On the Radius of Nonsplit Graphs and Information Dissemination in Dynamic Networks

Matthias Függer, Thomas Nowak, Kyrill Winkler

▶ To cite this version:

HAL Id: hal-02946849
https://hal.archives-ouvertes.fr/hal-02946849

Submitted on 23 Sep 2020

On the Radius of Nonsplit Graphs and Information Dissemination in Dynamic Networks*

Matthias Függer*

CNRS, LSV, ENS Paris-Saclay, Université Paris-Saclay, Inria

Thomas Nowak*

Université Paris-Saclay, CNRS

Kyrill Winkler**

TU Wien

Abstract

A nonsplit graph is a directed graph where each pair of nodes has a common incoming neighbor. We show that the radius of such graphs is in $O(\log \log n)$, where n is the number of nodes. This is an exponential improvement on the previously best known upper bound of $O(\log n)$. We then generalize the result to products of nonsplit graphs.

The analysis of nonsplit graph products has direct implications in the context of distributed systems, where processes operate in rounds and communicate via message passing in each round: communication graphs in several distributed systems naturally relate to nonsplit graphs and the graph product concisely represents relaying messages in such networks. Applying our results, we obtain improved bounds on the dynamic radius of such networks, i.e., the maximum number of rounds until all processes have received a message from a common process, if all processes relay messages in each round. We finally connect the

dynamic radius to lower bounds for achieving consensus in dynamic networks.

1. Introduction

Consider a distributed system of $n \geq 1$ processes that operate in lock-step synchronous rounds. Let $[n] = \{1, \ldots, n\}$ be the set of processes. In a round, each process broadcasts a message and receives messages from a subset of other processes, specified by the directed communication graph $G = ([n], E)$ whose nodes are the processes and there is an edge (i, j) in E if and only if process j receives the message sent by process i.

The radius of communication graph G is the minimum number of rounds until all processes have (transitively) received a message from a common process. Its value thus poses a lower bound on the number of rounds until information, originating at a single process, can be spread over the entire network. Related applications are from disease spreading and opinion dynamics.

Of particular interest in distributed computing are networks that potentially change during the execution of an algorithm, be it due to faulty processes, faulty links, mobility of the involved agents, etc.; see, e.g., [1] for a comprehensive overview. We thus generalize the investigation of the radius of a communication graph G to the dynamic radius of a sequence of communication graphs $G_1, G_2, \ldots$. Here, it is assumed that in the above scenario of broadcasting distributed processes, the communication graph for round $t \geq 1$ is G_t. The *dynamic radius of the sequence $G_1, G_2, \ldots$* is the minimum number of rounds until all processes have (transitively) received a message from a common process.

1.1. Radius of Nonsplit Digraphs

A nonsplit digraph is a directed graph where each pair of nodes has *at least one* common incoming neighbor. In this work, we study the radius of nonsplit digraphs: with $\ell(i, j)$ denoting the length of the shortest path from node i to node j, the radius is $\min_i \max_j \ell(i, j)$.

In the undirected case, the radius is trivially bounded by the diameter of the graph, which is 2 in the case of nonsplit graphs. Undirected graphs where each pair of nodes has *exactly one* common neighbor, have been studied by Erdős et al. [2], who showed that they are exactly the windmill graphs, consisting of triangles that share a common node. Thus, their radius is 1.

As demonstrated by the example in Figure 1 with radius 3, these bounds do not hold for nonsplit digraphs. We will show the following upper bound:

**Theorem 1.** *The radius of a nonsplit digraph with n nodes is in $O(\log \log n)$.*

Figure 1: Nonsplit digraph with radius 3. For example, node 1 and 6 have common incoming neighbor 6, while nodes 1 and 5 have node 4 as common incoming neighbor.

1.2. Communication over Nonsplit Digraphs

Nonsplit digraphs naturally occur as communication graphs in classical fault-models and as models for dynamic networks.

In fact, it was shown in [3] that, in the more general case where all communication graphs are rooted, i.e., are required to contain a rooted spanning tree, one can in fact simulate nonsplit communication graphs. Since, conversely, every nonsplit communication graph is rooted as well, nonsplit communication graphs are a convenient and concise abstraction for the technically more cumbersome rooted communication graphs. Furthermore, several classical fault-models were shown to lead to nonsplit communication graphs [4], among them link failures, as considered in [5], and asynchronous message passing systems with crash failures [1]. Nonsplit digraphs thus represent a convenient abstraction to

3

these classical fault-models as well. We will show in Section 4 that nonsplit digraphs arising from the classical model of asynchronous messages and crashes have dynamic radius at most 2.

