



HAL
open science

Distributed computation of temporal twins in periodic undirected time-varying graphs

Lina Azerouk, Binh-Minh Bui-Xuan, Camille Palisoc, Maria Potop-Butucaru, Massinissa Tighilt

► **To cite this version:**

Lina Azerouk, Binh-Minh Bui-Xuan, Camille Palisoc, Maria Potop-Butucaru, Massinissa Tighilt. Distributed computation of temporal twins in periodic undirected time-varying graphs. 2024. hal-04559327

HAL Id: hal-04559327

<https://hal.science/hal-04559327v1>

Preprint submitted on 25 Apr 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distributed computation of temporal twins in periodic undirected time-varying graphs

Lina Azerouk¹, Binh-Minh Bui-Xuan¹, Camille Palisoc¹, Maria Potop-Butucaru¹, and Massinissa Tighilt¹

LIP6 (CNRS – Sorbonne Université),

[lina.azerouk,buixuan,camille.palisoc,maria.potop-butucaru,massinissa.tighilt]@lip6.fr

Abstract. Twin nodes in a static network capture the idea of being substitutes for each other for maintaining paths of the same length anywhere in the network. In dynamic networks, we model twin nodes over a time-bounded interval, noted (Δ, d) -twins, as follows. A periodic undirected time-varying graph $\mathcal{G} = (G_t)_{t \in \mathbb{N}}$ of period p is an infinite sequence of static graphs where $G_t = G_{t+p}$ for every $t \in \mathbb{N}$. For Δ and d two integers, two distinct nodes u and v in \mathcal{G} are (Δ, d) -twins if, starting at some instant, the outside neighbourhoods of u and v has non-empty intersection and differ by at most d elements for Δ consecutive instants. In particular when $d = 0$, u and v can act during the Δ instants as substitutes for each other in order to maintain journeys of the same length in time-varying graph \mathcal{G} . We propose a distributed deterministic algorithm enabling each node to enumerate its (Δ, d) -twins in $2p$ rounds, using messages of size $O(\delta_{\mathcal{G}} \log n)$, where n is the total number of nodes and $\delta_{\mathcal{G}}$ is the maximum degree of the graphs G_t 's. Moreover, using randomized techniques borrowed from distributed hash function sampling, we reduce the message size down to $O(\log n)$ w.h.p.

Keywords: time-varying graph, twin, twin distributed computing.

1 Introduction

Maintaining connectivity is an important topic for dynamic networks. For instance, in the 1-interval-connected network model [9,7,8], if moreover the graph is periodic and each node's local view can be updated entirely every round on contact of its neighbours, then it is proven that a temporally optimal broadcast tree can be maintained in real time [3]. This is important for connectivity because of the real time updates, namely nodes do not need to process a posteriori the pieces of information they received about their neighbourhood. Here, the 1-interval-connected model allows for a fixed set of nodes to communicate through a temporal edge set changing at discrete time units called rounds. The structure as a whole is called a time-varying graph [4], where at every round the graph is also usually assumed to be connected.

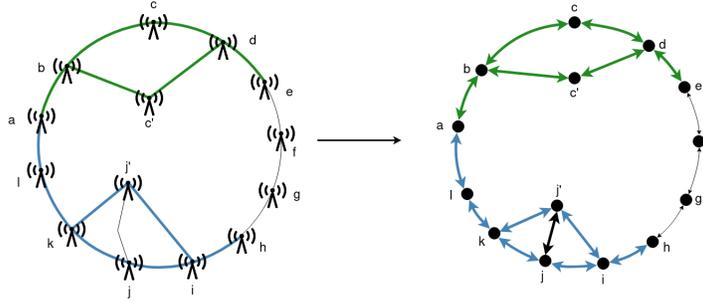


Fig. 1. Network modelled as an undirected graph. Arrows are in double direction to stress that communications are allowed both ways around. The two pairs of twins in the graph are $\{c, c'\}$ and $\{j, j'\}$. A packet traveling between node a and e can choose (a, b, c, d, e) as shortest path. If node c switches off or if it is faulty, the path of same length (a, b, c', d, e) can still be used. Between a and h , the two paths (a, l, k, j, i, h) and (a, l, k, j', i, h) are shortest, however, (a, l, k, j, j', i, h) is not a shortest path.

