



ELSEVIER

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs



Distinguishing views in symmetric networks: A tight lower bound

Dariusz Dereniowski^{a,*}, Adrian Kosowski^b, Dominik Pająk^c^a Faculty of Electronics, Telecommunications and Informatics, Gdańsk University of Technology, Poland^b Inria Paris and LIAFA, Université Paris Diderot, France^c Computer Laboratory, University of Cambridge, United Kingdom

ARTICLE INFO

Article history:

Received 22 October 2013

Received in revised form 29 January 2015

Accepted 8 March 2015

Available online 14 March 2015

Communicated by D. Peleg

Keywords:

Anonymous network

Port-labeled network

View

Quotient graph

ABSTRACT

The view of a node in a port-labeled network is an infinite tree encoding all walks in the network originating from this node. We prove that for any integers $n \geq D \geq 1$, there exists a port-labeled network with at most n nodes and diameter at most D which contains a pair of nodes whose (infinite) views are different, but whose views truncated to depth $\Omega(D \log(n/D))$ are identical.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

The notion of a *view* was introduced and first studied by Yamashita and Kameda in [19] in the context of distributed message passing algorithms. In so-called anonymous networks (without unique identifiers accessible to a distributed algorithm), the view is a fundamental concept which allows for identification of the network topology and for breaking of symmetries between nodes. Different views for a pair of nodes guarantee that the corresponding nodes are distinguishable, which is useful in, e.g., leader election algorithms. View-based approaches have been successfully used when designing algorithms for various network problems, including map construction [3,9], leader election [4,6,8,12,17,20], rendezvous [5,7,13], and other tasks [10,18].

The view from a node of a network is by definition (cf. Section 2) an infinite rooted tree, and therefore distributed algorithms (both for agents exploring the network or for the nodes in message passing models) can only know a finite subtree of the view. This motivates the question about the minimum integer l such that the view truncated to depth l contains all crucial information an algorithm may need.

Yamashita and Kameda proved that if views of two nodes truncated to depth n^2 are identical, then their infinite views are identical [19], where n is the number of nodes of the network. The bound has been improved to $n - 1$ by Norris [15]. Although this bound is asymptotically tight [1,15], it is far from being accurate for many networks. Hence, one may ask for bounds expressed as a function of different graph invariants. Fraigniaud and Pelc proved in [11] that if two nodes have the same views to depth $\hat{n} - 1$ then their views are the same, where \hat{n} is the number of nodes having different views

* Corresponding author.

E-mail addresses: deren@eti.pg.gda.pl (D. Dereniowski), adrian.kosowski@inria.fr (A. Kosowski), dsp39@cl.cam.ac.uk (D. Pająk).¹ Partially supported by National Science Centre of Poland grant DEC-2011/02/A/ST6/00201 and by ANR project DISPLEXITY.

(or equivalently, \hat{n} is the size of the quotient graph [19]). For some works on view computation see, e.g., [2,16]. Recently, Hendrickx [14] proved (for simple graphs with symmetric port labeling) an upper bound of $O(D \log(n/D))$ on the depth to which views need to be checked in order to be distinguished, where D is the diameter of the network, leaving the tightness of this bound as an open problem.

In this work we provide a corresponding lower bound of $\Omega(D \log(n/D))$. In particular, for each $D' \geq 3$ and $n' \geq D' \cdot 2^{12}/3$, we construct an n' -node graph G' with diameter at most D' such that taking truncations of the view to depth $\frac{D'-5}{6} \log_2 \frac{n'}{D'} - 0.41D'$ does not guarantee distinguishing a pair of nodes of this graph, which do in fact have different (infinite) views. Our construction is done in two steps. First, a list of graphs G_l , $l \geq 1$, is defined with the following properties: (a) $\text{diam}(G_l) = 3$ for each $l \geq 1$, and (b) G_l contains two nodes a_l and b_l such that the views from them to depth $l = \Theta(\log n)$ are identical but their (infinite) views are different, where n is the size of G_l . Next, in order to extend the bound for arbitrarily large diameter D' we then modify G_l by subdividing each of its edges roughly $D'/3$ times so that the new graph: (a) has diameter roughly D' , and (b) contains two nodes a_l and b_l such that their views are the same till depth $\Theta(D' \log_2(n'/D'))$ but their views are different, where n' is the size of the subdivided graph.

We remark that very recently [12], a construction of a class of labeled graphs has been put forward in the context of lower bounds for the leader election problem on anonymous graphs, which can also be used to obtain a separation of node views at distance $\Theta(\log n)$ in a graph of diameter $D = O(1)$. The analysis of that class appears somewhat more involved than for our construction.

2. Preliminaries

In this work we consider anonymous port labeled networks (the terms graph and network are used interchangeably throughout) in which the nodes do not have identifiers and each edge $\{u, v\}$ has two integers assigned to its endpoints, called the *port numbers* at u and v , respectively. The port numbers are assigned in such a way that for each node v they are pairwise different and they form a consecutive set of integers $\{1, \dots, k\}$, where k is the number of neighbors of v in G . The number of neighbors of v in G is called the *degree* of v and is denoted by $\text{deg}_G(v)$. To simplify some statements we introduce a *port labeling* function λ for G defined in such a way that for each pair u, v of adjacent nodes, $\lambda(u, v)$ is the port label at u of the edge $\{u, v\}$. For each node v of G and for each $p \in \{1, \dots, \text{deg}_G(v)\}$, $\text{next}_p(v)$ is the node u such that $\lambda(v, u) = p$, whereas $\text{end}_p(v) = \lambda(\text{next}_p(v), v)$ is the port label at the other end of the edge.

