Latency-optimized Broadcast in Mobile Ad Hoc Networks without Node Coordination

Bin Tang
State Key Lab for Novel Software Technology
Nanjing University, China
tb@nju.edu.cn

Baoliu Ye*
State Key Lab for Novel Software Technology
Nanjing University, China
yebl@nju.edu.cn

Sanglu Lu
State Key Lab for Novel Software Technology
Nanjing University, China
sanglu@nju.edu.cn

Song Guo
School of Computer Science and Engineering
The University of Aizu, Japan
sguo@u-aizu.ac.jp

Ivan Stojmenovic
School of Information Technology and Engineering
University of Ottawa, Canada
ivan@site.uottawa.ca

ABSTRACT

We consider the problem of broadcasting a message in a mobile ad hoc network (MANET) with the objective of minimizing the broadcast latency. Due to the mobility of network nodes, the coordination among nodes is hard and expensive. Thus it is much desired to design efficient, one-sided broadcast protocols where each node acts according to its own state solely. Although random scheduling is a popular and effective one-sided approach for leveraging the broadcast nature of wireless medium while coping with transmission collisions, both critical for reducing the broadcast latency, in this paper, we show that when nodes move very fast, the performance of pure random scheduling must be sub-optimal, no matter how the forwarding probabilities are specified. Furthermore, we propose a novel one-sided broadcast protocol named R^2, which first splits the message into a certain number of mini-messages and then couples a fine-grained random scheduling with random linear network coding for broadcasting the mini-messages. Theoretical analyses demonstrate that R^2 performs optimally in order sense, no matter how fast network nodes move around, although different mobility has distinct effect on the speed of message broadcast.

Categories and Subject Descriptors

C.2.2 [**Computer-Communication Networks**]: Network Protocols

Keywords

Mobile Ad Hoc Networks, Broadcast, Random Linear Network Coding, Random Scheduling

*Corresponding author

MobiHoc'14, August 11–14, 2014, Philadelphia, PA, USA.

http://dx.doi.org/10.1145/2632951.2632982 .

1. INTRODUCTION

With the fast development of wireless and social networking technologies, there has been increasing interest on the use of mobile ad hoc networks (MANETs) for data communication. A typical MANET is self-configured and formed by a collection of wireless nodes which act not only as hosts but also as relays, storing and forwarding data for other nodes in the network. The nodes in a MANET can move around at their will, which on one hand, brings the potential to improve the network performance, *e.g.*, capacity [18] and throughput-delay tradeoff [39], but on the other hand, raises great challenges for protocol design and performance optimization, as the network topology may change rapidly and unpredictably.

Network-wide broadcast, *i.e.*, delivering a message from a source node to all other network nodes, is a fundamental operation in MANETs. For many important applications of MANETs, like military communication, it is essential to design an efficient broadcast scheme with low *broadcast latency* (*i.e.*, the time it takes for all the network nodes to receive a copy of the message), so that the end-to-end delay for higher-level applications could be stringently guaranteed.

Towards minimizing the broadcast latency in MANETs, two key features of wireless communications should be carefully addressed. One is the broadcast nature of wireless medium, *i.e.*, a transmission of a message could be heard simultaneously by multiple neighboring nodes, which can speed up the broadcast process. The other is wireless interference/collision, *i.e.*, a node cannot receive a message if more than one neighboring nodes are transmitting at the same time, which can thus limit the diffusion of the message. In static wireless ad hoc networks, the two features are usually incorporated by constructing a broadcast tree together with a collision-free broadcast schedule [15, 25, 33, 14]. However, in a MANET, the network topology is dynamic over time, making the construction and maintenance of a broadcast tree prohibitive and also the coordination among nodes expensive, especially when the network nodes move very fast. Therefore, it is much desired if we could design an efficient *one-sided* broadcast protocol where each node acts according to its own state solely, avoiding the coordination among nodes.

The simplest one-sided broadcast scheme is flooding, where each node that has received the message keeps on forwarding the message to all its current neighboring nodes. Although the flooding scheme can enjoy the broadcast nature of wireless medium, it can also lead to severe collisions, making the broadcast latency intolerably high. One popular and effective approach to cope with these collisions while preserving the one-sided property of flooding is *random scheduling*, *i.e.*, each node that has received the message forwards the message only with a certain probability [3, 41, 21]. In general, the forwarding probability at each time is chosen mainly according to the number of nearby nodes that have received the message. However, due to the node mobility as well as the underlying randomness of broadcast process, such number may vary significantly over different nodes, and can also be hard to estimate, making the optimization of forwarding probabilities difficult.

Instead of focusing on the optimization of forwarding probabilities for random scheduling based broadcast directly, we theoretically demonstrate that, when nodes move very fast, no matter how the forwarding probabilities are specified, *i.e.*, varying over different nodes and time instances, the performance under pure random scheduling must be sub-optimal. The underlying reason is that the pure random scheduling suffers from the bottleneck in completing the broadcast process due to the very late receptions of the message by a few nodes.

In contrast to the above pessimistic result about pure random scheduling based broadcast, we further propose a novel one-sided broadcast protocol named R^2, which achieves the optimal broadcast latency asymptotically by coupling a fine-grained *Random scheduling* with *Random linear network coding (RLNC)*. Specifically, in the R^2 protocol, the message to be broadcast is firstly split into a certain number of mini-messages, such that random scheduling can be applied on a mini-message level and a uniform forwarding probability is enough for the efficiency of random scheduling. The mini-messages are then transmitted in an RLNC fashion, *i.e.*, all nodes forward coded mini-messages formed from random linear combinations of all previously received mini-messages, which makes nodes easier to accumulate useful information, and thus mitigates the bottleneck due to the application of random scheduling. As demonstrated by our theoretical arguments, R^2 can achieve an order-optimal broadcast latency no matter how fast nodes move around the network, even though different mobility of network nodes has distinct effect on the speed of message broadcast.

The contribution of the paper is two-fold:

- We first show that pure random scheduling is insufficient to achieve the optimal broadcast latency when nodes move very fast;
- We further propose a novel one-sided broadcast protocol R^2 which couples a fine-grained random scheduling with random linear network coding, and show that R^2 performs optimally no matter how fast nodes move around.

The reminder of the paper is organized as follows. In Sec. 2, we introduce the network model and the lower bounds of the broadcast latency. In Sec. 3, we present the protocols and main results, followed by proofs given in Sec. 4-6. Literature reviews are provided in Sec. 7. Finally, the conclusion and open questions are presented in Sec. 8.

2. THE MODEL

We consider a mobile ad-hoc network (MANET) consisting of n mobile nodes distributed uniformly on a $\sqrt{n} \times \sqrt{n}$ square. A source node, say s, wants to disseminate a message $\mathbf{m}$ with l bits to all the other nodes in the network. The *broadcast latency* is defined as the length of time interval between the starting time of the broadcast operation and the first time when all network nodes receive a copy of $\mathbf{m}$. Our main task is to minimize the broadcast latency by leveraging the broadcast nature of wireless medium as well as node mobility, while overcoming the interference due to concurrent transmissions. Keeping this in mind, we begin with the detailed introduction of the node mobility model and the communication model.

2.1 Mobility Model

The whole network region is partitioned into sub-squares of side length ρ each, resulting in a $\frac{\sqrt{n}}{\rho} \times \frac{\sqrt{n}}{\rho}$ *torus*. For ease of presentation, we refer to the formed sub-squares as ρ-cells and ρ as the velocity of each mobile node. The time is divided into slots of equal duration, and at the beginning of each time slot, each node independently moves to a new position inside its current ρ-cell or eight adjacent ρ-cells uniformly at random. Here we mean two ρ-cells (or two r-cells which are defined later) are adjacent if they touch each other by a side or by a corner. This mobility model can be viewed as a discrete version of the random walk mobility model, and has been widely used in literatures, *e.g.*, [9, 5]. In particular, when $\rho = \sqrt{n}/3$, this model degenerates to be the well-known *i.i.d.* mobility model [36].

