, it can be argued that constant space complexity is a reasonable assumption also from the practical point of view. Networked devices often need to be able to operate in a network of any size, while it is not necessarily easy to expand their memory resources afterwards. On the other hand, also nature provides us with plenty of phenomena that exhibit distributed behaviour. Natural organisms are usually of constant size, independent of the size of the swarm or flock they are part of. Therefore, the study of constant space may result in new avenues for applying distributed computing to advance the understanding of nature.

The line of research on anonymous models of distributed computing was initiated by Angluin [2], who introduced the well-known port numbering model. Our model of computation can be seen as a further restriction of the port-numbering model, with port numbers stripped out. While port numbers are a natural assumption in wired networks, the weaker variant makes more sense when applying distributed computing to a wireless setting. While our model has not been studied that much in prior work, a so-called *beeping model* [1, 7] is essentially similar. In the hierarchy of seven models defined by Hella et al. [12], our model is the weakest one; they call it the SB model. On the other hand, the case where nodes receive messages in a multiset instead of a set is discussed more often in the literature [12].

From the constant-space point of view, our setting bears similarities to the field of cellular automata [9, 17, 18]. In a cellular automaton, each cell can be in one of a constant number of states, and each cell updates its state synchronously using the same rule. However, in the case of cellular automata, one is usually interested in the kind of patterns an automata converges in, while we require each node to eventually stop and produce an output.

On the side of distributed computing, Emek and Wattenhofer [8] have considered a model where the network consists of finite state machines—hence making the space complexity constant. However, their model is asynchronous and randomised, while we study a fully synchronous and deterministic setting.

One of the main ingredients of our work, the Thue–Morse sequence, was used previously by Kuusisto [14], who proved that there exists a distributed algorithm that always halts in the class of graphs of maximum degree two but features a non-constant running time. However, his algorithm has also a non-constant space complexity—this is where our work provides a significant improvement.

## 2 Preliminaries

In this section we define our model of computation and introduce notions needed later in the proofs.

### 2.1 Model of computation

Our model of computation is a weaker variant of the standard LOCAL model, with the following restrictions:

- Nodes are anonymous, that is, they do not have unique identifiers.
- Nodes broadcast the same message to all their neighbours.
- Nodes receive the messages in a set, that is, they do not know which neighbour sent which message or how many identical messages they received.

In the following, we define distributed algorithms as state machines. Given an input graph, each node of the graph is equipped with an identical state machine. Machines in adjacent nodes can communicate with each other. In this work, we study only deterministic state machines and synchronous communication. The graph is always assumed to be simple, finite, connected and undirected, unless stated otherwise.

At the beginning, each state machine is only aware of its own local input (taken from some fixed finite set) and the degree of the node on which it sits. Then, computation is executed in synchronous rounds. In each round, each machine

- (1) broadcasts a message to its neighbours,
- (2) receives a set of messages from its neighbours,
- (3) moves to a new state based on the received messages and its previous state.

Each machine is required to eventually reach one of special halting states and stop execution. The local output is then the state of the node at the time of halting.

Note that while our model of computation is rather weak, it only makes our results stronger by limiting the capabilities of algorithms. Unique identifiers or unrestricted local inputs would not make much sense in the constant-space setting. Crucially, we require nodes to always halt in any input graph—nodes being allowed to run indefinitely or having the ability to continue passing messages after announcing an output would make engineering constant-space non-constant-time problems quite straightforward.

Next, we give a more formal definition of the model of computation used in this work by defining algorithms and graph problems.

### 2.2 Notation and terminology

For  $k \in \mathbb{N}$ , we denote by  $[k]$  the set  $\{1, 2, \dots, k\}$ . Given a graph  $G = (V, E)$ , the set of neighbours of a node  $v \in V$  is denoted by  $N(v) = \{u \in V : \{v, u\} \in E\}$ . For  $r \in \mathbb{N}$ , the *radius- $r$  neighbourhood* of a node  $v \in V$  is  $\{u \in V : \text{dist}(u, v) \leq r\}$ , that is, the set of nodes  $u$  such that there is a path of length at most  $r$  between  $v$  and  $u$ . We call any induced subgraph of  $G$  that contains node  $v$  simply a *neighbourhood* of  $v$ .

We will also work with strings of letters. Given a finite alphabet  $\Sigma$ , we denote elements of the alphabet by lowercase symbols such as  $x \in \Sigma$  or  $y \in \Sigma$ . On the other hand, *words* (finite sequences of letters) are denoted by uppercase symbols, for example  $X = x_1x_2 \dots x_i \in \Sigma^*$ . Given words  $X = x_1x_2 \dots x_i$  and  $Y = y_1y_2 \dots y_j$ , we write their concatenation simply  $XY = x_1x_2 \dots x_iy_1y_2 \dots y_j$ . A word  $X$  is a *subword* of  $Y$  if  $Y = Y_1XY_2$  for some (possibly empty)

words  $Y_1$  and  $Y_2$ . We identify words of length 1 with the letter they consist of. For any letter  $x$  and  $i \in \mathbb{N}$ ,  $x^i$  denotes the word consisting of  $i$  consecutive letters  $x$ . We say that a word  $X$  is of the form  $x^+$  if  $X = x^i$  for some  $i \in \mathbb{N}_+$ .