We recall the definition of a view [19]. Let G be a graph, v be a node of G and let λ be a port labeling for G . Given any $l \geq 0$, the (*truncated*) *view up to level l* , $\mathcal{V}_l(v)$, is defined as follows. $\mathcal{V}_0(v)$ is a tree consisting of a single node x_0 . Then, $\mathcal{V}_{l+1}(v)$ is the port-labeled tree rooted at x_0 and constructed as follows. For every node v_i , $i \in \{1, \dots, \text{deg}_G(v)\}$, adjacent to v in G there is a child x_i of x_0 in $\mathcal{V}_{l+1}(v)$ such that the port number at x_0 corresponding to edge $\{x_0, x_i\}$ equals $\lambda(v, v_i)$, and the port number at x_i corresponding to edge $\{x_0, x_i\}$ equals $\lambda(v_i, v)$. For each $i \in \{1, \dots, \text{deg}_G(v)\}$ the node x_i is the root of the truncated view $\mathcal{V}_l(v_i)$.

The *view* from v in G is the infinite port-labeled rooted tree $\mathcal{V}(v)$ such that $\mathcal{V}_l(v)$ is its truncation to level l , for each $l \geq 0$.

We remark that by adopting the above definitions, we are considering so-called *symmetric* networks in the sense that the port-labeled network corresponds to an unlabeled graph which is undirected, and that the encoding of port numbers at both endpoints of each edge appears in the labeling of the edges of the view.

A path in G is denoted as a sequence of nodes, $P = (v_0, v_1, \dots, v_k)$, such that $\{v_0, \dots, v_k\} \subseteq V(G)$ and $\{v_i, v_{i+1}\}$ is an edge in G for each $i \in \{0, \dots, k-1\}$. Note that nodes may repeat in a path, i.e., we do not assume that $v_i \neq v_j$ for $i \neq j$. We say that two paths $P_1 = (u_0, u_1, \dots, u_k)$ and $P_2 = (v_0, v_1, \dots, v_k)$ in G are *isomorphic* if $\lambda(u_i, u_{i+1}) = \lambda(v_i, v_{i+1})$ and $\lambda(u_{i+1}, u_i) = \lambda(v_{i+1}, v_i)$ for each $i \in \{0, \dots, k-1\}$. We will call a path *non-backtracking*² if it never follows the same edge twice on end in opposite directions, i.e., $\lambda(v_i, v_{i-1}) \neq \lambda(v_i, v_{i+1})$ for all $i \in \{1, \dots, k-1\}$.

Claim 2.1. (See [19].) *Let G be a graph, let u, v be two nodes of G , and let $l \geq 0$ be an integer. We have $\mathcal{V}_l(u) = \mathcal{V}_l(v)$ if and only if, for any path of length l starting at u , there exists an isomorphic path of length l starting at v , and vice versa. The claim also holds when restricting considerations to non-backtracking paths.*

We write $\text{diam}(G)$ to denote the *diameter* of G , i.e., the maximum (taken over all pairs of nodes u and v) length of a shortest path between u and v in G .

3. The lower bound

For each $l > 1$ we define the graph G_l which consists of nodes laid out on a regular grid with $l+2$ levels and 2^l columns, where the node in level $i \in \{0, 1, \dots, l+1\}$ and column $j \in \{0, 1, \dots, 2^l - 1\}$ is denoted by $v_i(j)$. Note that all levels are of size 2^l , and $n_l = |V(G_l)| = (l+2)2^l$.

² Boldi and Vigna used in [1] the term “non-stuttering” to denote such paths.

The construction of the edge set of G_l proceeds in four stages. Before giving a formal construction, we first provide some intuitions regarding the purpose served by edges introduced in different stages. The edges added to G_l in Stages 2 and 3 ensure that the graph is connected and has diameter of fixed size. The aim of Stage 3 is to add edges between consecutive levels in such a way that if one wants to detect a difference between some pairs of nodes in level $l+1$ (e.g., $v_{l+1}(0)$ and $v_{l+1}(2^{l-1})$), then two paths of sufficient length from those nodes need to be selected. In particular, the paths first need to go through all levels and reach level 0 (in the mentioned case, these are the nodes $v_0(0)$ and $v_0(1)$). The edges added to G_l in Stage 1 ensure that nodes in level 0 from two consecutive columns have different views truncated to depth 2.

Stage 1. Edges within level 0. In level 0, the edges form a matching between nodes $v_0(j)$ and $v_0(j \oplus 1)$, $j \in \{0, \dots, 2^l - 1\}$, with ports with labels $\{1, 2\}$, given as follows:

```
for j := 0, ..., 2l - 1 do
  λ(v0(j), v0(j ⊕ 1)) := 1 + ((j + 1) mod 2);
```

In the above, \oplus denotes the xor operation (bitwise modulo-2 addition of non-negative integers).