2.2 Communication Model

All network nodes communicate over a common wireless channel, and the network topology in time slot t is modeled by a random geometric graph [37] $\mathcal{G} = (\mathcal{V}, \mathcal{A})$, where $\mathcal{V} = \{1, 2, \ldots, n\}$ is the set of nodes and $\mathcal{A}_t = \{(i, j) : d_t(i, j) \leq r\}$ is the set of arcs. Here $d_t(i, j)$ denotes the Euclidean distance between nodes i and j in time slot t, and r is the common transmission radius. In other words, a pair of nodes are neighbors if their distance is no larger than r. To further characterize a successful reception of a transmission, we consider the classical *protocol model* [20], i.e., a transmission from node i to node j is successful if for any time t within the duration of the transmission, $d_t(i, j) \leq r$ and each node i' except i with $d_t(i', j) \leq (1 + \Delta)r$ is not transmitting simultaneously, where $\Delta > 0$ is a transmission guard constant independent with n.

Regarding the communication model, we make the following assumptions.

- We assume that the capacity of the wireless channel is l bits per time slot, so that the whole message $\mathbf{m}$ can be transmitted completely over a time slot. When the channel capacity is larger than l, our protocol still works and remains to be optimal asymptotically. In the opposite case where the channel capacity is smaller than l, the message to be broadcast can be considered as multiple messages each with length equal to the capacity. One might transmit the messages with our proposed protocol in a sequential manner, but how to transmit them efficiently in a concurrent manner requires further exploration.

- We assume that the transmission radius r is above the critical radius [19], which is on the order of $\sqrt{\log n}$. In this case, the network topology at any time is connected almost surely, and thus the broadcast process can be complete even if the network is static, *i.e.*, $\rho = 0$. On the other hand, r should be kept small enough for the purpose of energy saving. Regarding these considerations, we always assume $r = c\sqrt{\log n}$ for a sufficiently large constant $c > 0$. Besides, we also assume that all the nodes have common knowledge about n and r.
- Since the processing of wireless signals takes time, we assume that the information received by any node in some time slot can only be used for transmissions from the node in later time slots other than the same time slot.

2.3 Lower Bounds

As demonstrated in [7, 8], node velocity plays an important role in the speed of message dissemination, thus we consider two cases accordingly. One is named *low mobility regime*, where $\rho = O(r) = O(\sqrt{n})$. For this case, we have

LEMMA 1. *When $\rho = O(\sqrt{\log n})$, the broadcast latency under any protocol is $\Omega\left(\sqrt{\frac{n}{\log n}}\right)$ with probability $1-O(n^{-1})$.*

The other is referred to as *high mobility regime*, where $\rho = \omega(r) = \omega(\sqrt{\log n})$. For this case, we have

LEMMA 2. *When $\rho = \omega(\sqrt{\log n})$, the broadcast latency under any protocol is $\Omega\left(\frac{\sqrt{n}}{\rho} + \frac{\log n}{\log\log n}\right)$ with probability $1 - O(n^{-1})$.*

The above lemmas could be proved similarly to [7, 40] in spirit, with a slight more effort with Shannon measures. The details are thus omitted. From Lemma 2, we can see that, when the velocity is very large, say $\rho = \omega\left(\frac{\sqrt{n}}{\log n}\right)$, the lower bound of the broadcast latency could be sub-logarithmic in n, *i.e.*, $o(\log n)$. This implies that many approaches cannot achieve the lower bound in the high mobility regime as they require at least $\Omega(\log n)$ time slots. For instance, in gossip-based protocols (e.g., [5]), each node forwards one message to only a single neighboring node which is chosen randomly in each time slot. Since the number of nodes that have received the message can grow by two times in each time slot, thus no matter how the gossip-based protocols are specified, it requires at least $\Omega(\log n)$ time slots to complete the broadcast process. In contrast, our proposed R^2 protocol can achieve the lower bounds in either mobility regime, implying that *the lower bounds given in Lemma 1 and Lemma 2 are tight.*

3. PROTOCOLS AND RESULTS

In this section, we introduce the broadcast protocols as well as the results about their performance. The proofs of these results are deferred to the next sections.

3.1 Pure Random Scheduling Protocol

Random scheduling is an efficient technique for exploiting the broadcast nature of wireless medium while coping with interference in wireless networks [3]. Using random scheduling solely, a general broadcast protocol can be formalized as follows:

Pure Random Scheduling (PRS) Protocol:
In every time slot t, each node v that has received message $\mathbf{m}$ broadcasts it to all its neighboring nodes independently with probability $p_t(v)$, while other nodes keep silent.

Though the description of the PRS protocol is rather simple, it is a difficult task to optimize $p_t(v)$ so as to minimize the broadcast latency. What is worse, no matter how probabilities $p_t(v)$ are specified, the performance of the PRS protocol must be sub-optimal in some cases, as revealed by the following theorem.

THEOREM 3. *When $\rho = \sqrt{n}/3$, the broadcast latency under the PRS protocol is $\Omega(\log n)$ with probability $1-O(n^{-1})$, no matter how probabilities $p_t(v)$ in the PRS protocol are specified.*

As we will see later, when $\rho = \frac{\sqrt{n}}{3}$, the optimal latency is $\Theta\left(\frac{\log n}{\log\log n}\right)$. Therefore, there exists a $\Omega(\log\log n)$ factor gap between the performance of PRS and the optimum.

3.2 The R^2 Broadcast Protocol

3.2.1 Inspirational Example

Before introducing our broadcast protocol R^2, we first show how the RLNC is inspired via a simplified scenario. In this scenario, there are k nodes, say v_i, $i = 1, 2, \ldots, k$, need to get message $\mathbf{m}$. Assume that with random scheduling, in each time slot, each v_i can succeed to receive $\mathbf{m}$ independently with a constant probability $0 < p < 1$. Now we analyze the time it takes for all v_i to receive a copy of $\mathbf{m}$ with a failure probability of ϵ. Denoting the time as $T^d(\epsilon)$, we have

$$1 - \epsilon \le \left(1 - (1-p)^{T^d(\epsilon)}\right)^k \le e^{-(1-p)^{T^d(\epsilon)}k},$$

which gives

$$T^d(\epsilon) \ge -\frac{\ln k - \ln(-\ln(1-\epsilon))}{\ln(1-p)} = \Omega(\ln k). \qquad (1)$$

Comparing to the expected time for a single v_i to receive $\mathbf{m}$ which is just a constant, the total latency is much higher. This is mainly due to the very late reception of $\mathbf{m}$ by a few nodes.

Now suppose that message $\mathbf{m}$ is split into a number (denoted by f) of mini-messages and each time slot is also split into the same number of mini-slots, so that each mini-message could be transmitted in a mini-slot. Then transmit these mini-messages in an RLNC fashion such that every coded mini-message received by v_i is a random linear combination of all original mini-messages. Once v_i has collected f coded mini-messages with linearly independent coding vectors, it can recover $\mathbf{m}$ by Gaussian elimination. Assume that in each mini-slot and for each v_i, the probability that v_i can receive a coded mini-message remains to be a constant p'. Also, for ease of illustration, every coding coefficient is assumed to be chosen from a binary field $\mathbb{F}_2$ uniformly at random (this assumption does not hold in general networks). Let $T^{nc}(\epsilon)$ be defined similarly to $T^d(\epsilon)$. If $f \ge 2(\ln(2k) - \ln\epsilon)$, then according to [43], when v_i has received $2f$ coded mini-messages, the probability that it can recover the original f mini-messages is

$$\prod_{j=0}^{f-1}\left(1-2^{j-2f}\right) \geq \left(1-2^{-f}\right)^f \geq 1-\frac{f}{2^f} \geq 1-\frac{\epsilon}{2k}. \quad (2)$$

Meanwhile, by a standard Chernoff bound argument [34], we can show that in $\lceil 4/p' \rceil$ time slots, the probability that v_i receives at least $2f$ coded mini-messages is at least

$$1-e^{-\frac{f}{2}} \geq 1-\frac{\epsilon}{2k}. \quad (3)$$

Combining (2) and (3) and applying the union bound, we can see that the probability that v_i can recover $\mathbf{m}$ in $\lceil 4/p' \rceil$ time slots is at least $1-\frac{\epsilon}{k}$. By applying the union bound again over all k pairs, we thus have

$$T^{nc}(\epsilon) \leq \lceil 4/p' \rceil. \quad (4)$$

Note that the latency achieved by the RLNC approach is just a constant. The key insight is that the use of message split and RLNC makes nodes easier to accumulate useful information, relieving the bottleneck problem that cannot be avoided by the former approach.

