Local Gossip and Neighbour Discovery in Mobile Ad Hoc Radio Networks

Avery Miller

Abstract

We propose a new task called δ-local gossip, which can be viewed as a variant of both gossiping and geocasting. We motivate its study by showing how the tasks of discovering and maintaining neighbourhood information reduce to solving δ-local gossip. We then provide a deterministic algorithm that solves δ-local gossip when nodes travel on a line along arbitrary continuous trajectories with bounded speed.

keywords: neighbour discovery, radio networks, deterministic algorithms, gossiping

1 Introduction

Designing deterministic algorithms for communication tasks in mobile radio networks is difficult due to the fact that the neighbourhoods of nodes can change frequently and unpredictably while an algorithm is executing. Overcoming such difficulties is an important step towards designing ad hoc networks of devices that can cooperate to perform tasks with little to no human supervision, such as self-driving cars or aerial robotic drones. When a radio-equipped device is turned on, it may know some initial information about itself and possibly its location, but it does not have any information about other nearby devices. To form an ad hoc network, devices need to coordinate communication amongst themselves, despite wireless signal interference and changing neighbourhoods, so that they can perform more complex and interesting tasks. Solutions to such tasks often assume that every device initially has information about all other devices within its communication range (or even slightly beyond). We are interested in filling the current gaps in what is known about collecting this initial information without making simplifying assumptions, e.g., without assuming that there is an initial period during which devices are stationary.

To formalize the notion of sharing and collecting information among nearby mobile nodes, we define a new task called δ-local gossip. This task can be viewed as a variant of both gossiping [16] and geocasting [3]: each node initially has a piece of information that it needs to share, but rather than sharing it with all other nodes (as in gossiping), it shares this information with all nodes that enter a certain region (as in geocasting). However, the region is defined as all points that are within a δ-multiple of p's transmission range, so it depends on p's movement (unlike in geocasting where the region is fixed).

In this paper, we restrict our study to deterministic solutions for δ-local gossip. Our motivation stems from the fact that the collection of information could be an initial step or a subroutine of an algorithm that performs a more complex task. Randomized solutions that allow small amounts of error may cause subsequent algorithms that depend on accurate initial information to fail. For example, algorithms for calculating transmission schedules, for determining routing paths, or for avoiding physical collisions, could all be sensitive to errors in information about nearby nodes. This is particularly important in applications where human life is at risk. Further, one cannot amplify the success probability of a randomized algorithm for information collection by using repetition, since the neighbourhoods of mobile nodes can change across different executions.

The organization of the paper is as follows. In Section 1.1, we define our models of mobile radio networks and define δ-local gossip. In Section 2, we motivate the study of δ-local gossip by showing its connections to neighbour discovery. In Section 3, we give a solution to δ-local gossip in a specific model that involves nodes moving on a line along arbitrary continuous trajectories with bounded speed.

1.1 Models and Definitions

We formally define a mobile radio network model called *Mobile-Rcv*, which originally appeared in [8]. The network consists of a fixed set $\mathcal{V}$ of n nodes with distinct identifiers. With each node p, we associate a *trajectory function* p: given a time τ as input, $p[\tau]$ is p's location in the environment at time τ. Each node moves along a continuous trajectory, and, at all times, its speed is bounded above by some known constant σ. Each node can accurately determine its location relative to a global origin. Unless specified otherwise, we assume that every node knows its entire trajectory function. We assume that nodes have synchronized clocks. Time is partitioned into slots, where each slot is an interval of fixed length. For all $i \geq 1$, all nodes execute slot i of their algorithm at the same time.

Each node in $\mathcal{V}$ possesses a radio with which it can perform a transmission on the shared channel in any slot t. A node can only start transmissions at slot boundaries, and the length of each transmission is exactly the length of one slot. If a node is not transmitting in a slot, we say that it is *listening*. Each radio has *transmission radius* R and *interference radius* R'. A node p receives a message from a node q if and only if all of the following hold: (1) p does not transmit during slot t; (2) q transmits during slot t, and, at all times τ in slot t, the distance between p and q at time τ is at most R; (3) for all nodes $q' \neq q$ that transmit during slot t, at all times τ in slot t, the distance between q' and p at time τ is greater than R'. If the first two of the above conditions are satisfied, but the third is not, then we say that a *transmission collision* occurs due to interference. We assume that nodes do not possess collision detectors, so, for each node p and each slot t, either p receives a message in $\{0,1\}^*$ from a node q during slot t, or, p receives $\perp$ during slot t (which represents "no message"). We assume that all nodes know the values of R, R', and σ. For two nodes u, v, we say that v *is a neighbour of* u *for slot* t, if, when u is the only transmitting node in slot t, v receives u's transmitted message in slot t. In particular, u and v are neighbours for slot t if the distance between them for the entirety of slot t is at most R. The *neighbour graph* G_{nbr} is a sequence of graphs $\{G_t\}_{t\geq 1}$ where, for each $t \geq 1$, we have $V(G_t) = \mathcal{V}$ and $\{u, v\} \in G_t$ if and only if u and v are neighbours for slot t.

Local Gossip Informally, the main task considered in this paper is to guarantee that every node obtains a piece of information from each node that is ever located within a δ-multiple of its transmission radius. For example, if $\delta > 1$, this could allow nodes to get advance warning before another node becomes its neighbour.

Definition 1 (δ-local gossip)**.** *Suppose that each node p in the network initially possesses a piece* *of information I_j. Given a constant $\delta > 0$, there is a known $T_\delta > 0$ such that each node terminates at the end of slot T_δ and, at termination, for all nodes p_j and p_k such that $d(p_j, p_k) \leq \delta R$ at some time before or at the end of slot T_δ, p_j knows I_k.*

The condition that $d(p_j, p_k) \leq \delta R$ at some time during the execution is quite weak. For example, the distance between p_j and p_k could be greater than δR up until the exact moment that they both terminate. Also, if $\delta > 1$, it is possible that nodes p_j and p_k are never within transmission range of one another. Further, the task does not even require that p_j and p_k receive at least one transmission from one another, even if they are neighbours for the entire execution. This could happen due to signal interference from simultaneous transmissions.

At a high level, δ-local gossip is similar to geocasting [3], in that information must be disseminated to all nodes within a specified region. However, there are several significant differences. First, the geocasting task specifies a single source with a single piece of information, whereas, in δ-local gossip, each node has a distinct piece of information that it must share. In this sense, δ-local gossip can be considered a 'multi-source' version of geocasting. Also, in geocasting, the region within which the source's information must be shared is fixed relative to the starting point of the source, whereas, in δ-local gossip, the region follows the node's trajectory. Finally, in geocasting, only nodes that stay within the specified region for a long time must receive the source's message, whereas, in δ-local gossip, for any node q that is within distance δR from a node p at *some* time during the execution, p must receive q's piece of information.

1.2 Related Work

Motived by real-world networks, there has been increased activity in the theoretical study of distributed algorithms for mobile and dynamic networks [2, 15]. Much of this study has focused on dynamic graph models, while some models opt to explicitly include device movement along continuous trajectories and must deal with the resulting geometric considerations. In this latter category, the existing work about information dissemination (i.e., broadcasting, gossiping, geocasting) is most closely related to our work [1, 3, 16].

A survey of results regarding neighbour discovery in mobile networks can be found in [12]. In the model we consider in this paper, some previous work solves the *neighbourhood maintenance* task, in which it is assumed that each node initially knows its neighbourhood exactly, and, at all times before termination, each node must maintain an accurate list of its neighbours. Two solutions are known for this task in the *Mobile-Rcv* model: one for the line network by Ellen, Welch and Subramanian [8], which we refer to as the `EWS` algorithm, and one for road networks by Chung, Viqar and Welch [4], which we refer to as the `CVW` algorithm. These algorithms work if a certain amount of neighbourhood information is initially known by each node. In [4], they state that this information can be acquired by running a gossiping algorithm for static networks, under the assumption that all nodes stay close to their starting position for the entire execution. In Section 2, we will show how a δ-local gossip algorithm can be used to ensure that each node acquires this initial information without making such an assumption. In the *Mobile-Rcv* model with $R' = R$, Cornejo et al. [6] provide an algorithm for a neighbour discovery task that we call "continual stable-neighbour discovery" in Section 2. Their solution assumes the existence of a MAC layer, as defined in [14] (without aborts), so they do not need to consider transmission collisions. There are currently no deterministic implementations of the MAC layer that they use, and the probabilistic implementations that they cite do not guarantee message delivery.

2 Neighbour Discovery and Local Gossip

Neighbour discovery has been well-studied in the case of static networks, i.e., where the devices do not move or fail [5, 10, 13, 17, 18]. However, in mobile networks, it is not even clear how to define the task: since the neighbour graph G_{nbr} is a dynamic graph that can change often during an algorithm's execution, which nodes should a node v consider as its neighbours? In this section, we formalize neighbour discovery tasks and prove reductions that show the natural connections between neighbour discovery and δ-local gossip.

To solve a neighbour discovery task, each node p must calculate, at the start of each slot t, a *neighbour list for slot* t, which we denote by $\texttt{List}(p, t)$. To define a neighbour discovery task, we will specify conditions on the relationship between $\texttt{List}(p, t)$ and p's actual set of neighbours in G_t, denoted by $NBRS(p, t)$. In all cases, we assume that nodes have no initial information about their neighbours.

We distinguish between neighbourhood discovery tasks in several ways. First, we consider the *permanence* of the task: do nodes determine their neighbourhoods for a specific slot, i.e., *one-time* discovery, or do nodes keep learning about their neighbourhoods as the neighbourhoods change, i.e., *continual* discovery? One-time algorithms are useful in order to satisfy the initial conditions of a subsequent algorithm. In this case, we want all nodes to terminate the neighbour discovery algorithm at the same time, and we want all nodes to determine who their neighbours are at termination. In contrast, when discovering neighbours to facilitate on-going communication tasks such as routing, we want nodes to continually learn about changes in their neighbourhood. Next, we consider the *accuracy* of the task. In some problems, each node's neighbour list must match its actual neighbourhood exactly. For other problems, each node's neighbour list will be a subset or superset of its neighbourhood. Depending on the application, algorithms that calculate inexact neighbour lists might be sufficient. For example, if nodes are keeping a list of neighbours in order to avoid physical collisions, then, as long as each node's neighbour list is always a superset of its actual neighbourhood, the system can accomplish its goal. However, such a solution might be less efficient than in the case where exact neighbourhoods are known.

Using the above distinctions, we give formal definitions for six neighbour discovery tasks and provide reductions from each task to δ-local gossip. The reductions are general, i.e., they do not depend on our specific network model. In what follows, suppose that we have an algorithm $\texttt{LG}(\delta)$ that solves δ-local gossip, and we denote its running time by $|\texttt{LG}(\delta)|$.

One-time Exact Neighbour Discovery. A solution to this task guarantees that, at termination, each node knows the ID of each of its neighbours for the next slot. Such an algorithm is useful as an initialization step before executing an algorithm that assumes that each node initially knows its exact neighbourhood. This is equivalent to how neighbour discovery is defined for static networks. A parameter T_{Init} specifies how long the discovery process takes.

Definition 2. *An algorithm solves* one-time exact neighbourhood discovery *if, for some known* $T_{Init} > 0$*, every execution terminates at the end of slot* T_{Init}*, and, at termination, each node* p *has* $\texttt{List}(p, T_{Init} + 1) = NBRS(p, T_{Init} + 1)$.

Define algorithm $\texttt{OENL}$ as the execution of $\texttt{LG}(1)$, where, for each p_j, the value of I_j is p_j's trajectory for slot $|\texttt{LG}(1)| + 1$.

Lemma 1. $\texttt{OENL}$ *solves one-time exact neighbour discovery with* $T_{Init} = |\texttt{LG}(1)| + 1$.

Proof. From the definition of 1-local gossip, at the end of slot $|\texttt{LG}(1)|$, for each node p and every node p_j within distance R from p, p has received I_j. Since I_j consists of p_j's trajectory for slot $|\texttt{LG}(1)|+1$, p can determine which nodes will be at most distance R away for the entirety of slot $|\texttt{LG}(1)|+1$. Therefore, `OENL` solves one-time exact neighbour discovery with $T_{\text{Init}} = |\texttt{LG}(1)|+1$. □

Continual Exact Neighbour Discovery. This is the strongest possible version of neighbour discovery: after some fixed number of slots, all nodes know the identity of all neighbours at all times.

Definition 3. *An algorithm solves* continual exact neighbour discovery *if, for some known $T_{Init} > 0$, during every execution,* $\texttt{List}(p,t) = NBRS(p,t)$ *for every node p and every slot $t \geq T_{Init}+1$.*

First, we consider models where each node knows its entire future trajectory. One approach is to use a known neighbourhood maintenance algorithm. For example, in the *Mobile-Rcv* model where nodes travel along arbitrary, continuous trajectories on a line, we could run the `EWS` algorithm [8]. Similarly, in the *Mobile-Rcv* model where nodes travel along arbitrary, continuous trajectories in a road network, we could run the `CVW` algorithm [4]. However, both of these algorithms make strong assumptions about each node initially knowing the entire future trajectories of all other nodes within a certain distance. Let `EWSINIT` be the algorithm that consists of executing a $\frac{4}{3}$-local gossip algorithm where, for each node p_j, we take I_j to be p_j's entire future trajectory. Let `CVWINIT` be the algorithm that consists of executing a $\frac{13}{11}$-local gossip algorithm where, for each node p_j, we take I_j to be p_j's entire future trajectory. The following results can be verified using the model constraints from [8] and [4], respectively.

Lemma 2. *In the Mobile-Rcv model where nodes travel along arbitrary continuous trajectories on the line, executing* `EWSINIT` *followed by* `EWS` *solves continual exact neighbour discovery with* $T_{Init} = |\texttt{LG}(\frac{4}{3})|$.

Proof. The initial condition required by `EWS` is that each node has received the entire trajectory of each node that is initially within distance $R+2(m-1)\sigma$. From the definition of $\frac{4}{3}$-local gossip, it follows that, at termination, each node in the network has received I_j for each j such that p_j is within distance $\frac{4}{3}R$. Note that, from constraints (C1) and (C3) specified in [8], $L \leq \frac{R-3(m-1)\sigma-3K}{2} < \frac{R-3(m-1)\sigma-3(m-1)\sigma}{2}$. Since $L \geq 0$, it follows that $6(m-1)\sigma < R$, and, hence, $\frac{4}{3}R > R+2(m-1)\sigma$. Therefore, the initialization is complete at the end of slot $|\texttt{LG}(\frac{4}{3})|$, which means that the continual exact neighbour discovery task is solved with $T_{\text{Init}} = |\texttt{LG}(\frac{4}{3})|$. □

Lemma 3. *In the Mobile-Rcv model where nodes travel along arbitrary continuous trajectories in road networks, executing* `CVWINIT` *followed by* `CVW` *solves continual exact neighbour discovery with* $T_{Init} = |\texttt{LG}(\frac{13}{11})|$.

