Token Dissemination in Geometric Dynamic Networks*

Sebastian Abshoff, Markus Benter, Andreas Cord-Landwehr, Manuel Malatyali, and Friedhelm Meyer auf der Heide

Heinz Nixdorf Institute & Computer Science Department,
University of Paderborn, Fürstenallee 11, 33102 Paderborn, Germany
Phone: +49-5251-606433
Fax: +49-5251-606482
{abshoff,benter,phoenixx,malatya,fmadh}@hni.upb.de

Abstract. We consider the k-token dissemination problem, where k initially arbitrarily distributed tokens have to be disseminated to all nodes in a dynamic network (as introduced by Kuhn et al., STOC 2010). In contrast to general dynamic networks, our dynamic networks are unit disk graphs, i.e., nodes are embedded into the Euclidean plane and two nodes are connected if and only if their distance is at most R. Our worst-case adversary is allowed to move the nodes on the plane, but the maximum velocity $v_{\max}$ of each node is limited and the graph must be connected in each round. For this model, we provide almost tight lower and upper bounds for k-token dissemination if nodes are restricted to send only one token per round. It turns out that the maximum velocity $v_{\max}$ is a meaningful parameter to characterize dynamics in our model.

Keywords: Geometric Dynamic Networks, Token Dissemination, Distributed Computing

1 Introduction

Dynamic networks appear in many scenarios like peer-to-peer networks, mobile wireless ad-hoc networks or swarms of mobile robots. The dynamics in such models is diverse and different. Kuhn et al. [KLO10] have introduced a very general model with the aim of understanding limitations and possibilities when coping with dynamics in networks, independent of specific application. In this paper, we look at special dynamics motivated by agents that move in the Euclidean plane and that are able to communicate with nearby agents only. More particularly, we look at dynamic unit disk graphs as they are often used to model ad-hoc networks or robotic networks. We are mainly interested in exploring the impact

* This work was partially supported by the German Research Foundation (DFG) within the Collaborative Research Centre "On-The-Fly Computing" (SFB 901), by the EU within FET project MULTIPLEX under contract no. 317532, and the International Graduate School "Dynamic Intelligent Systems".

of a velocity limit of the agents on the time required to perform fundamental tasks such as token dissemination.

The nodes of our geometric dynamic network are embedded into the Euclidean plane and two nodes are connected if and only if their distance is at most a constant R, which models the limited range of a wireless communication device. We consider a worst-case dynamic that is able to move the nodes within this plane. This worst-case dynamic is restricted by a maximal velocity parameter $v_{\max}$ and it must preserve connectivity of the network. In our model, we can prove lower and upper bounds for the k-token dissemination problem, which has also been studied by Kuhn et al. in a general model. In the k-token dissemination problem, k initially arbitrarily distributed tokens have to be disseminated by the nodes of the dynamic network such that each node receives all tokens and also decides that it has received all k tokens since we assume k is not known by the nodes beforehand. Note that solving the all-to-all token dissemination problem, where each node starts with exactly one token, implicitly solves the counting problem if the nodes' unique IDs are considered as tokens. While solving the token dissemination problem, a distributed algorithm must cope with the dynamic of the network, i.e., the changes of edges as induced by a worst-case dynamic that moves the nodes.

In our model, we restrain the dynamic network model by Kuhn et al. by introducing a geometry that gives a natural restriction of the power of a worst-case dynamic by geometric means. From this, we expect new insights into the complexity of distributed computational problems by using different techniques that exploit the geometry of the dynamic network. As a first step, both our lower and upper bounds for the k-token dissemination problem contain the maximal velocity parameter $v_{\max}$, i.e., they are bound by the characteristic value of the network dynamic. More precisely, we define a dynamic unit disk graph with maximal node velocity $v_{\max}$ and communication radius R and require connectivity w.r.t. a unit disk graph with radius 1. Our algorithm terminates after $\mathcal{O}(n(n+k) \cdot \min\{v_{\max}, R\} \cdot R^{-2})$ rounds if $R > 1$. Moreover, we present a lower bound of $\Omega(n \cdot k \cdot \min\{v_{\max}, R\} \cdot R^{-3})$ for randomized knowledge-based token-forwarding algorithms. Note that for $k = \Omega(n)$, the upper bound simplifies to $\mathcal{O}(n \cdot k \cdot \min\{v_{\max}, R\} \cdot R^{-2})$ and the upper and the lower bound become almost tight.

2 The Geometric Dynamic Network Model

In this paper, we consider the following dynamic network model adapted from Kuhn et al. [KLO10,KO11,Osh12]: we assume a dynamic graph with a fixed set V of n nodes, and a discrete, synchronous time model. Each node v is identified by a unique ID, assigned by some injective function $\mathrm{id} : V \to \{1, \ldots, \mathrm{poly}(n)\}$. In round r, the dynamic graph has some edge set E_r, forming the graph $G_r = (V, E_r)$. We assume local broadcast communication, i.e., a message sent by node u in round r is delivered to u's neighbors in round $r+1$. Therefore, when sending a message in round r, a node usually does know to which neighbors the message will be delivered. In this paper, the message each node can send via local broadcast communication is limited to one token per round. Kuhn et al. introduced the concept of T-interval connectivity as a reasonable restriction of the dynamic: For each time interval I of length $T \geq 1$, there must be some stable and connected subgraph in all graphs G_r with $r \in I$. If $T = 1$, this just means the graph must be connected in each round r.