3.2.2 The R^2 Protocol and Performance

Now we formally introduce R^2, a broadcast protocol that couples a fine-grained random scheduling with RLNC. In R^2, the message $\mathbf{m}$ is partitioned into c^2 sub-messages with equal length, and the sub-messages are then broadcast in a sequential manner. (The reason for such partition will be explained at the end of Sec. 6.1.) In order to broadcast a sub-message,

- the sub-message, say $\mathbf{sm}$, is split by the source node into $\tau = \log n$ mini-messages, denoted by $\mathbf{mm}_1, \ldots, \mathbf{mm}_\tau$, each with $l/(c^2\tau)$ bits, and
- every time slot is also split into $\gamma = r^2 = c^2\tau$ mini-slots with equal duration, thus the transmission of a mini-message can be completed in a mini-slot.

Based on the above basic operations, the approach as the core of R^2 for broadcasting sub-message $\mathbf{sm}$ can then be described as follows:

R^2 for broadcasting sm:
Source node:

- In the first time slot that contains γ mini-slots, it disseminates the mini-message $\mathbf{mm}_i$, $i = 1, 2, \ldots, \tau$, to all its neighboring nodes in the i-th mini-slot. It remains silent in the left mini-slots.
- In later time slots, it works the same as other nodes.

Any other node in each time slot:

- If it has not received any mini-message, then it keeps silent in the whole time slot.
- Otherwise, in each mini-slot, it does the following with probability $1/\gamma$ (random scheduling):
  - (RLNC) it first picks out every mini-message received in previous time slots independently with probability 1/2, and then combines them into a coded mini-message by XORing them all.
  - it then broadcasts the generated mini-message to all its neighboring nodes.

In the above procedure, each generated mini-message is a linear combination of $\mathbf{mm}_1, \ldots, \mathbf{mm}_\tau$. For the sake of decoding, τ linear coefficients (referred to as the coding vector), each of which costs one bit, should also be appended into the mini-message for transmission. We can omit the coding vector overhead here to ease the presentation under a quite loose condition of $l = \Omega(\log^2 n)$ which can hold for large message. To see this under the assumption, we can adjust the number of sub-messages by a constant factor so that in the broadcast of a sub-message with R^2, a coded mini-message with coding vector appended has exactly $\frac{l}{c^2\tau}$ bits and thus can be transmitted in one mini-slot. This does not affect the performance in order sense (some remarks on the chosen of parameters in R^2 are also given at the end of Sec. 6.1).

Once a node receives enough coded mini-messages, say with τ linearly independent coding vectors, it can recover the whole sub-message $\mathbf{sm}$ by Gaussian elimination. Note that the generation of coded mini-messages in R^2 is well-known as RLNC with operations executed over a binary field $\mathbb{F}_2$. While the use of larger finite field may have the potential to improve the performance, it also incurs higher computational cost for encoding and decoding operations. On the other hand, the use of a binary field employs the simple XOR as a basic arithmetic operation, and also leads to minimum overhead due to coding vectors. Moreover, the binary field $\mathbb{F}_2$ is sufficient for the R^2 protocol to achieve asymptotic optimality in terms of broadcast latency, as shown by the following theorems.

THEOREM 4 (LOW MOBILITY REGIME). *When $\rho = O\left(\sqrt{\log n}\right)$ and $l = \Omega(\log^2 n)$, the R^2 protocol achieves a broadcast latency of $O\left(\sqrt{\frac{n}{\log n}}\right)$ with probability $1-O(n^{-1})$.*

THEOREM 5 (HIGH MOBILITY REGIME). *When $\rho = \omega\left(\sqrt{\log n}\right)$ and $l = \Omega(\log^2 n)$, the R^2 protocol achieves a broadcast latency of $O\left(\frac{\sqrt{n}}{\rho} + \frac{\log n}{\log\log n}\right)$ with probability $1 - O(n^{-1})$.*

Comparing Theorem 4 with Lemma 1 and Theorem 5 with Lemma 2, we can see that, as long as $l = \Omega(\log^2 n)$, R^2 performs optimally in order sense no matter how large the velocity ρ is.

4. PRELIMINARIES

Before proving Theorem 3, Theorem 4 and Theorem 5, in this section, we introduce some preliminaries.

4.1 Network Region Partition

We partition the whole square region into small sub-squares (referred to as r-cells), each with side length ar, where a is a suitable constant such that $\frac{1}{4\sqrt{2}} \leq a \leq \frac{1}{2\sqrt{2}}$ and $\frac{\sqrt{n}}{ar}$ is an integer. For an r-cell $\mathcal{R}$, we use $\mathcal{N}(\mathcal{R})$ to denote the set of all r-cells that are adjacent to $\mathcal{R}$ including R itself. Clearly, any pair of nodes in the same r-cell are neighbors. Also, for a node v lying in $\mathcal{R}$, every node in $\mathcal{N}(\mathcal{R})$ is within the transmission range of v.

A well-known result is that every r-cell in the region contains $\Theta(r^2)$ nodes throughout the broadcast process almost surely [2]. Formally, let $\sharp_{\mathcal{R}}(t)$ denote the number of nodes lying in r-cell $\mathcal{R}$ during time slot t, then

Fact 1. There exists positive constants η_1 and η_2 which only depend on a, such that

$$\Pr\{\mathcal{D}_n\} \geq 1 - n^{-4},$$

where $\mathcal{D}_t$ denotes the event of $\eta_1 r^2 \leq \sharp_{\mathcal{R}}(t') \leq \eta_2 r^2$ for every r-cell $\mathcal{R}$ and every time slot $t' \leq t$.

Hence, to prove a result with probability $1 - O(n^{-1})$, it is sufficient to show the same result conditioning on $\mathcal{D}_n$. In the rest of the paper, unless otherwise specified, it is implicitly assumed that $\mathcal{D}_n$ holds.

Now we introduce a simple but very useful concept about r-cells. Two r-cells are said to be *isolated* if the distance between any pair of points each in one r-cell is larger than $2(2+\Delta)r$, and a collection of r-cells are said to be isolated if any pair of r-cells are isolated. Under this definition, the success or not of any mini-message transmissions from/to nodes in one r-cell $\mathcal{R}$ are irrespective of those events in other r-cells that are isolated with $\mathcal{R}$ at any certain time.

Fact 2. For any set $\mathcal{S}$ of r-cells, there exists some subset $\mathcal{S}' \subseteq \mathcal{S}$ such that $|\mathcal{S}'| \geq |\mathcal{S}|/\phi$ and all the r-cells in $\mathcal{S}'$ are isolated, where $\phi = (2(2+\Delta)/a + 1)^2$.

Proof. We construct $\mathcal{S}'$ from an empty set by the following iterative algorithm: in each iteration, an arbitrary r-cell $\mathcal{R}$ is picked out from $\mathcal{S}$ and included into $\mathcal{S}'$, and all r-cells in $\mathcal{S}$ that are not isolated with $\mathcal{R}$ are removed away from $\mathcal{S}$; the iteration terminates when $\mathcal{S}$ becomes empty. Evidently, all r-cells in $\mathcal{S}'$ are isolated. Also, by some simple geometric arguments, it is easily shown that there are at most $\phi - 1$ r-cells each of which is not isolated to $\mathcal{R}$. Therefore, $|\mathcal{S}|$ is decreased by at most ϕ in each iteration, which implies that the number of iterations is at least $|\mathcal{S}|/\phi$ and thus $|\mathcal{S}'| \geq |\mathcal{S}|/\phi$. □

4.2 Projection Analysis

Projection analysis [22] is a newly developed technique for analyzing RLNC-based protocol. For the sake of completeness, we summarize its core concept and result as follows.

Definition 1. We say that node v knows about $\vec{\mu} \in \mathbb{F}_2^\tau$ or v is $\vec{\mu}$-informed, if $\vec{\mu} = \vec{0}$, or there is a vector $\vec{\nu}$ in the coding vector subspace of v such that $\vec{\mu} \cdot \vec{\nu} \neq 0$, where the coding vector subspace is induced by the coding vectors of all coded mini-messages in hand.

Lemma 6. *If a node knows about $\vec{\mu} \in \mathbb{F}_2^\tau$ and broadcasts a coded mini-message to a set of nodes successfully, then the probability that all nodes in the set know about $\vec{\mu}$ afterwards is at least $1/2$. Moreover, a node can recover all the τ mini-messages once it knows about all vectors in $\mathbb{F}_2^\tau$.*

With the help of projection analysis, we can claim that the latency of broadcasting a sub-message is upper bounded by T with probability $1 - O(n^{-1})$, if for any $\vec{\mu} \in \mathbb{F}_2^\tau$, the probability that all nodes know about $\vec{\mu}$ at time T is at least $1 - O\left(n^{-1}2^{-\tau}\right) = 1 - O(n^{-2})$. This is because under the assumption, we can apply the union bound over all $2^\tau = n$ vectors in $\mathbb{F}_2^\tau$ to show that at time T, the probability that all nodes know about all vectors in $\mathbb{F}_2^\tau$, or equivalently all nodes can recover the original τ mini-messages according to Lemma 6, is $1 - O(n^{-1})$.