Proof. The initial condition required by `CVW` is that each node has received the entire trajectory of each node that is within distance $R+2mu\sigma$. From the definition of $\frac{13}{11}$-local gossip, it follows that, at termination, each node in the network has received I_j for each j such that p_j is within distance $\frac{13}{11}R$. Note that, from constraints (D2) and (D4) specified in [4], $L \leq \frac{R-8mu\sigma-3K}{2} < \frac{R-8mu\sigma-3mu\sigma}{2}$. Since $L \geq 0$, it follows that $11mu\sigma < R$, and, hence, $\frac{13}{11}R > R+2mu\sigma$. Therefore, the initialization is complete at the end of slot $|\texttt{LG}(\frac{13}{11})|$, which means that the continual exact neighbour discovery task is solved with $T_{\text{Init}} = |\texttt{LG}(\frac{13}{11})|$. □

We can weaken the assumption about entire trajectory knowledge by using δ-local gossip to communicate trajectory updates. In each phase of our algorithm, called CENL, each node shares enough future trajectory information to ensure that all nodes can correctly calculate their neighbourhoods for all slots up until the completion of the next phase. More specifically, for some fixed δ (to be specified later), let phase $i \geq 0$ consist of the slots $i|\texttt{LG}(\delta)| + 1, \ldots, (i+1)|\texttt{LG}(\delta)|$. During each phase i, all nodes execute $\texttt{LG}(\delta)$, where each p_j sets I_j to be its trajectory for slots $(i+1)|\texttt{LG}(\delta)|+1, \ldots, (i+2)|\texttt{LG}(\delta)|$. So, by the end of phase i, each neighbour of p_j during phase $i+1$ receives p_j's trajectory for all times in phase $i+1$. The value of δ is chosen so that, for any node p_j, every node that is a neighbour of p_j during at least one slot in phase $i+1$ is located within distance δR of p_j at some time during phase i. It is sufficient to choose δ such that $\delta R \geq R + 2\sigma|\texttt{LG}(\delta)|$, which implies the following result.

Lemma 4. CENL *solves continual exact neighbour discovery with* $T_{Init} = |\texttt{LG}(\delta)|$*, where* δ *satisfies* $\delta R \geq R + 2\sigma|\texttt{LG}(\delta)|$.

One-Time Stable-Neighbour Discovery. In some applications, it might be useful to learn about neighbours that, in a sense, can be considered to be more 'dependable'. That is, node p cares only about neighbours that will stay nearby for a while, and ignores nodes that enter its neighbourhood and then leave shortly afterward. For example, in the context of routing algorithms, such a neighbour discovery algorithm could be used to periodically update routing tables at every node.

Definition 4. *An algorithm solves* one-time stable-neighbour discovery *if, for some fixed* $T_{Init} > 0$ *and some fixed* $T_{Stable} > 0$*, every execution terminates at the end of slot* T_{Init}*, and, for all nodes* p*,* $q \in \texttt{List}(p, T_{Init}+1)$ *if and only if* $q \in NBRS(p,t)$ *for all* $t \in \{T_{Init}+1, \ldots, T_{Init}+T_{Stable}\}$.

Suppose that, for some $F_{\text{traj}} > |\texttt{LG}(1)|$, each node initially knows its trajectory for slots $|\texttt{LG}(1)| + 1, \ldots, F_{\text{traj}}$. Define algorithm OSNL as the execution of $\texttt{LG}(1)$ where, for each node p_j, the value of I_j is defined as its trajectory for slots $|\texttt{LG}(1)|+1, \ldots, F_{\text{traj}}$. At the end of slot $|\texttt{LG}(1)|$, let $\texttt{List}(p, |\texttt{LG}(1)|+1) = \{p_k \mid I_k$ has been received, and, p_k will be a neighbour for all slots $|\texttt{LG}(1)|+1, \ldots, F_{\text{traj}}\}$.

Lemma 5. OSNL *solves one-time stable-neighbour discovery with* $T_{Init} = |\texttt{LG}(1)|$ *and* $T_{Stable} = F_{traj} - |\texttt{LG}(1)|$.

Proof. Consider any node q that is in $NBRS(p,t)$ for all $t \in \{|\texttt{LG}(1)|+1, \ldots, F_{\text{traj}}\}$. It must be within distance R of p at the end of slot $|\texttt{LG}(1)|$. By the definition of 1-local gossip, p receives q's trajectory for slots $|\texttt{LG}(1)|+1, \ldots, F_{\text{traj}}$ by the end of slot $|\texttt{LG}(1)|$. Therefore, at the beginning of slot $|\texttt{LG}(1)|+1$, p has received enough trajectory information for every node that it must include in $\texttt{List}(p, |\texttt{LG}(1)|+1)$. □

Continual Stable-Neighbour Discovery. In the continual version of stable-neighbour discovery, we want node p to consider a node q as a neighbour if and only if q enters p's neighbourhood and stays for a while. Informally, if q is a neighbour of p for all slots during some sufficiently long interval, then q must be contained in p's neighbour list for some suffix of the interval. Conversely, if p includes q in its neighbour list for some slot t, it must be the case q is a neighbour of p for a sufficiently long interval that includes t. A parameter T_{Stable} specifies how many slots must elapse before two neighbours are considered "stable", while a parameter T_{Delay} specifies an upper bound on the amount of delay before a node adds a stable neighbour to its neighbour list.

Definition 5. *An algorithm solves* continual stable-neighbour discovery *if there exist $T_{Stable} > 0$ and $T_{Delay} \geq 0$ such that, for every node p, every slot t, and every $T \geq t + T_{Stable} - 1$: if $q \in NBRS(p, t')$ for all $t' \in \{t, \ldots, T\}$, then $q \in \texttt{List}(p, t'')$ for all $t'' \in \{t + T_{Delay}, \ldots, T\}$, and, if $q \in \texttt{List}(p, t)$, then there exists $t' \in \{t - T_{Stable} + 1, \ldots, t\}$ such that $q \in NBRS(p, t')$ for all $t'' \in \{t', \ldots, t' + T_{Stable} - 1\}$.*

We now describe an algorithm, CSNL, for continual stable-neighbour discovery. At a high level, our solution proceeds in phases, each consisting of an execution of $\texttt{LG}(1)$. Initially, each node's neighbour list is empty for all slots. When a node p receives trajectory information about a node q that will be its neighbour for a while, then p adds q to its neighbour list for the appropriate future slots. More specifically, we ensure that, by the end of each phase i, each node p receives q's trajectory information for phases i and $i+1$, for each node q within distance R of p at some time during phase i. Using this trajectory information, p can determine if q is a neighbour for a suffix of phase i and a prefix of phase $i+1$ whose lengths a and b, respectively, add up to at least $|\texttt{LG}(1)| + 1$ slots. In this case, p adds q to its neighbour list for the first b slots of phase $i+1$. Our algorithm proceeds as follows at each node p_j and for each phase $i \geq 0$ consisting of slots $i|\texttt{LG}(1)| + 1, \ldots, (i+1)|\texttt{LG}(1)|$:

- Set I_j to be p_j's trajectory for slots $i|\texttt{LG}(1)| + 1, \ldots, (i+2)|\texttt{LG}(1)|$. Run $\texttt{LG}(1)$ at the start of slot $i|\texttt{LG}(1)| + 1$.
- At the end of each slot during which a message containing some I_k is received: find the largest positive integer $a \leq |\texttt{LG}(1)|$ such that p_j and p_k are neighbours for every slot in $\{(i+1)|\texttt{LG}(1)| + 1 - a, \ldots, (i+1)|\texttt{LG}(1)|\}$, and find the largest positive integer $b \leq |\texttt{LG}(1)|$ such that p_j and p_k are neighbours for every slot in $\{(i+1)|\texttt{LG}(1)| + 1, \ldots, (i+1)|\texttt{LG}(1)| + b\}$. If $a + b \geq |\texttt{LG}(1)| + 1$, put p_k in $\texttt{List}(p_j, t)$ for each $t \in \{(i+1)|\texttt{LG}(1)| + 1, \ldots, (i+1)|\texttt{LG}(1)| + b\}$.

Lemma 6. CSNL *solves the continual stable-neighbour discovery task with $T_{Stable} = |\texttt{LG}(1)| + 1$ and $T_{Delay} = |\texttt{LG}(1)|$.*

Proof. Note that, if two nodes p_j and p_k become neighbours during a phase i, they may not add each other to their neighbour lists until the end of phase i. In the worst case, p_k is p_j's neighbour for $|\texttt{LG}(1)| + 1$ slots starting at the beginning of a phase i, but p_j does not receive I_k until the end of phase i. In this case, p adds q as a neighbour $|\texttt{LG}(1)|$ slots after they became neighbours. Therefore, the worst-case delay, T_{Delay}, is at least $|\texttt{LG}(1)|$.

We now set out to prove that our algorithm solves the continual stable-neighbour discovery task. First, by the definition of the algorithm, if p_j adds p_k to $\texttt{List}(p_j, t)$ for some slot t, there must be an interval containing t of length at least $|\texttt{LG}(1)| + 1$ such that p_j and p_k are neighbours. Next, we show that if p_j and p_k are neighbours for an interval of at least $|\texttt{LG}(1)| + 1$ consecutive slots, p_j adds p_k to its neighbour list within $|\texttt{LG}(1)|$ slots from the start of the interval. The following result considers any subinterval that spans at most two consecutive phases of the algorithm.

Claim 1. *Consider an arbitrary node p_j and an arbitrary slot $t_1 = i|\texttt{LG}(1)| + \ell$ with $i \geq 0$ and $\ell \in \{1, \ldots, |\texttt{LG}(1)|\}$. If there exists a slot $t_2 \in \{t_1 + |\texttt{LG}(1)|, \ldots, (i+2)|\texttt{LG}(1)|\}$ and a node p_k such that, for all slots $t' \in \{t_1, \ldots, t_2\}$, $p_k \in NBRS(p_j, t')$, then, for all $t'' \in \{(i+1)|\texttt{LG}(1)| + 1, \ldots, t_2\}$, $p_k \in \texttt{List}(p_j, t'')$.*

To prove the claim, note that slot t_1 is contained in phase i. It follows that p_j and p_k are neighbours for some slot in phase i. Also, since $t_2 \geq t_1 + |\texttt{LG}(1)| \geq (i+1)|\texttt{LG}(1)| + 1$, it follows that slot t_2 is contained in phase $i+1$. In particular, this means that slot $(i+1)|\texttt{LG}(1)| + 1$ is contained in $\{t_1, \ldots, t_2\}$. By the definition of 1-local gossip, by the end of phase i, p_j has received I_k, which consists of p_k's trajectory for all slots in phases i and $i+1$. Namely, p_j has received p_k's trajectory for all slots in $\{t_1, \ldots, t_2\}$. Therefore, by the end of phase i, p_j has determined that p_k is its neighbour for all slots in $\{t_1, \ldots, t_2\}$ (which consists of at least $|\texttt{LG}(1)|+1$ slots), so p_j adds p_k to $\texttt{List}(p_j, t'')$ for all $t'' \in \{(i+1)|\texttt{LG}(1)|+1, \ldots, t_2\}$. This concludes the proof of the claim.

Now, let $t_2 \geq t_1 + |\texttt{LG}(1)|$ and suppose that $\{t_1, \ldots, t_2\}$ is a maximal set of slots for which nodes p_j and p_k are neighbours. For an arbitrary $t''' \in \{t_1 + |\texttt{LG}(1)|, \ldots, t_2\}$, we show that $p_k \in \texttt{List}(p_j, t''')$.

For some $i \geq 0$, $t_1 \in \{i|\texttt{LG}(1)|+1, \ldots, (i+1)|\texttt{LG}(1)|\}$. If $t_2 \leq (i+2)|\texttt{LG}(1)|$, then, by Claim 1, it follows that $p_k \in \texttt{List}(p_j, t''')$. So, in what follows, we assume that $t_2 \geq (i+2)|\texttt{LG}(1)|+1$. Suppose that $t''' \in \{(i+1)|\texttt{LG}(1)|+1, \ldots, (i+2)|\texttt{LG}(1)|\}$. By Claim 1 with $t_2 = (i+2)|\texttt{LG}(1)|$, it follows that $p_k \in \texttt{List}(p_j, t''')$. So, we proceed with the assumption that $t''' \geq (i+2)|\texttt{LG}(1)|+1$. This means that t''' is contained in a phase i' with $i' \geq i+2$. Since t_1 is contained in phase i, it follows that p_j and p_k are neighbours for every slot in phase $i'-1$. By Claim 1 with $t_1 = (i'-1)|\texttt{LG}(1)|+1$ and $t_2 = t'''$, it follows that $p_k \in \texttt{List}(p_j, t''')$. □

One-Time Delayed Neighbour Discovery. Another version of neighbourhood discovery involves each node discovering which nodes are its neighbours, but allowing a certain amount of delay before this occurs. An upper bound on the amount of delay is specified by a T_{Del} parameter. A node's neighbour list at termination could be a subset or a superset of its actual neighbourhood, but the learned neighbourhood information can still provide useful estimates. For example, if there is a known upper bound σ on the speed of nodes, then, knowing that a node q was a neighbour t' slots ago can help provide a range of possible locations for q in the current slot.

Definition 6. *An algorithm solves* one-time delayed neighbour discovery *if, for some fixed* $T_{Init} > 0$ *and some fixed* $T_{Del} > 0$, *every execution terminates at the end of slot* T_{Init}, *and, for all nodes* p, *node* $q \in \texttt{List}(p, T_{Init}+1)$ *if and only if* $q \in NBRS(p,t)$ *for some* $t \in \{T_{Init} - T_{Del} + 1, \ldots, T_{Init}\}$.

Define algorithm $\texttt{ODNL}$ as the execution of $\texttt{LG}(1)$, where, for each p_j, the value of I_j is defined as its trajectory for slots $1, \ldots, |\texttt{LG}(1)|$. For each I_j received, p can determine whether or not p_j was a neighbour at some time during the previous $|\texttt{LG}(1)|$ slots. This implies the following result.

Lemma 7. $\texttt{ODNL}$ *solves one-time delayed neighbour discovery with* $T_{Init} = T_{Del} = |\texttt{LG}(1)|$.