In our study, we consider all the above assumptions, except for that we do not require the time-varying graph to be connected at every round. This is convenient to model adhoc networks with both dense and sparse areas with predictable periodicity such as satellite networks. We would like to capture by the notion of twin nodes the possibility of backup routes or of a resource saving strategy during the times when some satellites are called twins. This is why we aim at a formalism where the connectivity in the rest of the network are maintained by journeys of the same length even though one of the twin satellites switches off for sleeping or maintenance during some time window of a given desired length, *cf.* Fig. 1.

In a static graph, such a concept can be formalised following the model of ϵ -modules [5] and d -contractions [1]: given an integer d , two nodes $u \neq v$ are d -twins if their outside neighbourhoods have non-empty intersection while the symmetric difference has size smaller than d , that is, the sum of $|N(u) \setminus N(v) \setminus \{u, v\}|$ and $|N(v) \setminus N(u) \setminus \{u, v\}|$ is at most d . We need the former assumption of non-empty intersection for obtaining efficient computational results. At the same time, isolated nodes or nodes without common neighbours are bad models of backup routes as in the above example of network of satellites. Hence, we discard them from our definition of d -twins which requires nodes to have at least one common neighbour. From this point, our extension to the dynamic case will follow that of Δ -twins [2]: given two integers Δ and d , two nodes $u \neq v$ are (Δ, d) -twins in a time-varying graph $\mathcal{G} = (G_t)_{t \in \mathbb{N}}$ if there exists $t_0 \in \mathbb{N}$ such that u and v are d -twins in every graph G_t , for $t_0 \leq t < t_0 + \Delta$.

Our contribution:

We give a distributed deterministic algorithm for every node in a periodic undirected time-varying graph of period p to compute its (Δ, d) -twins after $2p$ rounds, using messages of size $O(\delta_G \log n)$, where δ_G is the maximum degree of the graphs G_t 's. When randomized by sampling techniques from [6], our messages can be reduced to $O(\log n)$ w.h.p.

The main idea in our algorithm can be divided into two steps. In the first step, we show how to solve the problem of finding d -twins in a static graph G after 2 rounds. Roughly, while in 2 rounds every node can receive messages sent by its 2-neighbourhood, the receiver node need to detect those sent by its d -twins. In particular, note that while a path over 3 vertices admits two 0-twins, a path over $n \geq 4$ vertices has none (because for every node, one vertex from its 3-neighbourhood breaks down the 0-twin definition). In order to avoid examining messages from k -neighbourhoods with $k \geq 3$, we exhibit a set property involving the number of paths of length 2 which allows to detect d -twins. The main idea is to use inclusion-exclusion properties on the neighbourhoods of nodes in order to restrict the message size to $O(\delta_G \log n)$. Furthermore, by exploiting very simple twist inspired from distributed hash function sampling [6], we reduce the message size down to $O(\log n)$ w.h.p.

In a second step, we address the dynamic case. The main idea here is for the receiver nodes to store information while waiting for the time-varying graph to repeat its edges, using periodicity. At the same time, a sender node must be detected Δ consecutive times as a d -twin in order to be detected as (Δ, d) -twin. A special attention is needed for all computations to end after $2p$ rounds (and not $2p + \Delta$).

Eventually, we remark that our algorithm requires the time-varying graph to be undirected, meaning communication on an edge uv must be allowed both way around between u and v . This is because of the step in the static setting where every node attempts to count the number of paths of length 2 linking itself to every of its d -twins.

The paper is organised as follows. The formal framework of periodic time-varying graphs and (Δ, d) -twins is defined in Section 2. In Section 3 we briefly present a property of twins in static graphs, before using it to show a distributed algorithm for every node to compute its (Δ, d) -twins and prove its correctness. Section 4 is devoted to reducing the message size used by our algorithm to $O(\log n)$ w.h.p. We conclude the paper in Section 5 along with some open remarks for future works.

