

1 Tight Bounds on Distributed Exploration of 2 Temporal Graphs

3 **Tsuyoshi Gotoh**

4 Osaka University, Japan

5 **Paola Flocchini**

6 University of Ottawa, Canada

7 **Toshimitsu Masuzawa**

8 Osaka University, Japan

9 **Nicola Santoro**

10 Carleton University, Canada

11 — Abstract —

12 Temporal graphs (or evolving graphs) are time-varying graphs where time is assumed to be discrete.
13 In this paper, we consider for the first time the problem of exploring temporal graphs of *arbitrary*
14 *unknown topology*. We study the feasibility of exploration, under both the FSYNC and SSYNC
15 schedulers, focusing on the number of agents necessary and sufficient to explore such graphs.

16 We first consider the minimal (i.e., less restrictive) assumption on the dynamics of the graph
17 under which exploration is still feasible: *temporal connectivity*. Let \mathcal{H} be the class of temporally
18 connected graphs; we show that for any temporal graph $\mathcal{G} \in \mathcal{H}$ the number of agents sufficient
19 to perform exploration is related to the number of its transient edges, a parameter $\eta(\mathcal{G})$ we call
20 evanescence of the graph. More precisely, any $\mathcal{G} \in \mathcal{H}$ can be explored by a team of $k \geq 2\eta(\mathcal{G}) + 1$
21 agents; this bound is tight as we prove there are $\mathcal{G} \in \mathcal{H}$ that cannot be explored by $2\eta(\mathcal{G})$ agents.

22 We then turn our attention to the well-known stronger assumption on the dynamics of the graph,
23 called *1-interval connectivity*: the graph is connected at any time step. Let $\mathcal{W} \subset \mathcal{H}$ be the class
24 of these always-connected temporal graphs. For this class, we prove the existence of a difference
25 between FSYNC and SSYNC when there is a bound ℓ on the number of edges missing at each time.
26 In fact, we show a tight bound of $2\ell + 1$ on the number of agents necessary and sufficient in SSYNC,
27 and a smaller tight bound of 2ℓ in FSYNC. As a corollary, we re-establish two recently published
28 bounds for 1-interval connected rings.

29 **2012 ACM Subject Classification** Theory of computation \rightarrow Distributed algorithms; Theory of
30 Computation \rightarrow Distributed computing models;

31 **Keywords and phrases** Distributed algorithm, Mobile agents, Exploration of dynamic networks,
32 Arbitrary footprint

33 **Digital Object Identifier** 10.4230/LIPIcs.OPODIS.2019.23

34 **Acknowledgements** This research was partly supported by NSERC through the Discovery Grant
35 program, by Prof. Flocchini's University Research Chair,

36 **1** Introduction

37 **1.1** Framework and Background

38 The *graph exploration* problem (EXPLORATION), first introduced by Shannon [34], is a
39 fundamental problem in theoretical computer science, in particular in the field of distributed
40 computing by mobile entities. It requires each node of the graph to be visited by one or more
41 entities, called agents, a finite number of times (exploration *with termination*) or infinitely
42 often (*perpetual* exploration). In addition to its theoretical importance, EXPLORATION is
43 relevant from a practical viewpoint in networks with mobile entities (e.g., software agents,



© Tsuyoshi G., Paola F., Toshimitsu M. and Nicola S.;
licensed under Creative Commons License CC-BY

23rd International Conference on Principles of Distributed Systems (OPODIS 2019).

Editors: Pascal Felber, Roy Friedman, Seth Gilbert, and Avery Miller; Article No. 23; pp. 23:1–23:16

Leibniz International Proceedings in Informatics



Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

44 vehicles, or robots): by visiting all nodes, agents can check whether there are some nodes
 45 with problems in the network, propagate some data across the network, or collect (or search)
 46 specific information from the whole network.

47 This problem has been extensively studied over a variety of assumptions and settings
 48 depending on whether the nodes have distinct labelings or are anonymous, on the type of
 49 communication mechanisms available to the agents, on the degree of synchronization of the
 50 network, on the level of knowledge the agents have about the graph, on their memory, etc.
 51 (e.g., see [1, 8, 7, 10, 13, 14, 21, 22, 33, 35], and [9] for a recent survey). In spite of all the
 52 differences, the existing literature has until very recently made a common assumption: the
 53 graph is *static*, i.e., the link structure does not change during the exploration.

54 Recently, researchers in the distributed computing community have started to investigate
 55 *highly dynamic graphs*, that is graphs where the topological changes are not sporadic or
 56 anomalous, but rather inherent in the nature of the network. Various models have been
 57 proposed to describe highly dynamic networks, under a variety of names. A model that
 58 describes them in a simple and natural way is the one of *time-varying graphs*, formally
 59 defined in [6], where main classes of systems studied in the literature and their computational
 60 relationship were identified. When time is assumed to be *discrete*, the evolution of these
 61 systems can be equivalently described as a sequence of static graphs, called *evolving graph* or
 62 *temporal graph*, a model suggested in [25], formalized in [17].

63 If the dynamics of the changes is arbitrary and unrestricted, clearly any non-trivial
 64 computation is unfeasible and any non-trivial problem is unsolvable. Hence, all the studies
 65 are carried out under some assumptions restricting the arbitrariness of the dynamics. The
 66 minimal (i.e., less restrictive) assumption is *temporal connectivity*: starting at any time, there
 67 is temporal reachability between any two nodes (e.g., [5]). Stronger assumptions include
 68 *1-interval connectivity*: the graph is always connected (e.g., [24, 30, 31]); and *T-interval*
 69 *connectivity*: the graph is always connected and every $T > 1$ consecutive rounds contain
 70 the same spanning-tree (e.g., [28, 30]). A classification of the most common assumption was
 71 done in [6].

72 While there are several studies on computations by mobile agents moving in temporal
 73 graphs (for a recent survey see [11]), the results on the *exploration of temporal graphs* are
 74 rather limited. On the probabilistic side, there is an early seminal work on random walks [2].
 75 On the deterministic side there are: the study of the complexity of computing a foremost
 76 exploration schedule under the 1-interval-connectivity assumption [32], generalized and
 77 extended in [15] and then in [16]; the computation of an exploration schedule for *rings* under
 78 the stronger T-interval-connectivity assumption [28]; the computation of an exploration
 79 schedule for *cactuses* under the 1-interval-connectivity assumption [26]. These studies are
 80 however *centralized* (or off-line); that is, they assume that the exploring agents have complete
 81 a priori knowledge of the topological changes and the times of their occurrence. *Distributed*
 82 approaches have been studied under particular constraints on the network connectivity and on
 83 its underlying topology. Exploration with termination by a single agent of periodic temporal
 84 networks, including *carrier networks*, has been studied in [18, 19, 27, 28]. Exploration with
 85 termination of 1-interval connected *rings* by two and three agents under both synchronous
 86 and semi-synchronous schedulers has been considered in [12]. Perpetual exploration by three
 87 agents on temporally connected *rings* has been studied in [4, 5]. Perpetual exploration by
 88 $O(n)$ agents of $n \times m$ dynamic *tori* ($n \leq m$), where each column and row is a 1-interval
 89 connected ring, has been investigated in [23].

90 All the existing results on distributed exploration of time-varying graphs have been
 91 obtained for temporal graphs with very specific topologies (rings, tori, or collections of cycles

92 in the case of carrier networks). In this paper we start the investigation of the exploration of
 93 temporal graphs with *arbitrary* and *unknown* topologies.

94 1.2 Contributions

95 In this paper we consider perpetual exploration of time varying graphs whose topology is
 96 arbitrary and unknown to the agents. We focus on solvability of the exploration of such
 97 dynamic graphs and we determine the number of agents that are necessary and sufficient for
 98 exploration under the FSYNC and SSYNC activation schedulers.