Our modifications address a dynamic model motivated by geometric mobility as it appears, e.g., in swarm robotics. Here, we assume that each node v in each round r has a position $p_r(v) \in \mathbb{R}^2$ (however, our results also hold for $\mathbb{R}^3$). The distance between two nodes u, v in round r is denoted by $d_r(u, v) = |p_r(u) - p_r(v)|$. Then, for each round r, we define G_r as the unit disk graph with communication radius R. We omit the round parameter r when the round is clear from context. In addition, $1 \leq R \leq n$ holds throughout the paper for technical reasons. Furthermore, we assume that the maximum velocity of each node is bounded by a parameter $v_{\max} > 0$, i.e., the position of a node changes at most by a distance $v_{\max}$ from round to round. Such a model was also considered by Bienkowski et al. [BBKM09], and it is often (implicitly or explicitly) assumed for designing local strategies for robotic formation problems (for a survey see [KM11]).

Our results require a somewhat stronger notion of connectivity than in the general model by Kuhn et al.: We demand that the graph in round r is connected even if we restrict the communication radius to 1 instead of R. To distinguish these graphs, we talk about the *communication graph* G_r if radius R is used, and about the *connectivity graph* G'_r if radius 1 is used. Thus, we require that the connectivity graph G'_r is connected in each round r. This geometric model gives rise to another natural restriction of dynamics. A graph G_r is called C-connected if at least C edges have to be removed to transform G_r into a disconnected graph.

The focus of our paper lies on the k-token dissemination problem. In this problem, each node u in the network receives as input $I(u)$ a possibly empty subset of tokens such that $\left|\bigcup_{v \in V} I(v)\right| = k$. Then, the nodes have to disseminate these tokens such that each node eventually knows all k tokens and then explicitly terminates (i.e., it outputs the result and does not send/receive any further messages). Here, k is not known by the nodes beforehand. Additionally, we examine the implications of our results for the problem of counting, which is to determine the exact number of nodes in the network.

We show a result for a restricted class of algorithms that is called *knowledge-based token-forwarding algorithms* (cf. Kuhn et al.): let $A_u(r)$ denote the set of messages node u has received by the beginning of round r including its input $I(u)$. A *token-forwarding algorithm* requires each node to send only one pure token from $A_u(r)$ (without modification and without annotation) or the empty message, and it must not terminate before it has received all k tokens. A token-forwarding algorithm is called *knowledge-based* if the distribution that determines which token is sent by u in round r is a function only of its unique ID $\mathrm{id}(u)$, $A_u(0), \ldots, A_u(r-1)$ and the sequence of u's coin tosses up to round r (including r). Many natural strategies can be found in this class of knowledge-based token-forwarding algorithms, e.g., strategies like sending a known token sampled uniformly at random.

3 Related Work

Dynamic networks, where the set of edges in the network may change arbitrarily and in an adversarial way from round to round as long as the graph is strongly connected in each round, were introduced by Kuhn et al. [KLO10,KO11,Osh12]. In each round, each node may send a message of size $\mathcal{O}(\log n)$ bits that is delivered to all neighboring nodes in the following round. Computation in their model requires termination. On the one hand, for the k-token dissemination problem in T-interval connected dynamic networks, Kuhn et al. present a deterministic $\mathcal{O}(n(n+k)/T)$ token-forwarding algorithm. This algorithm can be used to obtain an $\mathcal{O}(n^2/T)$ algorithm for the counting problem. On the other hand, they give a $\Omega(nk/T)$ lower bound for the restricted class of knowledge-based token-forwarding algorithms and they provide an $\Omega(n \log k)$ lower bound for deterministic centralized token-forwarding algorithms.

Dutta et al. [DPR+13] improved the latter lower bound by Kuhn et al. to $\Omega(nk/\log n + n)$ for any randomized (even centralized) token-forwarding algorithm and showed for a weakly-adaptive adversary that k-token dissemination can be done in $\mathcal{O}((n+k)\log n \log k)$ w.h.p. Furthermore, they provide two polynomial time, randomized and centralized offline algorithms, one returns an $\mathcal{O}(n, \min\{k, \sqrt{k \log n}\})$ schedule w.h.p. and another one an $\mathcal{O}((n+k)\log^2 n)$ schedule w.h.p. if nodes can send a token along each edge per round. Using similar techniques, Haeupler and Kuhn [HK12] showed lower bounds if nodes are allowed to forward $b \leq k$ tokens or if they are only required to obtain a δ-fraction in T-interval connected dynamic networks and dynamic networks that are c-vertex connected in every round.