5. SUB-OPTIMALITY OF PRS: PROOF OF THEOREM 3

In the proof of Theorem 3, we only use the condition that $\mathcal{D}_{t_0}$ holds for some suitable $t_0 \leq n$, which will be clear soon. For any time $t \leq t_0$ and any node v, the communication range of v, *i.e.*, the circle centered at v with radius r, is covered by at most $\xi = \left(\left\lceil \frac{2r}{ar} \right\rceil + 2\right)^2 = \left(\left\lceil \frac{2}{a} \right\rceil + 2\right)^2$ r-cells. Since each r-cell contains at most $\eta_2 r^2$ nodes, each node is adjacent to at most $\xi\eta_2 r^2$ nodes at one time, which implies that the number of nodes that receive $\mathbf{m}$ can increase by at most $\xi\eta_2 r^2$ times in one time slot. Let $\mathcal{U}(t)$ be the set of nodes that have not received $\mathbf{m}$ in the first t time slots. Then there must be some time t_0, such that

$$\frac{n}{2\xi\eta_2 r^2} \leq n - |\mathcal{U}(t_0)| < n/2.$$

Lemma 7. *For any $t > t_0$ such that $|\mathcal{U}(t-1)| \geq \sqrt{n}$, there exists some constant $\chi < 1$ such that*

$$\Pr\{|\mathcal{U}(t)| \geq \chi|\mathcal{U}(t-1)|\} \geq 1 - O(n^{-2}). \tag{5}$$

Proof. A key observation is that in time slot t, the movement of a node $v \in \mathcal{V} \setminus \mathcal{U}(t-1)$ and whether v broadcasts $\mathbf{m}$ or not are both stochastic, and independent with each other. Thus, it is equivalent if the chronological order between the exposure on the randomness of movements of nodes in $\mathcal{V} \setminus \mathcal{U}(t-1)$ and the exposure on the randomness of whether these nodes broadcast $\mathbf{m}$ or not are exchanged. Let $\mathcal{B} \subseteq \mathcal{V} \setminus \mathcal{U}(t-1)$ be the set of nodes that broadcast $\mathbf{m}$ in time slot t. We will show that, for any possible value B of $\mathcal{B}$, there exists some constant $\chi < 1$ independent with B such that

$$\Pr\{|\mathcal{U}(t)| \geq \chi|\mathcal{U}(t-1)| \,|\, \mathcal{B} = B\} \geq 1 - O(n^{-2}). \tag{6}$$

Then by the law of total probability, we can get (5) straightforwardly.

Now we prove (6). Assume that $\mathcal{B} = B$ holds. We consider some set of isolated r-cells in the network region, say $\mathcal{S}$, such that $|\mathcal{S}| \geq \frac{n}{\phi(ar)^2}$. According to Fact 2, such $\mathcal{S}$ exists. For any r-cell $\mathcal{R} \in \mathcal{S}$, let $I_{\mathcal{R}}$ be an indicator variable such that $I_{\mathcal{R}} = 1$ if in time slot t, all nodes lying in $\mathcal{R}$ and belonging to $\mathcal{U}(t-1)$ do not receive $\mathbf{m}$ successfully, and $I_{\mathcal{R}} = 0$ if otherwise. Define $X_{\mathcal{R}}$ to be the number of nodes in B and lying in $\mathcal{R}$ in time slot t. According to the *i.i.d.* mobility model, $X_{\mathcal{R}}$ follows a binomial distribution, *i.e.*,

$$X_{\mathcal{R}} \sim \text{Binom}\left(|B|, \frac{(ar)^2}{n}\right)$$

with mean value $\lambda = \frac{|B|(ar)^2}{n}$. To cope with the dependency among $X_{\mathcal{R}}$, we apply the technique of Poisson approximation [34], where each $X_{\mathcal{R}}$ is approximated by an independent Poisson random variable $Y_{\mathcal{R}}$ with the same mean value λ. Also define $I_{\mathcal{R}}^P$ similarly to $I_{\mathcal{R}}$ in the Poisson case.

We first consider the Poisson case. Let $\mathcal{E}_1$ be the event that there is not any node in B lying in $\mathcal{R}$ or r-cells that are not isolated to $\mathcal{R}$. Then,

$$\begin{aligned}\Pr\{\mathcal{E}_1\} &= \Pr\{Y_{\mathcal{R}} = 0\} \prod_{\mathcal{R}':\text{not isolated with } \mathcal{R}} \Pr\{Y_{\mathcal{R}'} = 0\} \\ &\geq \left(e^{-\lambda}\right)^\phi = e^{-\lambda\phi}.\end{aligned}$$

Let $\mathcal{E}_2$ be the event that $\mathcal{R}$ contains at least two nodes in B in time slot t. Then,

$$\Pr\{\mathcal{E}_2\} = 1 - e^{-\lambda} - \lambda e^{-\lambda}.$$

If either $\mathcal{E}_1$ or $\mathcal{E}_2$ holds, then $I^P_{\mathcal{R}} = 1$. Therefore,

$$\begin{aligned} \Pr\{I^P_{\mathcal{R}} = 1\} &\geq \Pr\{\mathcal{E}_1\} + \Pr\{\mathcal{E}_2\} \\ &\geq e^{-\phi\lambda} + 1 - e^{-\lambda} - \lambda e^{-\lambda} \\ &\geq \begin{cases} e^{-\phi} & \text{if} \quad \lambda < 1 \\ 1 - \frac{2}{e} & \text{if} \quad \lambda \geq 1 \end{cases} \\ &\geq \text{some constant } \beta, \end{aligned}$$

where $\beta < 1$. Due to the mutual independence of all $I^P_{\mathcal{R}}$ for $\mathcal{R} \in \mathcal{S}$, we can use the Chernoff bound to obtain

$$\Pr\left\{\sum_{\mathcal{R}\in\mathcal{S}} I^P_{\mathcal{R}} \leq \beta|\mathcal{S}|/2\right\} \leq e^{-\beta|\mathcal{S}|/8} = O(n^{-3}).$$

According to [34], this implies that, in the actual binomial case,

$$\Pr\left\{\sum_{\mathcal{R}\in\mathcal{S}} I_{\mathcal{R}} \leq \beta|\mathcal{S}|/2\right\} \leq (|n - \mathcal{U}(t-1)|)\cdot O(n^{-3}) = O(n^{-2}).$$

In other words, with probability $1 - O(n^{-2})$, there are at least a $\frac{\beta}{2\phi}$ fraction of all the r-cells in which nodes belonging to $\mathcal{U}(t-1)$ can not receive $\mathbf{m}$ in time slot t. Meanwhile, by some standard Chernoff bound argument, we can show that the number of such nodes is at least $\frac{\beta}{4\phi}|\mathcal{U}(t-1)|$ with probability $1 - O(n^{-2})$. By setting $\chi = \frac{\beta}{4\phi} < 1$, we get (6). The proof is accomplished. □

Finally, according to Lemma 7, we can apply the union bound to see that, starting from time t_0 when $|\mathcal{U}(t_0)| \geq n/2$, the number of time slots required so that $|\mathcal{U}(t)|$ becomes less than $\sqrt{n}$ is at least $\log_\chi \frac{\sqrt{n}}{n/2} = \Omega(\log n)$ with probability $1 - O(n^{-1})$, which implies Theorem 3 directly.

6. OPTIMALITY OF R²: PROOFS OF THEOREMS 4 AND 5

In this section, we present the proofs of Theorem 4 and Theorem 5. Note that the broadcast latency under the R^2 protocol has the same order as the latency for broadcasting a sub-message as the number of sub-messages is a constant. In the following, we focus on the latency for broadcasting a sub-message.

Let $\vec{\mu}$ be an arbitrary non-zero vector in $\mathbb{F}_2^\tau$. We will analyze the knowledge spreading process of $\vec{\mu}$ to derive the time T it takes for all nodes to know about $\vec{\mu}$ with probability $1 - O(n^{-2})$, which implies that T is a upper bound on the latency for broadcasting a sub-message according to the theory of projection analysis.