Theorem 2. *The dynamic radius of a network on n nodes whose communication graphs are all nonsplit is $O(\log \log n)$.*

The study of nonsplit digraphs is also motivated by the study of a central problem in distributed computing: Agreeing on a common value in a distributed system is a problem that lies at the heart of many distributed computing problems, occurs in several flavors, and thus received considerable attention in distributed computing. However, even modest network dynamics already prohibit solvability of exact consensus, where agents have to decide on a single output value that is within the range of the agents' initial values [5]. For several problems, e.g., distributed control, clock synchronization, load balancing, etc., it is sufficient to asymptotically converge to the same value (asymptotic consensus), or decide on values not too far from each other (approximate consensus). Charron-Bost et al. [3] showed that both problems are solvable efficiently in the case of communication graphs that may vary arbitrarily, but are required to be nonsplit.

Motivated by this work on varying communication graphs, we will show that the following generalization of Theorem 1 holds:

Theorem 2. *The dynamic radius of a network on n nodes whose communication graphs are all nonsplit is $O(\log \log n)$.*

Traditionally, information dissemination, also called rumor spreading, is studied w.r.t. either all-to-all message relay or the time it takes for a fixed process to broadcast its message to everyone [6, 7]. In dynamic networks with nonsplit communication graphs, however, such strong forms of information dissemination are impossible. This can easily be seen by constructing appropriate sequences of star graphs (with self-loops), which are nonsplit graphs with radius 1. One possibility is to analyze information dissemination in dynamic networks that (probabilistically) guarantee some stability from time k to time $k + 1$; see, e.g., the work by Clementi et al. [8].

In this work, we follow an alternative route: Indeed, one-to-all broadcast

4

of *some* process is readily achieved in dynamic networks without any stability guarantees, which is why we focus on this characteristic here. While it is certainly not as universal as the previously mentioned primitives, it turns out that this type of information dissemination is crucial for the termination time of certain consensus algorithms based on vertex-stable root components [9]. Furthermore, we show the following theorem, relating the dynamic radius and the termination time of arbitrary consensus algorithms:

Theorem 3. *If the dynamic radius of a sequence of communication graphs is k, then, in every deterministic consensus algorithm, some process has not terminated before time k.*

Finally, we note that the dynamic radius is also an upper bound for the number of rounds until a single process aggregates the data of all other processes, when we use the dual interpretation of an edge (i, j) in a communication graph as a message sent by j and received by i. Even though this might not be the desired form of data aggregation in a standard setting, in a scenario where the communication is so constrained that aggregation by an *a priori* selected process is simply unobtainable, such a weak form might still be useful to transmit the collected data to a dedicated sink at regular intervals, for example.

We give a brief overview on related work in the next section.

1.3. Related Work

Information dissemination among an ensemble of n participants is a fundamental question that has been studied in a grand variety of settings and flavors (see [10, 6, 7, 11] for various reviews on the topic). While traditional approaches usually assume a static underlying network topology, with the rise of pervasive wireless devices, more recently, focus has shifted to dynamically changing network topologies [12]. A useful way of viewing the distribution of information is to denote the pieces of information that should be shared among the participants as tokens. For instance, the all-to-all token dissemination problem investigates the complete dissemination of n initially distributed tokens. This

5

problem was studied in [12] with a focus on bounds for the time complexity of the problem, i.e., how long it takes at least, resp. at most, until n tokens have been received by everyone. Here, the participants employed a token-forwarding algorithm mechanism, where tokens are stored and forwarded but not altered.

In the model of [12], it was assumed that the communication graphs are connected and undirected. For this, a lower bound of $\Omega(n \log n)$ and an upper bound of $O(n^2)$ for all-to-all token dissemination was established in the case where n is unknown to the participants, they have to terminate when the broadcast is finished, and the system is 1-interval connected, i.e., the communication graphs are completely independent of each other. In contrast, if the communication graphs are weakly connected, directed, and rooted, in the worst case only one of the tokens may ever be delivered to all participants. This can be seen, for example, when considering a dynamic graph that produces the same directed path for every round. We note that this example also provides a trivial lower bound of $\Omega(n)$ rounds until one token is received by everyone for the first time. As far as we are aware, the best lower bound for directed paths of varying linear order was established in [13, Theorem 4.3] to be $\lceil (3n-1)/2 - 2 \rceil$ rounds. Studying directed graphs is desirable as they represent a weaker, more general model and wireless communication is often inherently directed, for example due to localized fading or interference phenomena [14, 15] such as the capture effect or near-far problems [16].