99 Clearly, if the graph is not *temporally connected*, perpetual exploration is trivially im-
 100 possible to achieve. We thus start our investigation with the class \mathcal{H} of temporally connected
 101 temporal graphs. We show that for the graphs $\mathcal{G} \in \mathcal{H}$, the number of agents sufficient to
 102 perform exploration is related to the evanescence $\eta(\mathcal{G})$ of the graph, that is the number of
 103 transient edges. More precisely, any $\mathcal{G} \in \mathcal{H}$ can be explored by a team of $k \geq 2\eta(\mathcal{G}) + 1$
 104 agents; this bound is tight as we prove there are $\mathcal{G} \in \mathcal{H}$ that cannot be explored by $2\eta(\mathcal{G})$
 105 agents. The impossibility holds under very strong conditions (FSYNC scheduler, agents and
 106 nodes with distinct IDs, knowledge on n and k). On the other hand, the proposed exploration
 107 algorithm, based on the rotor router technique, works under very weak conditions (SSYNC
 108 scheduler, anonymous agents, no knowledge of topological parameters).

109 We then turn our attention to the stronger assumption on the dynamics of the graph,
 110 *1-interval connectivity*: the graph is always connected. Let $\mathcal{W}(\ell) \subset \mathcal{H}$ be the class of these
 111 always-connected temporal graphs where the number of missing edges at each time is at most
 112 ℓ . For this class, we first show a tight bound of $2\ell + 1$ under the SSYNC scheduler on the
 113 number of agents. We then prove the existence of a difference between FSYNC and SSYNC
 114 if the network size and the number of agents are known. In fact, in this case, while the
 115 bound for SSYNC remains unchanged, we prove a tight bound of 2ℓ for FSYNC. Moreover,
 116 we show that if $2\ell + 1$ agents are available in SSYNC, the exploration with termination
 117 is possible. As a corollary of these results, we re-establish a recently published bound for
 118 temporally-connected rings [5] and one for 1-interval connected rings [12].

119 Note that, when considering the class $\mathcal{H}(\ell)$ of temporally connected graphs with at most ℓ
 120 transient edges and the class $\mathcal{W}(\ell) \subset \mathcal{H}(\ell)$ of ℓ -bounded 1-interval connected graph, we have
 121 that the bound on the number of agents for $\mathcal{H}(\ell)$ is the same as the one for $\mathcal{W}(\ell)$ for SSYNC,
 122 while the two differs in the case of FSYNC, showing that the stronger connectivity assumption
 123 of \mathcal{W} does not influence the solvability bound in case of semi-synchronous schedulers, but
 124 does have an impact for fully synchronous ones.

125 2 The Model

126 2.1 The Network

127 The system is modeled as a time-varying graph (TVG), $\mathcal{G} = (V, E, \mathbb{T}, \rho)$, where V is a set of
 128 nodes, E is a set of edges, \mathbb{T} is the temporal domain, and $\rho : E \times \mathbb{T} \rightarrow \{0, 1\}$, called *presence*
 129 *function*, indicates whether a given edge is available at a given time. The graph $G = (V, E)$
 130 is called *underlying* graph (or *footprint*) of \mathcal{G} , with $|V| = n$ and $|E| = m$. Let $E(v)$ denote
 131 the set of edges incident on node v in the footprint, let $\delta_v = |E(v)|$ be the degree of node v
 132 in the footprint, and let $\Delta = \text{Max}_v \{\delta_v\}$ be the maximum degree of G .

133 In this paper we consider *discrete* time; that is, $\mathbb{T} = \mathbb{Z}^+$. Since time is discrete, the
 134 dynamics of the system can be viewed also in terms of a sequence of static graphs: $\mathcal{S}_{\mathcal{G}} =$
 135 $G_0, G_1, \dots, G_t, \dots$, where $G_t = (V_t, E_t)$ is the graph of the edges present at time t (also

23:4 Tight Bounds on Distributed Exploration of Temporal Graphs

136 called *snapshot* at time t). The TVG in this case is called *temporal graph* (or *evolving graph*).
 137 We denote by $\bar{E}_t = E \setminus E_t (\subseteq E)$ the set of edges that do not appear in the snapshot at time
 138 t .

139 In a temporal graph, the edge set E can be partitioned into the set of recurrent edges
 140 E^* , and the one of transient edges E^- . Formally, a *recurrent edge* $e^* \in E^*$ is such that
 141 $\forall t \in \mathbb{Z}^+, \exists t' > t : \rho(e^*, t') = 1$. In other words, a recurrent edge appears infinitely often. On
 142 the other hand, a *transient edge* $e^- \in E^-$ is such that $\exists t \in \mathbb{Z}^+, \forall t' \geq t : \rho(e^-, t') = 0$. In
 143 other words, a transient edge eventually ceases to exist forever.

144 The *solidity* of \mathcal{G} is defined as the number $\sigma(\mathcal{G})$ of recurrent edges, and the *evanescence*
 145 of \mathcal{G} , denoted by $\eta(\mathcal{G})$, as the number of transient edges (i.e., $\eta(\mathcal{G}) = |E| - \sigma(\mathcal{G})$).

146 A *journey* is a temporal walk in \mathcal{G} and it is defined as a sequence of couples $\mathcal{J} = \{(e_1, t_1),$
 147 $(e_2, t_2) \dots, (e_k, t_k)\}$, such that $\{e_1, e_2, \dots, e_k\}$ is a walk in G and $\forall i, 1 \leq i < k, \rho(e_i, t_i) = 1$
 148 and $t_{i+1} > t_i$. Let $J(u, v, t)$ denote the set of journeys from u to v starting at time $t' \geq t$.

149 A particularly important class of temporal graphs are *temporally connected* ones:

150 **► Definition 1 (Temporally connected).** *A TVG \mathcal{G} is temporally connected (or connected over*
 151 *time) if $\forall t \in \mathbb{Z}^+, \forall u, v \in V, J(u, v, t) \neq \emptyset$.*

152 Note that temporal connectivity is the minimal condition to be able to perform any global
 153 tasks; in particular, perpetual exploration (i.e., requiring every node to be visited infinitely
 154 often) is trivially impossible if the graph is not temporally connected. Let \mathcal{H} denote the class
 155 of temporally connected TVGs.

156 A variety of stronger assumptions have been studied in the literature. In this paper we
 157 are interested in a particular temporally connected graph, where connectivity is actually
 158 guaranteed at every time (*always connected* or *1-interval connected* temporal graphs); in
 159 particular, when the number of missing edges at any given time is bounded.

160 **► Definition 2 (ℓ -Bounded 1-Interval Connected).** *A temporal graph \mathcal{G} is 1-interval connected*
 161 *(or always connected) if $\forall G_i \in \mathcal{S}_{\mathcal{G}}, G_i$ is connected. Moreover, \mathcal{G} is ℓ -bounded 1-interval*
 162 *connected if it is always connected and $|\bar{E}_t| \leq \ell$.*

163 Let $\mathcal{W}(\ell) \subset \mathcal{H}$ denote the class of ℓ -bounded 1-interval connected temporal graphs.

164 The nodes of \mathcal{G} are anonymous (i.e., they have no IDs) and each node provides a constant
 165 amount of local memory called *whiteboard*. Each edge incident to node v is locally labeled by
 166 a bijection $\lambda_v : E \rightarrow \{0, \dots, \delta_v - 1\}$; no other assumptions are made about the labels. Every
 167 node v has ports p_i for $0 \leq i \leq \delta_v$ which are used to store at most one agent trying to move
 168 through e such that $\lambda_v(e) = i$.

169 2.2 Mobile agents

170 A set $A = \{a_0, a_1, \dots, a_{k-1}\}$ of k agents operate on the network, initially occupying arbitrary
 171 positions. Agents are anonymous and have access to their private notebook (local memory)
 172 and to whiteboards (memory of nodes).

173 The agents operate in synchronous rounds, and each round is composed by three phases:
 174 LOOK, COMPUTE, and MOVE, during which they execute the following actions [20]:

175 **LOOK:** Agent a_i observes the content of its own notebook and of the whiteboard of the
 176 node it occupies, and it checks, for each port of the node, if there are other agents at the
 177 same node.