Stage 2. Edges within level $l+1$. The edges in level $l+1$ form a clique on all 2^l nodes of the level, with port labels corresponding to the difference of identifiers of the connected nodes, computed modulo 2^l .

```
for j := 0, ..., 2l - 1 do
  for p := 1, ..., 2l - 1 do
    λ(vl+1(j), vl+1((j + p) mod 2l)) := p.
```

Stage 3. Edges connecting level $l+1$ with all lower levels. Each node $v_{l+1}(j)$ from level $l+1$ is connected to all nodes lying in lower levels, in the same column. The port numbers at node $v_{l+1}(j)$ leading to successive levels are successive integers starting from 2^l , and the port numbers at the other end of such edges are always equal to 1, except for level 0, where the port label is either 1 or 2 (depending on which port was not used at the considered node in Stage 1 of the construction):

```
for j := 0, ..., 2l - 1 do
  for i := 0, ..., l do
    λ(vl+1(j), vi(j)) := 2l + i;
  if i > 0 then
    λ(vi(j), vl+1(j)) := 1.
  else
    λ(v0(j), vl+1(j)) := 1 + (j mod 2).
```

Stage 4. Edges connecting adjacent levels. Each node belonging to a level $i \in \{0, \dots, l-1\}$ is connected by an edge to exactly one node of the level $i+1$ directly above, so that the set of edges between such two adjacent levels is a matching. Specifically, we introduce a permutation π_i on the set of integers $\{0, \dots, 2^l - 1\}$, defined for $i=0$ as the identity permutation $\pi_0(j) = j$, and for $i > 0$ as the involution (a function that is its own inverse) which swaps the values of the i -th and $(i-1)$ -th rightmost bits in the binary notation of its argument:

$$\pi_i(j) = (j - 2^i b_i(j) - 2^{i-1} b_{i-1}(j)) + 2^i b_{i-1}(j) + 2^{i-1} b_i(j), \quad (1)$$

where for $k \geq 0$, $b_k(j) = 1$ if $(j \bmod 2^{k+1}) \geq 2^k$, and $b_k(j) = 0$, otherwise. For each node at level $i \in \{1, \dots, l-1\}$, the port label used on the edge leading to level $i-1$ is always 2, and the port label leading to level $i+1$ is always 3, as follows:

```
for j := 0, ..., 2l - 1 do
  for i := 0, ..., l - 1 do
    λ(vi(j), vi+1(πi(j))) := 3;
    λ(vi+1(πi(j)), vi(j)) := 2.
```

The graph G_4 with some edges omitted is shown in Fig. 1. In particular, the edges between nodes in level $l+1$ and level i , $i \leq l$, are given only in column 0, and edges from the clique in level $l+1$ are omitted.

Claim 3.1. For each $l \geq 6$ it holds that $|E(G_l)| < 2^{2l}$ and $\text{diam}(G_l) \leq 3$.

Proof. The number of edges of G_l can be bounded by counting the number of edges added in Stages 1 to 4 and bounding for $l \geq 6$.

To bound the diameter, note that any node of G_l either belongs to level $l+1$ or is within distance 1 from a node in level $l+1$. Also, any two nodes in level $l+1$ are adjacent. \square

For a pair of integers $0 \leq j_1, j_2 < 2^l$, we will denote by $\delta(j_1, j_2)$ the number of rightmost bits in their binary representations which are all identical, i.e., $\delta(j_1, j_2)$ is the largest integer $\delta \in \{0, \dots, l\}$ such that $(j_1 \equiv j_2) \bmod 2^\delta$ (or equivalently, such that $b_k(j_1) = b_k(j_2)$ for all $0 \leq k < \delta$). The function $\delta(j_1, j_2)$ has several important properties with respect to transformations of its parameters.




Fig. 1. The construction of G_l for $l = 4$.

Lemma 3.1. Let $j_1, j_2, d \in \{0, \dots, 2^l - 1\}$ and $i \in \{1, \dots, l - 1\}$ be arbitrarily chosen. Then:

- (i) $\delta(j_1 \oplus d, j_2 \oplus d) = \delta(j_1, j_2)$.
- (ii) $\delta((j_1 + d) \bmod 2^l, (j_2 + d) \bmod 2^l) = \delta(j_1, j_2)$.
- (iii) $\delta(\pi_i(j_1), \pi_i(j_2)) \geq \delta(j_1, j_2) - 1$, where involution π_i is defined by (1).

Proof. Claims (i) and (ii) can be attributed to folklore. To prove claim (iii), note that the involution π_i consists in swapping adjacent bits at positions i and $i - 1$, only. Consequently, we have by definition of $\delta(j_1, j_2)$ that if $\delta(\pi_i(j_1), \pi_i(j_2)) \neq \delta(j_1, j_2)$ then either $\delta(j_1, j_2) = i - 1$ or $\delta(j_1, j_2) = i$. In both cases, we have $\delta(\pi_i(j_1), \pi_i(j_2)) \geq i - 1$, and claim (iii) follows. \square

In the following, for two nodes $v_i(j_1)$ and $v_i(j_2)$ belonging to the same level i of G_l , we will use the notation: $\delta(v_i(j_1), v_i(j_2)) \equiv \delta(j_1, j_2)$.