In [4], it was shown that the dynamic radius of a sequence of arbitrary nonsplit communication graphs is $O(\log n)$. Later, it was shown in [3] that the product of any $n-1$ rooted communication graphs is nonsplit. Put together, this means that the dynamic radius of a sequence of arbitrary rooted graphs is $O(n \log n)$. More recently, [13] provided an alternative proof for this fact that does not rely on the reduction to nonsplit graphs but instead uses a notion of influence sets. In addition to this, [13] provided linear $O(n)$ bounds in sequences of rooted trees with a constant number of leaves or inner nodes, established a dependency on the size of certain subtrees in sequences of rooted trees where the root remains the same, and investigated sequences of undirected trees.

6

2. Model and Definitions

We start with some definitions motivated by the study of information dissemination within a distributed system of n processes that operate in discrete, lock-step synchronous communication rounds. Starting with information being available only locally to each process, processes broadcast and receive information tokens in every round. We are interested in the earliest round where all processes have received an information token from a common process.

Clearly, the dissemination dynamics depends on the dynamics of the underlying network. For this purpose we define: A *communication graph on n nodes* is a directed graph $G = (V, E)$ with self-loops and the set of nodes $V = [n] = \{1, 2, \ldots, n\}$. For $i \in [n]$, let $\text{In}_i(G) = \{j \in [n] \mid (j, i) \in E\}$ denote the set of in-neighbors of i in G and $\text{Out}_i(G) = \{j \in [n] \mid (i, j) \in E\}$ denote its set of out-neighbors. Intuitively, communication graphs encode successful message reception within a round: an edge from i to j states that j received the message broadcast by i in this round.

A node $i \in [n]$ is called a *broadcaster in G* if it has an edge to all nodes, i.e., $\forall j \in [n] \colon (i, j) \in E$.

A communication graph $G = ([n], E)$ is *nonsplit* if every pair of nodes has a common incoming neighbor, i.e.,

$$\forall i, j \in [n]\ \exists k \in [n] \colon (k, i) \in E \wedge (k, j) \in E \ .$$

Given two communication graphs $G = ([n], E_G)$ and $H = ([n], E_H)$ on n nodes, define their *product graph* as $G \circ H = ([n], E_{G \circ H})$ where

$$(i, j) \in E_{G \circ H} \iff \exists k \in [n] \colon (i, k) \in E_G \wedge (k, j) \in E_H \ .$$

The empty product is equal to the communication graph $([n], E_\perp)$ which contains the self-loops (i, i) for all nodes i and no other edges. The graph product we use here is motivated by information dissemination within distributed systems of processes that continuously relay information tokens that they received: if k received i's information token in a round, and j received k's information to-

ken in the next round, then j received i's information token in the macro-round
formed by these two successive rounds.

Motivated by modeling communication networks that potentially change in
each round, we call each infinite sequence $\mathcal{G} = (G_1, G_2, G_3, \dots)$ of communica-
tion graphs on n nodes a *communication pattern on n nodes*. For every node
$i \in [n]$, define the *broadcast time $T_i(\mathcal{G})$ of node i in $\mathcal{G}$* as the minimum t such
that i is a broadcaster in the product of the first t communication graphs of $\mathcal{G}$.
If no such t exists, then $T_i(\mathcal{G}) = \infty$. The *dynamic radius $T(\mathcal{G})$ of $\mathcal{G}$* is the
minimal broadcast time of its nodes, i.e., $T(\mathcal{G}) = \min_{i \in [n]} T_i(\mathcal{G})$. Note that $T(\mathcal{G})$
is the earliest time, in terms of rounds, until that all nodes have received an
information token from a common node, given that the communication pattern
is $\mathcal{G}$.

A *network on n nodes* is a nonempty set of communication patterns on n
nodes; modeling potential uncertainty in a dynamic communication network. A
network's *dynamic radius* is defined as the supremum over all dynamic radii of its
communication patterns, capturing the worst-case of information dissemination
within this network.

3. The Dynamic Radius of Nonsplit Networks

In this section we show an upper bound on the dynamic radius of nonsplit
networks.

During this section, let $\mathcal{G} = (G_1, G_2, G_3, \dots)$ be a communication pattern
on n nodes in which every communication graph G_t is nonsplit.

In order to prove an upper bound on the dynamic radius of $\mathcal{G}$, we will prove
the existence of a relatively small set of $O(\log n)$ nodes that "infects" all other
nodes within only $O(\log \log n)$ rounds. Iteratively going back in time, it remains
to be shown that any such set is itself "infected" by an exponentially smaller
set within $O(\log \log n)$ rounds, until we reach a single node. It follows that this
single node has "infected" all nodes with its information token after $O(\log \log n)$
rounds.

Note that the strategy to follow "infection" back in time rather than considering the evolution of infected sets over time is essential in our proofs: it may very well be that a certain set of infected nodes cannot infect other nodes from some time on, since it only has incoming edges from nodes not in the set in all successive communication graphs. Going back in time prevents us to run into such dead-ends of infection.