178 COMPUTE: On the basis of the information obtained in the LOOK phase, agent a_i decides
 179 whether to move or not. It can write information on the whiteboard¹ and, if it decides
 180 to move, it places itself in correspondence of the selected port (if it is not occupied by
 181 another agent).

182 MOVE: If a_i occupies a port, it tries to move. If the corresponding edge exists, a_i reaches
 183 the other side, otherwise it stays on the port to try again at the next round. If a_i does
 184 not occupy a port, it does not move.

185 We distinguish between the *fully-synchronous* activation scheduler (FSYNC), when all
 186 the agents are activated in every round, and the *semi-synchronous* one (SSYNC), when an
 187 arbitrary subset of the agents is activated at each round. In SSYNC, the scheduler is an
 188 adversary which knows the algorithm of the agents, has infinite computing capacity, and
 189 tries to prevent agents from completing their task; however, it must activate every agent
 190 infinitely often. An agent which is not activated at round t is said to be *sleeping* at that
 191 round. The length of the sleeping time is finite but unbounded.

192 Under the semi-synchronous scheduler, we need to specify the behavior of the agents
 193 that fall asleep on a port when the corresponding edge is missing. In this paper, we assume
 194 the weakest rule, called *eventual transport rule* [12], in which the agent sleeping at a port
 195 will eventually be activated at a time when the edge corresponding to the port is present.
 196 This prevents the adversary from using semi-synchronicity to block an agent forever on a
 197 recurrent edge.

198 2.3 Configuration and execution

199 A configuration C_t is defined by: the contents of the whiteboards, the local memory of the
 200 agents, the locations of the agents, and the snapshot G_t of the temporal graph in the sequence
 201 \mathcal{S}_G , at round t . An execution $\mathcal{E}^A = C_0C_1\dots$ of an algorithm \mathcal{A} is an infinite sequence of
 202 configurations such that C_0 is an initial configuration (i.e., a configuration at round 0) and
 203 C_{t+1} is obtained from C_t by executing one round of algorithm \mathcal{A} . This execution is subject
 204 to two types of adversarial actions: those by the activation scheduler deciding which agents
 205 are activated in that round, and those of the topological scheduler deciding which edges are
 206 missing in that round. When no ambiguity arises, we use \mathcal{E} instead of \mathcal{E}^A .

207 2.4 The Exploration problem

208 We say that a node v is visited by round t if there exists a round t' ($0 \leq t' < t$) such that an
 209 agent occupies v at time t' . We say that the network is explored by round t if every node
 210 has been visited by round t .

211 A *perpetual* exploration algorithm is one where, in every execution, every node is visited
 212 infinitely often. An exploration *with termination* algorithm is one where the agents terminate
 213 when all nodes have been visited at least once. In this paper we are concerned with *perpetual*
 214 *exploration*.

215 3 Exploration of temporally connected TVGs

216 In this section, we show that the feasibility of exploration of temporally connected TVGs is
 217 related to their evanescence.

¹ Access to the whiteboard is done in fair mutual exclusion

218 **3.1 Impossibility**


219 Let $\mathcal{H}(\ell) = \{\mathcal{G} \in \mathcal{H} : \eta(\mathcal{G}) \leq \ell\}$ be the class of temporally connected TVGs with evanescence
 220 at most ℓ . In this section we show that it is impossible to perform perpetual exploration
 221 of all $\mathcal{G} \in \mathcal{H}(\ell)$ with 2ℓ agents. The result is quite strong as it applies also to TVGs
 222 that are connected at every time step, with uniquely labeled nodes and agents, under a
 223 fully-synchronous scheduler, and in presence of topological knowledge.

224 **► Theorem 3.** *There exist temporally connected time-varying graphs $\mathcal{G} \in \mathcal{H}(\ell)$ that cannot*
 225 *be explored by $k = 2\ell$ agents. The result holds even if nodes and/or agents have distinct IDs,*
 226 *the network is always connected, the agents have some topological knowledge (n , m or k),*
 227 *and the scheduler is fully-synchronous.*

228 **Proof.** We show the theorem by constructing a graph $\mathcal{G} \in \mathcal{H}(\ell)$ that cannot be explored by
 229 2ℓ agents by any algorithm. The main point of this proof is that an agent can eventually have
 230 only one of these two behaviors when wishing to traverse an edge that is missing: (i) the
 231 agent stays permanently on the chosen port, waiting for the appearance of the continuously
 232 missing edge; (ii) the agent eventually chooses a different edge. The former type of agents
 233 are called (with respect to the number of changes of a selected edge) *finite* and the latter
 234 *infinite*.

235 The components for constructing the graph are as follows. For $0 \leq i \leq 2\ell - 1$ ($= k - 1$),
 236 let S_i^{inf} be a star with center node c_i^{inf} and 3 leaf nodes $\{b_{(i,0)}^{\text{inf}}, b_{(i,1)}^{\text{inf}}, b_{(i,2)}^{\text{inf}}\}$ and S_i^{fin} be a star
 237 with center node c_i^{fin} and 3 leaf nodes $\{b_{(i,0)}^{\text{fin}}, b_{(i,1)}^{\text{fin}}, b_{(i,2)}^{\text{fin}}\}$. We construct the graph using S_i^{inf} ,
 238 S_i^{fin} and an additional node u .

239 Each component is connected as follows. For S_i^{inf} ($0 \leq i \leq 2\ell - 1$) and u , each $b_{(i,j)}^{\text{inf}}$
 240 ($0 \leq j \leq 2$) is connected with u by edge $(b_{(i,j)}^{\text{inf}}, u)$. For S_i^{fin} ($0 \leq i \leq 2\ell - 1$) and u , each $b_{(i,j)}^{\text{fin}}$
 241 ($j = 0$ or 1) is connected with u by edge $(b_{(i,j)}^{\text{fin}}, u)$. In addition to that, for $0 \leq i \leq \ell - 1$,
 242 $b_{(2i,2)}^{\text{fin}}$ and $b_{(2i+1,2)}^{\text{fin}}$ are connected by $(b_{(2i,2)}^{\text{fin}}, b_{(2i+1,2)}^{\text{fin}})$. A graph for $\ell = 2$ ($k = 4$) is depicted
 243 in Figure 1.



■ **Figure 1** Example of a graph for $\ell = 2$ and $k = 2\ell = 4$. There are four stars S_i^{fin} (S_i^{inf}) for $0 \leq i \leq 3$ on the top (bottom) of the figure. Each star S_i^{fin} (S_i^{inf}) has one center node c_i^{fin} (c_i^{inf}) and three leaf nodes $\{b_{(i,0)}^{\text{fin}}, b_{(i,1)}^{\text{fin}}, b_{(i,2)}^{\text{fin}}\}$ ($\{b_{(i,0)}^{\text{inf}}, b_{(i,1)}^{\text{inf}}, b_{(i,2)}^{\text{inf}}\}$).

244 For the constructed graph, we first show that, given any exploration algorithm using 2ℓ
 245 agents, the adversary can construct an execution for the algorithm such that in the execution

246 \mathcal{G} cannot be explored while the adversary may violate the restriction of $\mathcal{H}(\ell)$, i.e., $\eta(\mathcal{G})$ may
 247 be more than ℓ . Then, we give a way to convert the execution into another execution such
 248 that $\eta(\mathcal{G})$ is at most ℓ in the new execution and the agents cannot distinguish these two
 249 executions and thus cannot explore \mathcal{G} also in the new execution.

250 We start by showing that, given any exploration algorithm, say \mathcal{A} , using 2ℓ agents, the
 251 adversary can construct an execution \mathcal{E}_1 of \mathcal{A} in which the agents cannot explore \mathcal{G} . The
 252 adversary puts agent a_i on c_i^{inf} for $0 \leq i \leq 2\ell - 1$ in the initial configuration of \mathcal{E}_1 . During
 253 execution \mathcal{E}_1 of \mathcal{A} , the adversary deletes edge $(b_{(i,j)}^{\text{inf}}, u)$ whenever a_i is on $b_{(i,j)}^{\text{inf}}$. Clearly, this
 254 prevents any agent executing \mathcal{A} to visit u and thus \mathcal{G} is not explored permanently while the
 255 adversary violates the restriction for the number of transient edges (it is at most 2ℓ in \mathcal{E}_1).