Continual Delayed Neighbour Discovery. In the continual version of delayed neighbour discovery, we want each node to discover its neighbours within a bounded number of slots, and, if a node p has a node q in its neighbour list, then q must have been a neighbour of p in the recent past.

Definition 7. *An algorithm solves* continual delayed neighbour discovery *if, for some fixed* $T_{Del} > 0$, *for every node* p *and every slot* t, *if* $q \in NBRS(p,t)$, *then* $q \in \texttt{List}(p,t')$ *for some* $t' \in \{t+1, \ldots, t+T_{Del}\}$. *Also, if* $q \in \texttt{List}(p,t)$, *then* $q \in NBRS(p,t')$ *for some* $t' \in \{t - T_{Del}, \ldots, t-1\}$.

Define algorithm $\texttt{CDNL}$ as a sequence of phases, each consisting of an execution of $\texttt{LG}(1)$. In each phase $i \geq 0$, each node p_j sets I_j to be its trajectory for slots $i|\texttt{LG}(1)|+1, \ldots, (i+1)|\texttt{LG}(1)|$. If a node p_j receives a message containing I_k during some slot t in phase i, it compares I_k with its own trajectory for all slots that occur during phase i. If it determines that, for some slot t' of phase i, p_k is its neighbour, it adds p_k to $\texttt{List}(p_j, t'+1), \ldots, \texttt{List}(p_j, (i+1)|\texttt{LG}(1)|+1)$.

Lemma 8. $\texttt{CDNL}$ *solves continual delayed neighbour discovery with* $T_{Del} = |\texttt{LG}(1)|$.

Proof. Consider any two nodes p_j and p_k. Suppose that, for some slot t during an arbitrary phase i, $p_k \in NBRS(p_j, t)$. This means that p_k is within distance R from p_j at some point during phase i, so $\texttt{LG}(1)$ ensures that p_j receives I_k by the end of phase i. Therefore, for some $t' \in \{i|\texttt{LG}(1)|+1, \ldots, (i+1)|\texttt{LG}(1)|\}$, p_j receives I_k and adds p_k to every slot in $\texttt{List}(p_j, t+1), \ldots, \texttt{List}(p_j, (i+1)|\texttt{LG}(1)|+1)$. In particular, $p_k \in \texttt{List}(p_j, (i+1)|\texttt{LG}(1)|+1)$, and $(i+1)|\texttt{LG}(1)|+1 \in \{t+1, \ldots, t+|\texttt{LG}(1)|\}$.

Next, for arbitrary $i \geq 0$, suppose that $p_k \in \texttt{List}(p_j, t)$ for some $t \in \{i|\texttt{LG}(1)|+2, \ldots, (i+1)|\texttt{LG}(1)|+1\}$. When p_j added p_k to $\texttt{List}(p_j, t)$, the algorithm was in the process of adding p_k to $\texttt{List}(p_j, t'+1), \ldots, \texttt{List}(p_j, (i+1)|\texttt{LG}(1)|+1)$ for some slot $t' \in \{i|\texttt{LG}(1)|+1, \ldots, (i+1)|\texttt{LG}(1)|\}$. Further, it did so because $p_k \in NBRS(p_j, t')$. Since $t \geq t'+1$ and $t \leq (i+1)|\texttt{LG}(1)|+1 \leq t'+|\texttt{LG}(1)|$, we have shown that $p_k \in NBRS(p_j, t')$ for some $t' \in \{t-|\texttt{LG}(1)|, \ldots, t-1\}$. □

3 Solving Local Gossip

For the *Mobile-Rcv* model where nodes travel along continuous trajectories on a line with speed bounded above by σ, we describe a transmission schedule RBSched and use it to solve δ-local gossip.

At a high level, each phase of the transmission schedule consists of a subset of the nodes running a particular schedule that is based on cover-free families of sets. By carefully specifying which nodes participate in which phases, we ensure that each node transmits frequently and that the contents of its transmitted message gets propagated quickly in both directions. The idea is that the contents of a message transmitted by a node p gets forwarded to other nodes in the network at a faster rate than p can travel. We assume that, initially, each node p_j in the network possesses some piece of information I_j that it will include in its messages.

3.0.1 Phase Schedule PS

In each phase of RBSched, we will use a schedule called PS based on cover-free families of sets [9, 11]. Informally, a cover-free family is a collection of sets such that, for an arbitrary set S in the family, there is no way to construct a set containing all of the elements of S by taking the union of a small number of sets in the family other than S. The formal definitions of covers and cover-free families are as follows. For any set S, an *r-cover for S* is a set family of size r that does not contain S and whose union does contain S. For $r \geq 1$, an *r-cover-free family* $\mathscr{F}$ is a family of sets such that, for each $S \in \mathscr{F}$, there is no r-cover for S consisting of sets from $\mathscr{F} - \{S\}$.

We restrict attention to families whose sets contain only positive integers. In what follows, we will refer to the number of sets in a family $\mathscr{F}$ as the *size* of $\mathscr{F}$, denoted by $|\mathscr{F}|$. The *length* T of a family $\mathscr{F}$ is the largest integer contained in at least one set in $\mathscr{F}$. For an r-cover-free family, it is known that $T \in O(r^2 \log |\mathscr{F}|)$ [9]. It was shown in [5] that an r-cover-free family with size $|\mathscr{F}|$ and length T can be converted into a transmission schedule with T slots for nodes with IDs in the range $\{1, \ldots, |\mathscr{F}|\}$ with the following property: for each set X of at most $r+1$ nodes, for each node $p \in X$, there is a slot in which p is scheduled to transmit and all nodes in $X - \{p\}$ listen.

Let PS denote the schedule obtained from a $((3+2\rho)(\Delta+1))$-cover-free family of size U, where $\rho = \left\lceil \frac{R'}{R} \right\rceil$. Let $|\texttt{PS}|$ denote the number of slots in the resulting schedule. From the discussion above, we know that $|\texttt{PS}| \in O(\rho^2 \Delta^2 \log U)$, and that the following observation holds.

Observation 1. *Consider any set X of at most $(3+2\rho)(\Delta+1)$ nodes. For each node $p \in X$, there is a slot in* PS *during which p is scheduled to transmit and all nodes in $X - \{p\}$ listen.*

When a node p_j transmits, its message contains all of the information that it knows, i.e., $I_j \cup \{I_k \mid \text{node } p_j \text{ has previously received a message containing } I_k\}$.

We note that a different implementation for PS could be used in order to increase the algorithm's robustness. For example, the $(r; d)$-cover-free families from Dyachkov et al. [7] could be used, with $r = (3 + 2\rho)(\Delta + 1)$, to construct a schedule that gives a stronger property than Observation 1: for each node $p \in X$, there are at least d slots during which p is scheduled to transmit and all nodes in $X - \{p\}$ listen. This could help in models with include some amount of unpredictable interference or some number of transient failures.

3.1 Full Schedule RBSched

We start with some definitions. In our model, we assume that each node can accurately determine its location relative to a global origin. Using this origin, we define a set of regions that is known by all nodes. These regions are of length R, each overlapping its neighbouring regions by $5\sigma|\texttt{PS}|$ units. More formally, for all $z \in \mathbb{Z}$, let *region* ψ_z be the set of points $[z(R - 5\sigma|\texttt{PS}|), z(R - 5\sigma|\texttt{PS}|) + R)$.

Next, partition the set of slots into phases of $|\texttt{PS}|$ slots each. We assume that all nodes begin executing the schedule at the same time, starting with slot 1, so that the phase boundaries are aligned across all nodes. More formally, for all $a \in \mathbb{N}$, let *phase* π_a be the slots $a|\texttt{PS}| + 1, \dots, (a + 1)|\texttt{PS}|$.

For each phase, our schedule chooses a subset of *participants* that will start executing PS. Essentially, the nodes in odd-indexed regions will participate in odd-numbered phases, and nodes in even-indexed regions will participate in even-numbered phases. Participants that stay within the same region for the entire phase are *survivors*. More formally, a node p is a *participant in phase* π_a *for region* ψ_z if p is located in region ψ_z at the start of phase π_a and $z \equiv a \mod 2$. A node p is a *survivor in phase* π_a *for region* ψ_z if p is located in region ψ_z at all times during phase π_a and $z \equiv a \mod 2$. Let $\mathcal{P}_{a,z}$ denote the set of participants in phase π_a for region ψ_z. Let $\mathcal{S}_{a,z}$ denote the set of survivors in phase π_a for region ψ_z.

We can now define the schedule RBSched. At the beginning of each phase π_a, each node p checks if its current location is in a region ψ_z such that $a \equiv z \mod 2$. If this is not the case, then node p listens for the entire phase. If located in such a region ψ_z, then p is a participant for region ψ_z in phase π_a and starts running PS at the start of phase. If p ever leaves region ψ_z during phase π_a, it immediately stops executing PS and listens for the remainder of the phase. This is done to reduce contention on the wireless channel and to limit the set of potential locations where a node can cause transmission collisions. Note that, if a node knows its trajectory for the entire phase π_a in advance, it can determine at the start of the phase if it will ever leave its region, and, therefore, can determine whether or not it will be a survivor. In this case, we can make all non-survivors listen for the entire phase rather than have them start executing PS only to drop out later when they actually leave the region. This could result in fewer transmissions, which would reduce energy consumption.

The nodes continue to execute the phases of the transmission schedule until a specified termination condition is reached. We will specify this termination condition in Section 3.3 when describing how to use the schedule to solve δ-local gossip.

3.2 Analyzing the Schedule

To analyze the schedule, our strategy is to provide guarantees about all nodes that are found within specific regions of the network at specific times, and to ensure that, despite the fact that the nodes are following arbitrary trajectories, every node is eventually found in one of these regions.

Informally, the proof technique consists of three parts, which we will subsequently discuss in more detail:

1. Divide the physical environment into equal-sized regions large enough so that nodes cannot pass through them quickly. These regions should overlap so that a node cannot 'wiggle' back and forth between two regions. The region size will depend on the known upper bound σ on the distance a processor can move in one slot.

2. Create one or more "windows", each the size of one region, and define how they jump from region to region. Then, prove that every node is eventually located within some window.

3. Prove a window "invariant", which is a statement that is true about all nodes that are found in the window.

For our analysis, we partition the line into *segments* of length $L_{\text{SEG}} = R - 5\sigma|\texttt{PS}|$. We then define each region by extending each segment in both directions by $L_{\text{REACH}} = \frac{5}{2}\sigma|\texttt{PS}|$. In particular, for all $z \in \mathbb{Z}$, segment SEG_z is defined as the interval $[zL_{\text{SEG}} + L_{\text{REACH}}, (z+1)L_{\text{SEG}} + L_{\text{REACH}}) = [z(R - 5\sigma|\texttt{PS}|) + \frac{5}{2}\sigma|\texttt{PS}|, z(R - 5\sigma|\texttt{PS}|) + R - \frac{5}{2}\sigma|\texttt{PS}|)$, and region ψ_z is defined as the interval $[z(R - 5\sigma|\texttt{PS}|), z(R - 5\sigma|\texttt{PS}|) + R)$. To motivate these choices, notice that each region is of size R, so that two nodes in the same region are within communication range. Further, the overlap between regions is large enough so that, for example, if a node crosses the rightmost border of region ψ_z frequently within some phase, then the node is actually found in region ψ_{z+1} for the entire phase. It is also the case that even if a node travels at maximum speed σ, it cannot pass through an entire region during a single phase.

We define two windows as follows. One window is a rightward-moving window that 'jumps' one region rightward at the beginning of every 3 consecutive phases, while a leftward-moving window 'jumps' one region leftward at the beginning of every 3 consecutive phases. More formally, we say that, for phase $\pi_{a+3\varphi}$, a node q is *located in the rightward-moving window* if it is located in region $\pi_{z+\varphi}$ for the entire phase. Similarly, for phase $\pi_{a+3\varphi}$, a node q is *located in the leftward-moving window* if it is located in region $\pi_{z-\varphi}$ for the entire phase. At a high level, these windows represent how fast information is propagated through the network. Each time a transmission occurs containing some information I, we consider windows that move leftward and rightward from the region where the transmission originated. We give an upper bound on the amount of time until an arbitrary node q is found within a window, which gives a bound on how soon q will receive I.

To prove certain useful guarantees about our schedule, we will depend on the following assumptions that relate the various model parameters.

(A1) **Density:** Between the leftmost and rightmost nodes, there is never a line segment of length $\frac{R}{2}$ that contains no nodes.

(A2) **Region Spacing:** $R \geq 10\sigma|\texttt{PS}|$.

Assumption (A2) ensures that each region ψ_z only overlaps with regions ψ_{z-1} and ψ_{z+1}.

Proposition 1. *For every $z \in \mathbb{Z}$, $\psi_z \cap \psi_{z+2} = \emptyset$.*

First, we prove constraints on the possible locations of nodes at different times. For example, given a node's location at one specific time t, using the upper bound on node speed, we specify the range of possible locations where the node can be located within certain time intervals before or after t. Using these results, we will be able to prove upper bounds on the number of phases that elapse before an arbitrary node is found within one of the windows.

Then, we prove guarantees on the speed of message propagation. Namely, if I is transmitted during some phase, we give an upper bound on the time that elapses before nodes in a given area receive I. Using these results, we will be able to prove a window invariant that says that every node that is found in a window will receive I soon afterward.

3.2.1 Node Location Bounds

Based on the upper bound σ on node speed and the length $|\mathtt{PS}|$ of each phase, we get the following results that limit how far a node can travel from a known location within a bounded amount of time.

Observation 2. *If, at some time during phase $\pi_{a'}$, node p is located at a point x, then, at all times from the beginning of phase $\pi_{a'-k}$ until the end of phase $\pi_{a'+k}$, p is located in $[x-(k+1)\sigma|\mathtt{PS}|, x+(k+1)\sigma|\mathtt{PS}|]$.*

A useful application of this observation allows us to guarantee that, if a node p is located close enough to the middle of a region ψ_z at some time during a phase π_a, then p is located in region ψ_z for the entire phase. Further, if $a \equiv z \mod 2$, then it follows that p must be a survivor in phase π_a for region ψ_z.

Proposition 2. *Suppose that $a' \equiv z' \mod 2$. If, at some time during phase π_a, node p is located in $[z'(R-5\sigma|\mathtt{PS}|)+\sigma|\mathtt{PS}|, z'(R-5\sigma|\mathtt{PS}|)+R-\sigma|\mathtt{PS}|)$, then $p \in \mathcal{S}_{a',z'}$.*

Next, we show that each node is a survivor at least once in every two consecutive phases. This fact gives us an upper bound on the number of slots that elapse before any given node transmits.