For that purpose we define: Let $U, W \subseteq [n]$ be sets of nodes. We say that U *covers W in communication graph $G = ([n], E)$* if for every $j \in W$ there is some $i \in U$ that has an edge to j, i.e., $\forall j \in W \; \exists i \in U \colon (i, j) \in E$.

Now let $0 < t_1 \leq t_2$. We say that U *at time t_1 covers W at time t_2* if U covers W in the product graph $G_{t_1} \circ G_{t_1+1} \circ \cdots \circ G_{t_2-1}$.

Note that U at time t covers U at time t for all sets $U \subseteq [n]$ and all $t \geq 1$, by definition of the empty product as the digraph with only self-loops.

We first show that the notion of covering is transitive:

Lemma 1. *Let $0 < t_1 \leq t_2 \leq t_3$ and let $U, W, X \subseteq [n]$. If U at time t_1 covers W at time t_2, and W at time t_2 covers X at time t_3, then U at time t_1 covers X at time t_3.*

Proof. By definition, for all $k \in W$ there is some $i \in U$ such that (i, k) is an edge of the product graph $G_{t_1} \circ \cdots \circ G_{t_2-1}$. Also, for all $j \in X$ there is some $k \in W$ such that (k, j) is an edge of the product graph $G_{t_2} \circ \cdots \circ G_{t_3-1}$.

But, by the associativity of the graph product, this means that for all $j \in X$ there exists some $i \in U$ such that (i, j) is an edge in the product graph

$$\left(G_{t_1} \circ \cdots \circ G_{t_2-1}\right) \circ \left(G_{t_2} \circ \cdots \circ G_{t_3-1}\right) = G_{t_1} \circ \cdots \circ G_{t_3-1} \ .$$

That is, U at time t_1 covers X at time t_3. $\qquad\square$

We continue with some basic technical lemmas that we prove here for completeness.

Lemma 2. *For all $x \geq 1$ we have $\left\lceil \log_2 x \right\rceil = \left\lceil \log_2 \lceil x \rceil \right\rceil$.*

9

Proof. We have $\lceil \log_2 x \rceil = \min\{k \in \mathbb{Z} \mid x \leq 2^k\}$ if $x \geq 1$. Now, noting that the inequality $x \leq p$ is equivalent to $\lceil x \rceil \leq p$ whenever p is an integer concludes the proof. □

Lemma 3. *Let m and n be positive integers such that $|m - n| \leq 1$. Then $\lceil \log_2(m + n) \rceil \geq \lceil \log_2 m \rceil + 1$.*

Proof. By assumption we have $n \geq m - 1$. We distinguish between two cases for the positive integer m:

(i) If $m = 1$ then $n \in \{1, 2\}$, and we immediately obtain the lemma from $\lceil \log_2 2 \rceil = \lceil \log_2 1 \rceil + 1$ and $\lceil \log_2 3 \rceil = \log_2 4 = \lceil \log_2 2 \rceil + 1$.

(ii) Otherwise, $m \geq 2$. From $n \geq m-1$ we deduce $m+n \geq 2m-1$. This implies

$$\lceil \log_2(m + n) \rceil \geq \lceil \log_2(2m - 1) \rceil = \left\lceil \log_2\left(m - \frac{1}{2}\right) \right\rceil + 1 \ .$$

We are hence done if we can show $\left\lceil \log_2(m - \frac{1}{2}) \right\rceil = \lceil \log_2 m \rceil$. But this is just Lemma 2 with $x = m - \frac{1}{2} \geq 1$. □

Lemma 4. *Let n and m be positive integers such that $n \geq m$. Then there exist positive integers $n_1, n_2, \ldots, n_m$ such that $n = n_1 + \cdots + n_m$ and $\lceil \log_2 \frac{n}{m} \rceil \geq \lceil \log_2 n_i \rceil$ for all $1 \leq i \leq m$.*

Proof. Let $n = km + r$ with $k, r \in \mathbb{Z}$ and $0 \leq r < m$ be the integer division of n by m. Set $n_1 = n_2 = \cdots = n_r = k + 1$ and $n_{r+1} = n_{r+2} = \cdots = n_m = k$.

By Lemma 2, we have

$$\lceil \log_2 n_i \rceil \leq \left\lceil \log_2\left(k + \left\lceil \frac{r}{m} \right\rceil\right) \right\rceil = \left\lceil \log_2\left(k + \frac{r}{m}\right) \right\rceil = \left\lceil \log_2 \frac{n}{m} \right\rceil$$

for all $1 \leq i \leq m$. □

We continue with the following generalization of a result by Charron-Bost and Schiper [4]. In particular ($m = 1$), it shows that any set of nodes can be "infected" by a single node in such a way that the set of infected nodes grows exponentially in size per round.