### 2.3 Algorithms as state machines

Let  $G = (V, E)$  be a graph. An *input* for  $G$  is a function  $f: V \rightarrow I$ , where  $I$  is a finite set. For each node  $v \in V$ , we call  $f(v) \in I$  the *local input* of  $v$ .

A *distributed state machine* is a tuple  $\mathcal{A} = (S, H, \sigma_0, M, \mu, \sigma)$ , where

- $S$  is a set of states,
- $H \subseteq S$  is a finite set of halting states,
- $\sigma_0: \mathbb{N} \times I \rightarrow S$  is an initialisation function,
- $M$  is a set of possible messages,
- $\mu: S \rightarrow M$  is a function that constructs the outgoing messages,
- $\sigma: S \times \mathcal{P}(M)$  is a function that defines the state transitions, so that  $\sigma(h, \mathcal{M}) = h$  for each  $h \in H$  and  $\mathcal{M} \in \mathcal{P}(M)$ .

Given a graph  $G$ , and input  $f$  for  $G$  and a distributed state machine  $\mathcal{A}$ , the *execution* of  $\mathcal{A}$  on  $(G, f)$  is defined as follows. The state of the system in round  $r \in \mathbb{N}$  is a function  $x_r: V \rightarrow S$ , where  $x_r(v)$  is the state of node  $v$  in round  $r$ . To begin the execution, set  $x_0(v) = \sigma_0(\deg(v), f(v))$  for each node  $v \in V$ . Then, let  $A_{r+1}(v) = \{\mu(x_r(u)) : u \in N(v)\}$  denote the set of messages received by node  $v$  in round  $r + 1$ . Now the new state of each node  $v \in V$  is defined by setting  $x_{r+1}(v) = \sigma(x_r(v), A_{r+1}(v))$ .

The *running time* of  $\mathcal{A}$  on  $(G, f)$  is the smallest  $t \in \mathbb{N}$  for which  $x_t(v) \in H$  holds for all  $v \in V$ . The output of  $\mathcal{A}$  on  $(G, f)$  is then  $x_t: V \rightarrow H$ , where  $t$  is the running time, and for each  $v \in V$ , the *local output* of  $v$  is  $x_t(v)$ . The *space usage* of  $\mathcal{A}$  on  $(G, f)$  is defined as

$$\lceil \log |\{s \in S : s = x_r(v) \text{ for some } r \in \mathbb{N}, v \in V\}| \rceil,$$

that is, the number of bits needed to encode all the states that the state machine visits in at least one node of  $G$  during the execution. In case the execution does not halt, the running time can be defined to be  $\infty$ , and if the number of visited states grows arbitrarily large, we take the space usage to also be  $\infty$ .

From now on, we will use the terms *algorithm*  $\mathcal{A}$  and *distributed state machine*  $\mathcal{A}$  interchangeably, implying that each algorithm can be defined formally as a state machine.

### 2.4 Graph problems

The computational problems that we consider are graph problems with local input—that is, the problem instance is identical to the communication network, but possibly with an additional input value given to each node.

More formally, let  $I$  and  $O$  be finite sets. A *graph problem* is a mapping  $\Pi_{I,O}$  that maps each graph  $G = (V, E)$  and input  $f: V \rightarrow I$  to a set  $\Pi_{I,O}(G, f)$  of valid solutions. Each solution  $S$  is a function  $S: V \rightarrow O$ . In case of *decision graph problems*, we set  $O = \{\text{yes}, \text{no}\}$ .

If  $\Pi_{I,O}$  is a graph problem,  $T, U: \mathbb{N} \rightarrow \mathbb{N}$  are functions and  $\mathcal{A}$  a distributed state machine, we say that  $\mathcal{A}$  *solves*  $\Pi_{I,O}$  *in time*  $T$  *and in space*  $U$  if for each graph  $G = (V, E)$  and each input  $f: V \rightarrow I$  we have that the running time of  $\mathcal{A}$  on  $(G, f)$  is at most  $T(|V|)$ , the space usage of  $\mathcal{A}$  on  $(G, f)$  is at most  $U(|V|)$  and the output of  $\mathcal{A}$  on  $(G, f)$  is in the set  $\Pi_{I,O}(G, f)$ . In that case, we also say that the *time complexity* of algorithm  $\mathcal{A}$  is  $T$  and the *space complexity* of  $\mathcal{A}$  is  $U$ .

## 2.5 Thue–Morse sequence

In this section we present a concept that will be central in the proof of our main result in Section 4. The Thue–Morse sequence is the infinite binary sequence defined recursively as follows:

**Definition 1.** The *Thue–Morse* sequence is the sequence  $(t_i)$  satisfying  $t_0 = 0$ , and for each  $i \in \mathbb{N}$ ,  $t_{2i} = t_i$  and  $t_{2i+1} = 1 - t_i$ .