Lemma 9. *Suppose that, at the beginning of phase $\pi_{a'}$, node p is located in segment $SEG_{z'}$. Then, $p \in \mathcal{S}_{a',z'} \cup \mathcal{S}_{a'+1,z'}$.*

Proof. At the beginning of phase $\pi_{a'}$, p is located in $[z(R-5\sigma|\mathtt{PS}|)+\frac{5}{2}\sigma|\mathtt{PS}|, z(R-5\sigma|\mathtt{PS}|)+R-\frac{5}{2}\sigma|\mathtt{PS}|)$. By Observation 2 with $k=1$, at all times from the beginning of phase $\pi_{a'}$ until the end of phase $\pi_{a'+1}$, p is located in

$$\begin{aligned} & [z'(R-5\sigma|\mathtt{PS}|)+\tfrac{5}{2}\sigma|\mathtt{PS}|-2\sigma|\mathtt{PS}|, z'(R-5\sigma|\mathtt{PS}|)+R-\tfrac{5}{2}\sigma|\mathtt{PS}|+2\sigma|\mathtt{PS}|) \\ = & [z'(R-5\sigma|\mathtt{PS}|)+\tfrac{1}{2}\sigma|\mathtt{PS}|, z'(R-5\sigma|\mathtt{PS}|)+R-\tfrac{1}{2}\sigma|\mathtt{PS}|) \\ \subseteq & [z'(R-5\sigma|\mathtt{PS}|), z'(R-5\sigma|\mathtt{PS}|)+R) \\ = & \psi_{z'} \end{aligned}$$

Therefore, at all times in $\pi_{a'} \cup \pi_{a'+1}$, p is located in region $\psi_{z'}$. Since either $a' \equiv z' \mod 2$ or $a'+1 \equiv z' \mod 2$, $p \in \mathcal{S}_{a',z'} \cup \mathcal{S}_{a'+1,z'}$. □

Other useful applications of Observation 2 allow us to narrow down in which regions a node p can be located based on the fact that p was a survivor. Namely, if $p \in \mathcal{S}_{a,z}$, then p cannot be located far away from region ψ_z at times shortly before and after phase π_a.

Proposition 3. *Consider any node $p \in \mathcal{S}_{a,z}$. From the beginning of phase π_{a-4} until the end of phase π_{a+4}, p is located in $\psi_{z-1} \cup \psi_z \cup \psi_{z+1}$.*

Proof. If $p \in \mathcal{S}_{a,z}$, then, at all times during phase π_a, p is located in region $\psi_z = [z(R-5\sigma|\mathtt{PS}|), z(R-5\sigma|\mathtt{PS}|)+R)$. By Observation 2 with $a' = a$ and $k = 4$, at all times from the beginning of phase π_{a-4} until the end of phase π_{a+4}, p is located in

$$\begin{aligned}
& [z(R-5\sigma|\mathtt{PS}|)-5\sigma|\mathtt{PS}|, z(R-5\sigma|\mathtt{PS}|)+R+5\sigma|\mathtt{PS}|) \\
= \; & [(z-1)(R-5\sigma|\mathtt{PS}|)+R-10\sigma|\mathtt{PS}|, (z+1)(R-5\sigma|\mathtt{PS}|)+10\sigma|\mathtt{PS}|) \\
\subseteq \; & [(z-1)(R-5\sigma|\mathtt{PS}|), (z+1)(R-5\sigma|\mathtt{PS}|)+R) \text{ (by assumption (A2))} \\
= \; & \psi_{z-1} \cup \psi_z \cup \psi_{z+1}
\end{aligned}$$

□

We now prove some results that will help us show that each node is eventually found within one of the windows. First, we show that nodes cannot 'sneak by' the windows. For example, if a node is located to the right of some window at time t, it cannot be located to the left of the window at a time $t' > t$ without first being located in the window.

Proposition 4. *Suppose that, at the beginning of phase π_a, node q is located in or to the right of segment SEG_z. For any fixed $\gamma \geq 0$, if, for all $\ell \in \{0,\ldots,\gamma\}$, there is a time during phase $\pi_{a+3\ell}$ such that q is not located in region $\psi_{z+\ell}$, then q is located in or to the right of region $\psi_{z+\gamma}$ throughout phase $\pi_{a+3\gamma}$.*

Proof. The proof is by induction on γ. First, consider the case when $\gamma = 0$. Since the leftmost point of segment SEG_z is $z(R-5\sigma|\mathtt{PS}|)+\frac{5}{2}\sigma|\mathtt{PS}|$, q is located to the right of $z(R-5\sigma|\mathtt{PS}|)+\frac{5}{2}\sigma|\mathtt{PS}|$ at beginning of phase π_a. By Observation 2 with $a' = a$ and $k = 0$, at all times during phase π_a, q is located to the right of $z(R-5\sigma|\mathtt{PS}|)+\frac{5}{2}\sigma|\mathtt{PS}|-\sigma|\mathtt{PS}| = z(R-5\sigma|\mathtt{PS}|)+\frac{3}{2}\sigma|\mathtt{PS}|$. But, the leftmost point of region ψ_z is $z(R-5\sigma|\mathtt{PS}|)$, so q is not located to the left of region ψ_z at any time during phase π_a.

Now, consider the case when $\gamma > 0$. Suppose that, for all $\ell \in \{0,\ldots,\gamma\}$, there is a time during phase $\pi_{a+3\ell}$ such that q is not located in region $\psi_{z+\ell}$. As induction hypothesis, assume that q is not located to the left of region $\psi_{z+\gamma-1}$ at any time during phase $\pi_{a+3\gamma-3}$. However, there is a time during $\pi_{a+3\gamma-3}$ such that q is not located in region $\psi_{z+\gamma-1}$. Therefore, at some time during phase $\pi_{a+3\gamma-3}$, q is located to the right of the point $(z+\gamma-1)(R-5\sigma|\mathtt{PS}|)+R$. By Observation 2 with $a' = a+3\gamma-3$ and $k = 3$, at all times from the beginning of phase $\pi_{a+3\gamma-6}$ until the end of phase $\pi_{a+3\gamma}$, q is located to the right of the point $(z+\gamma-1)(R-5\sigma|\mathtt{PS}|)+R-4\sigma|\mathtt{PS}| > (z+\gamma-1)(R-5\sigma|\mathtt{PS}|)+(R-5\sigma|\mathtt{PS}|) = (z+\gamma)(R-5\sigma|\mathtt{PS}|)$. In particular, at all times during phase $\pi_{a+3\gamma}$, q is located in or to the right of region $\psi_{z+\gamma}$. □

By symmetry, the analogous result holds for the left-moving window.

Proposition 5. *Suppose that, at the beginning of phase π_a, node q is located in or to the left of segment SEG_z. For any fixed $\gamma \geq 0$, if, for all $\ell \in \{0,\ldots,\gamma\}$, there is a time during phase $\pi_{a+3\ell}$ such that q is not located in region $\psi_{z-\ell}$, then q is located in or to the left of region $\psi_{z-\gamma}$ throughout phase $\pi_{a+3\gamma}$.*

The following result combines Propositions 4 and 5 so that we will be able to prove that all nodes found within a certain area of the environment have all been located within a window at some point during an algorithm's execution.

Theorem 1. *Suppose that* $a \equiv z \mod 2$ *and* $\gamma \geq 0$. *If node* q *is located in* $((z-\gamma)(R-5\sigma|\texttt{PS}|)+\sigma|\texttt{PS}|, (z+\gamma)(R-5\sigma|\texttt{PS}|)+R-\sigma|\texttt{PS}|)$ *at the end of phase* $\pi_{a+3\gamma}$, *then, for some* $\varphi \in \{-\gamma, \ldots, \gamma\}$, $q \in \mathcal{S}_{a+3|\varphi|, z+\varphi}$.

Proof. Suppose that, for all $\varphi \in \{-\gamma, \ldots, \gamma\}$, $q \notin \mathcal{S}_{a+3|\varphi|, z+\varphi}$. Since $a \equiv z \mod 2$, it follows that $a + 3|\varphi| \equiv z + \varphi \mod 2$, so $q \notin \mathcal{S}_{a+3|\varphi|, z+\varphi}$ is equivalent to saying that there is a time during phase $\pi_{a+3|\varphi|}$ such that $q \notin \psi_{z+\varphi}$. Therefore, for each $\ell \in \{-\gamma, \ldots, \gamma\}$, there is a time during phase $\pi_{a+3\ell}$ such that q is not located in region $\psi_{z+\ell}$. Without loss of generality, assume that, at the beginning of phase π_a, q is located in or to the right of segment SEG_z. By Proposition 4, q is located in or to the right of region $\psi_{z+\gamma}$ at all times during phase $\pi_{a+3\gamma}$. Therefore, there exists a time during phase $\pi_{a+3\gamma}$ at which q is located at or to the right of $(z+\gamma)(R-5\sigma|\texttt{PS}|)+R$. By Observation 2 with $a' = a + 3\gamma$ and $k = 0$, at all times during phase $\pi_{a+3\gamma}$, q is located at or to the right of $(z+\gamma)(R-5\sigma|\texttt{PS}|)+R-\sigma|\texttt{PS}|$. Therefore, at the end of phase $\pi_{a+3\gamma}$, q is not located in $((z-\gamma)(R-5\sigma|\texttt{PS}|)+\sigma|\texttt{PS}|, (z+\gamma)(R-5\sigma|\texttt{PS}|)+R-\sigma|\texttt{PS}|)$. ☐

3.2.2 Message Propagation Bounds

Suppose that p is a survivor in phase π_a for region ψ_z, and that p transmits a message containing I during phase π_a. Note that, since $\mathcal{S}_{a,z} \neq \emptyset$, it follows that $a \equiv z \mod 2$. In this section, our goal is to show that, shortly thereafter, I gets re-transmitted by nodes in ψ_z and nearby regions. Essentially, we will show that I gets spread at a rate of one region per every three phases. More specifically, we show that:

(1) all other survivors in phase π_a for region ψ_z receive I during phase π_a,

(2) in phase π_{a+2}, if region ψ_z has at least one survivor, I will be transmitted by at least one such survivor during phase π_{a+2}, and,

(3) in phase π_{a+3}, if region ψ_{z+1} has at least one survivor, I will be transmitted by at least one such survivor during phase π_{a+3}.

(4) in phase π_{a+3}, if region ψ_{z-1} has at least one survivor, I will be transmitted by at least one such survivor during phase π_{a+3}.

To prove (1), we determine which nearby regions might contain nodes that cause transmission collisions at nodes in ψ_z during phase π_a, we consider the set X of all nodes that are located in these regions, and then we bound $|X|$ from above. Then, we note that PS is a schedule that ensures that, for each node $p' \in X$, there is a slot in phase π_a such that p' transmits and all nodes in $X - \{p'\}$ listen. We can then conclude that, for each $p \in \mathcal{S}_{a,z}$, there is a slot during which p transmits and that p's transmission is received by all other nodes that are within radius R of p for the entire slot (which includes all of $\mathcal{S}_{a,z} - \{p\}$).

Proposition 6. *For any* $z' \in \mathbb{Z}$, *the distance between any point in region* $\psi_{z'}$ *and the leftmost point in region* $\psi_{z'+2\rho+2}$ *is greater than* R'.

Proof. By definition, every point in region $\psi_{z'}$ is to the left of $z'(R-5\sigma|\texttt{PS}|)+R$, and the leftmost point of region $\psi_{z'+2\rho+2}$ is $(z'+2\rho+2)(R-5\sigma|\texttt{PS}|)$. Thus, the distance between any point in region

$\psi_{z'}$ and the leftmost point of region $\psi_{z'+2\rho+2}$ is greater than

$$\begin{aligned}
& (z'+2\rho+2)(R-5\sigma|\mathtt{PS}|)-(z'(R-5\sigma|\mathtt{PS}|)+R) \\
= \; & (2\left\lceil \tfrac{R'}{R} \right\rceil + 2)(R-5\sigma|\mathtt{PS}|)-R \\
\geq \; & (2\tfrac{R'}{R}+1)R-(2\tfrac{R'}{R}+2)(5\sigma|\mathtt{PS}|) \\
\geq \; & 2R'+R-(2\tfrac{R'}{R}+2)\tfrac{R}{2} \text{ (by assumption (A2))} \\
= \; & R'
\end{aligned}$$

□

Proposition 7. *Suppose that there is a node q located in $\psi_{z-1} \cup \psi_z \cup \psi_{z+1}$ at all times during phase π_a. Further, suppose that, for some $z'' \in \mathbb{Z}$, there exist $p \in \mathcal{P}_{a,z}$ and $q' \in \mathcal{P}_{a,z''}$ such that a transmission collision occurs at node q due to simultaneous transmissions by nodes p and q'. Then $z'' = z + 2w$ for some $w \in \{-\rho-1, \ldots, \rho+1\}$.*

Proof. Since $\mathcal{P}_{a,z''}$ is non-empty, we must have $z'' \equiv z \mod 2$, so $z'' = z + 2w$ for some $w \in \mathbb{Z}$. If $w \geq \rho + 2$, then at all times in phase π_a up to and including its transmission, q' is located in region $\psi_{z''}$, which is to the right of the leftmost point of region $\psi_{z+2\rho+3}$. Also, by assumption, q is located in or to the left of region ψ_{z+1} at all times in phase π_a. Therefore, by Proposition 6 with $z' = z + 1$, the distance between q' and q is greater than R' during the transmission by q'. This contradicts the fact that the transmission by q' caused a collision at node q. Hence, $w \leq \rho + 1$. Similarly, $w \geq -\rho - 1$. □

Proposition 8. *Suppose that there is a node q located in $\psi_{z-1} \cup \psi_z \cup \psi_{z+1}$ at all times during phase π_a. Further, suppose that there exists $p \in \mathcal{S}_{a,z}$ such that the distance between p and q is at most R at all times during phase π_a. Then, q receives a transmission by p during phase π_a.*

Proof. By Proposition 7, the only nodes that can cause a collision at a node in $\psi_{z-1} \cup \psi_z \cup \psi_{z+1}$ during phase π_a are in $X = \bigcup_{w=-\rho-1}^{\rho+1} \mathcal{P}_{a,z+2w}$. Since there are $(3 + 2\rho)$ sets in this union, and, since the size of each region is R, each set has size at most $(\Delta + 1)$, it follows that $|X| \leq (3 + 2\rho)(\Delta + 1)$.

Now, suppose that all participating nodes for phase π_a execute PS. By Observation 1, for every $p \in \mathcal{S}_{a,z}$, there exists a slot $t_p \in \pi_a$ during which p transmits and all nodes in $X - \{p\}$ listen. In particular, during slot t_p, a transmission collision cannot happen at node q. Since the distance between p and q is at most R at all times during phase π_a, it follows that, during slot t_p, q will receive p's transmission. □

Finally, since all survivors for region ψ_z in phase π_a are located in ψ_z for the entire phase, Proposition 8 implies that they all receive a message from each other during phase π_a.