2 Model and Problem definition

We consider distributed systems which are fault-free, message-passing, synchronous, with a unique ID for each process. This will be formalised as a time-varying graph [4] defined by $\mathcal{G} = (V, E, T, \rho)$, where V is a finite set of nodes representing processes, $E \subseteq \binom{V}{2}$ is the underlying set of all possible edges and $\rho : E \times T \rightarrow \{0, 1\}$ is the presence function defining whether an edge exists in a given round. Temporal edges are those bound to a specific time instant: $E_t = \{e \in E : \rho(e, t) = 1\}$. Alternatively, a time-varying graph can also be seen as a sequence of static graphs $\mathcal{G} = (G_t)_{t \in T}$, where $G_t = (V, E_t)$.

Inspired by the models of dynamic networks in [4,7,8,9], we suppose furthermore that the time-varying graph is characterized by a periodicity, non-anonymity, synchronicity, and local knowledge. However, we do not require it to be connected at every time instant.

- The **periodicity** means the existence of a specific value p such that, every edge present in the graph at a particular time t has its presence in the graph repeated periodically at $t + p$. Because of periodicity, we restrict our domain of study down to the cyclic group of p elements and abusively refer to it as $T = \mathbb{Z}/p\mathbb{Z} = \{0, 1, \dots, p-1\} = \llbracket 0, p-1 \rrbracket$, and subsequently restrict the input time-varying graph to the first p instants, $\mathcal{G} = (G_t)_{0 \leq t < p}$.
- The **non-anonymity** of the graph means that each node v has a unique $ID \in \mathbb{N}$, we assume that each ID is represented using at most $O(\log n)$ bits, where $n = |V|$.
- The **synchronicity** within the graph implies the presence of a global clock that ticks regularly for each $t \in T$, beginning with $t = 0$. During each tick of this clock, every node v in the graph performs three sequential actions: sending a message to each neighbouring node, receiving the message from each neighbour, and executing computational tasks. Each occurrence of the clock ticking is called *round*.
- By **local knowledge**, we assume that each node v in the graph has the knowledge of the number of its neighbours at any given time $t \in T$, but not their IDs . The set of neighbours of node v at time t is denoted as $N_t(v)$.

Remark 1. If the system does not allow for local knowledge, nodes can still pass their ID to all neighbours at every round t . Then, at round $t+1$, the system satisfies the retro-property of local knowledge of every past round. In other words, it is possible to equip any periodic time-varying

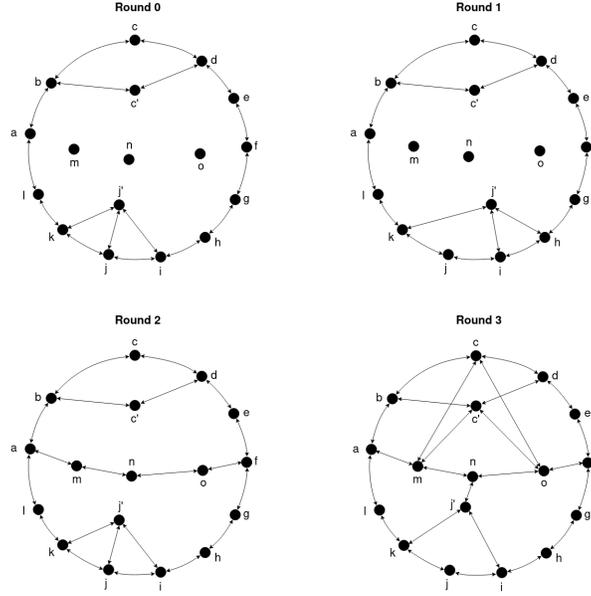


Fig. 2. A periodic time-varying graph with a period of 4, where nodes c and c' are $(4, 0)$ -twins.

graph of period p with the local knowledge property after a preprocessing using p rounds and messages of size $O(\log n)$.