O'Dell and Wattenhofer [OW05] analyzed information dissemination problems in slightly different but worst-case adversarial models. Das Sarma et al. [SMP12] developed randomized token-forwarding algorithms based on random walks on dynamic networks. Here, an oblivious adversary that is not aware of the random choices of the algorithm modifies the network. Haeupler and Karger go beyond the class of token-forwarding algorithms and send linear combinations of tokens. With this technique, they are able to solve the k-token dissemination problem in $\mathcal{O}(nk/\log n)$ rounds w.h.p. [HK11]. Brandes and Meyer auf der Heide [BM12] develop algorithms for counting if in addition every edge in the network fails with some probability. Michail et al. [MCS12b] studied computation in possibly disconnected dynamic networks and introduced temporal connectivity conditions. The same authors [MCS12a] looked into naming and counting in the absence of unique IDs in dynamic networks. Here, naming refers to the problem of generating unique IDs. Interestingly, they introduce a different communication model where the nodes in the network can send different, individual messages to their neighbors but without any information about their states.

The unit disk model has been extensively studied in the area of routing in wireless ad-hoc and sensor networks, in particular, in geographic routing algorithms. Geographic routing takes advantage of the availability of position information to decide which node becomes the next hop. Those algorithms assume that a node can get its own position using a location service such as GPS. A worst-case optimal and average-case efficient geographic routing algorithm has been proposed by Kuhn et al. [KWZ03,KWZZ03]. However, geographic routing focuses on single source routing while a well-designed token dissemination algorithm has to take congestion into account. For a broad overview about further routing algorithms in this area, we refer to Frey et al. [FRS09].

Gossiping algorithms are a class of algorithms for distributed computation in arbitrary graphs following a simple principle: the nodes are initialized with some values, they continuously exchange these values and calculate new values based on the ones they received together with a problem specific function.

Boyd et al. [BGPS04,BGPS05] analyzed the mixing time of the averaging problem. Later on, Dimakis et al. [DSW08] showed that the mixing time can be significantly improved in grid graphs and random geometric graphs, two communication models for realistic sensor networks. In a random geometric graph, the n sensor locations are chosen uniformly and independently in the unit square, and each pair of nodes is connected if their Euclidean distance is smaller than some constant transmission radius R. They proposed an algorithm that computes true average to accuracy $1/n^\alpha$ using $\mathcal{O}(n^{1.5}\sqrt{\log n})$ radio transmissions. This reduces the energy consumption by a factor of $\sqrt{n/\log n}$ compared to standard gossip algorithms.

4 Lower Bound on Token Dissemination in Geometric Dynamic Networks

In this chapter, we show that any knowledge-based token-forwarding algorithm needs $\Omega(n \cdot k \cdot \min\{v_{\max}, R\} \cdot R^{-3})$ rounds for solving the k-token dissemination problem in geometric dynamic networks. To do so, we follow a similar analysis like the one by Kuhn et al. for an $\Omega(nk)$ lower bound for dynamic networks with arbitrarily changing edges [KLO10].

For the sake of a simple presentation, we first introduce our construction for the special case $R = 1$, i.e., the communication graph is equal to the connectivity graph. Later, this result will be generalized to the case $R \geq 1$.

Theorem 1. *If $R = 1$, then any knowledge-based token-forwarding algorithm for k-token dissemination requires $\Omega(n \cdot k \cdot \min\{v_{\max}, 1\})$ rounds to succeed with probability $> \frac{1}{2}$.*

Proof. We create the setting as follows: Initially, some node v_0 knows all k tokens and all other nodes do not know any token. As the token-forwarding algorithm is knowledge-based, the probability distribution of the tokens sent by v_0 does not depend on the dynamic graph. Let $r^* := \left\lfloor \frac{(n-4)k}{2L} \right\rfloor - 1$ for $L := \left\lceil \frac{1+\varepsilon}{2v_{\max}} \right\rceil$, which is used as the length of a row of nodes in our construction. ϵ is a suitably small chosen value. Then, by linearity of expectation and Markov's inequality, there is some infrequently sent token t that is sent $< \frac{n-4}{L}$ times by v_0 until round r^* with probability of at least $\frac{1}{2}$. For this, we define a dynamic such that v_0 cannot terminate by round r^* since some node must be unaware of t at this time.