Thus, the beginning of the Thue–Morse sequence is

$$01101001100101101001\dots$$

For our purposes, the following two recursive definitions will be very useful.

**Definition 2.** Let  $T_0 = 0$ . For each  $i \in \mathbb{N}_+$ , let  $T_i = T_{i-1}\mathcal{C}(T_{i-1})$ , where  $\mathcal{C}$  denotes the Boolean complement.

Note that for each  $i \in \mathbb{N}$ , the word  $T_i$  is the prefix of length  $2^i$  of the Thue–Morse sequence.

**Definition 3.** Let  $T'_0 = 0$ . For each  $i \in \mathbb{N}_+$ , let  $T'_i$  be obtained from  $T'_{i-1}$  by substituting each occurrence of 0 with 01 and each occurrence of 1 with 10.

Again, we have  $T'_0 = 0$ ,  $T'_1 = 01$ ,  $T'_2 = 0110$ ,  $T'_3 = 01101001$  and so on. A straightforward induction shows that the above two definitions are equivalent:  $T_i = T'_i$  for each  $i \in \mathbb{N}$ . We call  $T_i$  the *Thue–Morse word of length  $2^i$* .

The Thue–Morse sequence contains lots of squares, that is, subwords of the form  $XX$ , where  $X \in \{0, 1\}^*$ . Interestingly, it does not contain any cubes—subwords of the form  $XXX$ . Note also that for each  $i \in \mathbb{N}_+$ ,  $T_{2i}$  is a palindrome.

## 3 Warm-up: graphs of maximum degree 3

In this section we present a graph problem that exhibits the constant space and non-constant time complexity, in case we do not restrict ourselves to paths and cycles. The proof is quite straightforward, which emphasises the fact that the degree-2 case considered in Section 4 is the most interesting one. We emphasise that in the following theorem, we do not need to make any additional assumptions about the graph; the described algorithm halts in all finite input graphs.

**Theorem 4.** *There exist a graph problem with constant space complexity and  $\Theta(\log n)$  time complexity in the class of graphs of maximum degree at most 3.*

*Proof.* Consider the following transformation: given a *directed* graph  $G' = (V', E')$ , replace each edge  $e = (u, v)$  by the gadget represented in Figure 1. The end result is an undirected graph  $G = (V, E)$ , where the gadgets encode the edge directions of the original graph. We say that an undirected graph  $G$  is *good* if it is obtained from some directed binary pseudotree  $G'$  (that is, a connected graph where each node has outdegree at most one and indegree either 0 or 2) by the above transformation.

Consider the following graph problem: if the input graph is good, each node has to output the parity of its distance to the nearest leaf node; if not, each node is allowed to output anything. In any case, all the nodes have to halt.

Define an algorithm  $\mathcal{A}$  as follows. First,  $\mathcal{A}$  checks locally whether the input graph  $G$  is good: Each node can detect in at most four communication rounds, whether it is contained in a valid gadget. Then, each node  $v$  corresponding to a node  $v'$  from the original directed graph can check that in the orientation obtained from the gadgets, the node either has two incoming edges and one outgoing edge, only one outgoing edge (a leaf node), or only two incoming edges (the root

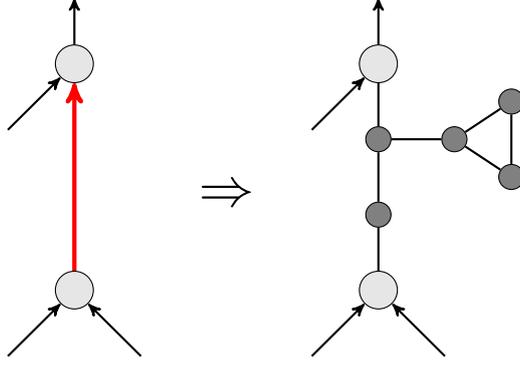


Figure 1: The red edge is replaced by the dark grey gadget consisting of five nodes. Similar encoding is performed to each edge of the original directed graph.

node in case of a tree). If either of the checks fails, the node broadcasts an instruction to halt to its neighbours and halts. Otherwise, it follows that the original graph  $G'$  is indeed a directed binary pseudotree with each edge directed away from the leaf nodes.

After verifying the goodness of the instance, algorithm  $\mathcal{A}$  starts to count the distances to leaf nodes. In the first round, each leaf node broadcasts message 0 to its neighbours, halts and outputs 0. When a node that has not yet halted receives message  $i$ , where  $i \in \{0, 1\}$ , it broadcasts  $i + 1 \bmod 2$  to its neighbours, halts and outputs  $i + 1 \bmod 2$ . Because the input graph is assumed to be finite and connected, each node eventually halts.

