Parallel Processing Letters Vol. 6 No. 1 (1996) 115–126

GOSSIPING ON A RING WITH RADIOS*

KRISHNAMURTHI RAVISHANKAR
Department of Math and Computer Science
SUNY at New Paltz
New Paltz, NY 12561, USA

SURESH SINGH
Department of Computer Science
University of South Carolina
Columbia, SC 29208, USA
E-mail: singh@cs.sc.edu

Received May 1995
Revised September 1995
Accepted by A. Liestman

ABSTRACT

We study the problem of gossiping where n nodes equipped with radios are placed on a *ring* of circumference L. Each radio has a transmission range of 1 and we assume that simultaneous transmissions by neighboring nodes results in garbled messages. We present an algorithm for gossiping and show that it works in asymptotically optimal time.

1. Introduction

Gossiping is a problem related to information dissemination in communication networks. In gossiping every node in the system has a piece of information (a *secret*) that needs to be communicated to everyone else. For given system configurations we are interested in developing *transmission schedules* that ensure gossiping completes in the shortest possible time.

Gossiping algorithms have been studied for a variety of models. In most models, communication between a pair of nodes is modeled as a *telephone call* during which the two nodes exchange all the information each of them has collected thus far. The edge set determines the set of allowable calls. Let us represent each *possible* call by an edge in a graph whose n vertices represent the n nodes in the system. In [4] the authors present results for gossiping in several graph-based models including the complete graph model. Other researchers have studied the problem of gossiping on hypergraphs [5], grid graphs [3] and trees [11]. Placing restrictions on the allowable *sequence of calls* yield further generalizations called the NODUP model (NO DUPlication) [12] and the NOHO model (No One Hears Own) [9,13].

*Partial funding for this work was provided by the NSF under grant number NCR-9410357.

In this paper we consider a network consisting of nodes that communicate via radio. One important application of gossiping in such networks is the dispersal of positional information for purposes of routing. In our model each node is equipped with a radio transmitter. Thus, when a node transmits, all others within range hear the message (if there is no other simultaneous transmission). Simultaneous transmissions result in collisions where all information contained in the messages is lost. We assume that the nodes are unaware of the system-wide topology and base transmission decisions on knowledge of local topological information only. In particular we assume that nodes know the location of all others within distance 1 (the transmission radius) and that they know their distance from some central spot (point 0 around the circle in a clockwise direction).

We have studied the problem of gossiping in *linear* configurations (i.e. all nodes lie on a line) in [7]. In this paper we extend those results to the case where nodes are placed randomly on a ring. This model is interesting because it is a first step to solving the far more complex problem of gossiping in two dimensions where nodes are placed on a plane (the ring is the simplest model where there are two distinct paths between any pair of nodes).

The algorithm presented in this paper is *asymptotically optimal* in the sense that for a given $\epsilon > 0$, for large enough n (number of nodes), we construct a gossiping algorithm whose gossiping time is within $(1+\epsilon)$ times the lower bound. The algorithm is distributed as well in the sense that every node bases transmission decisions on positional information local to it and on the history of transmissions thus far.We assume that the system topology remains fixed for the duration of the gossip. This is a reasonable assumption because transmission times are much faster as compared to time for nodes to move about.

2. The Model

In this paper, we consider a model in which nodes are placed on a *ring* and every node has a transmitter with a transmission radius 1. Nodes i and j are within range of each other if the distance between them is less than 1. We study only *connected configurations* where each node has a *path* to every other — i.e. each node is reachable from every other via a series of transmissions. Furthermore, we only consider the case where there are *two* node-disjoint paths (excluding the starting and terminating nodes) between all pairs of nodes. Let us call the set of all such connected configurations C_n. The other case, where there are no two node-disjoint paths between any pair of nodes, is equivalent to gossiping on a line and has been studied in [7].

An important aspect of our model deals with the problem of simultaneous transmissions. If a node receives two or more transmissions at the same time then we say that a collision has occurred. This means that the node does not receive the message(s). It is important to note that the transmitting nodes may or may not be aware that a collision has occured. Consider the two collision scenarios illustrated in Fig. 1. In the first kind of collision two or more nodes within transmission range of each other (i.e. the distance between them is less than 1) transmit simultaneously, see Fig. 1(a). In this case all three nodes hear noise. The second kind of collision is possible when three nodes 1, 2 and 3 are placed in such a way that node 3 is within range of both 1 and 2. If 1 and 2 transmit simultaneously, node 3 hears noise, however neither 1 nor 2 is aware that 3 heard noise because they cannot hear each others transmission, see Fig. 1(b).