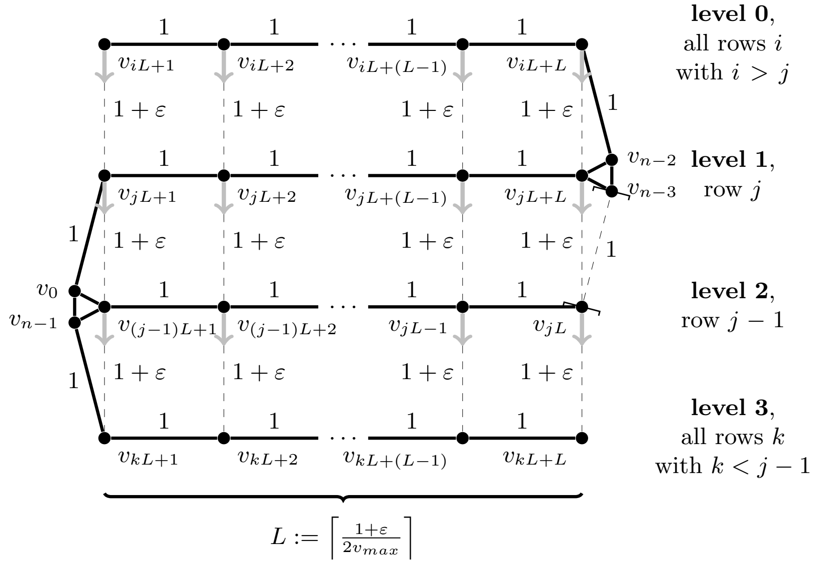

Fig. 1. Construction for $R = 1$ showing the positions of the nodes for a fixed j.

All nodes are positioned as shown in Figure 1: Except for the four nodes v_0, v_{n-1}, v_{n-2}, and v_{n-3}, all other nodes are assigned to horizontal rows where each row consists of L nodes that are positioned on four levels. On each level, the nodes have a distance of exactly 1, which maximizes the distance between the nodes such that the row is still connected. The distance between two levels is at least $1 + \epsilon$.

Node v_0 is connected to $v_{(j-1)L+1}$, v_{jL+1}, and v_{n-1}, and positioned such that the distance between v_0 and v_{jL+1} is exactly 1. Analogously, node v_{n-1} is connected to $v_{(j-1)L+1}$, v_{kL+1}, and v_0, and positioned such that the distance between v_n and v_{kL+1} is exactly 1. We will see later that these nodes are essential for preserving connectivity during the movement of the nodes.

Similarly, node v_{n-2} is connected to v_{jL+L}, v_{iL+L}, and v_{n-3}, and positioned such that the distance between v_{n-2} and v_{iL+L} is exactly 1. In contrast to that, node v_{n-3} is only connected to v_{jL+L} and v_{n-2} but the distance between v_{n-3} and v_{jL} is slightly larger than 1. As we will see later, this is important to ensure that the infrequently sent token t cannot be learned by the nodes on level 0.

Initially, one row $j = 0$ is at level 1 and all other rows $i > j$ are stacked at level 0. Level 2 and level 3 are not occupied by rows at this time.

When v_0 sends the token t, three rows start moving down. In particular, one row $j+1$ at level 0, row j at level 1, and row $j-1$ at level 2 start moving down with maximal relative velocity $2v_{\max}$ [1] for the next L rounds until they reach level 1, 2, and 3, respectively. Once a row reaches level 3, it does not move any further and all rows $k < j-1$ stack again. As soon as rows $j+1, j, j-1$ reach the next level, j can be incremented and Figure 1 shows the current situation. If v_0 again sends token t, the described procedure repeats and three rows move down.

The crucial property of our construction is that the graph is always connected while the infrequently sent token t never reaches the rows stacked on level 0. Let us first argue why the graph is always connected. Initially, level 0 and level 1 are occupied and the graph is connected. During movement, the location of node v_{n-2} ensures that the row moving between level 0 and level 1 is connected. Analogously, node v_{n-1} connects the row moving in between of level 2 and level 3. The placement of node v_0 and v_{n-3} ensures that the left side of the graph is connected to the right side of the graph via the row in between of level 1 and level 2.

Let us now consider the second property: level 0 never gets the infrequently sent token t. According to the definition of the dynamic graph, the only possibility for level 0 to get t is via the row moving between level 1 and level 2 via node v_{n-3} or v_{n-2}. Since t is sent less than $\lfloor \frac{n-4}{L} \rfloor$ times, the token needs more than L rounds to cross one row. Thus, according to the definition of the movement, the nodes v_{jL+L} and v_{n-3} become disconnected at least one round before t can be sent from v_{jL+L} to v_{n-3}. This is when the row arrives at level 2.

Since there are $\lfloor \frac{n-4}{L} \rfloor$ rows, the nodes on at least one row are unaware of the token. Hence, $r^* = \Omega(n \cdot k \cdot \min\{v_{\max}, 1\})$ rounds are required. □

Next, we extend this construction for an arbitrary communication radius $R \geq 1$:

Theorem 2. *If $R \geq 1$, then any knowledge-based token-forwarding algorithm for k-token dissemination requires $\Omega(n \cdot k \cdot \min\{v_{\max}, R\} \cdot R^{-3})$ rounds to succeed with probability $> \frac{1}{2}$.*