Lemma 3.2. Consider a pair of nodes $v_i(j_1), v_i(j_2)$ of G_l with $\delta(v_i(j_1), v_i(j_2)) > 0$. Then:

- (i) Nodes $v_i(j_1)$ and $v_i(j_2)$ are of the same degree d .
- (ii) For any port $p \in \{1, \dots, d\}$, nodes $next_p(v_i(j_1))$ and $next_p(v_i(j_2))$ belong to the same level in G_l .
- (iii) For any port $p \in \{1, \dots, d\}$, $\delta(next_p(v_i(j_1)), next_p(v_i(j_2))) \geq \delta(v_i(j_1), v_i(j_2)) - 1$.
- (iv) For any port $p \in \{1, \dots, d\}$, $end_p(v_i(j_1)) = end_p(v_i(j_2))$.

Proof. By the construction of G_l , all nodes in the same level are of the same degree, and claim (i) follows. Claim (ii) also follows directly from the construction of G_l .

The construction of the port labeling in G_l is such that the Stage $a \in \{1, 2, 3, 4\}$, during which an edge along any port p is added to a vertex $v_i(j)$, depends only on the value of its level i and the parity $j \bmod 2$ of its column number (this parity is only relevant for the case of $i = 0$ and $p = 2$, distinguishing edges added in Stage 1 and Stage 3). The nodes $v_i(j_1)$ and $v_i(j_2)$ belong to the same level. Moreover, since $\delta(j_1, j_2) > 0$, we have that

$$(j_1 \equiv j_2) \bmod 2. \tag{2}$$

It follows that the edges $e_1 = \{v_i(j_1), next_p(v_i(j_1))\}$ and $e_2 = \{v_i(j_2), next_p(v_i(j_2))\}$, corresponding to a traversal of the same port p starting from nodes $v_i(j_1)$ and $v_i(j_2)$, must necessarily have been defined in the same Stage a of the construction of the edge set of G_l . To complete the proofs of claims (iii) and (iv), we consider the corresponding four cases of $a \in \{1, 2, 3, 4\}$.

- Edges e_1 and e_2 were defined in Stage 1. Then, $i = 0, p \in \{1, 2\}$, and we have:

$$\begin{aligned} next_p(v_i(j_1)) &= v_i(j_1 \oplus 1), & end_p(v_i(j_1)) &= 1 + (j_1 \bmod 2), \\ next_p(v_i(j_2)) &= v_i(j_2 \oplus 1), & end_p(v_i(j_2)) &= 1 + (j_2 \bmod 2). \end{aligned}$$

By Lemma 3.1(i), we have:

$$\delta(next_p(v_i(j_1)), next_p(v_i(j_2))) = \delta(v_i(j_1), v_i(j_2)).$$

Moreover, taking into account (2), we obtain $end_p(v_i(j_1)) = end_p(v_i(j_2))$. This completes the proof of claims (iii) and (iv) for this case.

- Edges e_1 and e_2 were defined in Stage 2. Then, $i = l + 1$, $p \in \{1, \dots, 2^l - 1\}$, and we have:

$$\begin{aligned} \text{next}_p(v_i(j_1)) &= v_i((j_1 + p) \bmod 2^l), & \text{end}_p(v_i(j_1)) &= (2^l - p) \bmod 2^l, \\ \text{next}_p(v_i(j_2)) &= v_i((j_2 + p) \bmod 2^l), & \text{end}_p(v_i(j_2)) &= (2^l - p) \bmod 2^l. \end{aligned}$$

We immediately have $\text{end}_p(v_i(j_1)) = \text{end}_p(v_i(j_2))$, and moreover, by [Lemma 3.1\(ii\)](#):

$$\delta(\text{next}_p(v_i(j_1)), \text{next}_p(v_i(j_2))) = \delta(v_i(j_1), v_i(j_2)).$$

- Edges e_1 and e_2 were defined in Stage 3. Then, we need to consider two cases: either $i = l + 1$, or $i \in \{0, \dots, l\}$. If $i = l + 1$, then $p = 2^l + i'$ for some $i' \in \{0, \dots, l\}$. We have for $i' > 0$:

$$\begin{aligned} \text{next}_p(v_i(j_1)) &= v_{i'}(j_1), & \text{end}_p(v_i(j_1)) &= 1, \\ \text{next}_p(v_i(j_2)) &= v_{i'}(j_2), & \text{end}_p(v_i(j_2)) &= 1, \end{aligned}$$

whereas for $i' = 0$:

$$\begin{aligned} \text{next}_p(v_i(j_1)) &= v_0(j_1), & \text{end}_p(v_i(j_1)) &= 1 + (j_1 \bmod 2), \\ \text{next}_p(v_i(j_2)) &= v_0(j_2), & \text{end}_p(v_i(j_2)) &= 1 + (j_2 \bmod 2). \end{aligned}$$

Claims (iii) and (iv) follow directly, taking into account [Eq. \(2\)](#) in the latter case.