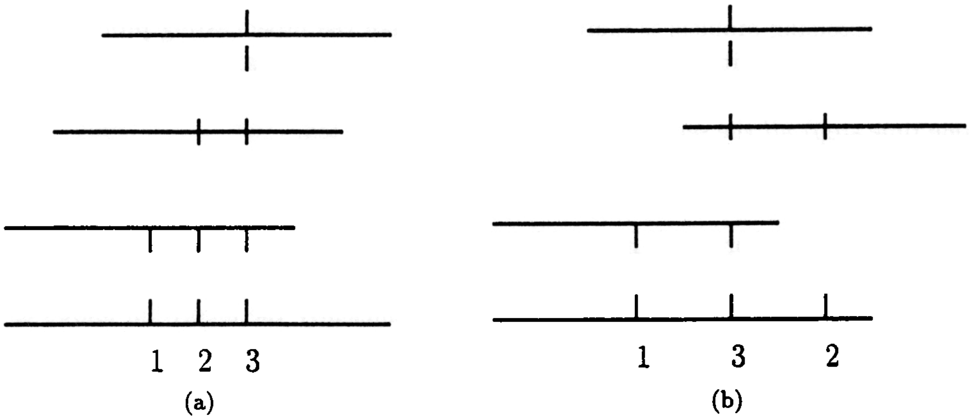

Fig. 1. Node 3 hears noise if 1 and 2 transmit simultaneously.

Finally, we assume that time is slotted and the system is synchronous. A transmission lasts exactly one slot and during a transmission a node can transmit all secrets it has collected thus far. We are interested in developing transmission algorithms that minimize the time taken to gossip.

3. Broadcasting on the Ring

Broadcasting refers to the restricted problem of gossiping when the secrets of only one node have to be disemminated to the rest of the system. We have studied the problem of broadcasting over a *line* [0, L_n] in [6] where we show that the time to broadcast, B_{L_n}, is L_n when,

$$L_n \sim n^\alpha, 0 < \alpha < 1$$

and is αL_n when,

$$L_n \sim \alpha n, \alpha > 0 .$$

For the case of the ring we observe that the time to broadcast is bounded below, for all cases, by,

$$B_{L_n} \geq \frac{L_n}{2}$$

where L_n is the circumference of the ring with n nodes such that there exist *two* node-disjoint paths between any pair of nodes. This is easy to see because,

- the broadcast originating at some node can proceed simultaneously in two directions,

- the maximum distance a broadcast can proceed in either direction during one time step (one transmission) is 1, the range of the radio transmitter.

This lower bound is used later to show that our gossiping algorithm is asymptotically optimal.

4. Gossiping

The gossiping problem is more complicated than the broadcasting problem because of the possibility of simultaneous transmissions by several nodes. As in many papers on gossiping, we assume that during a transmission by a node (which lasts one time step) *all* the secrets it possesses are transmitted.

4.1. Overview

Our algorithm proceeds in two phases. During the first phase, the idea is to allow all nodes to transmit their secrets at least once so that a smaller set of nodes ($w_{n,L}$) collectively know all secrets. To do this efficiently the algorithm allows as many nodes as possible to transmit simultaneously. In the second phase, $w_{n,L} \ll n$ broadcasts are initiated (in both directions). These broadcasts collect all secrets as they proceed. When all these broadcasts terminate, gossiping is complete.

4.2. Gossiping Algorithm

Given a *connected* configuration of nodes on the ring of circumference L, let us divide it into clusters of size $1+1/m$. Each cluster is further divided into smaller subclusters of size $1/m$ (for some m to be chosen appropriately, as discussed in the next section), see Fig. 2. Let these subclusters be labeled $c_{11}, c_{12}, \ldots, c_{1,m+1}, c_{21}, \ldots,$ etc. Let us assume that the kth cluster contains $q \leq m+1$ subclusters, i.e.

$$mL = (m+1)(k-1) + q\,.$$

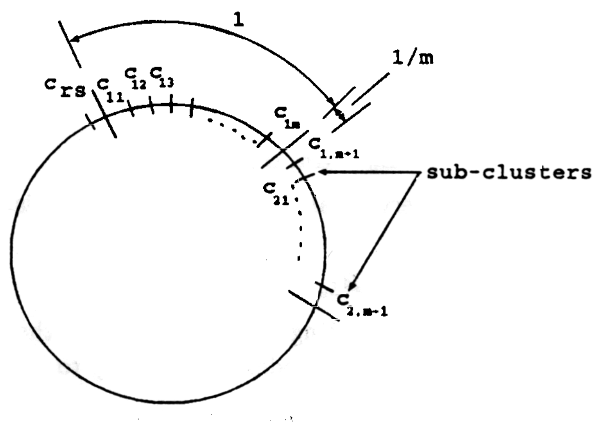

Fig. 2. Subcluster structure on the ring.