6.1 Basic Property of R²

We first introduce an elementary property of the R^2 protocol.

Lemma 8. *Let $\mathcal{W}$ be a set of $\vec{\mu}$-informed nodes lying in a same r-cell at the beginning of time slot t. Then there exists some constant $\theta > 1$ which does not depend on c, such that the probability that there exists some node in $\mathcal{W}$ making all its neighboring nodes $\vec{\mu}$-informed during time slot t is at least $1 - \theta^{-|\mathcal{W}|}$.*

Proof. According to the protocol model, we can show by some simple geometric argument that all nodes that can cause interference with nodes in $\mathcal{W}$ must be contained by at most $\lceil \frac{2+\Delta}{a} \rceil^2 + 1 \triangleq \phi_2$ r-cells. This implies that the number of such nodes is at most $\phi_2\eta_2 r^2$. Therefore, in each mini-slot of t, the probability that there exists some node in $\mathcal{W}$ that can successfully broadcast a coded mini-message to all its neighboring nodes is at least

$$|\mathcal{W}|\frac{1}{\gamma}\left(1 - \frac{1}{\gamma}\right)^{\phi_2\eta_2 r^2} \geq \left(\frac{1}{4}\right)^{\phi_2\eta_2} \frac{|\mathcal{W}|}{\gamma}$$

since $r^2 = \gamma$, which implies that the probability that there exists some node in $\mathcal{W}$ that can make all its neighboring nodes $\vec{\mu}$-informed is at least $\left(\frac{1}{2}\right)\left(\frac{1}{4}\right)^{\phi_2\eta_2}\frac{|\mathcal{W}|}{\gamma}$ according to Lemma 6. As there exists γ mini-slots in time slot t, therefore, the probability that no node in $\mathcal{W}$ can make all its neighboring nodes $\vec{\mu}$-informed during time slot t is at most

$$\left(1 - \left(\frac{1}{2}\right)\left(\frac{1}{4}\right)^{\phi_2\eta_2}\frac{|\mathcal{W}|}{\gamma}\right)^\gamma \leq e^{-\left(\frac{1}{2}\right)\left(\frac{1}{4}\right)^{\phi_2\eta_2}|\mathcal{W}|} = \theta^{-|\mathcal{W}|},$$

where $\theta \triangleq e^{\left(\frac{1}{2}\right)\left(\frac{1}{4}\right)^{\phi_2\eta_2}} > 1$, which is independent with c. The desired result follows immediately. □

Here we present some remarks on the setting of some parameters in R^2. According to the proof of Lemma 8, the scheduling probability is chosen to be $1/\gamma$ which guarantees that the above constant θ is independent with c. Meanwhile, RLNC operations are executed over mini-messages in a sub-message other than all mini-messages of the original message, so that the upper bound of the failure probability required by the means of projection analysis is also independent with c. Both independencies allow us to choose a large enough constant c for the convenience of analysis. From now on, all constants defined in the rest of the section, if not specified explicitly, are independent with c.

6.2 Low Mobility Regime

Now we prove Theorem 4. For an r-cell $\mathcal{R}$, we say it is completely $\vec{\mu}$-informed at time t if all nodes within $\mathcal{R}$ at t are $\vec{\mu}$-informed. Regarding this concept, a straightforward result is as follows.

Lemma 9. *Let $\mathcal{R}^*$ be the r-cell that contains the source s in the first time slot. Then, $\mathcal{R}^*$ is completely $\vec{\mu}$-informed at the end of the first time slot.*

Proof. According to the R^2 protocol, all nodes except for the source s keep silent when $t = 1$. Thus s can succeed to broadcast the whole sub-message $\mathbf{sm}$ to all nodes in $\mathcal{R}$ successfully, making all of them $\vec{\mu}$-informed. □

A key result is that a completely $\vec{\mu}$-informed r-cell can make all its adjacent r-cells completely $\vec{\mu}$-informed in one time slot almost surely. It is formally presented as follows.

Lemma 10. *Let $\mathcal{R}$ be an r-cell which is completely $\vec{\mu}$-informed at the end of time slot t. Then, any r-cell $\mathcal{R}' \in \mathcal{N}(\mathcal{R})$ will become completely $\vec{\mu}$-informed at the end of time slot $t+1$ with probability $1 - O(n^{-4})$.*

PROOF. We will show that $\mathcal{R}'$ fails to become completely $\vec{\mu}$-informed with probability $O(n^{-4})$. Note that the movements of nodes are dependent on condition of $\mathcal{D}_n$, which is assumed to be always true. However, due to the fact that $\Pr\{\cdot|\mathcal{D}_n\} \leq 2\Pr\{\cdot|\mathcal{D}_t\}$ [8, pp. 615], it is sufficient to show the same result on condition that $\mathcal{D}_t$ holds but the movements of nodes in time slot $t+1$ are independent.

For each node v lying in $\mathcal{R}$ during time slot t, let α be the probability that it will remain in $\mathcal{R}$ in $t+1$. By considering the circle centered at v with radius ρ within which v will stay in time slot $t+1$ with probability $\pi/9$, we can show that

$$\alpha \geq \begin{cases} \frac{\pi}{9} \cdot \frac{1}{4} = \frac{\pi}{36} & \text{if } \rho \leq r, \\ \frac{\pi r^2/4}{\pi \rho^2} \cdot \frac{\pi}{9} = \frac{\pi r^2}{36\rho^2} & \text{if } \rho > r. \end{cases}$$

Since $\rho = O(r)$, we conclude that $\alpha \geq \beta$ for some constant $\beta > 0$. Let X be the number of $\vec{\mu}$-informed nodes in $\mathcal{R}$ at the beginning of time slot $t+1$. Thus,

$$\mathbf{E}[X] \geq \beta\eta_1 r^2.$$

By the Chernoff bounds, we further have

$$\Pr\{X \leq \beta\eta_1 r^2/2\} \leq e^{-\beta\eta_1 r^2/8} = O(n^{-4}),$$

as long as c is sufficiently large. Besides, according to Lemma 8, conditioning on $X \geq \beta\eta_1 r^2/2$, the probability that all these X $\vec{\mu}$-informed nodes fail to make their respective neighboring nodes $\vec{\mu}$-informed, which include all the nodes in $\mathcal{R}'$, is at most $\theta^{-\beta\eta_1 r^2/2} = O(n^{-4})$. We thus get the result as desired. □

Starting with one completely $\vec{\mu}$-informed r-cell as given in Lemma 9, we can apply Lemma 10 iteratively and also the union bound over r-cells and time slots to show that at time $\frac{\sqrt{n}}{ar} = O\left(\sqrt{\frac{n}{\log n}}\right)$, all r-cells are completely $\vec{\mu}$-informed with probability $1-O(n^{-2})$. This finishes the proof of Theorem 4.

6.3 High Mobility Regime

Now we prove Theorem 5. Since $\rho = \omega(r)$, we can choose a suitable constant $\frac{1}{4\sqrt{2}} \leq a \leq \frac{1}{2\sqrt{2}}$, so that $\frac{\rho}{ar}$ is an integer and thus each ρ-cell is composed exactly by $\frac{\rho}{ar} \times \frac{\rho}{ar}$ r-cells. In the following, we focus on the evolution process of the number of $\vec{\mu}$-informed nodes in each ρ-cell.

For a ρ-cell $\mathcal{C}$, we denote the set of $\mathcal{C}$ and all its adjacent ρ-cells as $\mathcal{N}(\mathcal{C})$ and the number of $\vec{\mu}$-informed nodes in $\mathcal{C}$ at the end of time slot t as $K_{\mathcal{C}}(t)$. We then have the following immediate result.

LEMMA 11. *Let $\mathcal{C}^*$ be the ρ-cell that contains the source s when $t=1$. Then $K_{\mathcal{C}^*}(1) \geq \eta_1 r^2$.*

PROOF. The proof is very similar to the one of Lemma 9, and is thus omitted. □

Regarding the increment of $K_{\mathcal{C}}(t)$, we have the following result.