Lemma 10. *Every $q \in \mathcal{S}_{a,z}$ receives a message from each $p \in \mathcal{S}_{a,z} - \{q\}$ during phase π_a.*

Next, we prove (2).

Proposition 9. *Suppose that $p \in \mathcal{S}_{a,z}$ and p transmits I in phase π_a. If $\mathcal{S}_{a+2,z} \neq \emptyset$, then there exists a node in $\mathcal{S}_{a+2,z}$ that transmits I during phase π_{a+2}.*

Proof. If $p \in \mathcal{S}_{a+2,z}$, then p will transmit I in phase π_{a+2}. So, in what follows, we assume that $p \notin \mathcal{S}_{a+2,z}$. Therefore, at some time during phase π_{a+2}, p is located either to the left of region ψ_z or to the right of ψ_z. Without loss of generality, we assume that, at some time during phase

π_{a+2}, p is located to the right of ψ_z, i.e., in $[z(R-5\sigma|\texttt{PS}|)+R,\infty]$. By Observation 2 with $a'=a+2$ and $k=1$, it follows that, at the beginning of phase π_{a+1}, p is located at or to the right of $z(R-5\sigma|\texttt{PS}|)+R-2\sigma|\texttt{PS}|$. Moreover, since $p\in\mathcal{S}_{a,z}$, p is located to the left of $z(R-5\sigma|\texttt{PS}|)+R$ at the end of phase π_a. Hence, at the beginning of phase π_{a+1}, p is located in $[z(R-5\sigma|\texttt{PS}|)+R-2\sigma|\texttt{PS}|, z(R-5\sigma|\texttt{PS}|)+R)$.

First, suppose that, at the end of phase π_a, there exists a node p' in $[z(R-5\sigma|\texttt{PS}|)+2\sigma|\texttt{PS}|, z(R-5\sigma|\texttt{PS}|)+R-2\sigma|\texttt{PS}|)$. In this case, we can show that p' received I during phase π_a and will be a survivor for region ψ_z in phase π_{a+2}. By Observation 2 with $a'=a$ and $k=0$, at all times during phase π_a, p' is located in $[z(R-5\sigma|\texttt{PS}|)+\sigma|\texttt{PS}|, z(R-5\sigma|\texttt{PS}|)+R-\sigma|\texttt{PS}|)\subseteq[z(R-5\sigma|\texttt{PS}|), z(R-5\sigma|\texttt{PS}|)+R)=\psi_z$. Namely, $p'\in\mathcal{S}_{a,z}$. By Lemma 10, p' receives I during phase π_a. So, p' will include I in any transmissions in phases after π_a. Also, by Observation 2 with $a'=a+1$ and $k=1$, at all times during phase π_{a+2}, p' is located in $[(z(R-5\sigma|\texttt{PS}|)+2\sigma|\texttt{PS}|)-2\sigma|\texttt{PS}|, (z(R-5\sigma|\texttt{PS}|)+R-2\sigma|\texttt{PS}|)+2\sigma|\texttt{PS}|)=[z(R-5\sigma|\texttt{PS}|), z(R-5\sigma|\texttt{PS}|)+R)=\psi_z$. Namely, $p'\in\mathcal{S}_{a+2,z}$. This proves that there exists a node in $\mathcal{S}_{a+2,z}$ that transmits I during phase π_{a+2}.

So, in what follows, we assume that there is no node located in $[z(R-5\sigma|\texttt{PS}|)+2\sigma|\texttt{PS}|, z(R-5\sigma|\texttt{PS}|)+R-2\sigma|\texttt{PS}|)$ at the end of phase π_a. Note that the length of the interval $[z(R-5\sigma|\texttt{PS}|)+2\sigma|\texttt{PS}|, z(R-5\sigma|\texttt{PS}|)+R-2\sigma|\texttt{PS}|]$ is $R-4\sigma|\texttt{PS}|$. By assumption (A2), it follows that $R-4\sigma|\texttt{PS}|\geq(\frac{R}{2}+5\sigma|\texttt{PS}|)-4\sigma|\texttt{PS}|>\frac{R}{2}$. Then, assumption (A1) implies that this interval does not lie between the leftmost and rightmost nodes. Since we assumed that p is located to the right of this range, it follows that there is no node to the left of the point $z(R-5\sigma|\texttt{PS}|)+R-2\sigma|\texttt{PS}|$. Namely, at the beginning of phase π_{a+1}, each node in the network is located at or to the right of the point $z(R-5\sigma|\texttt{PS}|)+R-2\sigma|\texttt{PS}|$. By Observation 2 with $a'=a+1$ and $k=0$, at all times during phase π_{a+1}, each node in the network is located at or to the right of the point $z(R-5\sigma|\texttt{PS}|)+R-3\sigma|\texttt{PS}|$, a fact that we denote by (*).

It suffices to show that p is a survivor for region ψ_{z+1} in phase π_{a+1}, and that every node in $\mathcal{S}_{a+2,z}$ receives I from p during phase π_{a+1}. This would imply that, if $\mathcal{S}_{a+2,z}\neq\emptyset$, then all nodes in $\mathcal{S}_{a+2,z}$ transmit I during phase π_{a+2}.

By Observation 2 with $a'=a+1$ and $k=0$, at all times during phase π_{a+1}, p is located in $[z(R-5\sigma|\texttt{PS}|)+R-3\sigma|\texttt{PS}|, z(R-5\sigma|\texttt{PS}|)+R+\sigma|\texttt{PS}|)$. We can conclude that, at all times during phase π_{a+1}, p must be located in

$$\begin{aligned}
& [z(R-5\sigma|\texttt{PS}|)+R-3\sigma|\texttt{PS}|, z(R-5\sigma|\texttt{PS}|)+R+\sigma|\texttt{PS}|) \\
\subseteq\ & [z(R-5\sigma|\texttt{PS}|)+(R-5\sigma|\texttt{PS}|), z(R-5\sigma|\texttt{PS}|)+R+(R-5\sigma|\texttt{PS}|)) \text{ (by assumption (A2))} \\
=\ & [(z+1)(R-5\sigma|\texttt{PS}|), (z+1)(R-5\sigma|\texttt{PS}|)+R) \\
=\ & \psi_{z+1}
\end{aligned}$$

So, $p\in\mathcal{S}_{a+1,z+1}$. Therefore, by Lemma 10, all nodes in $\mathcal{S}_{a+1,z+1}$ receive I during phase π_{a+1}.

Finally, consider any survivor $q\in\mathcal{S}_{a+2,z}$. By the definition of $\mathcal{S}_{a+2,z}$, q must be located in $[z(R-5\sigma|\texttt{PS}|), z(R-5\sigma|\texttt{PS}|)+R)$ at the end of phase π_{a+1}. By Observation 2 with $a'=a+1$ and $k=0$, at all times during phase π_{a+1}, q is located in $[z(R-5\sigma|\texttt{PS}|)-\sigma|\texttt{PS}|, z(R-5\sigma|\texttt{PS}|)+R+\sigma|\texttt{PS}|)$. Thus, by (*), at all times during phase π_{a+1}, q is located in $[z(R-5\sigma|\texttt{PS}|)+R-3\sigma|\texttt{PS}|, z(R-5\sigma|\texttt{PS}|)+R+\sigma|\texttt{PS}|)\subseteq\psi_{z+1}$. Thus, $q\in\mathcal{S}_{a+1,z+1}$. □

Next, we set out to prove (3). At a high level, the proof proceeds by considering three disjoint cases:

1. Suppose that, at all times during phase π_a, there is never a node to the right of region ψ_z. If there are any survivors for region ψ_{z+1} in either phase π_{a+1} or π_{a+3}, then they will include I in their transmissions. (Proposition 10)

2. Suppose that the rightmost node of the network is to the right of region ψ_z at some time during phase π_a, and that, at some time during phase π_a, p is not located near the left border of region ψ_z. We show that there exists a survivor for region ψ_{z+1} in phase π_{a+1} that re-transmits I during phase π_{a+1} (Proposition 14). To prove this last statement, we suppose that p is not a survivor for region ψ_{z+1} in phase π_{a+1} that re-transmits I during phase π_{a+1} and proceed as follows:

   (a) p is not close to the right border of region ψ_z during phase π_a. (Proposition 11)

   (b) Since p is not close to the right border of region ψ_z and there exists a node to the right of region ψ_z, we apply the density assumption to prove that there is a node q that is in ψ_{z+1} at the end of phase π_a. (Proposition 12)

   (c) Such a node q received I during phase π_a and transmits I as a survivor for region ψ_{z+1} in phase π_{a+1}. (Proposition 13)

   Then, by Proposition 9, $\mathcal{S}_{a+3,z+1} \neq \emptyset$ implies that there exists a node in $\mathcal{S}_{a+3,z+1}$ that transmits I during phase π_{a+3}.

3. Suppose that the rightmost node of the network is to the right of region ψ_z at some time during π_a, and, at all times during phase π_a, p is located near the left border of region ψ_z. We start by showing that there must be some node q'' that is not near the left border of region ψ_z at the end of phase π_a which received I during phase π_a. (Proposition 15). If q'' is a survivor for region ψ_{z+1} in phase π_{a+1}, then, by Proposition 9, $\mathcal{S}_{a+3,z+1} \neq \emptyset$ implies that there exists a node in $\mathcal{S}_{a+3,z+1}$ which transmits I during phase π_{a+3}. Otherwise, we prove:

   (a) Node q'' is not near either border of region ψ_z during phase π_{a+1}. (Proposition 16)

   (b) Node q'' is a survivor for region ψ_z during phase π_{a+2} and q'' is not near the left border of region ψ_z at the beginning of phase π_{a+2}. Then, depending on whether or not there is a node located to the right of region ψ_z at some point during phase π_{a+2}, we essentially repeat the argument from cases 1 or 2 above with q'' instead of p and phase π_{a+2} instead of phase π_a. More specifically, if at all times during phase π_{a+2}, there is no node to the right of ψ_z, then, by Proposition 10, any survivor for region ψ_{z+1} in phase π_{a+3} will include I in its transmission. Otherwise, at some time during phase π_{a+2}, there is a node to the right of region ψ_z, and, at the beginning of phase π_{a+2}, q'' is not located near the left border of region ψ_z. Then, by applying Proposition 14, it follows that there exists a survivor for region ψ_{z+1} in phase π_{a+3} that re-transmits I during phase π_{a+3}. (Proposition 17)

Finally, combining these three cases gives us Proposition 19.

We now fill in the proof details. First, consider what happens at the right edge of the network. Namely, consider the case where, at all times during phase π_a, no node is located to the right of region ψ_z. In this case, using the upper bound on node speed, we will show that any survivors for region ψ_{z+1} in either phase π_{a+1} or π_{a+3} are also survivors for region ψ_z in phase π_a. Namely, these nodes received I during phase π_a, and, hence, will transmit I as part of all of their future transmissions.

Proposition 10. *Suppose that $p \in \mathcal{S}_{a,z}$ and that p transmits I during phase π_a. Further, suppose that, at all times during phase π_a, no node is located to the right of region ψ_z. Then, for each $\ell \in \{1,3\}$, each node in $\mathcal{S}_{a+\ell,z+1}$ will transmit I during phase $\pi_{a+\ell}$.*

Proof. For any $\ell \in \{1,3\}$, consider any node $q \in \mathcal{S}_{a+\ell,z+1}$. By definition, at all times during phase $\pi_{a+\ell}$, q is located in region ψ_{z+1}, so at or to the right of the point $(z+1)(R-5\sigma|\mathtt{PS}|)$. By Observation 2 with $a' = a+\ell$ and $k = 3$, at all times from the beginning of phase $\pi_{a+\ell-3}$ until the end of phase $\pi_{a+\ell+3}$, q is located at or to the right of the point $(z+1)(R-5\sigma|\mathtt{PS}|) - 4\sigma|\mathtt{PS}| = z(R-5\sigma|\mathtt{PS}|) + R - 9\sigma|\mathtt{PS}|$. By assumption (A2), $R \geq 10\sigma|\mathtt{PS}|$, so $z(R-5\sigma|\mathtt{PS}|) + R - 9\sigma|\mathtt{PS}| > z(R-5\sigma|\mathtt{PS}|)$. Therefore, at all times from the beginning of phase $\pi_{a+\ell-3}$ until the end of phase $\pi_{a+\ell+3}$, q is not located to the left of region ψ_z. Since $0 < \ell \leq 3$, $\pi_a \subseteq \pi_{a+\ell-3} \cup \cdots \cup \pi_{a+\ell+3}$. Thus, at all times during phase π_a, q is not located to the left of region ψ_z. By assumption, at all times during phase π_a, q is not located to the right of region ψ_z. Hence, $q \in \mathcal{S}_{a,z}$. By Lemma 10, q receives I during phase π_a. Since $q \in \mathcal{S}_{a+\ell,z+1}$, q transmits I during phase $\pi_{a+\ell}$ as a survivor for region ψ_{z+1}.

□

So, suppose that there is a time during phase π_a at which there is a node located to the right of region ψ_z. Since $p \in \mathcal{S}_{a,z}$, p is located in region ψ_z at all times during phase π_a. We consider two cases: either p is located near the left border of ψ_z for the entire phase, or, at some time during the phase, p is not located near the left border. We first consider the latter of these two cases. We will show that there exists a survivor for region ψ_{z+1} in phase π_{a+1} that re-transmits I during phase π_{a+1}. Then, by Proposition 9, $\mathcal{S}_{a+3,z+1} \neq \emptyset$ implies that there exists a node in $\mathcal{S}_{a+3,z+1}$ that transmits I during phase π_{a+3}.

If $p \in \mathcal{S}_{a+1,z+1}$, then p itself is a survivor for region ψ_{z+1} in phase π_{a+1} that re-transmits I during phase π_{a+1}. So, suppose that $p \notin \mathcal{S}_{a+1,z+1}$. The next result gives a constraint on the location of p at the beginning of phase π_{a+1}.