256 We now show how the adversary converts \mathcal{E}_1 into another execution, say \mathcal{E}_2 , so that
 257 the agents cannot distinguish \mathcal{E}_1 and \mathcal{E}_2 and $\eta(\mathcal{G})$ is at most ℓ in \mathcal{E}_2 . To decide the initial
 258 configuration of \mathcal{E}_2 , the adversary first separates the agents into two groups: *finite agents*
 259 and *infinite agents* depending on their behavior when faced with a missing edge during \mathcal{E}_1 .
 260 Let f ($0 \leq f \leq k - 1$) be the number of *finite agents*. In the following, *finite agents* are
 261 denoted by $a_0^{\text{fin}}, \dots, a_{f-1}^{\text{fin}}$, and the *infinite agents* are denoted by $a_0^{\text{inf}}, \dots, a_{k-f-1}^{\text{inf}}$. W.l.o.g.,
 262 we assume that $a_i^{\text{fin}} = a_i$, i.e., a_i^{fin} is the agent starting from c_i^{inf} in \mathcal{E}_1 .

263 The adversary decides the initial configuration of \mathcal{E}_2 as follows: each a_i^{inf} ($0 \leq i \leq k - f - 1$)
 264 is put on the same node as in the initial configuration of \mathcal{E}_1 , while each a_i^{fin} ($0 \leq i \leq f - 1$)
 265 is put on c_i^{fin} .

266 Then, the adversary changes the assignment of the port labels and the node ID (if any)
 267 of c_i^{fin} , $b_{(i,0)}^{\text{fin}}$, $b_{(i,1)}^{\text{fin}}$, and, $b_{(i,2)}^{\text{fin}}$ in S_i^{fin} so that a_i^{fin} cannot distinguish \mathcal{E}_1 and \mathcal{E}_2 . Let $v_i = b_{(i,x)}^{\text{fin}}$
 268 be the node where $a_i = a_i^{\text{fin}}$ finally waits a missing edge permanently in \mathcal{E}_1 . For $b_{(i,2)}^{\text{fin}}$, the
 269 assignment of the port labels and the node ID (if any) are copied from v_i . The ones of c_i^{fin}
 270 are copied from c_j^{inf} . The ones of $b_{(i,0)}^{\text{fin}}$ and $b_{(i,1)}^{\text{fin}}$ are copied from each of $b_{(i,y)}^{\text{inf}}$ for $y \neq x$.

271 Execution \mathcal{E}_2 with the initial configuration, the node ID, and, the assignment of port
 272 labels is constructed similarly to \mathcal{E}_1 : the adversary deletes the edge leading to u (resp, u or
 273 $S_{i'}^{\text{fin}}$ for $i' \neq i$) when $a_{i'}^{\text{inf}} = a_i$ (resp, a_i^{fin}) exists on $b_{(i,j)}^{\text{inf}}$ (resp, $b_{(i,j)}^{\text{fin}}$). Obviously, every agent
 274 cannot distinguish \mathcal{E}_1 and \mathcal{E}_2 , since the difference between \mathcal{E}_1 and \mathcal{E}_2 is only the order of the
 275 port labeling (and the node degrees are fixed). Thus, \mathcal{G} cannot be explored since u is not
 276 visited by any agent in \mathcal{E}_2 .

277 Finally, we show that, in \mathcal{E}_2 , $\eta(\mathcal{G})$ is at most ℓ . To prevent infinite agents, no transient
 278 edge is necessary; in fact, an infinite agent eventually changes its selected edge if it is kept
 279 missing, and no two infinite agents wait on the same edge (otherwise, the edge may be
 280 transient). For finite agents, by construction, a_{2i}^{fin} and a_{2i+1}^{fin} for $0 \leq i \leq (f - 1)/2$ eventually
 281 wait for the same edge $(b_{(2i,2)}^{\text{fin}}, b_{(2i+1,2)}^{\text{fin}})$ (when f is odd, only a_{f-1} waits for $(b_{(f-1,2)}^{\text{fin}}, b_{(f,2)}^{\text{fin}})$).
 282 Since f is at most $k = 2\ell$, at most ℓ edges are necessary to prevent finite agents. ◀

283 3.2 Semi Synchronous Exploration by $2\eta(\mathcal{G}) + 1$ agents

284 In this section, we show that every temporally connected time-varying network $\mathcal{G} \in \mathcal{H}$ can be
 285 explored by $2\eta(\mathcal{G}) + 1$ anonymous agents that do not know the topology. In fact, we propose
 286 an exploration algorithm for $2\eta(\mathcal{G}) + 1$ anonymous agents in an anonymous network, which
 287 works under the semi-synchronous scheduler with eventual transport.

288 The strategy is simple and it is based on the classical *rotor router* mechanism, which was
 289 introduced as a deterministic alternative to random walk and was studied in a variety of
 290 contexts, including static graph exploration (e.g., [3, 29, 35]).

291 In rotor router, each node v has a variable written on its whiteboard, pointer_v , indicating
 292 one of its incident ports. When an agent a visits node v , a checks each port in ascending

23:8 Tight Bounds on Distributed Exploration of Temporal Graphs

293 order from the port pointed by pointer_v . If a finds some unoccupied port p , a moves to that
 294 port and sets pointer_v to $p + 1$. If a finishes to check all the ports and they all are occupied,
 295 a does nothing.

Algorithm 1 Computation at node v

```

1: if not on a port then
2:    $i \leftarrow 0$ 
3:    $p \leftarrow \text{pointer}_v$ 
4:   while  $i < \delta_v \wedge$  port  $p$  is occupied do
5:      $p \leftarrow (p + 1) \bmod \delta_v$ 
6:      $i \leftarrow i + 1$ 
7:   if  $i < \delta_v$  then
8:      $\text{pointer}_v \leftarrow (p + 1) \bmod \delta_v$ 
9:     move to port  $p$ 

```

296 We first show that in any round, there exists at least one agent succeeding to move within
 297 finite time (Lemma 4). We then show that, $2l + 1$ agents achieve perpetual exploration using
 298 Algorithm 1 (Theorem).

299 ► **Lemma 4.** *For any round t , if $2\eta(\mathcal{G}) + 1$ agents execute Algorithm 1 in a temporally*
 300 *connected temporal graph \mathcal{G} , at least one of them eventually moves within finite time after t .*

301 **Proof.** By contradiction, assume that there exists a round t such that every agent never
 302 succeeds to move after t . We consider two cases: (i) there exists a node v containing more
 303 than $\delta_v - 1$ agents, and (ii) there does not exist such a node.

304 In the first case, every agent on v is activated within finite time after t because of the
 305 fairness of the scheduler, which means that every port of v is eventually occupied by an agent.
 306 Since at least one of the edges incident to v is a recurrent edge, say e , the agent sleeping on
 307 the corresponding port of e eventually succeeds to move because of the eventual transport
 308 rule. This is a contradiction.

309 Also in the second case, every agent on v is activated within finite time after round t
 310 because of the fairness of the scheduler. Since there is no node containing more agents than
 311 its degree, every agent eventually stays on a port. When this happens, at least one of the
 312 agents is sleeping at the port of a recurrent edge since the number of agents is $2\eta(\mathcal{G}) + 1$
 313 and there exist at most $2\eta(\mathcal{G})$ ports corresponding to transient edges. This means that, by
 314 the eventual transport rule, the agent sleeping at the port of a recurrent edge eventually
 315 succeeds to move after t ; a contradiction. ◀

316 Then, the following theorem holds.

317 ► **Theorem 5.** *Any $\mathcal{G} \in \mathcal{H}$ can be explored by $2\eta(\mathcal{G}) + 1$ anonymous agents under the*
 318 *semi-synchronous scheduler.*

319 **Proof.** Consider Algorithm 1. By definition of transient edges, there exists a time step t_e
 320 such that, for any transient edge e , $\rho(e, t) = 0$ for all $t > t_e$. Let t_E be $\max_{e \in E} t_e$, i.e., a
 321 time when all the transient edges have ceased to exist and all the edges that appear from
 322 this moment are recurrent. Let $x(t)$ be the sum of the number of visits over all the nodes
 323 from the beginning of the execution up to time t .