Lemma 5. *Let $W \subseteq [n]$ be nonempty and m be a positive integer. If $t_2 - t_1 \geq \log_2 \frac{|W|}{m}$, then there exists some $U \subseteq [n]$ with $|U| \leq m$ such that U at time t_1 covers W at time t_2.*

Proof. Using Lemma 4, we can assume without loss of generality that $m = 1$: For $m > 1$, we would need to show that for every $i \in U$, there are n_i distinct processes covered by i, given that $t_2 - t_1 \geq \log_2 n_i$ and $\sum_{i \in U} n_i = |W|$. This, however, is equivalent to the claim of the lemma for $m = 1$. We proceed by induction on $t_2 - t_1 \geq 0$.

Base case: If $t_2 - t_1 = 0$, i.e., $t_1 = t_2$, then $|W| = 1$ and the statement is trivially true since we can choose $U = W$.

Inductive step: Now let $t_2 - t_1 \geq 1$. Let $W = W_1 \cup W_2$ such that $W_1, W_2 \neq \emptyset$ and $\big| |W_1| - |W_2| \big| \leq 1$. Using Lemma 3, we see that $t_2 - (t_1 + 1) \geq \lceil \log_2 |W_s| \rceil$ for $s \in \{1, 2\}$. By the induction hypothesis, there hence exist nodes j_1 and j_2 that at time $t_1 + 1$ cover W_1 and W_2, respectively. But now, using the nonsplit property of communication graph G_{t_1}, we see that there exists a node i that covers $\{j_1, j_2\}$ in G_{t_1}. An application of Lemma 1 concludes the proof. $\square$

Note that Lemma 5, by choosing $W = [n]$ and $m = 1$, immediately provides an upper bound on the dynamic radius of $O(\log n)$. To show an upper bound of $O(\log \log n)$, we will apply this lemma only for the early infection phase of $O(\log \log n)$ rounds, and use a different technique, by the next two lemmas, for the late phase. Note that, for a set N, we use $\binom{N}{k}$ to denote all subsets of N with cardinality k and for an integer n, we use $\binom{n}{k}$ to denote the binomial coefficient n choose k.

Lemma 6. *Let U and W be finite sets with $|U| = k$, $|W| = n$, and $f : \binom{U}{\lfloor \log n \rfloor} \to W$. If $n \geq 8$, then there exists some $w \in W$ such that $\big| \bigcup f^{-1}[\{w\}] \big| \geq k/e^4$.*

Proof. By the pigeonhole principle and Stirling's formula, we get the existence of some $w \in W$ with

$$\big| f^{-1}[\{w\}] \big| \geq \frac{\binom{k}{\lfloor \log n \rfloor}}{n} \geq \frac{k^{\lfloor \log n \rfloor}}{n \lfloor \log n \rfloor^{\lfloor \log n \rfloor}} \ . \tag{1}$$

11

Write $M = \bigcup f^{-1}[\{w\}]$ and $m = |M|$. Since $S \in \binom{M}{\lfloor \log n \rfloor}$ for all $S \in f^{-1}[\{w\}]$, we have

$$|f^{-1}[\{w\}]| \leq \binom{m}{\lfloor \log n \rfloor} \leq \frac{m^{\lfloor \log n \rfloor} e^{\lfloor \log n \rfloor}}{\lfloor \log n \rfloor^{\lfloor \log n \rfloor}} \leq \frac{m^{\lfloor \log n \rfloor} n}{\lfloor \log n \rfloor^{\lfloor \log n \rfloor}} \ . \qquad (2)$$

Combining (1) and (2), we get

$$m \geq \frac{k}{n^{2/\lfloor \log n \rfloor}} = \frac{k}{e^{2 \log n / \lfloor \log n \rfloor}} \geq \frac{k}{e^{2 \log n / (\log n - 1)}}$$

$$= \frac{k}{e^{2/(1 - 1/\log n)}} \geq \frac{k}{e^{2/(1 - 1/2)}} = \frac{k}{e^4} \qquad (3)$$

where we used $\log n \geq \log 8 \geq 2$. This concludes the proof. $\qquad \square$

The next lemma shows that nodes are infected quickly in the late phase:

260 **Lemma 7.** *There exists some $C > 0$ such that for all $t \geq 1$ there exists a set of at most $C \log n$ nodes that at time t covers the set $[n]$ of all nodes at time $t + \lceil \log_2 \log n \rceil$.*

Proof. For every set $A \in \binom{V}{\lfloor \log n \rfloor}$ of $\lfloor \log n \rfloor$ nodes, let $f(A) \in V$ be a node that at time t covers A at time $t + \lceil \log_2 \log n \rceil$, which exists by Lemma 5.