Proposition 11. *Suppose that, at some time during phase π_a, p is located in $[z(R-5\sigma|\mathtt{PS}|) + \frac{R}{2} - 2\sigma|\mathtt{PS}|, z(R-5\sigma|\mathtt{PS}|) + R)$, and $p \notin \mathcal{S}_{a+1,z+1}$. At the beginning of phase π_{a+1}, p is located to the left of $(z+1)(R-5\sigma|\mathtt{PS}|) + \sigma|\mathtt{PS}|$.*

Proof. Since there exists a time during phase π_a at which p is located to the left of $z(R-5\sigma|\mathtt{PS}|)+R$, Observation 2 with $a' = a$ and $k = 0$ implies that, at all times during phase π_a, p is located to the left of the point

$$\begin{aligned} & z(R-5\sigma|\mathtt{PS}|) + R + \sigma|\mathtt{PS}| \\ = \ & (z+1)(R-5\sigma|\mathtt{PS}|) + 6\sigma|\mathtt{PS}| \\ < \ & (z+1)(R-5\sigma|\mathtt{PS}|) + 9\sigma|\mathtt{PS}| \\ \leq \ & (z+1)(R-5\sigma|\mathtt{PS}|) + R - \sigma|\mathtt{PS}| \text{ (by assumption (A2))}. \end{aligned}$$

By Proposition 2 with $a' = a+1$ and $z' = z+1$, it follows that, at the beginning of phase π_{a+1}, p is located to the left of the point $(z+1)(R-5\sigma|\mathtt{PS}|) + \sigma|\mathtt{PS}|$.

□

Next, we prove that there is some node q located in the range $[(z+1)(R-5\sigma|\mathtt{PS}|) + \sigma|\mathtt{PS}|, (z+1)(R-5\sigma|\mathtt{PS}|) + \sigma|\mathtt{PS}| + \frac{R}{2}]$ at the end of phase π_a. This does not immediately follow from the density assumption, since the assumption only holds between the leftmost and rightmost nodes of the network. So, using p's location and the fact that there is a node located to the right of region

ψ_z at some point during phase π_a, we can prove that there is at least one node to the left and one node to the right of the desired range at the end of phase π_a.

Proposition 12. *Suppose that there is a time during phase π_a at which there is a node p' located to the right of region ψ_z, and, at the beginning of phase π_{a+1}, p is located to the left of $(z+1)(R-5\sigma|\mathtt{PS}|)+\sigma|\mathtt{PS}|$. Then, there exists a node q (possibly p' itself) in $[(z+1)(R-5\sigma|\mathtt{PS}|)+\sigma|\mathtt{PS}|,(z+1)(R-5\sigma|\mathtt{PS}|)+\sigma|\mathtt{PS}|+\frac{R}{2}]$ at the end of phase π_a.*

Proof. Consider any node p' that is located to the right of region ψ_z, i.e., at or to the right of the point $z(R-5\sigma|\mathtt{PS}|)+R=(z+1)(R-5\sigma|\mathtt{PS}|)+5\sigma|\mathtt{PS}|$, at some time during phase π_a. By Observation 2 with $a'=a$ and $k=0$, at all times during phase π_a, p' is located at or to the right of point $(z+1)(R-5\sigma|\mathtt{PS}|)+4\sigma|\mathtt{PS}|$. If p' is in $[(z+1)(R-5\sigma|\mathtt{PS}|)+\sigma|\mathtt{PS}|,(z+1)(R-5\sigma|\mathtt{PS}|)+\sigma|\mathtt{PS}|+\frac{R}{2}]$ at the end of phase π_a, we are done. Otherwise, p' is located to the right of the point $(z+1)(R-5\sigma|\mathtt{PS}|)+\sigma|\mathtt{PS}|+\frac{R}{2}$. Since p is located to the left of the point $(z+1)(R-5\sigma|\mathtt{PS}|)+\sigma|\mathtt{PS}|$ at the end of phase π_a, then, by assumption (A1), there exists a node in the range $[(z+1)(R-5\sigma|\mathtt{PS}|)+\sigma|\mathtt{PS}|,(z+1)(R-5\sigma|\mathtt{PS}|)+\sigma|\mathtt{PS}|+\frac{R}{2}]$ at the end of phase π_a. □

Finally, we can show that q receives I during phase π_a and will transmit it in phase π_{a+1} as a survivor for region ψ_{z+1}.

Proposition 13. *Suppose that $p\in\mathcal{S}_{a,z}$, and, at the end of phase π_a, p is located to the left of $(z+1)(R-5\sigma|\mathtt{PS}|)+\sigma|\mathtt{PS}|$. Also, suppose that there is a time during phase π_a at which p is located at or to the right of $z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-2\sigma|\mathtt{PS}|$. Further, suppose that at the end of phase π_a, there exists a node q in $[(z+1)(R-5\sigma|\mathtt{PS}|)+\sigma|\mathtt{PS}|,(z+1)(R-5\sigma|\mathtt{PS}|)+\sigma|\mathtt{PS}|+\frac{R}{2}]$. Then $q\in\mathcal{S}_{a+1,z+1}$ and q transmits I during phase π_{a+1}.*

Proof. By Observation 2 with $a'=a$ and $k=0$, at all times during phase π_a, p is located in $\mathcal{P}=[z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-3\sigma|\mathtt{PS}|,(z+1)(R-5\sigma|\mathtt{PS}|)+2\sigma|\mathtt{PS}|)$. By Observation 2 with $a'=a$ and $k=0$, at all times during phases π_a and π_{a+1}, q is located in $\mathcal{Q}=[(z+1)(R-5\sigma|\mathtt{PS}|),(z+1)(R-5\sigma|\mathtt{PS}|)+2\sigma|\mathtt{PS}|+\frac{R}{2}]$.

Every point in $\mathcal{P}$ is located to the left of the rightmost point in $\mathcal{Q}$, and, the leftmost point in $\mathcal{Q}$ is to the right of the leftmost point in $\mathcal{P}$. This is because $(z+1)(R-5\sigma|\mathtt{PS}|)+2\sigma|\mathtt{PS}|<(z+1)(R-5\sigma|\mathtt{PS}|)+2\sigma|\mathtt{PS}|+\frac{R}{2}$ and $(z+1)(R-5\sigma|\mathtt{PS}|)\geq z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}>z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-3\sigma|\mathtt{PS}|$, by assumption (A2). Therefore, at all times during phase π_a, the distance between p and q is at most

$$\begin{aligned}
&\left((z+1)(R-5\sigma|\mathtt{PS}|)+2\sigma|\mathtt{PS}|+\tfrac{R}{2}\right)-\left(z(R-5\sigma|\mathtt{PS}|)+\tfrac{R}{2}-3\sigma|\mathtt{PS}|\right)\\
=\;&\left(z(R-5\sigma|\mathtt{PS}|)-3\sigma|\mathtt{PS}|+\tfrac{3R}{2}\right)-\left(z(R-5\sigma|\mathtt{PS}|)+\tfrac{R}{2}-3\sigma|\mathtt{PS}|\right)\\
=\;&R
\end{aligned}$$

Further, by assumption (A2), $R\geq 10\sigma|\mathtt{PS}|$, so $2\sigma|\mathtt{PS}|+\frac{R}{2}<R$. It follows that $\mathcal{Q}\subseteq[(z+1)(R-5\sigma|\mathtt{PS}|),(z+1)(R-5\sigma|\mathtt{PS}|)+R)=\psi_{z+1}$. So, by Proposition 8, q receives a transmission by p during phase π_a. In particular, q receives I during phase π_a. Furthermore, since q is located in $\mathcal{Q}$ at all times during phase π_{a+1}, $q\in\mathcal{S}_{a+1,z+1}$. □

Propositions 11, 12 and 13 can be combined to give us the following useful fact.

Proposition 14. *Suppose that $p\in\mathcal{S}_{a,z}$ and $p\notin\mathcal{S}_{a+1,z+1}$. Further, suppose that there is a time during phase π_a at which there is a node located to the right of region ψ_z. If, at some time during* *phase π_a, p is located in $[z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-2\sigma|\mathtt{PS}|, z(R-5\sigma|\mathtt{PS}|)+R)$, then there exists a node $q \in \mathcal{S}_{a+1,z+1}$ that transmits I during phase π_{a+1}.*

Assume that, at some time during phase π_a, p is not located near the left border of region ψ_z, i.e. it is located in $[z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-2\sigma|\mathtt{PS}|, z(R-5\sigma|\mathtt{PS}|)+R)$. Since $p \in \mathcal{S}_{a,z}$, $p \notin \mathcal{S}_{a+1,z+1}$, and there exists a time during phase π_a at which there is a node located to the right of region ψ_z, Proposition 14 implies that there exists a node $q \in \mathcal{S}_{a+1,z+1}$ that transmits I during phase π_{a+1}. Then, by Proposition 9, $\mathcal{S}_{a+3,z+1} \neq \emptyset$ implies that there exists a node in $\mathcal{S}_{a+3,z+1}$ that transmits I during phase π_{a+3}.

Finally, we complete the proof of (3) by assuming that, at all times during phase π_a, p is located near the left border of region ψ_z, i.e., it is located in $[z(R-5\sigma|\mathtt{PS}|), z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-2\sigma|\mathtt{PS}|)$. We begin by proving an analogue of Proposition 12. As in the case of Proposition 12, we have to make sure that the desired range lies between the leftmost and rightmost nodes of the network.

Proposition 15. *Suppose that $p \in \mathcal{S}_{a,z}$ and $p \notin \mathcal{S}_{a+1,z+1}$. Further, suppose that there is a time during phase π_a at which there is a node located to the right of region ψ_z, and, at all times during phase π_a, p is located in $[z(R-5\sigma|\mathtt{PS}|), z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-2\sigma|\mathtt{PS}|)$. Then, there exists a node q'' that receives I in phase π_a and is located in $[z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-\sigma|\mathtt{PS}|, z(R-5\sigma|\mathtt{PS}|)+R-\sigma|\mathtt{PS}|)$ at the end of phase π_a.*

Proof. Consider any node p' that is located to the right of region ψ_z, i.e., at or to the right of the point $z(R-5\sigma|\mathtt{PS}|)+R > z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}$, at some time during phase π_a. By Observation 2 with $a'=a$ and $k=0$, at all times during phase π_a, p' is located to the right of the point $z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-\sigma|\mathtt{PS}|$. If p' is located in $[z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-\sigma|\mathtt{PS}|, z(R-5\sigma|\mathtt{PS}|)+R-\sigma|\mathtt{PS}|)$ at the end of phase π_a, we are done. Otherwise, p' is located at or to the right of the point $z(R-5\sigma|\mathtt{PS}|)+R-\sigma|\mathtt{PS}|$ at the end of phase π_a, and, by assumption, p is located to the left of $z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-2\sigma|\mathtt{PS}| < z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-\sigma|\mathtt{PS}|$. Thus, by assumption (A1), at the end of phase π_a, there is a node q'' located in $[z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-\sigma|\mathtt{PS}|, z(R-5\sigma|\mathtt{PS}|)+R-\sigma|\mathtt{PS}|)$.

To prove that q'' receives I during phase π_a, we apply Observation 2 with $a'=a$ and $k=0$ to conclude that, at all times during phase π_a, q'' is located in $[z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-2\sigma|\mathtt{PS}|, z(R-5\sigma|\mathtt{PS}|)+R)$. By assumption (A2), $[z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-2\sigma|\mathtt{PS}|, z(R-5\sigma|\mathtt{PS}|)+R) \subseteq [z(R-5\sigma|\mathtt{PS}|), z(R-5\sigma|\mathtt{PS}|)+R)$, so $q'' \in \mathcal{S}_{a,z}$. By Lemma 10, q'' receives I from p during phase π_a. ☐

Since there is a time during phase π_a at which there is a node located to the right of region ψ_z, Proposition 15 implies that there exists a node q'' that received I during phase π_a and is located in $[z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-\sigma|\mathtt{PS}|, z(R-5\sigma|\mathtt{PS}|)+R-\sigma|\mathtt{PS}|)$ at the end of phase π_a. First, assume that $q'' \notin \mathcal{S}_{a+1,z+1}$. Then, we can prove the following restriction on the location of q'' during phase π_{a+1}.

Proposition 16. *At all times during phase π_{a+1}, q'' is located in $[z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-2\sigma|\mathtt{PS}|, z(R-5\sigma|\mathtt{PS}|)+R-4\sigma|\mathtt{PS}|)$.*

Proof. On one hand, by Observation 2 with $a'=a+1$ and $k=0$, at all times during phase π_{a+1}, q'' is located in $[z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-2\sigma|\mathtt{PS}|, z(R-5\sigma|\mathtt{PS}|)+R)$. By assumption (A2), $R \geq 10\sigma|\mathtt{PS}|$, so $2R-6\sigma|\mathtt{PS}| > R$. This implies that $[z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-2\sigma|\mathtt{PS}|, z(R-5\sigma|\mathtt{PS}|)+R) \subseteq [z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-2\sigma|\mathtt{PS}|, (z+1)(R-5\sigma|\mathtt{PS}|)+R-\sigma|\mathtt{PS}|)$.

On the other hand, by Proposition 2 with $a'=a+1$ and $z'=z+1$, at all times during phase π_{a+1}, q'' is not located in $[(z+1)(R-5\sigma|\mathtt{PS}|)+\sigma|\mathtt{PS}|, (z+1)(R-5\sigma|\mathtt{PS}|)+R-\sigma|\mathtt{PS}|] = [z(R-5\sigma|\mathtt{PS}|)+R-4\sigma|\mathtt{PS}|, (z+1)(R-5\sigma|\mathtt{PS}|)+R-\sigma|\mathtt{PS}|]$.

Thus, at all times during phase π_{a+1}, q'' is located in $[z(R-5\sigma|\texttt{PS}|)+\frac{R}{2}-2\sigma|\texttt{PS}|, z(R-5\sigma|\texttt{PS}|)+R-4\sigma|\texttt{PS}|)$. ☐

We now show that, if $\mathcal{S}_{a+3,z+1} \neq \emptyset$, then there exists a node in $\mathcal{S}_{a+3,z+1}$ that transmits I during phase π_{a+3}. We use Proposition 16 to conclude that q'' is a survivor for region ψ_z for phase π_{a+2} and that, at the beginning of phase π_{a+2}, q'' is not located near the left border of region ψ_z. Then, we either apply Proposition 10 or Proposition 14, depending on whether or not there is a node located to the right of region ψ_z at some point during phase π_{a+2}.