324 We now show that, from an arbitrary initial configuration, $2\eta(\mathcal{G}) + 1$ agents following
 325 Algorithm 1 visit all the nodes infinitely often.

326 First, note that there exists a node, say v , that is visited infinitely often (for $t \rightarrow \infty$)
 327 because $x(t)$ goes to infinity for $t \rightarrow \infty$) by Lemma 4.

328 We now show that every neighbor of v connected by a recurrent edge is also visited
 329 infinitely often. We prove it by contradiction. Suppose that a neighbor u of v connected by
 330 a recurrent edge is visited only a finite number of times and let t' be the last round when u
 331 is visited. Since v is visited infinitely often and the agents execute Algorithm 1 perpetually,
 332 some agent a visiting v eventually chooses (v, u) as the edge from which a moves out of
 333 v after time t' . Recall that (v, u) is a recurrent edge and the agents are activated by the
 334 eventual transport rule. It follows that a eventually visits u after round t' ; a contradiction.

335 Since G_r is temporally connected, we can apply inductively the claim (e.g., the neighbors
 336 of a neighbor of v is also visited infinitely often) to all the nodes, proving the theorem. ◀

337 From Theorems 3 and 5, the following Theorem holds.

▶ **Theorem 6.** *Exploration of all temporal graphs in $\mathcal{H}(\ell)$ is possible iff*

$$k \geq 2\ell + 1$$

338 Note that, if a graph is temporally connected, then its rigidity $\sigma(\mathcal{G}) \geq n - 1$; as a
 339 consequence, we have:

340 ▶ **Theorem 7.** *Every temporally connected temporal graph can be explored by $2(m - n) + 3$*
 341 *agents.*

342 **4 Exploration of 1-interval connected temporal graphs with bounded** 343 **missing edges**

344 In this Section, we turn our attention to the class $\mathcal{W}(\ell)$ of 1-interval connected temporal
 345 graphs where the number of missing edges is bounded in each round by a constant ℓ . In
 346 other words, at any time t the TVG is connected, and no more than ℓ edges are missing. We
 347 establish tight bounds for the exploration of this class of temporal graphs, in SSYNC and in
 348 FSYNC.

349 **4.1 Semi-synchronous model**

350 We first consider ℓ -bounded, 1-interval connected TVGs operating under a semi-synchronous
 351 scheduler and we show that there exists TVGs that cannot be explored by 2ℓ agents.

352 ▶ **Theorem 8.** *There exist 1-interval connected time-varying graphs $\mathcal{G} \in \mathcal{W}(\ell)$ that cannot*
 353 *be explored by $k = 2\ell$ anonymous agents. The result holds even if nodes and/or agents have*
 354 *distinct IDs and the agents have some topological knowledge (n , m or k).*

355 **Proof.** We use the same graph \mathcal{G} constructed for the proof of Theorem 3. The construction
 356 is omitted in this proof.

357 We first show that, given any exploration algorithm, say \mathcal{A} , using 2ℓ agents, the adversary
 358 can construct an execution \mathcal{E}_1 of \mathcal{A} , possibly violating the eventual transport rule, in which
 359 the agents cannot explore \mathcal{G} . We then show that it is always possible to convert this execution
 360 into another execution \mathcal{E}_2 that does not violate the eventual transport rule, and where the
 361 agents cannot explore \mathcal{G} .

362 In execution \mathcal{E}_1 , the adversary puts agent a_i on c_i^{inf} for $0 \leq i \leq k - 1 = 2\ell - 1$ in initial
 363 configuration of \mathcal{E}_1 . During \mathcal{E}_1 , exactly one agent is activated at each round: a_i is activated

at round t when $t \equiv i \pmod{k}$. When the adversary activates a_i and a_i exists on $b_{(i,j)}^{\text{inf}}$, the adversary deletes $(b_{(i,j)}^{\text{inf}}, u)$ whereas all the other edges are present. Note that the agents and the nodes are anonymous and thus either they are all *finite* (i.e., every agent permanently waits for appearance of its selected edge if the edge is permanently missing) or they are all *infinite* (i.e., every agent eventually changes its selected edge if the edge remains missing) in \mathcal{E}_1 . If the agents are *infinite*, since the eventual transport rule is not violated even in \mathcal{E}_1 , the adversary can prevent the agents from completing the exploration in \mathcal{E}_1 . If the agents are *finite*, the adversary converts \mathcal{E}_1 into another execution, say \mathcal{E}_2 , as follows. The adversary first puts a_i ($0 \leq i \leq k-1$) on c_i^{fin} in the initial configuration of \mathcal{E}_2 . Then, the adversary changes the assignment of the port labels and the node ID (if any) of c_i^{fin} , $b_{(i,0)}^{\text{fin}}$, $b_{(i,1)}^{\text{fin}}$, and, $b_{(i,2)}^{\text{fin}}$ in the same way explained in the proof of Theorem 3 (also omitted in this proof). In \mathcal{E}_2 , the adversary activates each agent in the same order as in \mathcal{E}_1 and deletes an edge leading to u or $S_{i'}$ for $i' \neq i$ whenever a_i is on $b_{(i,j)}^{\text{fin}}$. After some round t from when every agent a_i does not change its selected edge at $b_{(i,2)}^{\text{fin}}$ for $0 \leq i \leq 2l$, the adversary deletes $(b_{(2j,2)}^{\text{fin}}, b_{(2j+1,2)}^{\text{fin}})$ for $0 \leq j \leq l-1$ at every round. Obviously, every agent cannot distinguish \mathcal{E}_2 from \mathcal{E}_1 and \mathcal{G} cannot be explored since u is not visited by any agent in \mathcal{E}_2 . It is also clear that the eventually transport rule is not violated in \mathcal{E}_2 . ◀

Clearly, $\mathcal{W}(\ell) \subset H(\ell)$, thus any $\mathcal{G} \in \mathcal{W}(\ell)$ can be explored by Algorithm 1; that is:

► **Theorem 9.** *Any $\mathcal{G} \in \mathcal{W}(\ell)$ can be explored by $2\ell + 1$ anonymous agents under the semi-synchronous scheduler.*

From Theorems 8 and 9 it follows that:

► **Theorem 10.** *Under a semi-synchronous scheduler, exploration of all ℓ -bounded 1-interval connected TVG \mathcal{G} is possible iff*

$$k \geq 2\ell + 1$$

4.2 Fully-synchronous model


In this section, we show that, if the network size and the number of agents are known, there exists a difference between FSYNC and SSYNC in the exploration of ℓ -bounded 1-interval TVGs. In fact, we show that, $\mathcal{G} \in \mathcal{W}(\ell)$ can be explored if $k \geq 2\ell$, while there exist graphs that cannot be explored with $2\ell - 1$ agents.

4.2.1 Impossibility

We now consider ℓ -bounded, 1-interval connected TVGs operating under a fully-synchronous scheduler and we show that there exists TVGs that cannot be explored by $2\ell - 1$ agents, even if the agents know n , m , and k .