Since verifying the goodness takes only a constant number of rounds, and in the counting phase we count only up to 2, algorithm  $\mathcal{A}$  needs only a constant number of states. Furthermore, as in any binary pseudotree the distance from any node to the closest leaf node is at most  $O(\log n)$ , each node halts after at most  $O(\log n)$  rounds. On the other hand, in a balanced binary tree, the distance from the root to leaf nodes is  $\Omega(\log n)$ , and thus  $\Omega(\log n)$  rounds are needed until all the nodes can output the parity of their distance to the closest leaf. This completes the proof.  $\square$

## 4 Graphs of maximum degree 2

We are now ready to state the main theorem of this work.

**Theorem 5.** *There exist a decision graph problem with constant space complexity and  $\Theta(n)$  time complexity in the class of graphs of maximum degree at most 2.*

To define our graph problem, we will make use of the Thue–Morse sequence via the following definition—compare this to Definition 3 earlier.

**Definition 6** (Valid words). Define a set of words over  $\{0, 1, \_ \}$  recursively as follows:

- (1)  $\_0\_$  is *valid*,
- (2) if  $X$  is valid and  $Y$  is obtained from  $X$  by applying the substitutions  $0 \mapsto 0\_1\_1\_0$  and  $1 \mapsto 1\_0\_0\_1$  to each occurrence of 0 and 1, then  $Y$  is *valid*.

That is, valid words are obtained from the sequences  $T_{2i}, i \in \mathbb{N}$ , by inserting an underscore at the beginning, in between each symbol and at the end.

Now, we will define a decision graph problem that we call **ThueMorse** as follows.

**Definition 7** (ThueMorse). The local inputs of the problem are taken from the set  $I = \{\alpha, \beta, \gamma\} \times \{0, 1, \_ \}$  and each local output is either *yes* or *no*. Given an input function  $f: V \rightarrow I$ , we write  $f = (f_1, f_2)$ . Now, an instance is a *yes-instance* of **ThueMorse** if and only if

- the graph is a path graph,
- the first parts  $f_1(\cdot)$  of the local inputs define a consistent orientation,
- the second parts  $f_2(\cdot)$  of the local input define a valid word.

#### 4.1 Definition of the algorithm

Next, we will give an algorithm that is able to solve the decision problem **ThueMorse** by using only constant space. The high-level idea is as follows: First, we check that the local neighbourhood of each node looks correct. Then, we repeatedly apply substitutions that roll the configuration of the path graph back to a shorter prefix of the Thue–Morse sequence, until we reach a trivial configuration and accept the input—or fail to apply the substitutions unambiguously and consequently reject the input.

**Lemma 8.** *There exists an algorithm  $\mathcal{A}$  that solves problem **ThueMorse** in space  $O(1)$  and time  $O(n)$ .*

First, let us introduce some terminology and notation. In an instance  $(G, f)$  of **ThueMorse** where  $G$  is either a path or a cycle graph, the sequence of symbols defined by  $f_2$  in a given node neighbourhood is called the *input word*—note that contrary to usual words, the input word is unoriented. Sometimes we identify the node neighbourhood with the corresponding input word; the meaning should be clear from the context. During the execution of the algorithm, each node  $v$  has a *current symbol*  $c(v)$  as part of its state. The vertical bar  $|$  shall denote the end of the path graph. We will make use of it by assuming that each degree-1 node  $v$  sees a “virtual” neighbour  $u$  with the symbol  $|$ , that is,  $f_1(u) = f_2(u) = |$  and  $c(u)$  is always equal to  $|$ . The underscore symbols  $_$  will be called *separators*. Denote the alphabet which contains the possible current symbols by  $\Sigma = \{0, 1, _, |\}$ .

We define the algorithm  $\mathcal{A}$  in three parts, which we denote by I, II and III. In each part, any node can *abort*, which means that it sends a special abort message to its neighbours and then moves to state **no** and thus halts. If a node receives the abort message at any time, it aborts—that is, it passes the message on to all its neighbours, and then moves to state **no** and halts. At initialisation, each node  $v$  sets its symbol  $c(v)$  to be equal to  $f_2(v)$ , the second part of the local input.

In **part I**, each node  $v$  verifies its degree and the orientation: if  $\deg(v) \in \{1, 2\}$  and three different symbols from  $\{\alpha, \beta, \gamma, |\}$  can be found in the local inputs within the radius-1 neighbourhood of  $v$  (that is, we have  $|\{f_1(u) : u \in N(v) \cup \{v\}\}| = 3$ ), continue; otherwise, abort. Recall that the graph is assumed to be finite and connected. If none of the nodes aborts, it follows that the graph is either a path or a cycle and that the local inputs given by  $f_1$  define a word of the form  $\dots \alpha\beta\gamma\alpha\beta\gamma\alpha\beta\gamma\alpha\dots$

In **part II**, each node  $v$  verifies the input word in its radius-1 neighbourhood: if the neighbourhood is in  $\{|\_0, 0\_|, 0\_0, 1\_1, 0\_1, 1\_0, \_0\_ , \_1\_ \}$ , continue; otherwise, abort. Note that here we do not care about the orientation of the word, and hence the set of symbols received from the neighbours is enough for this step. If none of the nodes aborts, the input word is *locally valid*: every other symbol is the separator  $_$  and every other symbol is either 0 or 1.