Proposition 17. *Suppose that, at the beginning of phase π_{a+2}, node q'' is located in $[z(R-5\sigma|\texttt{PS}|)+\frac{R}{2}-2\sigma|\texttt{PS}|, z(R-5\sigma|\texttt{PS}|)+R-4\sigma|\texttt{PS}|)$, and q'' receives I during phase π_a. If $\mathcal{S}_{a+3,z+1} \neq \emptyset$, then there exists a node in $\mathcal{S}_{a+3,z+1}$ that transmits I during phase π_{a+3}.*

Proof. By Observation 2 with $a' = a+2$ and $k = 0$, at all times during phase π_{a+2}, q'' is located in

$$\begin{aligned} & [z(R-5\sigma|\texttt{PS}|)+\tfrac{R}{2}-3\sigma|\texttt{PS}|, z(R-5\sigma|\texttt{PS}|)+R-3\sigma|\texttt{PS}|) \\ \subseteq\ & [z(R-5\sigma|\texttt{PS}|)+2\sigma|\texttt{PS}|, z(R-5\sigma|\texttt{PS}|)+R-3\sigma|\texttt{PS}|) \text{ (by assumption (A2))} \\ \subseteq\ & [z(R-5\sigma|\texttt{PS}|), z(R-5\sigma|\texttt{PS}|)+R) \\ =\ & \psi_z \end{aligned}$$

Therefore, $q'' \in \mathcal{S}_{a+2,z}$. Since q'' receives I during phase π_a, q'' will transmit I during phase π_{a+2}. If $q'' \in \mathcal{S}_{a+3,z+1}$, we are done. So, suppose that $q'' \notin \mathcal{S}_{a+3,z+1}$.

If, at all times during phase π_{a+2}, no node is located to the right of region ψ_z, then, by Proposition 10 with p replaced by q'', a replaced by $a+2$ and $\ell = 1$, all nodes in $\mathcal{S}_{a+3,z+1}$ transmit I during phase π_{a+3}.

Otherwise, there exists a time during phase π_{a+2} at which a node is located to the right of region ψ_z. At the beginning of phase π_{a+2}, q'' is located in $[z(R-5\sigma|\texttt{PS}|)+\frac{R}{2}-2\sigma|\texttt{PS}|, z(R-5\sigma|\texttt{PS}|)+R-4\sigma|\texttt{PS}|) \subseteq [z(R-5\sigma|\texttt{PS}|)+\frac{R}{2}-2\sigma|\texttt{PS}|, z(R-5\sigma|\texttt{PS}|)+R)$. By Proposition 14 with p replaced by q'' and a replaced by $a+2$, there exists a node in $\mathcal{S}_{a+3,z+1}$ that transmits I during phase π_{a+3}. ☐

If $q'' \in \mathcal{S}_{a+1,z+1}$, then q'' transmits I in phase π_{a+1}, and, by Proposition 9, $\mathcal{S}_{a+3,z+1} \neq \emptyset$ implies that there is a node in $\mathcal{S}_{a+3,z+1}$ that transmits I during phase π_{a+3}, as desired. If $q'' \notin \mathcal{S}_{a+1,z+1}$, then, by Proposition 16, at all times during phase π_{a+1}, q'' is located in $[z(R-5\sigma|\texttt{PS}|)+\frac{R}{2}-2\sigma|\texttt{PS}|, z(R-5\sigma|\texttt{PS}|)+R-4\sigma|\texttt{PS}|)$. By Proposition 17, if $\mathcal{S}_{a+3,z+1} \neq \emptyset$, then there exists a node in $\mathcal{S}_{a+3,z+1}$ that transmits I during phase π_{a+3}. This gives us the following useful result.

Proposition 18. *Suppose that $p \in \mathcal{S}_{a,z}$ and $p \notin \mathcal{S}_{a+1,z+1}$. Further, suppose that there is a time during phase π_a at which there is a node located to the right of region ψ_z, and, at all times during phase π_a, p is located in $[z(R-5\sigma|\texttt{PS}|), z(R-5\sigma|\texttt{PS}|)+\frac{R}{2}-2\sigma|\texttt{PS}|)$. If $\mathcal{S}_{a+3,z+1} \neq \emptyset$, then there exists a node in $\mathcal{S}_{a+3,z+1}$ that transmits I during phase π_{a+3}.*

We are now ready to prove (3).

Proposition 19. *Suppose that $p \in \mathcal{S}_{a,z}$ and p transmits I during phase π_a. If $\mathcal{S}_{a+3,z+1} \neq \emptyset$, then there exists a node in $\mathcal{S}_{a+3,z+1}$ that transmits I during phase π_{a+3}.*

Proof. First, suppose that $p \in \mathcal{S}_{a+1,z+1}$. Then p is a survivor for region ψ_{z+1} in phase π_{a+1} that re-transmits I during phase π_{a+1}. So, by Proposition 9 with a replaced with $a+1$ and z replaced with $z+1$, if $\mathcal{S}_{a+3,z+1} \neq \emptyset$, then there exists a node in $\mathcal{S}_{a+3,z+1}$ that transmits I during phase π_{a+3}. So, in what follows, we assume that $p \notin \mathcal{S}_{a+1,z+1}$.

Next, suppose that, at all times during phase π_a, no node is located to the right of region ψ_z. Then, by Proposition 10 with $\ell = 3$, each node in $\mathcal{S}_{a+3,z+1}$ will transmit I during phase π_{a+3}. So, in what follows, we assume that there is a time during phase π_a at which there is a node located to the right of region ψ_z.

Next, suppose that, at some time during phase π_a, p is located in $[z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-2\sigma|\mathtt{PS}|, z(R-5\sigma|\mathtt{PS}|)+R)$. By Proposition 14, there exists a node $q \in \mathcal{S}_{a+1,z+1}$ that transmits I during phase π_{a+1}. So, by Proposition 9 with a replaced with $a+1$ and z replaced with $z+1$, if $\mathcal{S}_{a+3,z+1} \neq \emptyset$, then there exists a node in $\mathcal{S}_{a+3,z+1}$ that transmits I during phase π_{a+3}. So, we assume that, at all times during phase π_a, p is located in it is located in $[z(R-5\sigma|\mathtt{PS}|), z(R-5\sigma|\mathtt{PS}|)+\frac{R}{2}-2\sigma|\mathtt{PS}|)$. Finally, by Proposition 18, if $\mathcal{S}_{a+3,z+1} \neq \emptyset$, then there exists a node in $\mathcal{S}_{a+3,z+1}$ that transmits I during phase π_{a+3}. □

By symmetry, the same result holds for $\mathcal{S}_{a+3,z-1}$, which proves (4).

Proposition 20. *Suppose that $p \in \mathcal{S}_{a,z}$ and p transmits I during phase π_a. If $\mathcal{S}_{a+3,z-1} \neq \emptyset$, then there exists a node in $\mathcal{S}_{a+3,z-1}$ that transmits I during phase π_{a+3}.*

Using induction, we prove that this propagation of I continues in both directions.

Theorem 2. *Suppose that $p \in \mathcal{S}_{a,z}$ and p transmits I during phase π_a. For every $\ell \in \mathbb{Z}$, if $\mathcal{S}_{a+3|\ell|,z+\ell} \neq \emptyset$, then there exists a node in $\mathcal{S}_{a+3|\ell|,z+\ell}$ that transmits I during phase $\pi_{a+3|\ell|}$.*

Proof. Without loss of generality, we prove the result for $\ell \geq 0$. We proceed by induction on ℓ. If $\ell = 0$, then p is a node in $\mathcal{S}_{a+3\ell,z+\ell}$ that transmits I during phase $\pi_{a+3\ell}$. The case $\ell = 1$ is proven in Proposition 19. So, let $\ell \geq 1$. As induction hypothesis, assume that, if $\mathcal{S}_{a+3\ell,z+\ell} \neq \emptyset$, then there exists a node in $\mathcal{S}_{a+3\ell,z+\ell}$ that transmits I during phase $\pi_{a+3\ell}$.

Suppose that $\mathcal{S}_{a+3(\ell+1),z+\ell+1} \neq \emptyset$. If $\mathcal{S}_{a+3\ell,z+\ell} \neq \emptyset$, then, by the induction hypothesis, it follows that a node in $\mathcal{S}_{a+3\ell,z+\ell}$ transmits I during phase $\pi_{a+3\ell}$. Therefore, by Proposition 19 with a replaced with $a+3\ell$ and z replaced with $z+\ell$, there exists a node in $\mathcal{S}_{a+3\ell+3,z+\ell+1}$ that transmits I during phase $\pi_{a+3\ell+3}$. Thus, to prove the induction step, it suffices to show that $\mathcal{S}_{a+3\ell,z+\ell} \neq \emptyset$.

Let $q \in \mathcal{S}_{a+3(\ell+1),z+\ell+1}$. By Proposition 3 with a replaced with $a+3\ell+3$ and z replaced with $z+\ell+1$, from the beginning of phase $\pi_{a+3\ell-1}$ until the end of phase $\pi_{a+3\ell+7}$, q is located in $\psi_{z+\ell} \cup \psi_{z+\ell+1} \cup \psi_{z+\ell+2}$. In particular, at all times during phase $\pi_{a+3\ell}$, q is located in $\psi_{z+\ell} \cup \psi_{z+\ell+1} \cup \psi_{z+\ell+2}$.

If $q \in \mathcal{S}_{a+3\ell,z+\ell}$, we are done. So, assume that, without loss of generality, at some time during phase $\pi_{a+3\ell}$, q is located to the right of region $\psi_{z+\ell}$. We will show that, at the beginning of phase $\pi_{a+3\ell}$, there is some node in $[(z+\ell)(R-5\sigma|\mathtt{PS}|)+R-\sigma|\mathtt{PS}|-\frac{R}{2}, (z+\ell)(R-5\sigma|\mathtt{PS}|)+R-\sigma|\mathtt{PS}|]$.

We consider p's location at the beginning of phase $\pi_{a+3\ell}$. Since $p \in \mathcal{S}_{a,z}$, p is located in $[z(R-5\sigma|\mathtt{PS}|), z(R-5\sigma|\mathtt{PS}|)+R)$ at all times during phase π_a. By Observation 2 with $a' = a$ and $k = 3\ell-1$, at all times during phase $\pi_{a+3\ell-1}$, p is located in $[z(R-5\sigma|\mathtt{PS}|)-(3\ell)\sigma|\mathtt{PS}|, z(R-5\sigma|\mathtt{PS}|)+R+(3\ell)\sigma|\mathtt{PS}|)$. In particular, at the beginning of phase $\pi_{a+3\ell}$, p is located to the left of

$$\begin{aligned}
& z(R-5\sigma|\mathtt{PS}|)+R+(3\ell)\sigma|\mathtt{PS}| \\
= \;& (z+\ell)(R-5\sigma|\mathtt{PS}|)+R-\ell R+8\ell\sigma|\mathtt{PS}| \\
\leq \;& (z+\ell)(R-5\sigma|\mathtt{PS}|)+R-10\ell\sigma|\mathtt{PS}|+8\ell\sigma|\mathtt{PS}| \text{ (by assumption (A2))} \\
= \;& (z+\ell)(R-5\sigma|\mathtt{PS}|)+R-2\ell\sigma|\mathtt{PS}|.
\end{aligned}$$

If $p \in [(z+\ell)(R-5\sigma|\texttt{PS}|)+R-\sigma|\texttt{PS}|-\frac{R}{2},(z+\ell)(R-5\sigma|\texttt{PS}|)+R-\sigma|\texttt{PS}|]$, then we just set $q'=p$. Otherwise, p is located to the left of the point $(z+\ell)(R-5\sigma|\texttt{PS}|)+R-\sigma|\texttt{PS}|-\frac{R}{2}$. However, at some time during phase $\pi_{a+3\ell}$, q is located to the right of region $\psi_{z+\ell}$, that is, at or to the right of $(z+\ell)(R-5\sigma|\texttt{PS}|)+R$. By Observation 2 with $a'=a+3\ell$ and $k=0$, at all times during phase $\pi_{a+3\ell}$, q is located at or to the right of $(z+\ell)(R-5\sigma|\texttt{PS}|)+R-\sigma|\texttt{PS}|$. It follows that the range $[(z+\ell)(R-5\sigma|\texttt{PS}|)+R-\sigma|\texttt{PS}|-\frac{R}{2},(z+\ell)(R-5\sigma|\texttt{PS}|)+R-\sigma|\texttt{PS}|]$ is between the leftmost and rightmost nodes in the network, so, by assumption (A1), there is a node q' located in this range at the beginning of phase $\pi_{a+3\ell}$.

By assumption (A2), $R-\sigma|\texttt{PS}|-\frac{R}{2}=\frac{R}{2}-\sigma|\texttt{PS}|\geq 5\sigma|\texttt{PS}|-\sigma|\texttt{PS}|>\sigma|\texttt{PS}|$, so, $[(z+\ell)(R-5\sigma|\texttt{PS}|)+R-\sigma|\texttt{PS}|-\frac{R}{2},(z+\ell)(R-5\sigma|\texttt{PS}|)+R-\sigma|\texttt{PS}|]\subseteq[(z+\ell)(R-5\sigma|\texttt{PS}|)+\sigma|\texttt{PS}|,(z+\ell)(R-5\sigma|\texttt{PS}|)+R-\sigma|\texttt{PS}|]$. By Proposition 2 with p replaced by q', $a'=a+3\ell$ and $z'=z+\ell$, it follows that $q'\in\mathcal{S}_{a+3\ell,z+\ell}$. □

Finally, we combine the major results of this section to demonstrate how the window technique provides useful guarantees about information dissemination amongst nodes following `RBSched`. We will use the following result to prove the correctness of various algorithms that are based on `RBSched`.

Theorem 3. *Suppose that there is a node $p\in\mathcal{S}_{a,z}$ for some phase π_a and region ψ_z. During phase π_a, suppose that p transmits a message containing information I. If, for some $\gamma\geq 0$, a node q is located in $((z-\gamma)(R-5\sigma|\texttt{PS}|)+\sigma|\texttt{PS}|,(z+\gamma)(R-5\sigma|\texttt{PS}|)+R-\sigma|\texttt{PS}|)$ at the end of phase $\pi_{a+3\gamma}$, then q receives I by the end of phase $\pi_{a+3\gamma}$.*

Proof. Since $\mathcal{S}_{a,z}\neq\emptyset$, $a\equiv z \mod 2$. By Theorem 1, for some $\varphi\in\{-\gamma,\ldots,\gamma\}$, $q\in\mathcal{S}_{a+3|\varphi|,z+\varphi}$. By Theorem 2, there exists a node $q'\in\mathcal{S}_{a+3|\varphi|,z+\varphi}$ that transmits I during phase $\pi_{a+3|\varphi|}$. By Lemma 10 with a replaced with $a+3|\varphi|$ and z replaced with $z+\varphi$, q receives I during phase $\pi_{a+3|\varphi|}$ via a transmission by q'. □

3.3 Our Local Gossip Algorithm

In this section, we consider the *Mobile-Rcv* model where nodes travel with bounded speed along continuous trajectories on a line. We use our `RBSched` transmission schedule to give a deterministic algorithm for δ-local gossip. Since our algorithms are based on `RBSched`, we assume that the assumptions from Section 3.2 hold.