265 We recursively define the following sequence of nodes v_i, ≥ 1 and sets of nodes V_i, $i \geq 0$:

- $V_0 = V$

- For $i \geq 1$, we choose v_i such that $|\bigcup f^{-1}[\{v_i\}]| \geq |V_{i-1}|/e^4$, which exists by Lemma 6, and $V_i = V_{i-1} \setminus \bigcup f^{-1}[\{v_i\}]$.

270 Note that, setting $r = 1 + \log n / \log \frac{e^4}{e^4 - 1}$, we have $V_r = \emptyset$. Hence the set $\{v_1, \ldots, v_r\}$ at time t covers all nodes at time $t + \lceil \log_2 \log n \rceil$. Noting $r = O(\log n)$ concludes the proof. $\qquad \square$

We are now ready to combine Lemma 5 for the early phase and Lemma 7 for the late phase to prove the main result of this section, Theorem 2.

275 *Proof of Theorem 2.* Let $t = \lceil \log_2(C \log n) \rceil$ where C is the constant from Lemma 7.

12

By Lemma 7, there is a set A of nodes with $|A| \leq C \log n$ that at time t covers all nodes at time $t + \lceil \log_2 \log n \rceil$. By Lemma 5, a single node at time 1 covers A at time t.

Combining both results via Lemma 1 shows that a single node at time 1 covers all nodes at time $\lceil \log_2(C \log n) \rceil + \lceil \log_2 \log n \rceil = O(\log \log n)$. $\qquad\square$

4. Nonsplit Networks from Asynchronous Rounds

We now show that in an important special case of nonsplit networks, namely those evolving from distributed algorithms that establish a round structure over asynchronous message passing in the presence of crashes, the dynamic radius is at most 2.

In the classic asynchronous message passing model with crashes, it is assumed that all messages sent have an unbounded but finite delay until they are delivered. Furthermore, processes do not operate in lock-step but may perform their computations at arbitrary times relative to each other. In addition, some processes may be faulty in the sense that they are prone to crashes, i.e., they may seize to perform computations at an arbitrary point in time.

This means that in a system where up to f processes may be faulty, in order to make progress in a distributed algorithm, a process may wait until it received a message from $n - f$ different processes but no more: If a process waits for a message from $> n - f$ different processes, but there were in fact f crashes, this process will wait forever. For this reason, algorithms for this asynchronous model often employ the concept of *asynchronous rounds*, sometimes realized as a local round counter variable t_i, which is held by each process $i \in [n]$ and appended to every message. A process i increments t_i only if it received a message containing a round counter $\geq t_i$ from $n - f$ different processes.

One may now ask how fast information can spread in this distributed computing model.

For this purpose, we consider a network whose communication patterns are induced by n processes communicating in asynchronous rounds. Here, an edge

13

(i, j) in the communication graph G_t represents that j, before incrementing its round counter from $t_j = t$ to $t_j = t + 1$, received a message from i and the round counter, appended to this message, was $t_i = t$. When deriving the communication graph G_t of such an asynchronous round t, we get a digraph where each correct process has at least $n - f$ incoming neighbors. In fact, when we slightly abuse notation and define the arrival of a message at a crashed process as a (virtual) reception that is represented in the communication graph G_t as well, we get that G_t is actually a digraph where all processes have at least $n - f$ incoming neighbors.

We restrict our attention to the case where $n > 2f$, i.e., a majority of the processes is correct. This implies that the sets of incoming neighbors of any two processes in a communication graph have a non-empty intersection, which means that the communication graph is nonsplit. Note that if $n \leq 2f$, then the network is not necessarily nonsplit. In fact, it can be disconnected into two disjoint sets of processes that do not receive messages from each other until termination of the algorithm. Below, we establish a constant upper bound on the dynamic radius of this important class of nonsplit graphs.

Theorem 4. *Let $f \geq 0$, $n > 2f$, and $(G_t)_{t \geq 1}$ be a sequence of communication graphs with $\mathrm{In}_i(G_t) \geq n - f$ for all t and all i. The dynamic radius of $(G_t)_{t \geq 1}$ is at most 2.*

Proof. To show the bound on the radius, we prove that there exists a center node m that realizes the dynamic radius, i.e.,

$$\exists m \in [n] \; \forall i \in [n] \; \exists j \in [n] : m \in \mathrm{In}_j(G_1) \wedge j \in \mathrm{In}_i(G_2). \tag{4}$$

Equation (4) now follows from

$$\exists m \in [n] : |\mathrm{Out}_m(G_1)| \geq f + 1, \tag{5}$$

by the following arguments: Equation (5) states that the information at m has been transmitted to at least $f + 1$ nodes. By assumption $\mathrm{In}_i(G_2) \geq n - f$ for all $i \in [n]$. Thus each i must have an incoming neighbor j in digraph G_2 such that $j \in \mathrm{Out}_m(G_1)$; equation (4) follows.