If $q = m + 1$ then the subclusters of the kth clusters are $c_{k1}, \ldots, c_{k,m+1}$ (see Fig. 3(a)). If $q < m + 1$ then the subclusters are numbered, $c_{k,m-q+1}, \ldots, c_{k,m+1}$, (see Fig. 3(b)).

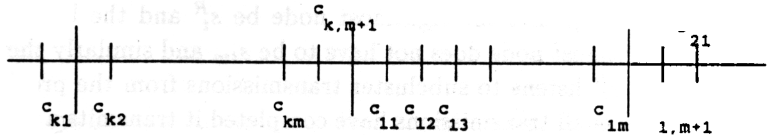

(a) $q = m + 1$

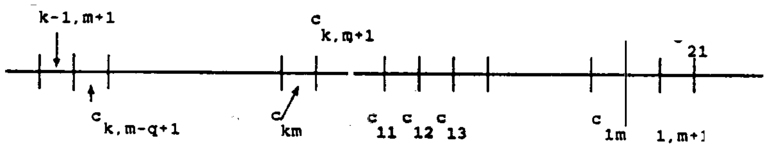

(b) $q < m + 1$

Fig. 3. Naming clusters.

Algorithm 1

Phase 1: Some steps in this phase of the algorithm are similar to Algorithm 2 presented in our earlier paper [7]. This phase achieves information *collection* on a local level (we assume that all nodes begin this phase of the algorithm simultaneously). That is, at the end of phase 1, in any arc of length l we can identify $l + c$ (for some constant c) nodes who collectively know the secrets of all nodes that lie within this arc.

1. The algorithm proceeds in $m + 1$ stages. During each stage j, all nodes in subcluster c_i, $\forall i$, transmit their secrets in a *left to right order* (i.e. in clockwise order). The rightmost node r_{ij} of subcluster c_{ij} (see Fig. 4) hears all transmissions made by nodes in c_{ij} correctly (note that sub-clusters may be empty in which case there is no node r_{ij} and this step is trivial). Observe that clusters c_{ij} and $c_{i+1,j}$ are distance $1 + 1/m$ apart and thus there is no interference.
   Transmissions within cluster c_{ij} are initiated when node r_{ij} transmits a SIGNAL. Node r_{ij} knows the total number of nodes to its right within distance 1. It counts the number of transmissions (by listening to successful transmissions and collisions) that occur in subclusters $c_{i+1,1}, \ldots, c_{i+1,j-1}$. When the number of transmissions heard is equal to the number of nodes in $c_{i+1,1}, \ldots, c_{i+1,j-1}$ (which are within distance 1 of r_{ij}) then it sends a signal. Note that nodes r_{kj} in the kth cluster count the transmission in subclusters of cluster c_1.
2. At the end of the above step, the secrets of each subcluster c_{ij} are contained in node r_{ij}. Next, these secrets are disseminated to all nodes to the *right* of r_{ij}s within distance 1.

Let us rename the nodes r_{ij}. If r_{ij} lies within arc $[t-1, t]$ and is the vth such node in that arc, name it s_{tv}. Thus all nodes r_{1j} in $[0,1]$ are renamed s_{1v}, $1 \leq v \leq m$. Nodes in [1, 2] are renamed s_{2v}, $1 \leq v \leq m$. Node $r_{1,m+1}$ is renamed s_{21}, r_{21} is renamed s_{22}, etc.
Consider arc $[t-1, t]$. Let the rightmost node be s_t^R and the leftmost be s_t^L (note that the rightmost node does not have to be s_{tm} and similarly the leftmost need not be s_{t1}). s_t^R listens to subcluster transmissions from the previous step. When it detects that all transmissions have completed it transmits a SIGNAL to s_t^L in an *odd* time step if t is odd or in an *even* time step if t is even. s_t^R detects completion of transmissions by a method similar to the one used in step 1 above. After $s_t^L = s_{tj}$ receives the signal, it waits until step 1 transmissions to its left have completed. It then transmits all its accumulated secrets in the next odd or even time step (depending on whether t is odd or even). Its transmission is received correctly by at least $s_{t,j+1}$. $s_{t,j+1}$ transmits all accumulated secrets in the next even or odd numbered time step, etc. At the end of this process, when s_t^R hears the transmission from its left, it transmits *all* accumulated secrets in the next (odd/even) time step. Thus, all secrets of nodes in interval $[t-1, t]$ are made known to nodes in an interval of length at least 1 around s_t^R.
Two problems arise in this algorithm:

(a) if k is odd (in Fig. 2) then, according to the algorithm above, step 2 transmissions in [0, 1] and $[L-1, L]$ could occur at the same time resulting in collisions and a violation of the stated result. To avoid this problem, we ensure that step 2 transmissions do not begin until all step 1 transmissions in [0, 1] have completed.

(b) if $\lfloor L \rfloor < L$ (i.e. L is not an integer) then even if $[\lfloor L \rfloor, L]$ is an even numbered arc, there is a potential for collisions in step 2 transmissions between transmissions from [0, 1] and $[\lfloor L \rfloor - 1, \lfloor L \rfloor]$. To alleviate this problem it is necessary to delay step 2 transmissions in $[\lfloor L \rfloor - 1, \lfloor L \rfloor]$ until all step 2 transmissions in [0, 1] have been completed.

In either of these two cases, this step of the algorithm requires at most m more steps to complete.

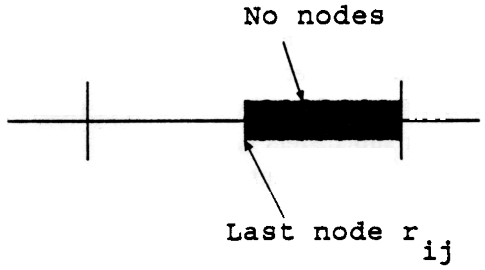

Fig. 4. Rightmost node in c_{ij}.

Phase 2: Let us divide the ring into w arcs of length $l = L/w$ each. These arcs are $[0, l]$, $[l, 2l], \ldots, [L - l, L]$. Identify w nodes $b_1, b_2, \ldots, b_w$ such that b_i is a node closest to the left endpoint of the ith arc. All nodes between b_i and b_{i+1} (inclusive) form the ith *sector* S_i.

Node b_i listens to transmissions on either side within distance 1. When it detects completion of phase 1 transmissions on both sides, it initiates a broadcast (i.e. it transmits all secrets it possesses) to its *left* (i.e. anti-clockwise) and to its *right* (i.e. clockwise). These broadcasts "collect" all information as they proceed. These broadcasts proceed around the ring. When b_i's *anti-clockwise* broadcast meets b_{i-1}'s *clockwise* broadcast in a sector other than sector S_{i-1}, they destroy each other (i.e. they are not forwarded any more).

At any time there may be as many as w broadcasts proceeding in a clockwise direction and an equal number proceeding in the anti-clockwise direction. When any two broadcasts moving in opposite directions cross one another it results in a collision and an added time complexity of at most c_1 time steps for each of the two broadcasts because one of the colliding broadcasts will have to 'wait' while the other passes by.

4.3. Complexity

The complexity of phase 1 can be divided into two parts — the time to complete local transmissions within each subcluster and the time to disemminate the secrets of each cluster to nodes within distance 1 of the rightmost node in the cluster. Let n_{ij} denote the number of nodes within subcluster c_{ij}. A bound on the time to complete local transmissions within subclusters is,

$$\Delta_{m,n} = (m + 1)M_{m,n}$$

where,

$$M_{m,n} = \max_{i,j} n_{ij}$$

The time to disemminate secrets within distance 1 of nodes s_t^R is $c_2 m$ for some constant c_2. Thus, the complexity for phase 1 is,

$$\Delta_{m,n} + c_2 m$$

The complexity of phase 2 is the time taken for each of the $2w$ broadcasts to end. This time is bounded by,

$$\Gamma_{w,L} \leq B_L + wc_1 + B_l$$

where B_L is the time to broadcast over the ring and B_l is the time to broadcast over a sector of length l. c_1 is the cost when two broadcasts moving in opposite directions collide.