For $u \neq v$ two nodes of \mathcal{G} , we say that u and v are (Δ, d) -twins if, starting from a specific time instant $t_0 \in T$, their outside neighbourhoods have a non-empty intersection while the difference between these two sets is less than or equal to a fixed value d for Δ consecutive rounds, that is, $|N_t(u) \setminus N_t(v) \setminus \{u, v\}| + |N_t(v) \setminus N_t(u) \setminus \{u, v\}| \leq d$, for every $t \in \llbracket t_0, t_0 + \Delta - 1 \rrbracket$ if $t_0 + \Delta \leq p$ and for every $t \notin \llbracket t_0 + \Delta - 1 \bmod p, t_0 \rrbracket$ otherwise. Fig. 2 exemplifies (Δ, d) -twins. We address the following problem.

PERIODIC (Δ, d) -TWINS

INPUT: $\mathcal{G} = (G_t)_{t \in T}$ with $T = \llbracket 0, p - 1 \rrbracket$ representing a time-varying graph of period p ; integers $\Delta \leq p$ and d .

OUTPUT: for every node u in \mathcal{G} , a list of all its (Δ, d) -twins, that is, a list of (v, t) 's with v being a node ID and t the starting time instant where u and v become d -twins for Δ instants.

Note that it is required that the period p be given among the input of every node, but not the graph \mathcal{G} . During each round t , every node has a local knowledge of \mathcal{G} , which is the size of its current neighbourhood, in

G_t . In the next Section 3, we prove that PERIODIC(Δ, d)-TWINS can be solved after $2p$ rounds, using messages of size $O(\delta_{\mathcal{G}} \log n)$, where n is the total number of nodes and $\delta_{\mathcal{G}} \leq n$ is the maximum degree of the graphs G_t 's.

3 Distributed computation of (Δ, d)-twins

The main idea of our solution is based on the following property of static graphs, which hints that every node can compare its 1-neighbourhood with the 1-neighbourhood of any other node in its 2-neighbourhood in order to determine whether that node is one of its d -twins. For $u \neq v$ two nodes of a static graph $G = (V, E)$, we say that u and v are d -twins if they have at least one common neighbour and that the difference between the sets of their outside neighbours is at most d , namely $|N(u) \setminus N(v) \setminus \{u, v\}| + |N(v) \setminus N(u) \setminus \{u, v\}| \leq d$. A path of length 2 between u and v is a path having 2 edges and 3 nodes, of the form (u, w, v) with $uw \in E$ and $vw \in E$, for some $w \neq u$ and $w \neq v$.

Property 1 *Let $G = (V, E)$ be an undirected static graph, let $u \neq v$ be two distinct nodes in V with a common neighbour. Then, u and v are d -twins if and only if $d \geq k$ or there are at least $k - d$ paths of length 2 between u and v , where $k = |N(u) \cup N(v) \setminus \{u, v\}|$.*

Proof. The case when $d \geq k$ is clear. Otherwise let $A = N(u) \setminus \{u, v\}$ and $B = N(v) \setminus \{u, v\}$. Then, $k = |A \cup B|$. The paths of length 2 between u and v are only paths passing through their common neighbours. Therefore, the number of such paths is $n_p = |A \cap B|$. We need to prove that u and v are d -twins if and only if $n_p \geq k - d$. From the definition of d -twins, we have $|A \setminus B| + |B \setminus A| \leq d$. Since $|A \cup B| - |A \cap B| = |A \setminus B| + |B \setminus A|$, this is equivalent to $|A \cup B| - |A \cap B| \leq d$, in other words, $k - n_p \leq d$. \square

Remark 2. Note that u and v could have very different neighbourhoods. For instance when they have no common neighbours, k could be large, while u and v are not d -twins until $d \geq k$.

Let $\mathcal{G} = (G_t)_{t \in T}$ be a p -periodic time-varying graph with $T = \llbracket 0, p-1 \rrbracket$. For every node to compute its list of d -twins, we would need 2 rounds per static graph G_t , so that the node collects information about its 2-neighbourhood. Since the time-varying graph is periodic, we can proceed a first phase of p rounds for passing every node's 1-neighbourhood information to its neighbours. Then, in a second phase consisting of p extra rounds, every node can forward all the pieces of information it has already

collected from its neighbours during the first phase. The information received in this second phase informs every node about its 2-neighbourhood. There are however two drawbacks of the above idea.