Otherwise, if $i < l + 1$, then $p = 2$ (if $i = 0$ and $j_1 \equiv j_2 \equiv 1 \pmod{2}$), or $p = 1$ (in all other cases). We have:

$$\begin{aligned} \text{next}_p(v_i(j_1)) &= v_{l+1}(j_1), & \text{end}_p(v_i(j_1)) &= 2^l + i, \\ \text{next}_p(v_i(j_2)) &= v_{l+1}(j_2), & \text{end}_p(v_i(j_2)) &= 2^l + i, \end{aligned}$$

and claims (iii) and (iv) immediately follow as well.

- Edges e_1 and e_2 were defined in Stage 4. Then, $p \in \{2, 3\}$ and $i \in \{0, \dots, l\}$.

We first consider the case of $p = 3$, i.e., when $i < l$ and port p leads up to level $i + 1$. We have:

$$\begin{aligned} \text{next}_p(v_i(j_1)) &= v_{i+1}(\pi_i(j_1)), & \text{end}_p(v_i(j_1)) &= 2, \\ \text{next}_p(v_i(j_2)) &= v_{i+1}(\pi_i(j_2)), & \text{end}_p(v_i(j_2)) &= 2. \end{aligned}$$

Claim (iv) follows directly, and so does claim (iii), taking into account that by [Lemma 3.1\(iii\)](#):

$$\delta(\text{next}_p(v_i(j_1)), \text{next}_p(v_i(j_2))) = \delta(\pi_i(j_1), \pi_i(j_2)) \geq \delta(j_1, j_2) - 1 = \delta(v_i(j_1), v_i(j_2)) - 1.$$

In the case of $p = 2$, i.e., when $i > 0$ and port p leads down to level $i - 1$, we have:

$$\begin{aligned} \text{next}_p(v_i(j_1)) &= v_{i-1}(\pi_{i-1}^{-1}(j_1)), & \text{end}_p(v_i(j_1)) &= 3, \\ \text{next}_p(v_i(j_2)) &= v_{i-1}(\pi_{i-1}^{-1}(j_2)), & \text{end}_p(v_i(j_2)) &= 3. \end{aligned}$$

We obtain the claims as in the previous case, this time noting that since π_{i-1} is an involution, we have $\pi_{i-1}^{-1} \equiv \pi_{i-1}$, and we can apply [Lemma 3.1\(iii\)](#) for π_{i-1} to show Claim (iii). \square

Lemma 3.3. Consider a pair of nodes $v_i(j_1)$, $v_i(j_2)$ of G_l with $\delta \equiv \delta(v_i(j_1), v_i(j_2)) > 0$. Then, the views of nodes $v_i(j_1)$ and $v_i(j_2)$ are equal at least up to depth δ , $\mathcal{V}_\delta(v_i(j_1)) = \mathcal{V}_\delta(v_i(j_2))$.

Proof. The proof proceeds by induction with respect to δ .

When $\delta = 1$, by [Lemma 3.2\(i\)](#), the nodes $v_i(j_1)$ and $v_i(j_2)$ have the same degree d , and by [Lemma 3.2\(iv\)](#), after traversing an edge labeled with any port $p \in \{1, \dots, d\}$ from either node, we enter the adjacent node by the same port: $\text{end}_p(v_i(j_1)) = \text{end}_p(v_i(j_2))$. Hence, $\mathcal{V}_1(v_i(j_1)) = \mathcal{V}_1(v_i(j_2))$.

Now, let $\delta > 1$ and suppose that the claim of the lemma holds for all $\delta' \leq \delta - 1$. Again, by [Lemma 3.2\(i\)](#) and (iv), the nodes $v_i(j_1)$ and $v_i(j_2)$ have the same degree, and after traversing an edge labeled with any port $p \in \{1, \dots, d\}$ from either node, we enter the adjacent node by the same port. Moreover, we have by [Lemma 3.2\(iii\)](#) that $\delta(\text{next}_p(v_i(j_1)), \text{next}_p(v_i(j_2))) \geq \delta - 1$, and, by [Lemma 3.2\(ii\)](#), $\text{next}_p(v_i(j_1))$ and $\text{next}_p(v_i(j_2))$ belong to the same level of G_l . Hence, by the inductive assumption, $\mathcal{V}_{\delta-1}(\text{next}_p(v_i(j_1))) = \mathcal{V}_{\delta-1}(\text{next}_p(v_i(j_2)))$. Since port p was arbitrarily chosen, it follows from the recursive definition of the view that $\mathcal{V}_\delta(v_i(j_1)) = \mathcal{V}_\delta(v_i(j_2))$, and so we have the claim. \square