LEMMA 12. *Let $\mathcal{C}$ be some ρ-cell such that $K_{\mathcal{C}}(t) \geq \eta_1 r^2$. Then there exists positive constants φ and ψ such that, with probability $1-O(n^{-3})$, every ρ-cell $\mathcal{C}' \in \mathcal{N}(\mathcal{C})$ satisfies*

$$K_{\mathcal{C}'}(t+1) \geq \min\{\varphi r^2 K_{\mathcal{C}}(t), \psi\rho^2\}.$$

PROOF. We distinguish two cases. We first consider the case $\rho = \Omega(r^2)$. Let Z be the number of r-cells that contain one or more $\vec{\mu}$-informed nodes after the move phase of time slot $t+1$. According to [9], there exists some constants φ_1 and ψ_1 such that

$$\Pr\left\{Z \geq \min\left\{\varphi_1 K_{\mathcal{C}}(t), \frac{\psi_1\rho^2}{r^2}\right\}\right\} \geq 1 - O(n^{-3}).$$

Now consider some subset $\mathcal{S}$ of these r-cells in which all the r-cells are isolated, and also

$$|\mathcal{S}| \geq Z/\phi.$$

According to Fact 2, such $\mathcal{S}$ exists. Let Y be the number of r-cells in $\mathcal{S}$ in which all nodes get $\vec{\mu}$-informed at the end of time slot $t+1$. According to Lemma 8, for any r-cell $\mathcal{R} \in \mathcal{S}$, the probability that all nodes in $\mathcal{R}$ get $\vec{\mu}$-informed at the end of time slot $t+1$ is at least $1-\theta^{-1} = \zeta$. This implies that

$$\mathbf{E}[Y] \geq \zeta|\mathcal{S}| \geq \frac{\zeta}{\phi}Z.$$

Finally, thanks to the isolation between r-cells in $\mathcal{S}$, we can apply the Chernoff bounds to show

$$\Pr\left\{Y \leq \frac{\zeta}{2\phi}Z\middle|\mathcal{E}\right\} \leq \exp\left(-\frac{\zeta}{8\phi}Z\middle|\mathcal{E}\right) = O(n^{-3}),$$

where condition $\mathcal{E}$ is defined as the event of $Z \geq \min\{\varphi_1 K_{\mathcal{C}}(t), \psi_1\rho^2/r^2\} = \Omega(r^2) = \Omega(\log n)$. The desired result then follows by the fact that each r-cell contains at least $\eta_1 r^2$ nodes.

Now consider the opposite case $\rho = o(r^2)$. We will show that $K_{\mathcal{C}'}(t+1) \leq \psi\rho^2$ with probability $O(n^{-3})$ for some constant ψ. Without loss of generality, we assume that $K_{\mathcal{C}}(t) = \eta_1 r^2$. Note that the movements of nodes are dependent on condition of $\mathcal{D}_n$. Again, according to the fact that $\Pr\{\cdot|\mathcal{D}_n\} \leq \Pr\{\cdot|\mathcal{D}_t\}$ [8, pp. 615], it is sufficient to show the lemma on condition that $\mathcal{D}_t$ holds but the movements of nodes in time slot $t+1$ are independent.

In the following, we say a node is of type 1 if it is $\vec{\mu}$-informed and was in $\mathcal{C}$ at time t, and of type 2 if otherwise. By a standard Chernoff bound argument, it is straightforward to show that the event that each r-cell contains $\Theta(r^2)$ nodes of type 2 in time slot $t+1$ holds with probability $1-O(n^{-3})$. Thus we can assume that this event holds exactly. Now define $X_{\mathcal{R}}$ to be the number of nodes of type 1 in $\mathcal{R}$ for each r-cell $\mathcal{R}$ in $\mathcal{N}(\mathcal{C})$. Clearly, $X_{\mathcal{R}}$ follows a binomial distribution, *i.e.*,

$$X_{\mathcal{R}} \sim \text{Binom}\left(K_{\mathcal{C}}(t), \frac{(ar)^2}{9\rho^2}\right).$$

Note that $X_{\mathcal{R}}$'s are dependent as their sum over all $\mathcal{R}$ in $\mathcal{N}(\mathcal{C})$ is equal to $K_{\mathcal{C}}(t)$. To deal with the dependency, we apply the Poisson approximation technique again by approximating each $X_{\mathcal{R}}$ with a Poisson random variable $Y_{\mathcal{R}}$ with the same mean value $\lambda = \frac{(ar)^2 K_{\mathcal{C}}(t)}{9\rho^2} = \frac{\eta_1 a^2 r^4}{9\rho^2}$. For r-cell $\mathcal{R}$ in $\mathcal{C}'$, we define $I_{\mathcal{R}}$ ($I^P_{\mathcal{R}}$) to be a 0-1 random variable so that it equals to 0 if if all nodes in $\mathcal{R}$ will become $\vec{\mu}$-informed at the end of time slot $t+1$ or equals to 1 if otherwise in the actual binomial (Poisson) case. To characterize $I^P_{\mathcal{R}}$, we first show that for each $\mathcal{R}$ in $\mathcal{C}'$,

$$\Pr\{Y_{\mathcal{R}} \leq \eta_2 r^2\} \geq 1 - O(n^{-4}). \quad (7)$$

Let Y' be a Poisson random variable with mean value $\lambda' = \frac{\eta_2}{2}r^2$. Since $\rho = \omega(r)$, we have $\lambda < \lambda'$. Using the coupling technique based on the Knuth's algorithm [27], we can easily infer that

$$\Pr\{Y_{\mathcal{R}} > \eta_2 r^2\} \leq \Pr\{Y' > \eta_2 r^2\}.$$

On the other hand,

$$\Pr\{Y' > \eta_2 r^2\} = \Pr\{Y' > 2\lambda'\} \leq \left(\frac{e}{4}\right)^{\lambda'} = O(n^{-4}),$$

where the second step follows due to a tail inequality for Poisson distribution [42] and the last step follows since $r = c\sqrt{\log n}$ for a sufficiently large c. By combining the above two inequalities, we get (7).

Now consider a set $\mathcal{S}$ of isolated r-cells in $\mathcal{C}'$ so that $|\mathcal{S}| \geq \frac{\rho^2}{\phi(ar)^2}$ whose existence is guaranteed by Fact 2. Let $\mathcal{E}_{\mathcal{R}}$ be the event that all r-cells that are not isolated to $\mathcal{R}$ contain no more than $\eta_2 r^2$ nodes of type 1. Thus according to (7),

$$\Pr\{\mathcal{E}_{\mathcal{R}}\} \geq 1 - O(n^{-4}).$$

Similarly to Lemma 8, we can show that there exists some constant $\theta_2 > 1$ such that for any $i = 1, 2, \ldots, \eta_2 r^2$,

$$\Pr\{I_R = 1 | \mathcal{E}_{\mathcal{R}}, Y_{\mathcal{R}} = i\} \leq \theta_2^{-i}.$$

Therefore,

$$\begin{aligned}
&\Pr\{I_{\mathcal{R}}^P = 1\} \\
&\leq \Pr\{\overline{\mathcal{E}_{\mathcal{R}}}\} + \Pr\{I_{\mathcal{R}}^P = 1 | \mathcal{E}_{\mathcal{R}}\} \\
&\leq \Pr\{\overline{\mathcal{E}_{\mathcal{R}}}\} + \sum_{i=0}^{\eta_2 r^2} \Pr\{I_{\mathcal{R}}^P = 1 | \mathcal{E}_{\mathcal{R}}, Y_{\mathcal{R}} = i\} \Pr\{Y_{\mathcal{R}} = i\} \\
&\quad + \Pr\{Y_{\mathcal{R}} > \eta_2 r^2\} \\
&= \sum_{i=0}^{\eta_2 r^2} \theta_2^{-i} \frac{e^{-\lambda}\lambda^i}{i!} + O(n^{-4}) \\
&\leq e^{-(1-1/\theta_2)\lambda} + O(n^{-4}) \\
&\leq 2e^{-(1-1/\theta_2)\lambda},
\end{aligned}$$

where the last step follows since $\rho = \omega(r)$ and thus $\lambda = o(r^2) = o(\log n)$. This implies

$$\mathbf{E}\left[\sum_{\mathcal{R}\in\mathcal{S}} I_{\mathcal{R}}^P\right] \leq 2|\mathcal{S}|e^{-(1-1/\theta_1)\lambda} = o(|\mathcal{S}|),$$

as $\rho = o(r^2)$ and thus $\lambda = \omega(1)$. By the Chernoff bound, we further have

$$\begin{aligned}
&\Pr\left\{\sum_{\mathcal{R}\in\mathcal{S}} I_{\mathcal{R}}^P > |\mathcal{S}|/2 + \mathbf{E}\left[\sum_{\mathcal{R}\in\mathcal{S}} I_{\mathcal{R}}^P\right]\right\} \\
&\leq \exp\left(-\frac{|\mathcal{S}|^2}{12\mathbf{E}\left[\sum_{\mathcal{R}\in\mathcal{S}} I_{\mathcal{R}}^P\right]}\right) \\
&\leq \exp\left(-|\mathcal{S}|e^{(1-1/\theta_2)\lambda}/24\right) \\
&\leq \exp\left(-|\mathcal{S}|(1-1/\theta_2)\lambda/24\right) \\
&\leq \exp\left(-\frac{\rho^2}{24\phi(ar)^2}\left(1-\frac{1}{\theta_2}\right)\lambda\right) \\
&= \exp\left(-\frac{(1-1/\theta_2)\eta_1}{216\phi} r^2\right) \\
&= O(n^{-4}).
\end{aligned}$$