Firstly, we need to control the size of the messages sent from each node, especially in the second phase. In Lemma 2 below, we prove that the information related to the entire neighbourhood is not necessary, but only the size of it. Roughly, we plan to only send the size of the neighbourhood each time we need to forward this information. More specifically, following Lemma 2 each node u in \mathcal{G} can determine whether a node v in \mathcal{G} is such that u and v are d -twins at round t , by adding the number of u 's neighbours, obtained at round t , to the number of v 's neighbours, obtained at round $t' = t + p$ and subtracting the number of common neighbours between them. Then, if the result is less than d , they are d -twins, according to the inequality given in Lemma 2.

Secondly, we aim at computing (Δ, d) -twins in \mathcal{G} , with Δ representing a continuous window of Δ consecutive time instants starting at $t_0 \in T$. For most t_0 this can be done using in the receiver node's internal storage a counter per sender node, which increments every time the receiver node detects a sender node as its d -twin at some G_t . The counter is reset to zero whenever that sender node is detected as a non- d -twin at some (other) G_t . However, cases involving extremities of interval T such as when $t_0 + \Delta \geq p$ will have their time window divided into two disjoint intervals: $\llbracket 0, t_0 + \Delta - 1 \pmod p \rrbracket$ and $\llbracket t_0, p - 1 \rrbracket$. If we proceed as with all other cases, the information will be ready at some round after the $2p$ -th round for this case. We can accelerate this to be as early as in the $2p$ -th round using internal storage of a table in each receiver node, instead of the previously mentioned incremental counters. Another positive consequence of using a table instead of using incremental counters is that, starting from round $p + \Delta$, every node has access in real time to its (Δ, d) -twins of previous Δ rounds. Then, at the $2p$ -th round, every node has access to the full list of its (Δ, d) -twins, where the computation terminates.

Algorithm 1 implements the above ideas, calling as subroutines both Algorithms 2 and 3 at every round. All tables in Algorithm 1 are dictionary data structures. The result for every node u will be stored in table `Delta_d_Twins`. For its computation, node u maintains furthermore two tables, called `Twins` and `Count`. For a node $v \neq u$ with identifier id , a strictly positive value of `Twins[id]` means that v and u have been twin vertices for the last `Twins[id]` number of rounds. Hence, when `Twins[id]` is at least Δ , we append id to the resulting list `Delta_d_Twins` along

Algorithm 1 Finding the (Δ, d) -twins, all tables are dictionary data structures

```

1: procedure PERIODIC( $\Delta, d$ )-TWINS( $p, \Delta, d$ )
                                     ▷ Initialisation
2:   Table Tab with  $p$  associating to each round the set of neighbours and the number
   of their neighbours
3:   Table Count associating to a given ID, the number of times it has been received
   in the current round
4:    $nbMsg \leftarrow 0$                                      ▷ The number of messages received.
5:   Table Neighbors associating to a given ID its number of neighbours in the
   current round
6:   Table Twins storing the consecutive  $d$ -twins of the node, updated at each round,
   initially empty
7:   Table dTwinsRound associating to each round the found  $d$ -twins of the node
8:   Set Delta_d_Twins initially empty
9:   for each round do
10:    Call Algorithm 2                                     ▷ Sending a message
11:    Call Algorithm 3                                     ▷ For received messages
12:    if  $round = 2p - 1$  then
13:       $i \leftarrow 0$ 
14:      while  $i < p$  and Twins  $\neq \emptyset$  do
15:        for  $id$  in Twins do
16:          if  $id$  in dTwinsRound[ $i$ ] then
17:            Twins[ $id$ ]  $\leftarrow$  Twins[ $id$ ] + 1
18:            if Twins[ $id$ ] =  $\Delta$  then
19:              Append  $\{id, p + i - \Delta + 1\}$  to Delta_d_Twins
20:              Remove  $id$  from Twins
21:            else
22:              Remove  $id$  from Twins
23:             $i \leftarrow i + 1$ 
24:    return Delta_d_Twins