14

It remains to show (5). Suppose that the equation does not hold, i.e.,

$$\forall j \in [n] : |\operatorname{Out}_j(G_1)| \le f \,. \tag{6}$$

By assumption on digraph G_1, we have

$$\sum_{i \in [n]} |\operatorname{In}_i(G_1)| \ge n(n - f) \tag{7}$$

By the handshake lemma,

$$\sum_{i \in [n]} |\operatorname{In}_i(G_1)| = \sum_{j \in [n]} |\operatorname{Out}_j(G_1)|$$

and, using (6),

$$\sum_{i \in [n]} |\operatorname{In}_i(G_1)| \le nf \,.$$

Together with (7), we have $n(n - f) \le nf$; a contradiction to the assumption that $n > 2f$. $\qquad\square$

5. A Lower Bound for Consensus in Dynamic Networks

We now show that the dynamic radius of a network provides a lower bound on the time complexity of a consensus algorithm for this network.

Let $[n] = \{1, \dots, n\}$ be a set of processes that operate in lock-step synchronous rounds $t = 1, 2, \dots$ delimited by times $t = 0, 1, \dots$ where, by convention, round t happens between time $t - 1$ and time t. Each round consists of a phase of communication, followed by a phase of local computation. Like in the previous sections, a communication pattern defines, for each round, which messages reach their destination.

In the (exact) consensus problem, every node $i \in [n]$ starts with an input value $x_i \in X$ from an arbitrary domain X and holds a unique write-once variable y_i, initialized to $y_i = \bot$, where $\bot$ denotes a special symbol s.t. $\bot \notin X$. Since we are concerned with an impossibility result here, we may restrict ourselves without loss of generality to the binary consensus problem, i.e., the case where

15

$X = \{0, 1\}$. An execution of a deterministic consensus algorithm is a sequence of state transitions according to the algorithm and determined by the input assignment and the communication pattern. An algorithm solves consensus if it satisfies in all of its executions:

(Termination) Eventually for every $i \in [n]$, $y_i \neq \bot$.

(Validity) If $y_i \neq \bot$ then $y_i = x_j$ for some $j \in [n]$.

(Agreement) For every $i, j \in [n]$, if $y_i \neq \bot$ and $y_j \neq \bot$ then $y_i = y_j$.

Theorem 5. *If the dynamic radius of the network is k, then, in every deterministic consensus algorithm, not all processes can have terminated before time k.*

Proof. Let $\mathcal{G}$ be a communication pattern with dynamic radius k, which occurs in the network by assumption. Suppose, in some deterministic consensus algorithm A, all $i \in [n]$ have terminated by time $k-1$ in every execution based on $\mathcal{G}$. Let C_0 be the input assignment where $x_i = 0$ for all $i \in [n]$ and C_1 be the input assignment where $x_i = 1$ for all $i \in [n]$. By validity, when running A under $\mathcal{G}$ and starting from C_0, all $i \in [n]$ have $y_i = 0$ by time $k-1$ and when starting from C_1, they have $y_i = 1$. Thus, there are input assignments C, C' that differ only in the input assignment x_j of a single process j and, for all $i \in [n]$, at time $k-1$, $y_i = 0$ when applying A under $\mathcal{G}$ when starting from C and $y_i = 1$ when starting from C'. Since there is no broadcaster in $\mathcal{G}$ before round k, there is some process i' that did not receive a (transitive) message from j and thus i' is in the same state in both executions. Therefore, i' decides on the same value in both executions, which is a contradiction and concludes the proof. $\square$

6. Conclusion

In this paper, we studied nonsplit networks, which are a convenient abstraction that arises naturally when considering information dissemination in a variety of dynamic network settings. Since classic information dissemination

problems are trivially impossible in these nonsplit dynamic networks, it made sense to study the more relaxed dynamic radius here. As we showed in Theorem 5, this is an important characteristic with respect to the impossibility of exact consensus. For our main technical contribution, we proved a new upper bound in Theorem 2, which shows that the dynamic radius of nonsplit networks is in $O(\log \log n)$. This is an exponential improvement of the best previously known upper bound of $O(\log n)$.

We also showed an upper bound of 2 asynchronous rounds for the dynamic radius in the asynchronous message passing model with crash failures. Thus, in this important class of nonsplit networks, information dissemination is considerably quicker than what is currently known for the general case.