► **Theorem 11.** *There exist 1-interval ℓ -bounded time-varying graphs $\mathcal{G} \in \mathcal{W}(\ell)$ that cannot be explored by $k = 2\ell - 1$ anonymous agents in FSYNC. The result holds even if the agents have some topological knowledge (n, m, k) .*

Proof. Let $K_{2\ell} = (V_{2\ell}, E_{2\ell})$ be the complete graph with 2ℓ nodes where $V_{2\ell} = \{v_0, v_1, \dots, v_{2\ell-1}\}$. It is well known that the edges of $K_{2\ell}$ can be colored with $2\ell - 1$ colors, that is, $E_{2\ell}$ can be partitioned into $2\ell - 1$ disjoint independent edge sets (or complete matchings): $E_{2\ell}^{(0)}, E_{2\ell}^{(1)}, \dots, E_{2\ell}^{(2\ell-2)}$. For example, the following separation leads to disjoint independent edge sets: each $E_{2\ell}^{(i)}$ has ℓ edges, $(v_i, v_{2\ell-1}), (v_{i-1}, v_{i+1}), (v_{i-2}, v_{i+2}), \dots, (v_{i-l+1}, v_{i+l-1})$, see Figure 2 (for simplicity, mod 2ℓ is omitted).



■ **Figure 2** Example of coloring for the proof of Lemma 11. The bold lines are the edges of $E_8^{(0)}$.

403 The execution where $v_{2\ell-1}$ remains unvisited is constructed as follows. For $0 \leq i \leq 2\ell - 1$,
 404 the adversary places each agent a_i on v_i and for $0 \leq j \leq 2\ell - 1$ assigns a label j to the port
 405 of v_i corresponding to e , if $e \in E_{2\ell}^{(j)}$. Note that, since agents and nodes are anonymous, all
 406 the agents select the port with the same label to move at each round. Thus, the adversary
 407 can prevent any agent from moving by deleting all the edges of $E_{2\ell}^{(i)}$ when the agent selects
 408 port i ; as a consequence, none of the agents can move out of its current nodes. This means
 409 that $v_{2\ell-1}$ remains unvisited forever.

410 In this execution, the number of missing edges is always ℓ and the network is obviously
 411 kept connected. Thus, the theorem holds. ◀

412 4.2.2 Bound on Exploration time

413 Let $\mathcal{G} \in \mathcal{W}(\ell)$. Since $\mathcal{W}(\ell) \subset H(\ell)$, we can clearly execute Algorithm 1 in graph \mathcal{G} .
 414 Interestingly, when executed on $\mathcal{G} \in \mathcal{W}(\ell)$, it can be shown that the time complexity of
 415 exploration can be bounded under the fully-synchronous scheduler. More specifically, we
 416 show that within $\Delta^n(\Delta + 1)^k(n - 1)^k$ rounds, all nodes of the graph have been visited at
 417 least once by a team of $2\ell + 1$ agents.

418 We prove the theorem by a sequence of lemmas. First of all, we can easily show that
 419 $2\ell + 1$ agents executing Algorithm 1 cannot be all prevented from moving at any given round.

420 ▶ **Lemma 12.** *If $2\ell + 1$ agents activated fully-synchronously execute Algorithm 1 in ℓ -bounded
 421 1-interval TVGs, at least one of them succeeds to move at every round.*

422 **Proof.** There exist two cases as in the proof of Lemma 4: at round t , (i) there exists a node
 423 v containing more than $\delta_v - 1$ agents, and (ii) there does not exist such a node.

424 In the first case, since there are more than $\delta_v - 1$ agents at v , every port is occupied
 425 by one agent at t since every agent is activated. In addition to that, v has at least one
 426 adjacent edge present at t by the connectivity of the TVG. This implies that at least one
 427 agent succeeds to move at round t .

428 In the second case, each agent occupies one port by assumption and by fully-synchronous
 429 activation, which means that $2\ell + 1$ ports are occupied. Moreover, at most ℓ edges are
 430 missing at each round, which means that at most 2ℓ ports are blocked at each round. It
 431 follows that at least one agent can move at round t also in this case.

432 ◀

23:12 Tight Bounds on Distributed Exploration of Temporal Graphs

433 To show the upper-bound on time complexity, we introduce the notions of *augmented*
434 *configuration* and *augmented execution*.

435 In an augmented configuration C_t^{aug} , a new variable visited_v written and read only by
436 an external observer, is added to each node v . The initial value of visited_v is 0. When v is
437 visited, visited_v is set to 1 by the external observer. An augmented configuration C_t^{aug} is
438 defined by configuration C_t and the value of visited_v of every node v at round t . We say that
439 an augmented configuration is *terminal* when $\text{visited}_v = 1$ for any node v .

440 An augmented execution $\mathcal{E}^{\text{aug}} = C_0^{\text{aug}} C_1^{\text{aug}} \dots C_r^{\text{aug}}$ is a sequence of augmented configura-
441 tions such that C_0^{aug} is an initial augmented configuration; C_{t+1}^{aug} is obtained from C_t^{aug} by
442 $2\ell + 1$ agents executing one round of Algorithm 1 fully-synchronously, with the action of the
443 adversary deciding which edges are missing; C_r^{aug} is a unique terminal configuration in \mathcal{E}^{aug} .
444 Note that the agents keep executing Algorithm 1 after round r , but augmented configurations
445 after round r are ignored in \mathcal{E}^{aug} . For \mathcal{E}^{aug} , the following lemma holds.

446 **► Lemma 13.** *In an augmented execution by $2\ell + 1$ agents, any two augmented configurations*
447 *are different.*

448 **Proof.** First note that Lemma 12 precludes the same two consecutive augmented configura-
449 tions C_t^{aug} and C_{t+1}^{aug} in an augmented execution where no agents move between C_t^{aug} and
450 C_{t+1}^{aug} . Suppose that there exist two augmented configurations C_t^{aug} and $C_{t'}^{\text{aug}}$ for $t < t'$ in an
451 augmented execution \mathcal{E}^{aug} . Let $\mathcal{E}_{t,t'}^{\text{aug}} = C_t^{\text{aug}} C_{t+1}^{\text{aug}} \dots C_{t'-1}^{\text{aug}}$ be a subsequence of \mathcal{E}^{aug} . In this
452 case, the adversary can create an infinite augmented execution from \mathcal{E}^{aug} by repeating $\mathcal{E}_{t,t'}^{\text{aug}}$,
453 which means that the adversary can create an (augmented) execution where $2\ell + 1$ agents
454 cannot complete the exploration forever. This contradicts Theorem 5. Thus, the lemma
455 holds. ◀

456 We are now ready to show an upper bound on the exploration time of Algorithm 1, which
457 is obtained by calculating the maximum length among all the augmented executions.

458 **► Lemma 14.** *The length of any possible augmented execution by $k = 2\ell + 1$ agents is*
459 *bounded by $\Delta^n (\Delta + 1)^k (n - 1)^k$.*

460 **Proof.** Let α be the maximum length among all the possible augmented executions. By
461 Lemma 13, α is bounded by the number of possible augmented configurations in an execution.

462 The number of possible configurations on a fixed node set $V' \subseteq V$ is bounded by
463 $\Delta^{|V'|} (|V'| (\Delta + 1))^k$, which corresponds to all the combinations of the directions of pointers
464 (i.e., $\Delta^{|V'|}$) and all of the the agents' locations (i.e., $(|V'| (\Delta + 1))^k$). Notice that only pointer_v
465 of each node v is used as a variable in Algorithm 1. Since the number of visited nodes is
466 not decreasing during the exploration, the exploration time is smaller than or equal to the
467 sum of $\Delta^{|V'|} (|V'| (\Delta + 1))^k$ for $1 \leq |V'| \leq n - 1$, i.e., $\alpha \leq \sum_{|V'|=1}^{n-1} \Delta^{|V'|} (|V'| (\Delta + 1))^k \leq$
468 $\Delta^n (\Delta + 1)^k (n - 1)^k$ rounds. ◀

469 It then follows that:

470 **► Theorem 15.** *Under a fully-synchronous scheduler, Algorithm 1 executed by $k = 2\ell + 1$*
471 *anonymous agents explores any ℓ -bounded 1-interval connected TVG within $\Delta^n (\Delta + 1)^k (n - 1)^k$*
472 *rounds.*

473 Note that, as a consequence, we obtain a terminating exploration algorithm for ℓ -bounded
474 1-interval connected TVGs.