5. Asymptotic Results

Theorem 1. *(Lower Bound for time needed to Gossip) Let G_n be the random variable (over the probability space of connected configurations) which denotes the minimum number of time steps required to gossip when there are n nodes in the system.*

(a) $\lim_{n\to\infty} P_n\left(G_n > \frac{n}{L_n} + B_L - \log n\right) = 1$ if $\lim_{n\to\infty} \frac{L_n}{n^\alpha} = c_\alpha > 0\,, 0 < \alpha < 1$

(b) $G_n \geq B_L$ in all cases

where B_n is the time required to broadcast.

Proof. Part (a): For any sequence of transmissions, observe that at the end of $n/L_n - 1$th step at least L_n nodes have not transmitted their information. At least one of these L_n nodes lies within a distance $\log n$ to the right of 0 with a probability approaching 1 (please see [7] for a detailed proof of this statement). The minimum number of steps required to get this node's secret all the way across the configuration is then $B_L - \log n$. This yields a total number of steps equal to $n/L + B_L - \log n$.

Part (b): Proof is trivial. □

Lemma 1. *Let $L_n/n^\alpha \to c > 0$ as $n \to \infty$, $0 \leq \alpha < 1$. Then,*

$$\tilde{P}_n(C_n) \to 1 \text{ as } n \to \infty$$

where C_n is the set of all connected configurations over the ring. If we identify the point L_n with 0, $\tilde{P}_n$ denotes the product measure on $[0, L_n)^n$ with uniform marginals (i.e. all possible placements of n nodes on the ring without regard for connectedness).

The proof of this lemma is easy and is similar to Lemma 1 in [7].

Theorem 2. *If $L \sim n^\alpha$, $0 \leq \alpha < 1$, then $\forall \epsilon > 0$,*

$$\lim_{m\to\infty} \lim_{n\to\infty} P_n\left(\frac{\Delta_{m,n} + c_2 m + B_L + B_l + c_1 w}{n/L_n + B_L} > 1 + \epsilon\right) \to 0$$

Proof. Since $\frac{B_L}{L_n} \xrightarrow{P} 1$, $\frac{B_l}{l_n} \xrightarrow{P} 1$ and w is a fixed constant independent of n it is sufficient to show that,

$$\lim_{m\to\infty} \lim_{n\to\infty} P_n\left(\frac{\Delta_{m,n} + c_2 m + B_L}{n/L_n + B_L} > 1 + \epsilon\right) \to 0$$

The proof of this limit follows from Theorem 2 in [7]. □

Theorem 3. *If $L_n \sim \alpha n$, $0 \leq \alpha \leq 1$, then $\forall \epsilon > 0$,*

$$\lim_{n\to\infty} P_n\left(\frac{\Delta_{1,n}+c_2+B_{L_n}}{B_{L_n}} > 1+\epsilon\right) \longrightarrow 0$$

Proof. The ring T (of circumference L), can be coordinatized by specifying the angle θ, $0 \leq \theta < 2\pi$ that the radius vector makes with a fixed direction as shown below.

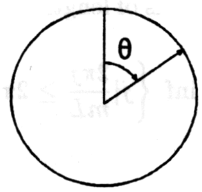

A configuration of n nodes in T is an element of the n-torus T^n, which is coordinatized by an ordered n-tuple $(\theta_1, \theta_2, \ldots, \theta_n)$, $\theta_i \in T$. We now define a family of bijections from T^n into $[0, L)^n$ parameterized by $\mathcal{R}$ (set of real numbers).

Let,

$$y \in \mathcal{R} =, U_y : T^n \longrightarrow [0, L)^n$$

where

$$U_y((\theta_1, \theta_2, \ldots, \theta_n)) = \left(\left(\frac{\theta_1}{2\pi}L\right) - y(\text{mod } L), \ldots, \left(\frac{\theta_n}{2\pi}L - y(\text{mod } L)\right)\right)$$

where $\mathcal{B}$ denotes the Borel sigma algebra, [1] (we define $y(\text{mod } L) = z$ where $y = kL + z$ for some integer k). Let $\tilde{P_n}$ and $\tilde{P'_n}$ denote the product measure with uniform marginals on T^n and $[0, L)^n$ respectively. Then $\forall y \in \mathcal{R}$, U_y is a measure preserving map from,

$$(T^n, \mathcal{B}(T^n), \tilde{P_n}) \text{ into } ([0, L)^n, \mathcal{B}([0, L)^n), \tilde{P'_n})$$

Let $C'_n \subseteq [0, L)^n$ denote the set of connected configurations in $[0, L)^n$. These are the configurations with the property that the distance between successive nodes is less than or equal to one. While the definition of connection in T^n is intuitively clear, the formal definition is more complicated since there are two ways to measure the distance between two points on T.