According to [34], this implies that in the actual binomial case we have

$$\begin{aligned}
\Pr\left\{\sum_{\mathcal{R}\in\mathcal{S}} I_{\mathcal{R}} > \left(\frac{1}{2} + o(1)\right)|\mathcal{S}|\right\} \leq & e\sqrt{K_{\mathcal{C}}(t)} \cdot O(n^{-4}) \\
= & O(n^{-3}).
\end{aligned}$$

In other words, there are at least $(1/2 - o(1))|\mathcal{S}| \geq \frac{(1-o(1))\rho^2}{2\phi(ar)^2}$ r-cells, within which all nodes will get $\vec{\mu}$-informed in time slot $t+1$ with probability $1 - O(n^{-3})$. Recall that each r-cell contains $\Theta(r^2)$ nodes of type 2. The proof is thus accomplished. □

The key observation from Lemma 12 is that with probability $1 - O(n^{-3})$, the number of $\vec{\mu}$-informed nodes in any ρ-cell will become at least $\eta_1 r^2$ in $\frac{\sqrt{n}}{2\rho} + 1$ time slots, and then increase to $\psi\rho^2$ in another no more than $\log_{\varphi r^2} \frac{\psi\rho^2}{\eta_1 r^2} = O\left(\frac{\log \rho}{\log r}\right)$ time slots. Therefore, at time $O\left(\frac{\sqrt{n}}{\rho} + \frac{\log \rho}{\log r}\right)$, all ρ-cells will have at least $\psi\rho^2$ $\vec{\mu}$-informed nodes with probability $1 - O(n^{-2})$. Once this event holds, the knowledge spreading process of $\vec{\mu}$ will be finished in one time slot almost surely, as stated in the following result.

Lemma 13. *For any ρ-cell $\mathcal{C}$, if $K_{\mathcal{C}}(t) \geq \psi\rho^2$, then the probability that all nodes in any $\mathcal{C}' \in \mathcal{N}(\mathcal{C})$ will get $\vec{\mu}$-informed at the end of time slot $t+1$ is at least $1 - O(n^{-3})$. Therefore, if $K_{\mathcal{C}}(t) \geq \psi\rho^2$ for every ρ-cell $\mathcal{C}$, then all nodes in the network will become $\vec{\mu}$-informed at time slot $t+1$ with probability $1 - O(n^{-2})$.*

Proof. Consider an arbitrary r-cell $\mathcal{R}$ within $\mathcal{C}'$. Let X be the number of $\vec{\mu}$-informed nodes that move from $\mathcal{C}$ to $\mathcal{C}'$ at the beginning of the time slot $t+1$, and $\mathcal{E}$ be the event that all nodes in $\mathcal{R}$ will get $\vec{\mu}$-informed at the end of $t+1$. It can be easily shown that there exists some constant ψ_1 depending on ψ such that

$$\Pr\{X < \psi_1 r^2\} < O(n^{-4}).$$

According to Lemma 8, we then have

$$\begin{aligned}
\Pr\{\mathcal{E}\} &\geq \Pr\{X \geq \psi_1 r^2\} \Pr\{\mathcal{E} | X \geq \psi_1 r^2\} \\
&\geq (1 - O(n^{-4}))(1 - \theta^{-\psi_1 r^2}) \\
&\geq 1 - O(n^{-4}).
\end{aligned}$$

By the union bound over all r-cells in $\mathcal{N}(\mathcal{C})$, we get the desired result. □

Finally, we have that at time $O\left(\frac{\sqrt{n}}{\rho} + \frac{\log \rho}{\log r}\right)$, all nodes in the network will become $\vec{\mu}$-informed with probability $1 - O(n^{-2})$. Note that when $\rho = \omega(r)$,

$$\frac{\sqrt{n}}{\rho} + \frac{\log \rho}{\log r} = \Theta\left(\frac{\sqrt{n}}{\rho} + \frac{\log n}{\log\log n}\right).$$

This completes the proof of Theorem 5.

7. RELATED WORK

As a fundamental operation in wireless/mobile ad hoc networks, broadcast has been very extensively studied in the literature. Here we mainly present a brief review of recent research much related to our work.

Broadcast in Static Wireless Networks: Starting with the pioneering work [6], there has been a large body of work in designing efficient broadcast protocols with low latency in static radio networks [3, 10, 30, 16, 31, 17]. While the interference of wireless communications is captured by the radio network model, the network topology of a radio network could be arbitrary. In contrast, the topology of a wireless network is usually captured by the proximity of network nodes. With this restriction on the network topology, the broadcast latency could be further improved. Specifically, a series of approximation algorithms have been proposed for wireless ad hoc networks [15, 25, 33, 14]. Most of these works followed a general approach by constructing a broadcast tree as well as a deterministic or randomized conflict-aware transmission scheduling. However, such approach is hardly to be applied in mobile ad hoc networks, since the network topology may change unpredictably over time, and the coordination among nodes are expensive.

Broadcast in MANETs: Some recent works have studied the impact of node mobility on the broadcast latency. They can be grouped into two categories according to whether the transmission radius is above the critical radius or not. In the sparse regime, where the transmission radius is below the critical radius, the broadcast latency was investigated over a rich set of mobility models [26], including the well-known *i.i.d.* mobility model [29], the random walk model [38], *etc.* In the dense regime where the transmission radius is above the critical radius, Clementi *et al.* did several work under an almost the same network model as studied in our work [7, 8, 9]. They established a lower bound on the broadcast latency, and analyzed the broadcast latency under the flooding approach without taking the interference into account. Tang *et al.* established a new lower bound [40], which is complementary to the one given by Clementi *et al.* when nodes move very fast. However, both the work of Clementi *et al.* and Tang *et al.* did not take the interference into account, albeit the lower bounds still work. Chen *et al.* [5] introduced a gossip-like approach, which deals with interference by power controlling; however, it cannot exploit the broadcast nature of wireless medium, resulting in a logarithmic factor gap from the optimum.

RLNC-based Information Dissemination: In the seminar work of Ahlswede, Cai, Li and Yeung [1], network coding was demonstrated to achieve the capacity bounds of different destinations simultaneously. Following this work, Li *et al.* [32] showed that linear coding is enough for multicast. Koetter and Medard [28] then presented polynomial time algorithms for encoding and decoding operations. Further, Ho *et al.* introduced RLNC [23, 24], making network coding easily applicable to networks with unknown topologies and with packet loss. RLNC was first introduced by Deb *et al.* [11, 13, 12] for multi-message information dissemination. Following this, some works focused on the characterization of the performance of RLNC-based protocols in different settings in terms of network parameters [35, 4]. Later, Haeupler [22] proposed a powerful tool named projection analysis for analyzing RLNC-based protocols. While these works investigated the power of RLNC in some sense, they mainly focused on wired networks, which are significantly different from MANETs. In contrast, it remains unknown how to efficiently incorporate RLNC and how RLNC performs in MANETs.

8. CONCLUDING REMARKS

In this paper, we sho that the pure random scheduling is insufficient to achieve optimal broadcast latency when nodes move very fast. We further proposed a novel broadcast protocol R^2, which couples a fine-grained random scheduling with random linear network coding. Theoretical analyses show that R^2 can achieve an optimal broadcast latency in order sense, no matter how fast nodes move around the network.