475 **► Theorem 16.** *With knowledge of n and k , exploration with termination of an arbitrary ℓ -*
476 *bounded 1-interval connected temporal graph $\mathcal{W}(\ell)$ can be achieved in $\Delta^n (\Delta + 1)^{2\ell+1} (n - 1)^{2\ell+1}$*
477 *rounds by $2\ell + 1$ agents under the fully synchronous scheduler.*

4.2.3 Exploration by 2ℓ agents

The result of the previous section can be used to obtain a perpetual exploration algorithm of ℓ -bounded 1-interval connected graphs by 2ℓ agents (which know n and k). The solution (Algorithm 2 below) is obtained by applying Algorithm 1 bounding the waiting time of an agent blocked on a missing edge.

In fact, while an agent keeps waiting for a missing edge forever in Algorithm 1, in Algorithm 2, an agent waits for a missing edge up to kT rounds where T is calculated on the basis of the results of Section 4.2.2.

Apart from the waiting time, the rest of the algorithm is the same as in Algorithm 1: each node has `pointerv` pointing to a port. When a visits v , a checks each port in ascending order from the port pointed by `pointerv`. If a finds some unoccupied port p , a moves to the port and sets `pointerv` to $p + 1$. If a finishes to check all the ports and they all are occupied, a does nothing.

Variable `Waiting` of an agent represents the elapsed time since the last round when the agent moved to the port.

Algorithm 2 Computation at node v

```

1: if on a port then
2:   Waiting  $\leftarrow$  Waiting + 1
3:   if Waiting >  $kT$  then
4:     exit the current port
5: if not on a port then
6:   Waiting  $\leftarrow$  0
7:    $i \leftarrow 0$ 
8:    $p \leftarrow \text{pointer}_v$ 
9:   while  $i < \delta_v \wedge$  port  $p$  is occupied do
10:     $p \leftarrow (p + 1) \bmod \delta_v$ 
11:     $i \leftarrow i + 1$ 
12:   if  $i < \delta_v$  then
13:     pointerv  $\leftarrow (p + 1) \bmod \delta_v$ 
14:     move to the port  $p$ 

```

► **Lemma 17.** *Let 2ℓ agents execute Algorithm 2. If an agent waits at u for a missing edge $e = (u, v)$ for kT rounds, during this time either another agent starts to wait for e at v , or the other $2\ell - 1$ agents complete the exploration.*

Proof. Suppose that an agent a at u starts to wait for a missing edge (u, v) at round t and (u, v) is kept missing for the next kT rounds (including t).

We first show that there exist T successive rounds in $[t, t + kT)$ during which all the agents but a keep waiting without satisfying predicate `Waiting` > kT at its chosen edge, if it remains missing.

We show the claim by contradiction. We assume that in any interval of T successive rounds in $[t, t + kT)$, there is an agent that satisfies `Waiting` > kT .

By assumption, at least k agents other than a must satisfy `Waiting` > kT , since $kT/T \geq k$. This means that at least one agent (different from a) satisfies the predicate twice since the number of the agents (excluding a) is $k - 1$. However, once an agent satisfies `Waiting` > kT at round $t' \in [t, t + kT)$, the agent never satisfies the predicate in $[t, t + kT)$ since the length of the interval is kT . This is a contradiction. Thus, there exist T successive rounds in $[t, t + kT)$

23:14 Tight Bounds on Distributed Exploration of Temporal Graphs

508 during which all the agents (except for a) keep waiting their chosen edge without satisfying
 509 `Waiting` $> kT$ if the edge is kept missing.

510 Now, we show the lemma, i.e., show that another agent at v starts to wait for $e = (u, v)$
 511 or the exploration is completed. Suppose that no agent at v starts to wait for e in these T
 512 rounds. Since e is missing during these T rounds, during that time the network (without
 513 e) can be considered as a $(\ell - 1)$ -bounded 1-interval connected TVG. By Theorem 15,
 514 $2(\ell - 1) + 1 = 2\ell - 1$ agents complete the exploration of the $(\ell - 1)$ -bounded TVGs in these
 515 T rounds. This means that the $2\ell - 1$ agents other than a complete the exploration of the
 516 network without e during those T rounds, because none of them starts to wait for e at v
 517 during that time by assumption. Thus, the lemma holds. \blacktriangleleft

518 **► Theorem 18.** *Any ℓ -bounded 1-interval connected temporal graph $\mathcal{G} \in \mathcal{W}(\ell)$ can be explored*
 519 *by $k = 2\ell$ anonymous agents with knowledge of n and k , under a fully-synchronous scheduler.*

520 **Proof.** The proof follows the same lines of Theorem 5. We first show that, executing
 521 Algorithm 2, there exists at least one node v which is visited infinitely often, and we then
 522 show that all the nodes are visited infinitely often. Let $x(t)$ be the sum of the number of
 523 visits over all the node from the beginning of the execution up to time t and $V_A^{(t)}$ be a node
 524 set such that there exists at least one agent on every $w \in V_A^{(t)}$ at round t .

525 We show that $x(t)$ goes to infinity (for $t \rightarrow \infty$), which leads to the existence of a node
 526 v visited infinitely often. We consider the configuration at round t and show that after t
 527 rounds, $x(t)$ eventually increases. Two cases are considered: *Case 1*) there exists a node
 528 $\hat{v} \in V_A^{(t)}$ with $\delta_{\hat{v}}$ or more agents and *Case 2*) there does not exist such a node.

529 *Case 1*) Suppose that there exists a node \hat{v} with $\delta_{\hat{v}}$ or more agents at round t . Note that
 530 at least one of the edges incident to \hat{v} exists at round t because the network is 1-interval
 531 connected. In this case, at least one of the agents on \hat{v} succeeds to move because all the
 532 ports of \hat{v} are occupied. Therefore, $x(t)$ increases.

533 *Case 2*) Suppose that there does not exist a node \hat{v} with $\delta_{\hat{v}}$ or more agents. We show
 534 that $x(t)$ increases within finite rounds from t by contradiction. We assume that no agent
 535 moves out of its current node after t . Clearly, there exists a node $\tilde{v} \in V_A^{(t)}$ which has a
 536 neighbor \tilde{u} with no agent (otherwise, the exploration would have been completed). An agent
 537 changes its port if it is blocked by the same missing edge for kT rounds by Algorithm 2;
 538 an agent \tilde{a} on \tilde{v} eventually chooses (\tilde{v}, \tilde{u}) to move from \tilde{v} . At this round, the adversary
 539 must prevent \tilde{a} from moving by deleting (\tilde{v}, \tilde{u}) . This means that the adversary must prevent
 540 $2(\ell - 1) + 1 = 2\ell - 1$ agents from moving by deleting $\ell - 1$ edges, which is impossible. This
 541 leads to a contradiction. Therefore, $x(t)$ increases and goes to infinity for $t \rightarrow \infty$, and thus a
 542 node (say v) visited infinitely often exists.

543 We now show that all the neighbors of v are also visited by agents infinitely often. We
 544 prove it by contradiction. Suppose that a neighbor u of v is visited only a finite number of
 545 times and let t' be the last round when u is visited. Since v is visited infinitely often and the
 546 agents execute Algorithm 2, some agent a visiting v eventually chooses (v, u) as the edge
 547 from which a moves after t' . If (v, u) appears by the kT -th round since a chose it, a visits u
 548 as soon as (v, u) appears. Otherwise, another agent visits u by Lemma 17. It follows that u
 549 is eventually visited after t' rounds, which is a contradiction.