Proof. As before, we want to find a token that is sent infrequently over some cut. Yet, for $R \geq 2$, the communication graph is $\lfloor R \rfloor$-connected in each round, i.e., there is no single cut vertex as in the construction before for $R = 1$. Therefore, multiple nodes $v_0, v_{n-1}, \ldots, v_{n-\lfloor R \rfloor+3}$ initially receive all k tokens such that the probability distribution of the tokens sent by them does not depend on the dynamic graph. Define $r^* := \left\lfloor \frac{(n-cR)k}{2L\lfloor R \rfloor} \right\rfloor - 1$ for $L := \left\lceil \frac{3(R+\varepsilon)}{2v_{\max}} \right\rceil \cdot (R+1)$ and some constant $c \in \mathbb{N}^+$ that is specified later. Then, by linearity of expectation and Markov's inequality, there is some infrequently sent token t which is sent $<$

[1] To upper bound the worst-case traveling distance for a fixed node pair u and v, we can w.l.o.g. assume that u is static while v moves with velocity of at most $2v_{\max}$.

$\frac{(n-cR)k}{L\lfloor R\rfloor}$ times by all nodes $v_0, v_{n-1}, \dots, v_{n-\lfloor R\rfloor+3}$ until round r^* with probability of at least $\frac{1}{2}$. We present a dynamic such that all nodes $v_0, \dots, v_{\lfloor R\rfloor-1}$ cannot terminate by round r^* since there still is a node that is unaware of t at this time.

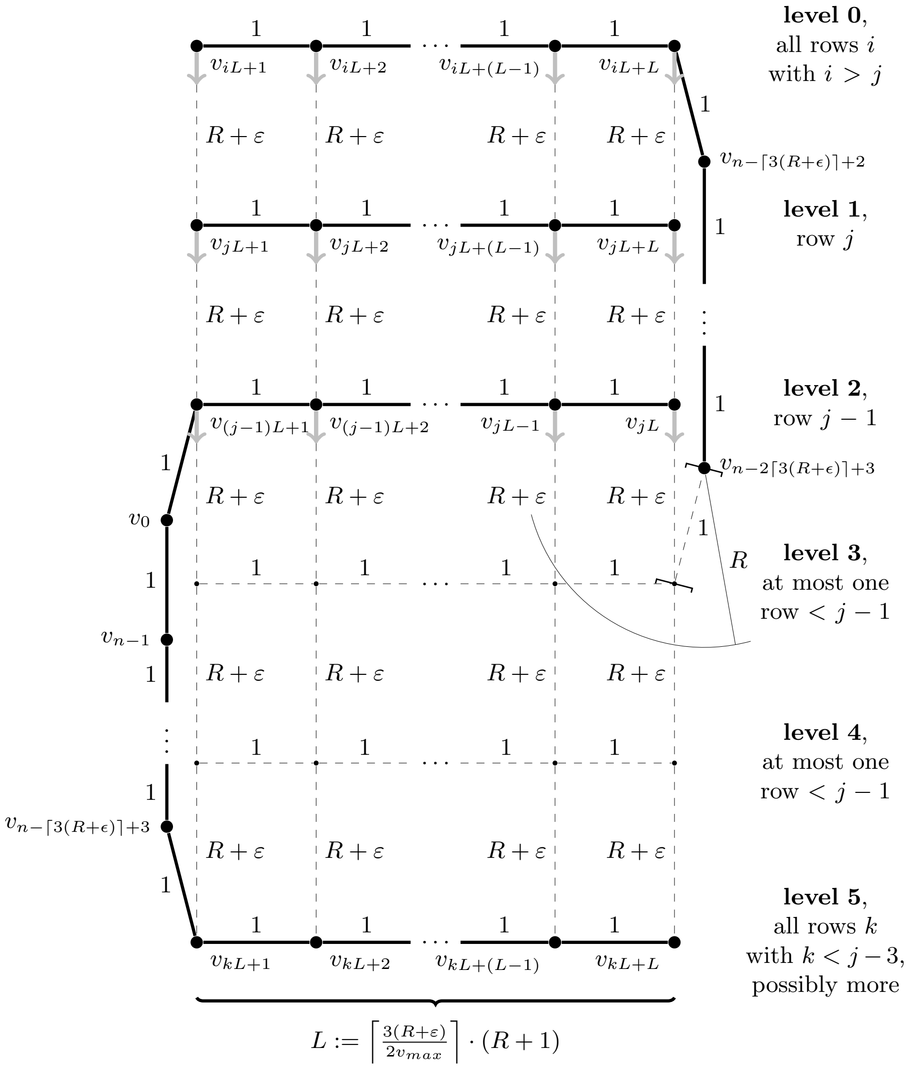

Fig. 2. Construction for $R \geq 1$ showing the positions of the nodes for a fixed j.