Several interesting issues regarding the broadcast problem are still open and to be further investigated. In the current work, the message length is assumed to be $\Omega(\log^2 n)$ so that the RLNC overhead due to the coding vectors would not affect the performance asymptotically. Thus the first question is, whether the overhead of R^2 could be reduced or eliminated so that it remains to be optimal even when the message length is $\Theta(\log n)$ which holds in general as the message usually contains the identity of the source node, which costs $\Omega(\log n)$. Another question is to derive similar results under more realistic models, *e.g.*, the physical communication model [20].

9. ACKNOWLEDGMENTS

This work was partially supported by NSFC Grants (No. 61170069, 61373014, 91218302 and 61321491); Science and Technology Pillar Program (Industry) of Jiangsu Province Grant (No. BE2013116); Collaborative Innovation Center of Novel Software Technology and Industrialization; EU FP7 IRSES MobileCloud Project Grant (No. 612212). B. Tang was also partially supported by the program A for outstanding PhD candidates of Nanjing University.

10. REFERENCES

[1] R. Ahlswede, N. Cai, S.-Y. Li, and R. Yeung. Network information flow. *IEEE Trans. Inf. Theory*, 46(4):1204–1216, 2000.

[2] C. Avin and G. Ercal. On the cover time and mixing time of random geometric graphs. *Theor. Comput. Sci.*, 380(1-2):2–22, 2007.

[3] R. Bar-Yehuda, O. Goldreich, and A. Itai. On the time-complexity of broadcast in multi-hop radio networks: An exponential gap between determinism and randomization. *J. Comput. Syst. Sci.*, 45(1):104–126, 1992.

[4] M. Borokhovich, C. Avin, and Z. Lotker. Tight bounds for algebraic gossip on graphs. In *Proc. IEEE ISIT*, pages 1758–1762, 2010.

[5] Y. Chen, S. Shakkottai, and J. G. Andrews. On the role of mobility for multimessage gossip. *IEEE Trans. Inf. Theory*, pages 3953–3970, 2013.

[6] I. Chlamtac and S. Kutten. On broadcasting in radio networks–problem analysis and protocol design. *IEEE Trans. Commun.*, 33(12):1240–1246, 1985.

[7] A. Clementi, A. Monti, F. Pasquale, and R. Silvestri. Information spreading in stationary markovian evolving graph. *IEEE Trans. Parallel Distrib. Syst.*, 22(9):1425–1432, 2011.

[8] A. Clementi, F. Pasquale, and R. Silvestri. Opportunistic manets: mobility can make up for low transmission power. *IEEE/ACM Trans. Netw.*, 21(2):610–620, 2013.

[9] A. Clementi and R. Silvestri. Parsimonious flooding in geometric random-walks. In *Proc. DISC*, 2011.

[10] A. Czumaj and W. Rytter. Broadcasting algorithms in radio networks with unknown topology. *J. Algorithm*, 60(2):115–143, 2006.

[11] S. Deb and M. Medard. Algebraic gossip: a network coding approach to optimal multiple rumor mongering. In *Proc. Allerton*, 2004.

[12] S. Deb, M. Medard, and C. Choute. On random network coding based information dissemination. In *Proc. IEEE ISIT*, pages 278–282, 2005.

[13] S. Deb, M. Medard, and C. Choute. Algebraic gossip: a network coding approach to optimal multiple rumor mongering. *IEEE Trans. Inf. Theory*, 52(6):2486–2507, 2006.

[14] R. Gandhi, Y.-A. Kim, S. Lee, J. Ryu, and P.-J. Wan. Approximation algorithms for data broadcast in wireless networks. *IEEE Trans. Mobile Comput.*, 11(7):1237–1248, 2012.

[15] R. Gandhi, A. Mishra, and S. Parthasarathy. Minimizing broadcast latency and redundancy in ad hoc networks. *IEEE/ACM Trans. Netw.*, 16(4):840–851, 2008.

[16] L. Gąsieniec, D. Peleg, and Q. Xin. Faster communication in known topology radio networks. *Distrib. Comput.*, 19(4):289–300, 2007.

[17] M. Ghaffari, B. Haeupler, and M. Khabbazian. Randomized broadcast in radio networks with collision detection. In *Proc. ACM PODC*, pages 325–334, 2013.

[18] M. Grossglauser and D. Tse. Mobility increases the capacity of ad hoc wireless networks. *IEEE/ACM Trans. Netw.*, 10(4):477–486, 2002.

[19] P. Gupta and P. R. Kumar. Critical power for asymptotic connectivity. In *Proc. CDC*, pages 1106–1110, 1998.

[20] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Trans. Inf. Theory*, 46(2):388–404, 2000.

[21] Z. J. Haas, J. Y. Halpern, and L. Li. Gossip-based ad hoc routing. *IEEE/ACM Trans. Netw.*, 14(3):479–491, 2006.

[22] B. Haeupler. Analyzing network coding gossip made easy. In *Proc. ACM STOC*, pages 293–302, 2011.

[23] T. Ho, R. Koetter, M. Medard, D. Karger, and M. Effros. The benefits of coding over routing in a randomized setting. In *Proc. IEEE ISIT*, page 442, 2003.

[24] T. Ho, M. Medard, R. Koetter, D. Karger, M. Effros, J. Shi, and B. Leong. A random linear network coding approach to multicast. *IEEE Trans. Inf. Theory*, 52(10):4413–4430, 2006.

[25] S.-H. Huang, P.-J. Wan, X. Jia, H. Du, and W. Shang. Minimum-latency broadcast scheduling in wireless ad hoc networks. In *Proc. IEEE INFOCOM*, pages 733–739, 2007.

[26] P. Jacquet, B. Mans, and G. Rodolakis. Information propagation speed in mobile and delay tolerant networks. In *Proc. IEEE INFOCOM*, pages 244–25, 2009.

[27] D. E. Knuth. *The Art of Computer Programming, Volume 2: Seminumerical Algorithms*. Addison Wesley, 1969.

[28] R. Koetter and M. Médard. An algebraic approach to network coding. *IEEE/ACM Trans. Netw.*, 11(5):782–795, 2003.

[29] Z. Kong and E. M. Yeh. On the latency for information dissemination in mobile wireless networks. In *Proc. ACM MOBIHOC*, pages 139–148, 2008.

[30] D. R. Kowalski and A. Pelc. Broadcasting in undirected ad hoc radio networks. *Distrib. Comput.*, 18(1):43–57, 2005.

[31] D. R. Kowalski and A. Pelc. Optimal deterministic broadcasting in known topology radio networks. *Distrib. Comput.*, 19(3):185–195, 2007.

[32] S.-Y. Li, R. Yeung, and N. Cai. Linear network coding. *IEEE Trans. Inf. Theory*, 49(2):371–381, 2003.

[33] R. Mahjourian, F. Chen, R. Tiwari, H. Zhai, and Y. Fang. An approximation algorithm for conflict-aware broadcast scheduling in wireless ad hoc networks. In *Proc. ACM MobiHoc*, pages 331–340, 2008.

[34] M. Mitzenmacher and E. Upfal. *Probability and computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2004.

[35] D. Mosk-Aoyama and D. Shah. Information dissemination via network coding. In *Proc. IEEE ISIT*, pages 1748–1752, 2006.

[36] M. J. Neely and E. Modiano. Capacity and delay tradeoffs for ad hoc mobile networks. *IEEE Trans. Inf. Theory*, 51(6):1917–1937, 2005.

[37] M. Penrose. *Random Geometric Graphs*. Oxford University Press, 2003.

[38] A. Pettarin, A. Pietracaprina, G. Pucci, and E. Upfal. Tight bounds on information dissemination in sparse mobile networks. In *Proc. ACM PODC*, pages 355–362, 2011.

[39] G. Sharma, R. Mazumdar, and N. B. Shroff. Delay and capacity trade-offs in mobile ad hoc networks: A global perspective. *IEEE/ACM Trans. Netw.*, 15(5):981–992, 2007.

[40] B. Tang, B. Ye, S. Guo, S. Lu, and D. O. Wu. Order-optimal information dissemination in manets via network coding. *IEEE Trans. Parallel Distrib. Syst.*, 2013. To appear.

[41] Y.-C. Tseng, S.-Y. Ni, Y.-S. Chen, and J.-P. Sheu. The broadcast storm problem in a mobile ad hoc network. *Wirel. Netw.*, 8(2-3):153–167, 2002.

[42] Wikipedia. Poisson distribution –wikipedia, the free encyclopedia, 2013.

[43] S. Yang and R. W. Yeung. Batched sparse codes, 2012. accepted by IEEE Trans. Inf. Theory.