Then, we proceed to **part III**, which contains the most interesting steps of the algorithm. Now we will make use of the orientation given as part of the input. Define  $\alpha < \beta$ ,  $\beta < \gamma$  and  $\gamma < \alpha$ . For each node  $v$ , we say that neighbour  $u$  of  $v$  is the *left neighbour* of  $v$  if  $f_1(u) < f_1(v)$  and the *right neighbour* of  $v$  if  $f_1(u) > f_1(v)$ . Thus, each node can essentially send a different message to each of its two neighbours: the orientation symbol of the recipient indicates which part of the message is intended for which recipient. From now on, we will also assume that each node attaches its own orientation symbol as part of every message sent, so that nodes can distinguish between incoming messages.



(2) Path graph, no-instance:

$$\begin{array}{c}
|_0_1_1_0_1_0_0_1_1_0_1_0_0_1_| \\
\downarrow \\
|_0000000_1111111_... \\
..._0000000_1111111_| \\
\downarrow \text{ (ambiguous substitutions)} \\
\text{abort}
\end{array}$$

(3) Cycle graph (the ends marked with ... are connected circularly):

$$\begin{array}{c}
..._0_1_1_0_1_0_0_1_1_0_0_1_0_1_1_0_... \\
\downarrow \text{ (unambiguous substitutions)} \\
..._0000000_1111111_1111111_0000000_... \\
\downarrow \text{ (no matches)} \\
\text{abort}
\end{array}$$

## 4.2 Proof of correctness

In this section we show that algorithm  $\mathcal{A}$  executes correctly and always halts with the desired output. Since parts I and II are quite trivial, we will start with part III of the algorithm.

We say that a word  $X = x_1x_2\dots x_i$  is a *compressed version* of word  $Y = y_1y_2\dots y_j$  if  $x_i \neq x_{i+1}$  for all  $i \in \{1, 2, \dots, i-1\}$  and there exist a surjective *compression mapping*  $f: [j] \rightarrow [i]$  such that  $f(1) = 1$ ,  $f(k) \in \{f(k-1), f(k-1) + 1\}$  for all  $k \in \{2, 3, \dots, j\}$  and  $y_k = x_{f(k)}$  for all  $k \in [j]$ .

**Lemma 10.** *After any node  $v$  has finished collecting the buffers,  $V(v)$  is the compressed version of the actual input word in the neighbourhood of  $v$ .*

*Proof.* We use induction on the number of rounds  $t$  after starting the phase. After the first round of the phase, each node  $v$  has received  $L(v) = c(u)$  from its left neighbour  $u$  and  $R(v) = c(w)$  from its right neighbour  $w$ . Hence  $L(v)$  and  $R(v)$  are compressed versions of the left and right 1-neighbourhoods of  $v$ , respectively.

Assume then that  $L(u)$  and  $R(w)$  are compressed versions of the left and right  $(t-1)$ -neighbourhoods of  $u$  and  $w$ , respectively, and let  $f_l$  and  $f_r$  be the corresponding mappings. Consider now the definition of the algorithm. The definition of the mapping  $r$  implies that what  $v$  receives from the left in round  $t$  of the phase is  $L(u)$ —extended by  $c(u)$  if and only if  $c(u)$  differs from the last symbol of  $L(u)$ . If it does differ, we extend  $f_l$  by defining  $f_l(t) = |L(u)| + 1$ , otherwise  $f_l(t) = |L(u)|$ . Hence the new value of the buffer  $L(v)$  is a compressed version of the left  $t$ -neighbourhood of  $v$ . The case of  $R(v)$  is handled analogously.

Suppose that  $v$  has finished collecting the buffers. Now  $L(v)$  and  $R(v)$  are the compressed versions of the left and right  $k$ -neighbourhoods of  $v$  for some  $k$ . Then, it follows from the definition of the function  $C$  that an appropriate compression mapping can be formed and  $V(v) = C(L(v), c(v), R(v))$  is the compressed version of the  $k$ -neighbourhood of  $v$ .  $\square$

We call the sequence of current symbols  $c(u)$  in the graph a *configuration* (in the case of a cycle graph, the sequence is infinite in both directions). A maximal sequence of adjacent nodes such that each node  $v$  has the same current symbol  $c(v)$  is called a *block*. We sometimes identify a block with the subword consisting of the current symbols of the block nodes.

In the next two lemmas, we show how updating the current symbol locally in nodes results in global substitutions in the configuration.