All nodes are positioned as shown in Figure 2: similar to the construction in Theorem 1, except for $c := 2\left(\lceil 3(R+\epsilon)\rceil - 2\right)$ nodes $v_0, v_{n-1}, \dots, v_{n-2\lfloor 3(R+\epsilon)\rfloor+3}$, all other nodes are assigned to horizontal rows where each row consists of L nodes that are positioned on six levels. The distance between $v_{n-2\lceil 3(R+\epsilon)\rceil+3}$ and the position below v_{jL} is slightly greater than one. Initially, one row $j = 0$ is at level 1 and all other rows $i > j$ are stacked at level 0. Levels $2, \ldots, 5$ are not occupied by rows at this time.

When one of the nodes $v_0, v_{n-1}, \ldots, v_{n-\lfloor R\rfloor+3}$ sends the token t, one row at level 0 and all rows on levels $1, \ldots, 4$ start moving down with maximal relative velocity $2v_{\max}$ for the next L rounds until they reach the next level. The row from level 0 stops at level 1, but all other rows continue moving until they reach level 5, where they do not move any further and stack again. As soon as row $j+1$ reaches level 0, j can be incremented and Figure 2 shows the current situation until t is sent again. If any of the nodes $v_0, v_{n-1}, \ldots, v_{n-\lfloor R\rfloor+3}$ sends the token t again, the described procedure repeats and further rows move down.

Observe that the graph is always connected and that token t cannot reach $v_{n-\lceil 3(R+\epsilon)\rceil+3}$ or any node above since it is sent less than $\left\lfloor \frac{(n-cR)k}{L\lfloor R\rfloor} \right\rfloor$ times and the token needs more than L rounds to cross one row. Since there are $\left\lfloor \frac{n-cR}{L} \right\rfloor$ rows, the nodes on at least one row are unaware of the token. Hence, $r^* = \Omega(n \cdot k \cdot \min\{v_{\max}, R\} \cdot R^{-3})$ rounds are required. □

Remark 1. Our results implies that the lower bound by Kuhn et al. [KLO10] already holds for a much more restricted model of dynamics: If we choose $R = 1$ and $v_{\max}$ constant, e.g. $v_{\max} = 1$, then we achieve the $\Omega(nk)$ lower bound for knowledge-based token-forwarding algorithms.

5 Upper Bound on Token Dissemination in Geometric Dynamic Networks

In this chapter, we present a k-token dissemination algorithm for geometric dynamic networks with bounded maximum velocity $v_{\max}$. The algorithm is basically an extension of the algorithm by Kuhn et al. which allows to solve k-token dissemination under arbitrary edge dynamics in $\Theta(n(n+k))$ rounds. Under the restriction of $R > 1$, it is possible to speed up the algorithm up to $\Theta(n(n+k) \cdot \min\{v_{\max}, R\} \cdot R^{-1})$. Moreover, if $R \geq 2$, then the $\Theta(R)$-connectivity of the communication graph can be exploited to get another speed-up of $\Theta(R)$, i.e., the algorithm needs $O(n(n+k) \cdot \min\{v_{\max}, R\} \cdot R^{-2})$ rounds in total.

Let us first sketch the dissemination algorithm by Kuhn et al. for $2T$-interval connected graphs. By definition of $2T$-interval connectivity, there is a spanning connected subgraph for at least $2T$ rounds. This subgraph is used to establish a pipelining effect such that at least the T smallest tokens are disseminated to all nodes in $\Theta(n)$ rounds. The algorithm proceeds in $\lfloor \frac{n}{T} \rfloor$ phases, where each phase consists of $2T$ rounds.[2] In each round of each phase, each node sends the smallest token it has not yet sent in this phase. To disseminate k tokens, this

[2] Note that n is not known by the nodes beforehand but as described by Kuhn et al. [KLO10] it can be determined involving the dissemination procedure itself using different estimates for n.

procedure can be repeated $\lceil \frac{k}{T} \rceil$ times. We restate the following results are either provided in the paper by Kuhn et al. or that directly follow from their results.

Theorem 3 **([KLO10,Osh12]).** *For $T \geq 1$, in a T-interval connected dynamic network with arbitrarily changing edges, the algorithms by Kuhn et al. for k-token dissemination and counting can be sped up by a factor of T, i.e., they need $\Theta(n(n+k) \cdot T^{-1})$ rounds for k-token dissemination and $\Theta(n^2 \cdot T^{-1})$ rounds for counting.*

Theorem 4 **([KLO10,Osh12]).** *For $T, C \geq 1$, in a T-interval connected dynamic network with arbitrarily changing edges where the stable subgraph is C-connected, the algorithms by Kuhn et al. for k-token dissemination and counting can be sped up by a factor of $T \cdot C$, i.e., they need $\Theta(n(n+k) \cdot T^{-1} \cdot C^{-1})$ rounds for k-token dissemination and $\Theta(n^2 \cdot T^{-1} \cdot C^{-1})$ rounds for counting.*

Note that it is assumed that T and C are known by the nodes. Furthermore, we would like to stress that it is not enough that $G(r)$ is C-connected in each round r. To make use of the pipelining effect it is also important that the stable subgraph is C-connected.