Combining our Theorem 2 with the result from [3] that established a $O(n)$ simulation of nonsplit networks in rooted networks, i.e., networks where every communication graph contains a rooted spanning tree, yields an improvement of the best previously known upper bound for the dynamic radius of rooted dynamic networks from $O(n \log n)$ to $O(n \log \log n)$:

Theorem 6. *The dynamic radius of a dynamic networks whose communication graphs are rooted is $O(n \log \log n)$.*

While this is another hint at the usefulness of the nonsplit abstraction for dynamic networks, the tightness of this bound remains an open question.

7. Acknowledgements

We want to thank Ulrich Schmid for thoroughly proof reading the article.

References

[1] N. A. Lynch, Distributed algorithms, Elsevier, 1996.

[2] P. Erdös, A. Rényi, V. T. Sós, On a problem of graph theory, Studia Sci. Math. Hungar. 1 (1966) 215–235.

[3] B. Charron-Bost, M. Függer, T. Nowak, Approximate consensus in highly dynamic networks: The role of averaging algorithms, in: M. M. Halldòrsson, K. Iwama, N. Kobayashi, B. Speckmann (Eds.), Automata, Languages, and Programming, Vol. 9135 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2015, pp. 528–539. `doi:10.1007/978-3-662-47666-6_42`.

[4] B. Charron-Bost, A. Schiper, The Heard-Of model: computing in distributed systems with benign faults, Distributed Computing 22 (1) (2009) 49–71. `doi:10.1007/s00446-009-0084-6`.

[5] N. Santoro, P. Widmayer, Time is not a healer, in: B. Monien, R. Cori (Eds.), 6th Symposium on Theoretical Aspects of Computer Science, Vol. 349 of LNCS, Springer, Heidelberg, 1989, pp. 304–313. `doi:10.1007/BFb0028994`.

[6] H. A. Harutyunyan, A. L. Liestman, J. Peters, D. Richards, Broadcasting and gossiping), in: J. Gross, P. Zhang (Eds.), The Handbook of Graph Theory, Chapman and Hall/CRC, 2013, pp. 1477–1494.

[7] S. M. Hedetniemi, S. T. Hedetniemi, A. L. Liestman, A survey of gossiping and broadcasting in communication networks, Networks 18 (4) (1988) 319–349. `doi:10.1002/net.3230180406`.

[8] A. Clementi, P. Crescenzi, C. Doerr, P. Fraigniaud, M. Isopi, A. Panconesi, F. Pasquale, R. Silvestri, Rumor spreading in random evolving graphs, in: European Symposium on Algorithms, Springer, 2013, pp. 325–336. `doi:10.1007/978-3-642-40450-4_28`.

[9] K. Winkler, M. Schwarz, U. Schmid, Consensus in rooted dynamic networks with short-lived stability, Distributed Computing (2019). `doi:10.1007/s00446-019-00348-0`.

[10] P. Fraigniaud, E. Lazard, Methods and problems of communication in usual

18

networks, Discrete Applied Mathematics 53 (1) (1994) 79 – 133. `doi:`
`10.1016/0166-218X(94)90180-5`.

[11] J. Hromkovič, R. Klasing, B. Monien, R. Peine, Dissemination of information in interconnection networks (broadcasting & gossiping), in: D.-Z. Du, D. F. Hsu (Eds.), Combinatorial Network Theory, Springer US, Boston, MA, 1996, pp. 125–212. `doi:10.1007/978-1-4757-2491-2_5`.

[12] F. Kuhn, N. Lynch, R. Oshman, Distributed computation in dynamic networks, in: Proceedings of the Forty-second ACM Symposium on Theory of Computing, STOC '10, ACM, New York, NY, USA, 2010, pp. 513–522. `doi:10.1145/1806689.1806760`.

[13] M. Zeiner, M. Schwarz, U. Schmid, On linear-time data dissemination in dynamic rooted trees, Discrete Applied Mathematics 255 (2019) 307 – 319. `doi:10.1016/j.dam.2018.08.015`.

[14] U. Schilcher, C. Bettstetter, G. Brandner, Temporal correlation of interference in wireless networks with rayleigh block fading, IEEE Transactions on Mobile Computing 11 (12) (2012) 2109–2120. `doi:10.1109/TMC.2011.244`.

[15] A. Goiser, S. Khattab, G. Fassl, U. Schmid, A new robust interference reduction scheme for low complexity direct-sequence spread-spectrum receivers: Performance, in: Proceedings 3rd International IEEE Conference on Communication Theory, Reliability, and Quality of Service (CTRQ'10), Athens, Greece, 2010, pp. 15–21. `doi:10.1109/CTRQ.2010.52`.

[16] C. Ware, J. Judge, J. Chicharo, E. Dutkiewicz, Unfairness and capture behaviour in 802.11 adhoc networks, in: 2000 IEEE International Conference on Communications. ICC 2000. Global Convergence Through Communications., 2000, pp. 159–163 vol.1. `doi:10.1109/ICC.2000.853084`.