**Lemma 11.** *Assume that in the current configuration, each maximal subword of the form  $0+$  or  $1+$  is of length  $\ell$  and each maximal subword of the form  $_+$  is of length 1. If the algorithm is executed for one phase and no node aborts, in the resulting configuration the lengths are  $4\ell + 3$  and 1, respectively. More precisely, the execution of one phase always results in substitutions of the following kinds:*

$$_0^\ell _1^\ell _1^\ell _0^\ell _1^\ell _0^\ell _0^\ell _1^\ell _ \mapsto _0^{4\ell+3} _1^{4\ell+3} _ , \quad (1)$$

$$_1^\ell _0^\ell _0^\ell _1^\ell _0^\ell _1^\ell _1^\ell _0^\ell _ \mapsto _1^{4\ell+3} _0^{4\ell+3} _ . \quad (2)$$

*Proof.* Since no node aborts, each node  $v$  is able to find an unambiguous new value for  $c(v)$ . Consider an arbitrary node  $u$ . Suppose that the pattern  $P_1 = _0_1_1_0_1_0_0_1_$  matches  $V(u)$  so that the  $i$ th symbol of the pattern is aligned with the  $q(v)$ th symbol of  $V(u)$ . Now it follows from Lemma 10 that  $P_1$  is the compressed version of an actual neighbourhood  $N$  of  $u$ . Due to the assumption, for each  $j$  such that the  $j$ th symbol of  $P_1$  is 0 or 1, in the neighbourhood there are exactly  $\ell$  consecutive symbols 0 or 1, respectively, that are mapped to the  $j$ th symbol of  $P_1$  by the compression mapping  $f$ .

Let us consider the case  $i = 6$  as an example. Now  $c(u) = 1$ . As the buffers are gathered until each of them contain eight separators  $_$ , the next 5 blocks to the left from the block of  $u$ , as well as the next 11 blocks to the right, gather views that are compressed versions of a neighbourhood containing  $N$ . Hence  $P_1$  matches also their views. For example, for all nodes  $w$  in the block two steps left from the block of  $u$ ,  $P_1$  matches  $V(w)$  at position  $i - 2 = 4$ . It follows from the definition of the algorithm that the new value for  $c(u)$ , as well as for  $c(w)$ , where  $w$  is in the next 4 blocks to the left from  $u$  or 2 blocks to the right from  $u$ , is 0. The new value for the node in the 5th block left from  $u$  as well as 3rd block right from  $u$  is  $_$ . Thus, after the phase, node  $u$  will be part of a 0-block of length  $\ell + 1 + \ell + 1 + \ell + 1 + \ell = 4\ell + 3$ .

The cases for all other values of  $i$ , as for as the matching the reverse pattern  $P_2$ , are analogous.  $\square$

We call a word  $_x_1^i x_2^i \dots x_p^i _$  a *padded Thue–Morse word of length  $p$*  if  $x_1 x_2 \dots x_p$  is a prefix of the Thue–Morse sequence.

**Lemma 12.** *Let  $W$  be a subword of a configuration  $C$  at the beginning of a phase. Let  $k \geq 3$ . If  $W$  is a (complement of a) padded Thue–Morse word of length  $2^k$  and the algorithm is executed for one phase on  $C$  without aborting, the subword  $W$  is transformed to a (complement of a) padded Thue–Morse word of length  $2^{k-2}$ .*

*Proof.* We use induction on  $k$ . If  $k = 3$ , we have  $W = _0^i _1^i _1^i _0^i _1^i _0^i _0^i _1^i _$  or  $W = _1^i _0^i _0^i _1^i _0^i _1^i _1^i _0^i _$  for some  $i$ . Then Lemma 11 implies that  $W$  is transformed to  $_0^j _1^j _$  or  $_1^j _0^j _$ , respectively, where  $j = 4i + 3$ . Hence the claim holds for  $k = 3$ .

Suppose then that  $k > 3$  and the claim holds for each subword that is a (complement of a) padded Thue–Morse word of length  $2^{k-1}$ . Let  $W$  be a padded Thue–Morse word of length  $2^k$ . Now the definition of the Thue–Morse sequence implies that we can write  $W = W_1 W_2$ , where  $W_1$  is a padded Thue–Morse word of length  $2^{k-1}$  and  $_W_2$  is a complement of a padded Thue–Morse word of length  $2^{k-1}$ . By the inductive hypothesis,  $W_1$  is transformed to a padded Thue–Morse word of length  $2^{k-3}$  and  $_W_2$  is transformed to a complement of a padded Thue–Morse word of length  $2^{k-3}$ . Now the definition of the Thue–Morse sequence again implies that  $W = W_1 W_2$  is transformed to a padded Thue–Morse word of length  $2^{k-2}$ . The case where  $W$  is a complement of a padded Thue–Morse word is completely analogous.  $\square$

Now we are ready to aggregate our previous lemmas to establish that algorithm  $\mathcal{A}$  actually works correctly.