In the following, we show that our geometric dynamic networks are $\Theta(R \cdot v_{\max}^{-1})$-interval connected if $R > 1$ and that the stable subgraphs are $\Theta(R)$-connected if $R \geq 2$.

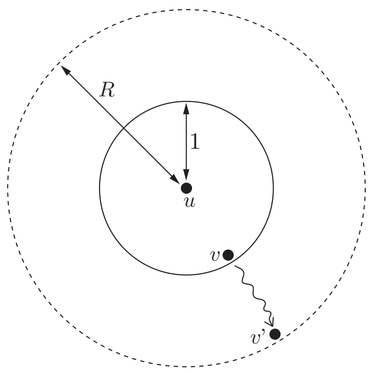

Fig. 3. After v moved to position v', the nodes u and v are still connected.

Lemma 1. *Assume the nodes of a geometric dynamic network move with maximum velocity $v_{\max}$. Then, the geometric dynamic network is $\left\lfloor \frac{(R-1)}{2 \cdot v_{\max}} \right\rfloor + 1$-interval connected.*

Proof. Consider a fixed node pair u and v which is connected in the connectivity graph of round r. Observe that the distance between two nodes can increase by at most $2v_{\max}$ per round. Thus, nodes that are connected in the connectivity graph (radius 1) stay connected in the communication graph (radius R) for at least $\left\lfloor \frac{(R-1)}{2 \cdot v_{\max}} \right\rfloor$ further rounds. This implies the lemma (cf. Figure 3). $\square$

Lemma 2. *Assume the nodes of a geometric dynamic network move with maximum velocity $v_{\max}$. Then, the geometric dynamic network contains a spanning $\left\lfloor \frac{1}{2} R \right\rfloor$-connected subgraph that is stable for $\left\lfloor \frac{R}{4 v_{\max}} \right\rfloor + 1$ rounds.*

Proof. Consider a path of length $\left\lfloor \frac{1}{2} R \right\rfloor$ in the connectivity graph of round r. The nodes on this path form a clique in the communication graph. Following a similar argument as in Lemma 1, these nodes stay connected in the communication graph for further $\left\lfloor \frac{\frac{1}{2} R}{2 v_{\max}} \right\rfloor$ rounds. This implies the lemma. $\square$

Theorem 5. *If $R > 1$, then k-token dissemination can be done in $\Theta(n(n+k) \cdot \min\{v_{\max}, R\} \cdot R^{-2})$ rounds and counting can be done in $\Theta(n^2 \cdot \min\{v_{\max}, R\} \cdot R^{-2})$ rounds.*

Proof. If $R > 1$, then according to Lemma 1 the geometric dynamic network is $\Theta(R \cdot v_{\max}^{-1})$-interval connected. Thus, by Theorem 3, the algorithms by Kuhn et al. need $\Theta(n(n+k) \cdot \min\{v_{\max}, R\} \cdot R^{-1})$ rounds for k-token dissemination and $\Theta(n^2 \cdot \min\{v_{\max}, R\} \cdot R^{-1})$ rounds for counting.

If in addition $R \geq 2$, then according to Lemma 2, the communication graph contains a spanning $\Theta(R)$-connected subgraph that is stable for $\Theta(R \cdot v_{\max}^{-1})$ rounds. Thus, by Theorem 4, the algorithms by Kuhn et al. need $\Theta(n(n+k) \cdot \min\{v_{\max}, R\} \cdot R^{-2})$ rounds for k-token dissemination and $\Theta(n^2 \cdot \min\{v_{\max}, R\} \cdot R^{-2})$ rounds for counting. $\square$

Comparing the lower and the upper bound, we can observe that the bounds are almost matching (despite of a factor of R^{-1}) if $k = \Theta(n)$. However, it should be pointed out that the graph model is a bit relaxed by introducing the connectivity graph in addition to the communication graph. It is an interesting question for further research to consider less relaxed models or even matching models.

6 Conclusion and Future Prospects

We showed that the k-token dissemination problem in geometric dynamic networks can be solved asymptotically faster than in traditional dynamic networks. For this, we utilized a communication radius larger than one, which is the connectivity radius. Specifically, by introducing natural conditions, the k-token dissemination problem can be solved in $\mathcal{O}(n(n+k) \cdot \min\{v_{\max}, R\} \cdot R^{-2})$ rounds in our model with $R > 1$ while the lower bound for arbitrary edge dynamics for knowledge-based token-forwarding algorithms is $\Omega(nk)$ [KLO10]. Additionally, these results can also be applied to count the number of nodes of the network in $\mathcal{O}(n^2 \cdot \min\{v_{\max}, R\} \cdot R^{-2})$ rounds if $R > 1$.

Our lower bound shows that even an optimal knowledge-based token-forwarding algorithm needs $\Omega(n \cdot k \cdot \min\{v_{\max}, R\} \cdot R^{-3})$ rounds to disseminate k tokens.