Suppose that, initially, each node p_j in the network has a piece of information I_j. Recall that, to solve δ-local gossip, all nodes must terminate at the same time, and, for all nodes p_j,p_k such that the distance between them is at most δR at some time during the execution of the algorithm, p_j has received I_k and p_k has received I_j.

Given $\delta>0$, our algorithm for δ-local gossip, denoted by $\texttt{LG}(\delta)$, consists of each node running the schedule `RBSched` from Section 3.1 for exactly $3\alpha+2$ phases. Since δ, R, σ and $|\texttt{PS}|$ are known values, it is clear that all nodes will terminate the algorithm at the same time.

Suppose that, at the beginning of some phase π_a, a node p is located in some region ψ_z. Let q be a node within distance δR from p during a phase π_b with $b\geq a$. The following result limits the number of regions that lie between region ψ_z and q's location during phases after π_b. In particular, it takes 3 phases for q to move a distance of one region away from region ψ_z, even if it moves away as fast as possible. This result will allow us to use the window technique to prove that p's message eventually 'catches up' with q.

Lemma 11. *Suppose that, at the beginning of phase π_a, an arbitrary node p is located in region ψ_z. For any $k\geq\frac{\delta R+2\sigma|\texttt{PS}|+1}{R-8\sigma|\texttt{PS}|}$, suppose that q is within distance δR of p at some time between* *the beginning of phase* π_a *and the end of phase* π_{a+3k}*. At the end of phase* π_{a+3k}*,* q *is located in* $((z-k)(R-5\sigma|\mathtt{PS}|)+\sigma|\mathtt{PS}|,(z+k)(R-5\sigma|\mathtt{PS}|)+R-\sigma|\mathtt{PS}|)$*.*

Proof. At the beginning of phase π_a, i.e., at time $a|\mathtt{PS}|$, p is located in $[z(R-5\sigma|\mathtt{PS}|), z(R-5\sigma|\mathtt{PS}|)+R)$. Suppose that, for some $\tau \leq (3k+1)|\mathtt{PS}|$, the distance between p and q is at most δR at time $a|\mathtt{PS}|+\tau$. From the beginning of phase π_a until time $a|\mathtt{PS}|+\tau$, the distance that p can travel is at most $\tau\sigma$. Therefore, at time $a|\mathtt{PS}|+\tau$, p is located in $[z(R-5\sigma|\mathtt{PS}|)-\tau\sigma, z(R-5\sigma|\mathtt{PS}|)+\tau\sigma)$, and, furthermore, q is located in $[z(R-5\sigma|\mathtt{PS}|)-\tau\sigma-\delta R, z(R-5\sigma|\mathtt{PS}|)+\tau\sigma+\delta R)$. The amount of time that elapses from time $a|\mathtt{PS}|+\tau$ until the end of phase π_{a+3k} is $(a+3k+1)|\mathtt{PS}|-(a|\mathtt{PS}|+\tau) = (3k+1)|\mathtt{PS}|-\tau$. The distance that q can travel during this time is at most $((3k+1)|\mathtt{PS}|-\tau)\sigma$. Therefore, at the end of phase π_{a+3k}, q is located at or to the right of $z(R-5\sigma|\mathtt{PS}|)-\tau\sigma-\delta R-((3k+1)|\mathtt{PS}|-\tau)\sigma$ and to the left of $z(R-5\sigma|\mathtt{PS}|)+\tau\sigma+\delta R+((3k+1)|\mathtt{PS}|-\tau)\sigma$, i.e., in $[z(R-5\sigma|\mathtt{PS}|)-(3k+1)\sigma|\mathtt{PS}|-\delta R, z(R-5\sigma|\mathtt{PS}|)+(3k+1)\sigma|\mathtt{PS}|+\delta R)$.

Since $k \geq \frac{\delta R+2\sigma|\mathtt{PS}|+1}{R-8\sigma|\mathtt{PS}|}$,

$$\begin{aligned} k(R-5\sigma|\mathtt{PS}|)-\sigma|\mathtt{PS}| &= kR-(5k+1)\sigma|\mathtt{PS}| \\ &= \delta R+(3k+1)\sigma|\mathtt{PS}|+k(R-8\sigma|\mathtt{PS}|)-(\delta R+2\sigma|\mathtt{PS}|) \\ &> \delta R+(3k+1)\sigma|\mathtt{PS}|. \end{aligned}$$

Adding $z(R-5\sigma|\mathtt{PS}|)+R$ to both sides shows that $(z+k)(R-5\sigma|\mathtt{PS}|)+R-\sigma|\mathtt{PS}| > z(R-5\sigma|\mathtt{PS}|)+R+(3k+1)\sigma|\mathtt{PS}|+\delta R$. Similarly, subtracting both sides from $z(R-5\sigma|\mathtt{PS}|)$ shows that $(z-k)(R-5\sigma|\mathtt{PS}|)+\sigma|\mathtt{PS}| < z(R-5\sigma|\mathtt{PS}|)-(3k+1)\sigma|\mathtt{PS}|-\delta R$. Therefore, $[z(R-5\sigma|\mathtt{PS}|)-(3k+1)\sigma|\mathtt{PS}|-\delta R, z(R-5\sigma|\mathtt{PS}|)+R+(3k+1)\sigma|\mathtt{PS}|+\delta R) \subseteq ((z-k)(R-5\sigma|\mathtt{PS}|)+\sigma|\mathtt{PS}|,(z+k)(R-5\sigma|\mathtt{PS}|)+R-\sigma|\mathtt{PS}|)$. □

We now prove the correctness of the algorithm. We first suppose that an arbitrary node p is a survivor for some region ψ_z in phase π_a when it first transmits its initial information I.

Proposition 21. *Suppose that, for some* $a \in \mathbb{Z}$ *and some* $z \in \mathbb{N}$*, there is a node* $p \in \mathcal{S}_{a,z}$ *with initial information* I*. Let* $\delta > 0$*, and suppose that the schedule* `RBSched` *is followed from the start of phase* π_a *until the end of phase* π_{a+3k}*, where* $k = \left\lceil \frac{\delta R+2\sigma|\mathtt{PS}|+1}{R-8\sigma|\mathtt{PS}|} \right\rceil$*. Then, for each node* q *that is within distance* δR *from* p *at some time between the beginning of phase* π_a *and the end of phase* π_{a+3k}*,* q *receives* I *by the end of phase* π_{a+3k}*.*

Proof. Suppose that we run `RBSched` for $3\left\lceil \frac{\delta R+2\sigma|\mathtt{PS}|+1}{R-8\sigma|\mathtt{PS}|} \right\rceil+1$ phases starting at the beginning of some phase π_a. Lemma 11 combined with Theorem 3 shows that I_j has been received by each node q that is found within distance δR from p_j at some time during the execution. □

To extend this result to all nodes in the network, it suffices to execute the schedule for one additional phase.

Theorem 4. *Consider any network node* p_j *with initial information* I_j*. For any* $\delta > 0$*, after executing* `LG`(δ)*, each node that is within distance* δR *from* p_j *at some time during the execution has received* I_j *before termination.*

Proof. By Lemma 9, for an arbitrary node p and any phase π_a, p is a survivor in either phase π_a or phase π_{a+1}. The result follows by applying Proposition 21. □

Next, we find upper and lower bounds for $\left\lceil \frac{\delta R+2\sigma|\texttt{PS}|+1}{R-8\sigma|\texttt{PS}|} \right\rceil$ as functions of δ. Since the number of phases used by $\texttt{LG}(\delta)$ is $3\alpha+2$, we get upper and lower bounds on the number of phases used by $\texttt{LG}(\delta)$.

Proposition 22. $\delta < \left\lceil \frac{\delta R+2\sigma|\texttt{PS}|+1}{R-8\sigma|\texttt{PS}|} \right\rceil \leq 5\delta + 3.$

Proof. First, notice that

$$\frac{\delta R + 2\sigma|\texttt{PS}| + 1}{R - 8\sigma|\texttt{PS}|} = \delta + \frac{(8\delta + 2)\sigma|\texttt{PS}| + 1}{R - 8\sigma|\texttt{PS}|}$$

It follows that $\left\lceil \frac{\delta R+2\sigma|\texttt{PS}|+1}{R-8\sigma|\texttt{PS}|} \right\rceil > \delta$. Next, by constraint (A2), $R \geq 10\sigma|\texttt{PS}|$, so $R - 8\sigma|\texttt{PS}| \geq 2\sigma|\texttt{PS}|$. Hence, $\left\lceil \frac{\delta R+2\sigma|\texttt{PS}|+1}{R-8\sigma|\texttt{PS}|} \right\rceil \leq \lceil 5\delta + 2 \rceil \leq 5\delta + 3$. □

Corollary 5. $\texttt{LG}(\delta)$ *uses at most* $15\delta + 11$ *phases, and no fewer than* $3\delta + 3$ *phases, where each phase consists of* $O(\rho^2 \Delta^2 \log U)$ *slots.*

4 Conclusion and Future Work

Local gossip is a task that captures the need for nodes in a mobile network to share information with other nearby nodes, even if they are not within communication range of one another for a long period of time. A solution to this task can be useful as a fundamental building block in algorithms for mobile networks, and we have demonstrated that this is the case with neighbour discovery via reductions that hold very generally. The resulting solutions to neighbour discovery need not make the simplifying assumptions made elsewhere in the literature, and they are deterministic, which means that they can be used as subroutines without introducing error. By solving the local gossip task in the one-dimensional *Mobile-Rcv* model, we obtain solutions to one-time exact neighbour discovery in the same model, which can be used to answer open questions about initializing the algorithms in [4, 8, 16].

One important direction for future work is to solve δ-local gossip in more general environments, such as road networks or the plane. An important generalization is to weaken the assumption that the transmission radius is the same for all nodes. We feel that the results in this paper could be generalized to the asymmetric case without too much difficulty, and is left for future work. One important decision in such a generalization is to decide whether to define a node's neighbours as the nodes to which it can send a message, or the nodes from which it can receive a message.

A challenging problem to consider would be to determine the trade-off between future trajectory knowledge and feasible values of the parameters for neighbour discovery tasks. Such a trade-off would have an impact on the design of real-world systems, where there can be varying degrees of future trajectory knowledge. For example, satellites or bus routes are fully specified in advance, while a self-driving car has a planned route until it reaches its destination, while a human-controlled car has no knowledge of its future trajectory.

References

[1] Antonio Fernández Anta, Alessia Milani, Miguel A. Mosteiro, and Shmuel Zaks. Opportunistic information dissemination in mobile ad-hoc networks: the profit of global synchrony. *Distributed Computing*, 25(4):279–296, 2012.

[2] John Augustine, Gopal Pandurangan, and Peter Robinson. Distributed algorithmic foundations of dynamic networks. *SIGACT News*, 47(1):69–98, 2016.

[3] Roberto Baldoni, Antonio Fernández Anta, Kleoni Ioannidou, and Alessia Milani. The impact of mobility on the geocasting problem in mobile ad-hoc networks: Solvability and cost. *Theor. Comput. Sci.*, 412(12-14):1066–1080, 2011.

[4] Hyun Chul Chung, Saira Viqar, and Jennifer L. Welch. Neighbor knowledge of mobile nodes in a road network. In *2012 IEEE 32nd International Conference on Distributed Computing Systems, Macau, China, June 18-21, 2012*, pages 486–495. IEEE Computer Society, 2012.

[5] Andrea E. F. Clementi, Angelo Monti, and Riccardo Silvestri. Selective families, superimposed codes, and broadcasting on unknown radio networks. In S. Rao Kosaraju, editor, *Proceedings of the Twelfth Annual Symposium on Discrete Algorithms, January 7-9, 2001, Washington, DC, USA.*, pages 709–718. ACM/SIAM, 2001.

[6] Alejandro Cornejo, Saira Viqar, and Jennifer L. Welch. Reliable neighbor discovery for mobile ad hoc networks. *Ad Hoc Networks*, 12:259–277, 2014.

[7] A. G. Dyachkov, V. V. Rykov, and A. M. Rashad. Superimposed distance codes. *Problems Control Inform. Theory/Problemy Upravlen. Teor. Inform.*, 18(4):237–250, 1989.

[8] Faith Ellen, Sivaramakrishnan Subramanian, and Jennifer L. Welch. Maintaining information about nearby processors in a mobile environment. In Soma Chaudhuri, Samir R. Das, Himadri S. Paul, and Srikanta Tirthapura, editors, *Distributed Computing and Networking, 8th International Conference, ICDCN 2006, Guwahati, India, December 27-30, 2006.*, volume 4308 of *Lecture Notes in Computer Science*, pages 193–202. Springer, 2006.

[9] P. Erdös, P. Frankl, and Z. Füredi. Families of finite sets in which no set is covered by the union of r others. *Israel J. Math*, 51:79–89, 1985.

[10] Leszek Gasieniec, Aris Pagourtzis, Igor Potapov, and Tomasz Radzik. Deterministic communication in radio networks with large labels. *Algorithmica*, 47(1):97–117, 2007.

[11] W. Kautz and R. Singleton. Nonrandom binary superimposed codes. *IEEE Trans. on Info. Theory*, 10(4):363–377, 1964.

[12] Athar Ali Khan, Mubashir Husain Rehmani, and Yasir Saleem. Neighbor discovery in traditional wireless networks and cognitive radio networks: Basics, taxonomy, challenges and future research directions. *J. Network and Computer Applications*, 52:173–190, 2015.

[13] Srinivasan Krishnamurthy, Mansi Ramakrishnan Thoppian, Srikant Kuppa, Ramaswamy Chandrasekaran, Neeraj Mittal, Subbarayan Venkatesan, and Ravi Prakash. Time-efficient distributed layer-2 auto-configuration for cognitive radio networks. *Computer Networks*, 52(4):831–849, 2008.

[14] Fabian Kuhn, Nancy A. Lynch, and Calvin C. Newport. The abstract MAC layer. *Distributed Computing*, 24(3-4):187–206, 2011.

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

[16] Avery Miller. Gossiping in one-dimensional synchronous ad hoc wireless radio networks. In Lélia Blin and Yann Busnel, editors, *4th Workshop on Theoretical Aspects of Dynamic Distributed Systems, TADDS '12, Roma, Italy, December 17, 2012*, pages 32–43. ACM, 2012.

[17] Neeraj Mittal, Srinivasan Krishnamurthy, R. Chandrasekaran, Subbarayan Venkatesan, and Yanyan Zeng. On neighbor discovery in cognitive radio networks. *J. Parallel Distrib. Comput.*, 69(7):623–637, 2009.

[18] Shailesh Vaya. Round complexity of leader election and gossiping in bidirectional radio networks. *Inf. Process. Lett.*, 113(9):307–312, 2013.