```

Algorithm 2 Send messages

```

1: procedure SENDMESSAGE( $p, \Delta, d, round$ )
2:   if  $round < p$  then
3:      $msg \leftarrow \langle myId, nbNeighbors \rangle$ 
4:     Send  $msg$  to all neighbours.
5:   else if  $p \leq round < 2 * p$  then
6:      $msg \leftarrow \emptyset$ 
7:      $index \leftarrow round - p$ 
8:     for  $elem \in$  Tab[ $index$ ] do
9:       Append  $elem$  to  $msg$ 
10:    Send  $msg$  to all neighbours

```

with the number of rounds where they start to be (Δ, d) -twins. This is implemented in Algorithm 3, lines 21-22.

Algorithm 3 Upon Receiving a Message **msg**

```

1: procedure RECEIVEMESSAGE( $p, \Delta, d, round$ )
2:   if  $round < p$  then
3:     Append  $msg$  to  $\text{Tab}[round]$ 
4:   else if  $p \leq round < 2 * p$  then
5:      $nbMsg \leftarrow nbMsg + 1$ 
6:     for  $\langle id, nb \rangle$  in  $msg$  do
7:        $\text{Count}[id] \leftarrow \text{Count}[id] + 1$ 
8:        $\text{Neighbors}[id] \leftarrow nb$ 
9:     if  $nbMsg = nbNeighbors$  then
10:      for each  $id$  in  $\text{Twins}$  do
11:        if  $id$  not in  $\text{Count}$  then
12:          Remove  $id$  from  $\text{Twins}$ 
13:      for each  $id$  in  $\text{Count}$  do
14:         $dTwin \leftarrow nbNeighbors + \text{Neighbors}[id] - 2 * \text{Count}[id]$ 
15:        if  $dTwin \leq d$  then
16:          Append  $id$  to  $d\text{TwinsRound}[round - p]$ 
17:          if  $id$  in  $\text{Twins}$  then
18:             $\text{Twins}[id] \leftarrow \text{Twins}[id] + 1$ 
19:          else
20:             $\text{Twins}[id] \leftarrow 1$ 
21:          if  $\text{Twins}[id] \geq \Delta$  then
22:            Append  $\{id, round - \Delta + 1\}$  to  $\text{Delta}_d\text{Twins}$ 
23:          else
24:            Remove  $id$  from  $\text{Twins}$ 

```

In case we conclude that id is not a twin with u , such as with Algorithm 3's lines 12 and 24, a quick way to save this information is to remove key id from the (dictionary) table Twins . Table Count is to be used in Algorithm 3 and related to the current round. We use an alternative computation for what is stored in table Count : rather than using Property 1 which would force us to count the number of paths of length 2, we use the equivalent quantities showed in Lemma 2 below instead. These quantities from Lemma 2 are encoded precisely at line 14 of Algorithm 3. The size of the message in the second phase is composed by two additive terms. The first term is the number of neighbours, multiplied by the size of the nodes ID's. The second term is the size of the neighbourhood of every neighbour, which is upper bounded by the first term. Thus, the maximum size of messages are bounded by $O(\delta_G \log n)$, where n is the total number of nodes and $\delta_G \leq n$ is the maximum degree of the graphs G_t 's, *cf.* Lemma 1 below. Fig. 3 exemplifies the messages received in the second phase of every node in the example given in Fig. 2.

Round 4		Round 5	
Sender node	Message	Sender node	Message
a	{<a, 2>, <b, 3>, <l, 2>}	a	{<a, 2>, <b, 3>, <l, 2>}
b	{<b, 3>, <a, 2>, <c, 2>, <c', 2>}	b	{<b, 3>, <a, 2>, <c, 2>, <c', 2>}
c	{<c, 2>, <b, 3>, <d, 3>}	c	{<c, 2>, <b, 3>, <d, 3>}
c'	{<c', 2>, <b, 3>, <d, 3>}	c'	{<c', 2>, <b, 3>, <d, 3>}
j	{<j, 3>, <i, 3>, <j', 3>, <k, 3>}	j	{<j, 2>, <i, 3>, <k, 3>}
j'	{<j', 3>, <i, 3>, <j, 3>, <k, 3>}	j'	{<j', 3>, <i, 3>, <k, 3>, <h, 3>}