**Lemma 13.** *In a yes-instance, each node eventually halts and outputs yes.*

*Proof.* Let  $(G, f)$  be a yes-instance of ThueMorse. By definition, algorithm  $\mathcal{A}$  executes parts I and II successfully. Notice then that since the input word is valid, the configuration at the beginning of part III is actually a padded Thue–Morse word. If the input word is either  $\_0\_$  or  $\_0\_1\_1\_0$ , the nodes move immediately to the yes state and halt. Otherwise, the input word is a padded Thue–Morse word of length at least 16, and by iterating Lemma 12 we obtain a shorter padded Thue–Morse word after every phase. Note that the new current symbol will be unambiguous for each node on each phase. Eventually the configuration will match with  $\_0\_1\_1\_0\_$ , and the instance will get accepted.  $\square$

**Lemma 14.** *In a no-instance, each node eventually halts and outputs no.*

*Proof.* Let  $(G, f)$  be a no-instance of ThueMorse. By assumption,  $G$  is finite and connected. If  $G$  contains a node of degree higher than 2, algorithm  $\mathcal{A}$  aborts in part I. Hence we can assume that  $G$  has maximum degree at most two. It follows immediately from Lemma 11 that algorithm  $\mathcal{A}$  halts on  $(G, f)$ : since the size of blocks of 0’s or 1’s grows in each phase, we will eventually run out of nodes.

Graph  $G$  is either a path or a cycle graph. If  $\mathcal{A}$  rejects in part I or part II, we are done. Hence, we can assume that the input  $f$  defines a consistent orientation and each 1-neighbourhood is of the correct form—that is, every second symbol on the input word is a separator  $\_$ . Now it follows from Lemma 11 that the computation proceeds synchronously: in part III, each node starts a new phase in the same round.

Consider the case that  $G$  is a cycle graph. Due to Lemma 11, the number of blocks decreases after each phase. Eventually, if no node rejects, the execution reaches a configuration with  $b \leq 4$  blocks of 0’s and 1’s. Then, each node sees the same sequence of  $b$  blocks repeating in its view. It follows that neither pattern  $P_1$  or  $P_2$  matches and the node rejects.

Assume then that  $G$  is a path graph. If there exists a yes-instance  $(G', f')$  with the same number of nodes as  $G$  has, we proceed as follows. Suppose for a contradiction that  $(G, f)$  gets accepted. Then there is a smallest  $i$  such that before the  $i$ th phase, the configurations  $C$  on  $(G, f)$  and  $C'$  on  $(G', f')$ , respectively, are different, but after the  $i$ th phase, they are identical,  $C''$ . It follows that  $C''$  contains a subword of the form  $\_0^i\_1^i\_$  for some  $i$ , such that  $C$  and  $C'$  differ on a position overlapping with the subword. But this is a contradiction, as it follows from Lemma 11 that  $C$  and  $C'$  cannot differ on such a position.

Finally, consider the case where  $(G, f)$  is a no-instance such that there does not exist a yes-instance with the same number of nodes. Suppose again for a contradiction that  $(G, f)$  gets accepted. Let  $(G', f')$  be the largest yes-instance no larger than  $(G, f)$  and let  $(G'', f'')$  be the yes-instance that is one step larger from  $(G', f')$ . Lemma 12 implies that the algorithm  $\mathcal{A}$  needs exactly one phase more on  $(G'', f'')$  than on  $(G', f')$ . The number of phases on  $(G, f)$  equals either the number of phases on  $(G', f')$  or on  $(G'', f'')$ . But either case is a contradiction, since the instances have a different amount of blocks in the beginning, and Lemma 11 implies that the size of blocks grows at a fixed rate.  $\square$

The last thing left to do is analysing the complexity of the algorithm. This is taken care of in the following lemmas.

**Lemma 15.** *The space usage of algorithm  $\mathcal{A}$  is in  $O(1)$ .*

*Proof.* Parts I and II of algorithm  $\mathcal{A}$  clearly use only a constant number of states. Consider then part III. When gathering the buffers, consecutive blocks of symbols in the neighbourhood are represented by only one symbol, and only a constant amount of blocks are gathered (eight separator symbols  $\_$  in each buffer). Hence a constant number of states is enough to represent the contents of the buffers. Furthermore, the buffers get erased after each phase has completed. Thus, algorithm  $\mathcal{A}$  can be implemented using a constant number of states, independent of the size of the input.  $\square$

**Lemma 16.** *The running time of algorithm  $\mathcal{A}$  is in  $O(n)$ .*

*Proof.* Executing each phase of part III of the algorithm can be done in  $8i + 8 = 8(i + 1)$  rounds, where  $i$  is the size of the blocks at the start of the phase. Recall that after each phase, the size of the blocks grows from  $\ell$  to  $4\ell + 3$ . Note also that  $\frac{1}{2} \log n$  phases are enough to grow the block size so large that the algorithm has to either accept or reject. It follows that

$$8(1 + 1) + 8(7 + 1) + 8(31 + 1) + \dots + 8(2^{2(\frac{1}{2} \log n) - 1}) \leq 8n$$

is an upper bound for the running time. □

On the other hand,  $\Omega(n)$  rounds are clearly necessary to solve ThueMorse: node at the end of the path has to receive information from the other end of the path to be able to verify that the instance is actually a yes-instance. This concludes the proof of Theorem 5.

## 5 Conclusions

We identified a model of distributed computing and a class of graphs, where the question on the existence of constant-space, non-constant-time algorithms is interesting and non-trivial. The answer turned out to be positive. This opens up the way to study constant space further—we can ask, for example, what other time complexities besides  $\Theta(\log n)$  and  $\Theta(n)$  can possibly be found and in which graph classes.