550 By the connectivity assumption, we can apply inductively the claim (e.g., the neighbors
 551 of a neighbor of v are also visited infinitely often) to all the nodes, proving the theorem. \blacktriangleleft

552 From Theorems 11 and 18, we have:

553 **► Theorem 19.** *Under the fully-synchronous scheduler, with knowledge of n and k , the*
 554 *exploration of all ℓ -bounded 1-interval connected TVGs is possible iff $k \geq 2\ell$.*

5 Conclusion

In this paper, we considered perpetual exploration of temporal graphs with arbitrary topology, focusing on the number of agents that are necessary and sufficient to perform the task. We considered two common dynamic models: temporally connected networks, and 1-interval connected (or always connected) networks with a bounded number of missing edges at each round. We derived tight bounds for both models under fully synchronous and semi-synchronous settings.

This is the first study on distributed exploration of temporal graphs with *arbitrary topology* and it has considered only temporally connected and 1-interval connected networks: the investigation of other connectivity classes of temporal graphs with arbitrary topology is the main research direction left open.

In this paper the focus was exclusively on feasibility of exploration; clearly, an important avenue of investigation is also the design of efficient solutions, whenever they exist.

References

- 1 S. Albers and M. Henzinger. Exploring unknown environments. *SIAM Journal on Computing*, 29(4):1164–1188, 2000.
- 2 C. Avin, M. Koucky, and Z. Lotker. How to explore a fast-changing world. In *Proc. of the 35th Int. Colloquium on Automata, Languages and Programming (ICALP)*, pages 121–132, 2008.
- 3 E. Bampas, L. Gasieniec, N. Hanusse, D. Ilcinkas, R. Klasing, A. Kosowski, and T. Radzik. Robustness of the rotor-router mechanism. *Algorithmica*, 78(3):869–895, 2017.
- 4 M. Bournat, A.K. Datta, and S. Dubois. Self-stabilizing robots in highly dynamic environments. *Theoretical Computer Science*, 772:88–110, 2019.
- 5 M. Bournat, S. Dubois, and F. Petit. Computability of perpetual exploration in highly dynamic rings. In *Proc. IEEE 37th Int. Conference on Distributed Computing Systems (ICDCS)*, pages 794–804, 2017.
- 6 A. Casteigts, P. Flocchini, W. Quattrociocchi, and N. Santoro. Time-varying graphs and dynamic networks. *Int. Journal of Parallel, Emergent and Distributed Systems*, 27(5):387–408, 2012.
- 7 J. Chalopin, P. Flocchini, B. Mans, and N. Santoro. Network exploration by silent and oblivious robots. In *Proc. of 36th International Workshop on Graph Theoretic Concepts in Computer Science (WG)*, pages 208–219, 2010.
- 8 R. Cohen, P. Fraigniaud, D. Ilcinkas, A. Korman, and D. Peleg. Label-guided graph exploration by a finite automaton. *ACM Trans. Algorithms*, 4(4):1–18, 2008.
- 9 S. Das. *Graph Exploration with Mobile Agents*. Chapter 16 of [20], pages 403–422, 2019.
- 10 X. Deng and C. H. Papadimitriou. Exploring an unknown graph. *Journal of Graph Theory*, 32(3):265–297, 1999.
- 11 G.A. Di Luna. *Mobile Agents on Dynamic Graphs*. Chapter 20 of [20], pages 549–584, 2019.
- 12 G.A. Di Luna, S. Dobrev, P. Flocchini, and N. Santoro. Live exploration of dynamic rings. In *Proc. IEEE 36th International Conference on Distributed Computing Systems (ICDCS)*, pages 570–579, 2016.
- 13 Y. Dieudonné and A. Pelc. Deterministic network exploration by anonymous silent agents with local traffic reports. *ACM Transactions on Algorithms*, 11(2):Article No. 10, 2014.
- 14 S. Dobrev, L. Narayanan, J. Opatrny, and D. Pankratov. Exploration of high-dimensional grids by finite automata. In *Proc. 46th Int. Colloquium on Automata, Languages, and Programming (ICALP)*, pages 1–16, 2019.
- 15 T. Erlebach, M. Hoffmann, and F. Kammer. On temporal graph exploration. In *Proc. of 42th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 444–455, 2015.

- 603 16 T. Erlebach and J. T. Spooner. Faster exploration of degree-bounded temporal graphs. In
604 *Proc. of 43rd International Symposium on Mathematical Foundations of Computer Science*
605 *(MFCS)*, pages 1–13, 2018.
- 606 17 A. Ferreira. Building a reference combinatorial model for MANETs. *IEEE Network*, 18(5):24–
607 29, 2004.
- 608 18 P. Flocchini, M. Kellett, P. Mason, and N. Santoro. Searching for black holes in subways.
609 *Theory of Computing Systems*, 50(1):158–184, 2012.
- 610 19 P. Flocchini, B. Mans, and N. Santoro. On the exploration of time-varying networks. *Theor-*
611 *etical Computer Science*, 469:53–68, 2013.
- 612 20 P. Flocchini, G. Prencipe, and N. Santoro (Eds). *Distributed Computing by Mobile Entities*.
613 Springer, 2019.
- 614 21 P. Fraigniaud, D. Ilcinkas, G. Peer, A. Pelc, and D. Peleg. Graph exploration by a finite
615 automaton. *Theoretical Computer Science*, 345(2–3):331–344, 2005.
- 616 22 P. Fraigniaud, D. Ilcinkas, and A. Pelc. Impact of memory size on graph exploration capability.
617 *Discrete Applied Mathematics*, 156(12):2310–2319, 2008.
- 618 23 T. Gotoh, Y. Sudo, F. Ooshita, H. Kakugawa, and T. Masuzawa. Group exploration of
619 dynamic tori. In *Proc. IEEE 38th International Conference on Distributed Computing Systems*
620 *(ICDCS)*, pages 775–785, 2018.
- 621 24 B. Haeupler and F. Kuhn. Lower bounds on information dissemination in dynamic networks.
622 In *Proc. 26th International Symposium on Distributed Computing (DISC)*, pages 166–180,
623 2012.
- 624 25 F. Harary and G. Gupta. Dynamic graph models. *Mathematical and Computer Modelling*,
625 25(7):79–88, 1997.
- 626 26 D. Ilcinkas, R. Klasing, and A.M. Wade. Exploration of constantly connected dynamic
627 graphs based on cactuses. In *Proc. 21st International Colloquium Structural Information and*
628 *Communication Complexity (SIROCCO)*, pages 250–262, 2014.
- 629 27 D. Ilcinkas and A.M. Wade. On the power of waiting when exploring public transportation sys-
630 tems. In *Proc. 15th International Conference on Principles of Distributed Systems (OPODIS)*,
631 pages 451–464, 2011.
- 632 28 D. Ilcinkas and A.M. Wade. Exploration of the T-interval-connected dynamic graphs: the
633 case of the ring. *Theory Computing Systems*, 62(4):1144–1160, 2018.
- 634 29 A. Kosowski and D. Pajak. Does adding more agents make a difference? a case study of cover
635 time for the rotor-router. *J. Comput. Syst. Sci.*, 106:80–93, 2019.
- 636 30 F. Kuhn, N. A. Lynch, and R. Oshman. Distributed computation in dynamic networks. In
637 *Proc. 42nd ACM Symposium on Theory of Computing (STOC)*, pages 513–522, 2010.
- 638 31 F. Kuhn and R. Oshman. Coordinated consensus in dynamic networks. In *Proc. 30th*
639 *Symposium on Principles of Distributed Computing (PODC)*, pages 1–10, 2011.
- 640 32 O. Michail and P. Spirakis. Traveling salesman problems in temporal graphs. *Theoretical*
641 *Computer Science*, 634:1–23, 2016.
- 642 33 P. Panaite and A. Pelc. Exploring unknown undirected graphs. *Journal of Algorithms*,
643 33:281–295, 1999.
- 644 34 C. Shannon. Presentation of a maze-solving machine. In *Proc. of the 8th Conf. of the Josiah*
645 *Macy Jr. Foundation (Cybernetics)*, pages 173–180, 1951.
- 646 35 V. Yanovsky, A. Bruckstein, and I. Wagner. A distributed ant algorithm for efficiently patrolling
647 a network. *Algorithmica*, 37(3):165–186, 2003.