Observe that the nodes $a_l = v_l(0)$ and $b_l = v_l(2^{l-1})$ have distinct views in G_l . Indeed, consider a sequence of l traversals along port 2, starting from nodes a_l and b_l . We argue, by induction on $i \in \{0, \dots, l-1\}$, that after i edge traversals the node reached from a_l is $v_{l-i}(0)$, and the node reached from b_l is $v_{l-i}(2^{i-1})$. For $i = 0$ the claim is trivial and hence assume that it holds for some $0 \leq i < l - 1$. The edge with port number 2 at $v_{l-i}(0)$ clearly leads to $v_{l-1-i}(0)$ as required. Hence,

it remains to argue that there is an edge between $v_{l-1-i}(2^{l-2-i})$ and $v_{l-i}(2^{l-1-i})$ in G . According to construction of edges between the levels $l-2-i$ and $l-1-i$ in Stage 4, we need to argue that

$$\pi_{l-1-i}(j) = 2^{l-1-i}, \quad \text{where } j = 2^{l-2-i}. \quad (3)$$

By (1),

$$\pi_{l-1-i}(j) = \left(j - 2^{l-1-i} b_{l-1-i}(j) - 2^{l-2-i} b_{l-2-i}(j) \right) + 2^{l-1-i} b_{l-2-i}(j) + 2^{l-2-i} b_{l-1-i}(j).$$

We have $b_{l-1-i}(j) = 0$ because $2^{l-2-i} \bmod 2^{l-1-i} = 0 < 2^{l-1-i}$, and $b_{l-2-i}(j) = 1$ because $2^{l-2-i} \bmod 2^{l-1-i} \geq 2^{l-2-i}$. Thus, $\pi_{l-1-i}(j) = j - 2^{l-2-i} + 2^{l-1-i} = 2^{l-1-i}$ as required, which completes the proof of (3). Thus, for $i = l-1$, we reach nodes $v_1(0)$ and $v_1(1)$, respectively. Then, after following port 2 for the l -th time, we reach nodes $v_0(0)$ and $v_0(1)$, respectively. Finally, after following port 2 for the $(l+1)$ -th time, we reach nodes $v_0(1)$ and $v_{l+1}(1)$, respectively.

In the last step of the traversal of this sequence of ports, node $v_0(1)$ is entered by port 1, while node $v_{l+1}(1)$ is entered by port 2^l . Hence, $\mathcal{V}(a_l) \neq \mathcal{V}(b_l)$. On the other hand, $\delta(v_l(0), v_l(2^{l-1})) = l-1$, so by Lemma 3.3, $\mathcal{V}_{l-1}(a_l) = \mathcal{V}_{l-1}(b_l)$. We obtain the following claim.

Proposition 3.1. *For any integer $l \geq 6$, there exists a graph G_l on $(l+2)2^l$ nodes, at most 2^{2l} edges, and diameter at most 3, which contains a pair of nodes a_l, b_l having distinct views and having the same views up to depth $l-1$.*

This result completes our proof for the case of graphs of diameter 3. Now, in order to obtain an asymptotic lower bound of $\Omega(D \log(n/D))$, where n and D are, respectively, the size and the diameter of a graph, we modify each of G_l 's to obtain graphs of arbitrarily large diameter.

Let D be an odd integer. For each G_l , $l \geq 1$, define $\xi_D(G_l)$ to be a graph constructed by replacing each edge $\{u, v\}$ from G_l by a path $P(\{u, v\})$ of length D with endpoints u and v . Note that $|V(\xi_D(G_l))| = |V(G_l)| + (D-1)|E(G_l)|$ and $|E(\xi_D(G_l))| = D|E(G_l)|$. Also, $\xi_1(G_l) = G_l$. We define the port labeling λ_D for $\xi_D(G_l)$ as follows. For each $\{u, v\} \in E(G_l)$ take the corresponding path $P(\{u, v\}) = (u, x_1, \dots, x_{D-1}, v)$ and set $\lambda_D(u, x_1) = \lambda(u, v)$, $\lambda_D(v, x_{D-1}) = \lambda(v, u)$. The remaining port labels of $P(\{u, v\})$ are assigned arbitrarily but in such a way that whenever two edges of G_l have the same port labels at the endpoints, then we select isomorphic port labelings for the two corresponding paths in $\xi_D(G_l)$. Formally, for any two edges $\{u, v\}$ and $\{u', v'\}$ of G_l satisfying $\lambda(u, v) = \lambda(u', v')$ and $\lambda(v, u) = \lambda(v', u')$, for the two corresponding paths $P(\{u, v\}) = (u = x_0, x_1, \dots, x_{D-1}, x_D = v)$ and $P(\{u', v'\}) = (u' = x'_0, x'_1, \dots, x'_{D-1}, x'_D = v')$ it holds that $\lambda_D(x_j, x_{j+1}) = \lambda_D(x'_j, x'_{j+1})$ and $\lambda_D(x_{j+1}, x_j) = \lambda_D(x'_{j+1}, x'_j)$ for each $j \in \{0, \dots, D-1\}$. The latter is possible for any D when $\lambda(u, v) \neq \lambda(v, u)$ and it is possible for odd D for 'symmetric' edges, i.e., when $\lambda(u, v) = \lambda(v, u)$. As an example of such labeling consider the following. If D is odd and $\lambda(u, v) = \lambda(v, u)$, then we set

$$\lambda_D(x_j, x_{j-1}) = 1 \text{ and } \lambda_D(x_j, x_{j+1}) = 2 \text{ for each } j \in \{1, \dots, \lfloor D/2 \rfloor\},$$

and

$$\lambda_D(x_j, x_{j-1}) = 2 \text{ and } \lambda_D(x_j, x_{j+1}) = 1 \text{ for each } j \in \{\lfloor D/2 \rfloor + 1, \dots, D-1\}.$$

If, on the other hand, $\lambda(u, v) \neq \lambda(v, u)$, then one can set

$$\lambda_D(x_j, x_{j-1}) = 1 \text{ and } \lambda_D(x_j, x_{j+1}) = 2 \text{ for each } j \in \{1, \dots, D-1\}.$$

We also have the following claim.