Round 6		Round 7	
Sender node	Message	Sender node	Message
a	{<a, 3>, <b, 3>, <l, 2>, <m, 2>}	a	{<a, 3>, <b, 3>, <l, 2>, <m, 4>}
b	{<b, 3>, <a, 3>, <c, 2>, <c', 2>}	b	{<b, 3>, <a, 3>, <c, 4>, <c', 4>}
c	{<c, 2>, <b, 3>, <d, 3>}	c	{<c, 4>, <b, 3>, <d, 3>, <m, 4>, <o, 4>}
c'	{<c', 2>, <b, 3>, <d, 3>}	c'	{<c', 4>, <b, 3>, <d, 3>, <m, 4>, <o, 4>}
j	{<j, 3>, <i, 3>, <j', 3>, <k, 3>}	j	{<j, 2>, <i, 3>, <k, 3>}
j'	{<j', 3>, <i, 3>, <j, 3>, <k, 3>}	j'	{<j', 3>, <i, 3>, <k, 3>, <n, 3>}

Fig. 3. Tables describing the messages sent to the direct neighbourhood of nodes in the example presented in Fig. 2. The messages described here-above are sent through the second period of the time-varying graph, corresponding to the information collected during the first period.

We now prove the correctness of the algorithm, as well as the maximum size of the messages used in the algorithm.

Lemma 1. *Given a p -periodic time-varying graph $\mathcal{G} = (G_t)_{t \in T}$, with n nodes, if each node sends the list of IDs of its neighbours associated with the number of their neighbours collected at a round t in Algorithm 3 in line 3, then the maximum size of a message sent at round $t + p$ as stated in Algorithm 2 lines 6-11 is $O(\delta_{\mathcal{G}} \log(n\delta_{\mathcal{G}}))$.*

Proof. A node u can have at most $\delta_{\mathcal{G}} = n - 1$ neighbours. For every neighbour w of u , u needs to store the ID of w , which requires a size of $O(\log n)$, and the number of neighbours of w , which can be at most $\delta_{\mathcal{G}}$, requiring a size $\log \delta_{\mathcal{G}}$. So for a neighbour w of u , u stores information of size $O(\log n) + \log \delta_{\mathcal{G}} = O(\log n\delta_{\mathcal{G}})$, and for all neighbours of u , it stores information of size $O(\delta_{\mathcal{G}} \log n\delta_{\mathcal{G}})$, and this information represents the size of the message that u will send in the second phase. \square

Lemma 2. *Given a p -periodic time-varying graph $\mathcal{G} = (G_t)_{t \in T}$, and two distinct nodes $u \neq v$, u and v are (Δ, d) -twins if $|N_t(u) \setminus \{u, v\}| + |N_t(v) \setminus \{u, v\}| - 2 \times |(N_t(u) \setminus \{u, v\}) \cap (N_t(v) \setminus \{u, v\})| \leq d$ for $t_0 \leq t < t_0 + \Delta$ that is, for Δ consecutive rounds.*

Proof. The proof is very similar to that of Property 1. Let $A = N_t(u) \setminus \{u, v\}$ and $B = N_t(v) \setminus \{u, v\}$. From the definition of d -twins, we have $|A \setminus B| + |B \setminus A| \leq d$. Since $|A| + |B| - 2 \times |A \cap B| = |A \setminus B| + |B \setminus A|$, this is equivalent to $|A| + |B| - 2 \times |A \cap B| \leq d$. \square

The latter lemma allows us to store the d -twins in Algorithm 3 and check the number of consecutive round two nodes are d -twins in Algorithm 3 line 26 and Algorithm 1 line 20. We have proved the following theorem.