## Acknowledgements

We have discussed this work with numerous people, including at least Juho Hirvonen, Janne H. Korhonen, Antti Kuusisto, Joel Rybicki and Przemysław Uznański. Many thanks to all of them, as well as to others that we may have forgotten.

## References

- [1] Yehuda Afek, Noga Alon, Ziv Bar-Joseph, Alejandro Cornejo, Bernhard Haeupler, and Fabian Kuhn. Beeping a maximal independent set. In *Proc. 25th International Symposium on Distributed Computing (DISC 2011)*, volume 6950 of *Lecture Notes in Computer Science*, pages 32–50. Springer, 2011. doi:10.1007/978-3-642-24100-0\_3.
- [2] Dana Angluin. Local and global properties in networks of processors. In *Proc. 12th Annual ACM Symposium on Theory of Computing (STOC 1980)*, pages 82–93. ACM Press, 1980. doi:10.1145/800141.804655.
- [3] Sebastian Brandt, Orr Fischer, Juho Hirvonen, Barbara Keller, Tuomo Lempiäinen, Joel Rybicki, Jukka Suomela, and Jara Uitto. A lower bound for the distributed Lovász local lemma. In *Proc. 48th Annual ACM Symposium on Theory of Computing (STOC 2016)*, pages 479–488. ACM Press, 2016. doi:10.1145/2897518.2897570. arXiv:1511.00900.
- [4] Sebastian Brandt, Juho Hirvonen, Janne H. Korhonen, Tuomo Lempiäinen, Patric R. J. Östergård, Christopher Purcell, Joel Rybicki, Jukka Suomela, and Przemysław Uznański. LCL problems on grids. In *Proc. 36th Annual ACM Symposium on Principles of Distributed Computing (PODC 2017)*, 2017. To appear. arXiv:1702.05456.
- [5] Yi-Jun Chang, Tsvi Kopelowitz, and Seth Pettie. An exponential separation between randomized and deterministic complexity in the LOCAL model. In *Proc. 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2016)*, pages 615–624. IEEE Computer Society Press, 2016. doi:10.1109/FOCS.2016.72. arXiv:1602.08166.

- [6] Yi-Jun Chang and Seth Pettie. A time hierarchy theorem for the LOCAL model, 2017. arXiv:1704.06297.
- [7] Alejandro Cornejo and Fabian Kuhn. Deploying wireless networks with beeps. In *Proc. 24th International Symposium on Distributed Computing (DISC 2010)*, volume 6343 of *Lecture Notes in Computer Science*, pages 148–162. Springer, 2010. doi:10.1007/978-3-642-15763-9\_15.
- [8] Yuval Emek and Roger Wattenhofer. Stone age distributed computing. In *Proc. 32nd Annual ACM Symposium on Principles of Distributed Computing (PODC 2013)*, pages 137–146. ACM Press, 2013. doi:10.1145/2484239.2484244. arXiv:1202.1186.
- [9] Martin Gardner. The fantastic combinations of John Conway’s new solitaire game “life”. *Scientific American*, 223(4):120–123, 1970.
- [10] Mohsen Ghaffari, Fabian Kuhn, and Yannic Maus. On the complexity of local distributed graph problems. In *Proc. 49th Annual ACM Symposium on Theory of Computing (STOC 2017)*, 2017. To appear. arXiv:1611.02663.
- [11] Mika Göös, Juho Hirvonen, and Jukka Suomela. Lower bounds for local approximation. *Journal of the ACM*, 60(5):39:1–23, 2013. doi:10.1145/2528405. arXiv:1201.6675.
- [12] Lauri Hella, Matti Järvisalo, Antti Kuusisto, Juhana Laurinharju, Tuomo Lempiäinen, Kerkko Luosto, Jukka Suomela, and Jonni Virtema. Weak models of distributed computing, with connections to modal logic. *Distributed Computing*, 28(1):31–53, 2015. doi:10.1007/s00446-013-0202-3. arXiv:1205.2051.
- [13] Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. Local computation: lower and upper bounds. *Journal of the ACM*, 63(2):17:1–17:44, 2016. doi:10.1145/2742012. arXiv:1011.5470.
- [14] Antti Kuusisto. Infinite networks, halting and local algorithms. In *Proc. 5th International Symposium on Games, Automata, Logics and Formal Verification (GandALF 2014)*, volume 161 of *Electronic Proceedings in Theoretical Computer Science*, pages 147–160, 2014. doi:10.4204/EPTCS.161.14. arXiv:1408.5963.
- [15] Nathan Linial. Locality in distributed graph algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992. doi:10.1137/0221015.
- [16] Moni Naor and Larry Stockmeyer. What can be computed locally? *SIAM Journal on Computing*, 24(6):1259–1277, 1995. doi:10.1137/S0097539793254571.
- [17] John von Neumann. *Theory of Self-Reproducing Automata*. University of Illinois Press, 1966.
- [18] Stephen Wolfram. *A New Kind of Science*. Wolfram Media, 2002.