Claim 3.2. *For each $l \geq 1$ and $D \geq 1$ it holds that $\text{diam}(\xi_D(G_l)) \leq 3D$.*

We now consider the nodes $a_l, b_l \in V(G_l)$ satisfying Proposition 3.1, and characterize their (truncated) views within graph $\xi_D(G_l)$.

Lemma 3.4. *For any $l \geq 1$, $i \leq l-1$, and odd $D \geq 1$, in graph $\xi_D(G_l)$ we have: $\mathcal{V}_{D_i}(a_l) = \mathcal{V}_{D_i}(b_l)$ and $\mathcal{V}(a_l) \neq \mathcal{V}(b_l)$.*

Proof. In order to prove that $\mathcal{V}_{D_i}(a_l) = \mathcal{V}_{D_i}(b_l)$, we will use the characterization from Claim 2.1. Let $P_j = (u_0^j, u_1^j, \dots, u_{kD}^j)$, $j \in \{1, 2\}$, be any two non-backtracking paths in $\xi_D(G_l)$ such that $u_0^1 = a_l$ and $u_0^2 = b_l$.

By construction, $P_j^i = (u_0^i, u_D^i, u_{2D}^i, \dots, u_{kD}^i)$ is a path in G_l for each $j \in \{1, 2\}$. By the definition of port labeling of $\xi_D(G_l)$, for paths ending at nodes within $V(G_l)$, the port labelings of P_1 and P_2 are identical if and only if the port labelings of P_1^i and P_2^i are identical. Thus, P_1 and P_2 are isomorphic in $\xi_D(G_l)$ if and only if P_1^i and P_2^i are isomorphic in G_l . Since $i \leq l-1$, by Claim 2.1 we obtain that $\mathcal{V}_{D_i}(a_l) = \mathcal{V}_{D_i}(b_l)$. The fact that $\mathcal{V}(a_l) \neq \mathcal{V}(b_l)$ follows from similar arguments. \square

Theorem 3.1. *Let $D' \geq 3$ and $n' \geq 1$ be arbitrary integers with $n' \geq D' \cdot 2^{12}/3$. There exists a graph G with at most n' nodes and diameter at most D' , which contains two nodes having distinct views which are identical when truncated up to depth $\frac{D'-5}{6} \log_2 \frac{n'}{D'} - 0.41D'$.*

Proof. Let D be the largest odd integer such that $3D \leq D'$. Note that $D \geq (D' - 5)/3$ and $1 \leq D \leq D'/3$. Take $G = \xi_D(G_l)$, $a = a_l$ and $b = b_l$, where l is selected so that $n = |V(G)| \geq n'$. Observe that, by Claim 3.1, the number of nodes of G satisfies:

$$n = |V(G_l)| + (D - 1)|E(G_l)| < D|E(G_l)| < D2^{2l} \leq D'2^{2l}/3.$$

Thus, $n \leq n'$ is satisfied if $D'2^{2l}/3 \leq n'$; we put $l = \lfloor \frac{1}{2} \log_2(3n'/D') \rfloor$. (Note that $l \geq 6$ by assumption.)

By Lemma 3.4, the views of a_l and b_l are different in G , but the same when truncated up to depth $D(l - 1)$. We have:

$$\begin{aligned} D(l - 1) &\geq \frac{D' - 5}{3} \cdot \left(\frac{1}{2} \log_2 \frac{3n'}{D'} - 2 \right) = \frac{D' - 5}{6} \log_2 \frac{n'}{D'} + \frac{D'}{3} \left(\frac{1}{2} \log_2 3 - 2 \right) - \frac{5}{3} \left(\frac{1}{2} \log_2 3 - 2 \right) > \\ &> \frac{D' - 5}{6} \log_2 \frac{n'}{D'} - 0.41D'. \quad \square \end{aligned}$$

4. Final remarks

We have shown a tight lower bound of $\Omega(D \log(n/D))$ on the depth to which the views of a pair of nodes of a symmetric anonymous network need to be checked in order to decide if their views in the graph are different. We remark that our problem of view distinction can be generalized in the following two directions:

- One may consider scenarios in which some information (labels) is also encoded at nodes of the network, and also appears as a node-labeling in the definition of the view. (Such an extended definition of views has appeared, e.g., in the context of leader election in networks where not all identifiers are distinct [20].)
- One may ask about the depth of the view which suffices not only to distinguish a pair of nodes of the same graph having distinct views, but also any pair of nodes of two arbitrary graphs, which have the same view. (This type of distinction is required in, e.g., in so-called map construction problems [3].)