Theorem 1. *In a p -periodic time-varying graph $\mathcal{G} = (G_t)_{t \in T}$ where integers p, Δ, d are given as input to every node, problem PERIODIC(Δ, d)-TWINS can be solved after $2p$ rounds, using messages of size $O(\delta_{\mathcal{G}} \log n)$, where n is the total number of nodes and $\delta_{\mathcal{G}} \leq n$ is the maximum degree of the graphs G_t 's.*

4 Twin sampling with $O(\log n)$ message size

In a static graph, we can reduce the message size as follows. According to Lemma 2 in [6], every node v can receive its neighbourhood in a first round, then apply a well selected hash function sampled from a universe \mathcal{U} of hash functions and forward this to every neighbour u in a second round. When receiving the second round message, node u can compute the value of $|N(u) \cap N(v)|$ within $\epsilon \max(n_u, n_v)$ with probability $1 - \nu$, where $n_u = |N(u)|$ and $n_v = |N(v)|$. The process uses $O(1)$ messages of size $O(\epsilon^4 \log(1/\nu) + \log \log |\mathcal{U}| + \log \max(n_u, n_v))$ bits. Whence, w.h.p. after 2 rounds the inequality in our Lemma 2 can be decided using messages of $O(\log n)$ bits. For the dynamic case, the extension is similar to construction proposed in Section 3.

5 Conclusion and perspectives

We introduce the problem of finding (Δ, d) -twins and propose a distributed algorithm to compute them in any p -periodic time-varying graph $\mathcal{G} = (G_t)_{t \in \mathbb{N}}$ under a distributed model similar to the 1-interval-connected network. After $2p$ rounds, every node can compute the nodes that are its (Δ, d) -twins using messages of size $O(\delta_{\mathcal{G}} \log n)$, where n is the total number of nodes and $\delta_{\mathcal{G}}$ is the maximum degree of the graphs G_t 's. Using techniques borrowed from [6], we reduce the message size down to $O(\log n)$ w.h.p. Finding (Δ, d) -twins can be useful in several ways. For instance, it could be used for alternately scheduling sleeping times of the twin nodes

to save resources, while maintaining connectivity for the rest of the network. Furthermore, it can be used in order to compute disjoint paths or disjoint broadcast trees. As for the next steps of research, it would be useful to extend our algorithm to compute ϵ -modules as defined in [5].

References

1. Bonnet, E., Kim, E., Thomassé, S., Watrigant, R.: Twin-width I: tractable FO model checking. In: 61st IEEE Annual Symposium on Foundations of Computer Science. pp. 601–612. IEEE (2020)
2. Bui-Xuan, B., Hourcade, H., Miachon, C.: Computing temporal twins in time logarithmic in history length. In: 9th International Conference on Complex Networks and their Applications. SCI, vol. 944, pp. 651–663 (2020)
3. Casteigts, A., Flocchini, P., Mans, B., Santoro, N.: Measuring temporal lags in delay-tolerant networks. *IEEE Trans. Computers* **63**(2), 397–410 (2014)
4. Casteigts, A., Flocchini, P., Quattrociocchi, W., Santoro, N.: Time-varying graphs and dynamic networks. *International Journal of Parallel, Emergent and Distributed Systems* **27**(5), 387–408 (2012)
5. Habib, M., Mouatadid, L., Zou, M.: Approximating modular decomposition is hard. In: 6th International Conference on Algorithms and Discrete Applied Mathematics. LNTCS/LNCS, vol. 12016, pp. 53–66 (2020)
6. Halldórsson, M., Nolin, A., Tonoyan, T.: Overcoming congestion in distributed coloring. In: 2022 ACM Symposium on Principles of Distributed Computing. pp. 26–36. ACM (2022)
7. Kuhn, F., Lynch, N., Oshman, R.: Distributed computation in dynamic networks. In: 42nd ACM Symposium on Theory of Computing. pp. 513–522. ACM (2010)
8. Luna, G.D., Viglietta, G.: Computing in anonymous dynamic networks is linear. In: 63rd Annual Symposium on Foundations of Computer Science. pp. 1122–1133. IEEE (2022)
9. O’Dell, R., Wattenhofer, R.: Information dissemination in highly dynamic graphs. In: DIALM-POMC Joint Workshop on Foundations of Mobile Computing. pp. 104–110. ACM (2005)