For the more interesting case of $k = \Omega(n)$, the upper bound simplifies to $\mathcal{O}(n \cdot k \cdot \min\{v_{\max}, R\} \cdot R^{-2})$ and becomes almost tight. It should be pointed out that our upper bound model is somehow relaxed by introducing the connectivity graph in addition to the communication graph. However, we think that this is a good starting point for analyzing, how we can improve the performance of k-token dissemination, counting, and other related problems by restricting general edge dynamics.

As a next step, it would be interesting to further restrict the model used for the upper bound. Particularly, is it possible to show similar round complexity results with matching models in the upper and lower bound such that we can omit the restriction $R > 1$? To gain more intuition about bounds in the geometric dynamic network model, it is another open question whether global knowledge is an advantage in this model and if so, to what degree. In other words, can a central online algorithm for k-token dissemination that is able to observe the positions of the nodes perform better when facing the network dynamic?

Moreover, more general geometric network models could be of interest. Those are for example asynchronous time models or different graph classes such as disk graphs or quasi unit disk graphs. Yet, one could also think about looking at different and specifically non-geometric restrictions to the network dynamic. A challenging but very interesting question is, whether it is possible to build up a hierarchy of dynamic network restrictions similar to hierarchies known from complexity theory.

References

BBKM09. Marcin Bienkowski, Jaroslaw Byrka, Miroslaw Korzeniowski, and Friedhelm Meyer auf der Heide. Optimal algorithms for page migration in dynamic networks. *J. Discrete Algorithms*, 7(4):545–569, 2009.

BGPS04. S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Analysis and optimization of randomized gossip algorithms. In *Decision and Control, 2004. CDC. 43rd IEEE Conference on*, volume 5, pages 5310–5315 Vol.5, 2004.

BGPS05. Stephen P. Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. Gossip algorithms: design, analysis and applications. In *INFOCOM*, pages 1653–1664, 2005.

BM12. Philipp Brandes and Friedhelm Meyer auf der Heide. Distributed computing in fault-prone dynamic networks. In *TADDS*, pages 9–14, 2012.

DPR+13. Chinmoy Dutta, Gopal Pandurangan, Rajmohan Rajaraman, Zhifeng Sun, and Emanuele Viola. On the complexity of information spreading in dynamic networks. In *SODA*, pages 717–736, 2013.

DSW08. Alexandros G. Dimakis, Anand D. Sarwate, and Martin J. Wainwright. Geographic gossip: Efficient averaging for sensor networks. *IEEE Transactions on Signal Processing*, 56(3):1205–1216, 2008.

FRS09. Hannes Frey, Stefan Rührup, and Ivan Stojmenovic. Routing in wireless sensor networks. In S. Misra, I. Woungang, and S. C. Misra, editors, *Guide to Wireless Sensor Networks*, pages 81–111. Eds. Springer, London, 2009.

HK11. Bernhard Haeupler and David R. Karger. Faster information dissemination in dynamic networks via network coding. In *PODC*, pages 381–390, 2011.

HK12. Bernhard Haeupler and Fabian Kuhn. Lower bounds on information dissemination in dynamic networks. In *DISC*, pages 166–180, 2012.

KLO10. Fabian Kuhn, Nancy A. Lynch, and Rotem Oshman. Distributed computation in dynamic networks. In *STOC*, pages 513–522, 2010.

KM11. Barbara Kempkes and Friedhelm Meyer auf der Heide. Local, self-organizing strategies for robotic formation problems. In *ALGOSENSORS*, pages 4–12, 2011.

KO11. Fabian Kuhn and Rotem Oshman. Dynamic networks: models and algorithms. *SIGACT News*, 42(1):82–96, March 2011.

KWZ03. Fabian Kuhn, Roger Wattenhofer, and Aaron Zollinger. Worst-case optimal and average-case efficient geometric ad-hoc routing. In *MobiHoc*, pages 267–278, 2003.

KWZZ03. Fabian Kuhn, Roger Wattenhofer, Yan Zhang, and Aaron Zollinger. Geometric ad-hoc routing: of theory and practice. In *PODC*, pages 63–72, 2003.

MCS12a. Othon Michail, Ioannis Chatzigiannakis, and Paul G. Spirakis. Brief announcement: Naming and counting in anonymous unknown dynamic networks. In *DISC*, pages 437–438, 2012.

MCS12b. Othon Michail, Ioannis Chatzigiannakis, and Paul G. Spirakis. Causality, influence, and computation in possibly disconnected synchronous dynamic networks. In *OPODIS*, pages 269–283, 2012.

Osh12. Rotem Oshman. *Distributed Computation in Wireless and Dynamic Networks*. PhD thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139, September 2012.

OW05. Regina O'Dell and Roger Wattenhofer. Information dissemination in highly dynamic graphs. In *DIALM-POMC*, pages 104–110, 2005.

SMP12. Atish Das Sarma, Anisur Rahaman Molla, and Gopal Pandurangan. Fast distributed computation in dynamic networks via random walks. In *DISC*, pages 136–150, 2012.