Since our lower bound concerns a more restricted scenario, it immediately applies to both of the above cases as well. Formally, when considering a pair of graphs, as n and D we take the maximum order and diameter of the two graphs.

At the same time, the techniques used by Hendrickx [14] to show a corresponding upper bound of $O(D \log_2(n/D))$ for distinguishing a pair of nodes of a connected graph can be adapted to apply to all of the above cases as well, including the scenario of distinguishing a pair of views in two different graphs. Indeed, suppose that there exist a graph G_1 on n_1 nodes with diameter D_1 containing a node v_1 , and a graph G_2 on n_2 nodes with diameter D_2 containing a node v_2 , such that nodes v_1 and v_2 have views in their respective graphs indistinguishable up to some distance $l > 1$. Then, one can construct a new connected graph G on $n = n_1 + n_2$ nodes with diameter $D \leq D_1 + D_2$, in which there exists a pair of nodes with views indistinguishable also up to distance l . To achieve this, denoting by d the degrees of v_1 in G_1 and of v_2 in G_2 , which are necessarily equal, we form G by taking the disjoint union of graphs G_1 and G_2 , and connecting vertices v_1 and v_2 by an edge labeled with port $d + 1$ at both ends.

Thus, we can say that the question of the necessary depth of view reconstruction with respect to the diameter of a symmetric port-labeled networks has been completely resolved.

References

- [1] P. Boldi, S. Vigna, Fibrations of graphs, *Discrete Math.* 243 (1–3) (2002) 21–66.
- [2] P. Boldi, S. Vigna, Universal dynamic synchronous self-stabilization, *Distrib. Comput.* 15 (3) (2002) 137–153.
- [3] J. Chalopin, S. Das, A. Kosowski, Constructing a map of an anonymous graph: applications of universal sequences, in: *OPODIS*, 2010, pp. 119–134.
- [4] J. Chalopin, Y. Métivier, An efficient message passing election algorithm based on Mazurkiewicz's algorithm, *Fund. Inform.* 80 (1–3) (2007) 221–246.
- [5] J. Czyzowicz, A. Kosowski, A. Pelc, How to meet when you forget: log-space rendezvous in arbitrary graphs, *Distrib. Comput.* 25 (2) (2012) 165–178.
- [6] S. Das, P. Flocchini, A. Nayak, N. Santoro, Effective elections for anonymous mobile agents, in: *ISAAC*, 2006, pp. 732–743.
- [7] S. Das, M. Mihalák, R. Srámek, E. Vicari, P. Widmayer, Rendezvous of mobile agents when tokens fail anytime, in: *OPODIS*, 2008, pp. 463–480.
- [8] D. Dereniowski, A. Pelc, Leader election for anonymous asynchronous agents in arbitrary networks, *Distrib. Comput.* 27 (1) (2012) 21–38.
- [9] D. Dereniowski, A. Pelc, Drawing maps with advice, *J. Parallel Distrib. Comput.* 72 (2) (2012) 132–143.
- [10] P. Flocchini, A. Roncato, N. Santoro, Computing on anonymous networks with sense of direction, *Theoret. Comput. Sci.* 301 (1–3) (2003) 355–379.
- [11] P. Fraigniaud, A. Pelc, Decidability classes for mobile agents computing, in: *Proc. Latin American Symposium on Theoretical Informatics, LATIN'12*, 2012, pp. 362–374.
- [12] E.G. Fusco, A. Pelc, Knowledge, level of symmetry, and time of leader election, *Distrib. Comput.* (2015), <http://dx.doi.org/10.1007/s00446-014-0237-0>, in press.
- [13] S. Guilbault, A. Pelc, Asynchronous rendezvous of anonymous agents in arbitrary graphs, in: *OPODIS*, 2011, pp. 421–434.
- [14] J.M. Hendrickx, Views in a graph: to which depth must equality be checked?, *IEEE Trans. Parallel Distrib. Syst.* 25 (7) (2014) 1907–1912.
- [15] N. Norris, Universal covers of graphs: isomorphism to depth $N - 1$ implies isomorphism to all depths, *Discrete Appl. Math.* 56 (1) (1995) 61–74.
- [16] S. Tani, Compression of view on anonymous networks – folded view, *IEEE Trans. Parallel Distrib. Syst.* 23 (2) (2012) 255–262.
- [17] S. Tani, H. Kobayashi, K. Matsumoto, Exact quantum algorithms for the leader election problem, *ACM Trans. Comput. Theory* 4 (1) (2012) 1.

- [18] M. Yamashita, T. Kameda, Computing functions on asynchronous anonymous networks, *Math. Syst. Theory* 29 (4) (1996) 331–356.
- [19] M. Yamashita, T. Kameda, Computing on anonymous networks, Part I: characterizing the solvable cases, *IEEE Trans. Parallel Distrib. Syst.* 7 (1) (1996) 69–89.
- [20] M. Yamashita, T. Kameda, Leader election problem on networks in which processor identity numbers are not distinct, *IEEE Trans. Parallel Distrib. Syst.* 10 (9) (1999) 878–887.