Let $w \in T^n$. We relabel $U_0(w)$ such that $x_1 < x_2 < \cdots < x_n$ (where the x_i denote the node positions). Now $U_0(w) \in C'_n$ or $U_0(w) \notin C'_n$. If $U_0(w) \notin C'_n$, let,

$$d = \min_{2 \leq i \leq n} \{i | x_i - x_{i-1} > 1\}$$

clearly $x_d \in [k-1, k]$ for some k, $2 \leq k \leq L$. If $U_{k-1}(w) \in C'_n$ we say that w is a connected configuration in T^n.

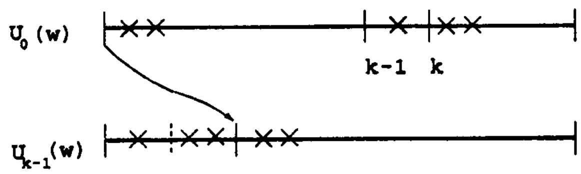

Definition. $C_n \subseteq T^n$, $C_n = \{w \in T^n | U_m(w) \in C'_n\}$ for some m, $0 \leq m \leq \lceil L \rceil - 1$.

Partition T and $[0, L)$ into subintervals of length $1/m$. Let $w \in T^n$ and $w' \in [0, L)^n$. Let,

$$l = \inf \left\{ j | \frac{2\pi j}{mL} \geq 2\pi \right\}$$

and $w = (\theta_1, \theta_2, \ldots, \theta_n)$,

$$\eta_k(w) = \begin{cases} \# \left\{ i | \theta_i \in \left[\frac{2\pi}{mL}(k-1), \frac{2\pi}{mL} k \right) \right\} & \text{if } 1 \leq k \leq l-1 \\ \# \left\{ i | \theta_i \in \left[\frac{2\pi}{mL}(l-1), 2\pi \right) \right\} & \text{if } k = l \end{cases}$$

where $\#\{\ldots\}$ denotes the number of elements in the set. Then,

$$M_{m,n}(w) = \max_{1 \leq k \leq l} \eta_k(w)$$

$$\Delta_{m,n} = (m+1) M_{m,n} \,.$$

Similarly let $w' = (x_1, x_2, \ldots, x_n)$

$$\eta'_k = \begin{cases} \# \left\{ i | x_i \in \left[\frac{k-1}{m}, \frac{k}{m} \right) \right\} & \text{if } 1 \leq k \leq l-1 \\ \# \left\{ i | x_i \in \left[\frac{l-1}{m}, L \right) \right\} & \text{if } k = l \,. \end{cases}$$

Then,

$$M'_{m,n}(w') = \max_{1 \leq k \leq l} \eta'_k(w)$$

and,

$$\Delta'_{m,n} = (m+1) M'_{m,n}$$

η_k and η'_k are the number of nodes in the kth subintervals in T and $[0, L)$ respectively. Let,

$$A_p = \{w \in C_n | U_p(w) \in C'_n\} \,.$$

Clearly,

$$\bigcup_{p=0}^{\lceil L \rceil - 1} A_p = C_n \,.$$

Let P_n and P'_n denote the restriction of $\tilde{P}_n$ and $\tilde{P'}_n$ to C_n and C'_n respectively. If $w \in A_p$ then $w' = U_p(w) \in C'_n$. We note that,

$$\Delta_{m,n}(w) = \Delta'_{m,n}(w')$$

$$P_n(\Delta_{m,n} \geq \lfloor \delta L_n \rfloor) \leq \sum_{p=0}^{\lceil L \rceil - 1} P_n(\Delta_{m,n} \geq \lfloor \delta L_n \rfloor, A_p)$$

$$= \sum_p P_n(\Delta'_{m,n}(U_p(w)) \geq \lfloor \delta L_n \rfloor, A_p).$$

Now using the fact that U_p is a bijection and is measure preserving we have,

$$P_n(\Delta'_{m,n}(U_p(w)) \geq \lfloor \delta L_n \rfloor, A_p) = P'_n(\Delta'_{m,n} \geq \lfloor \delta L_n \rfloor, U_p A_p).$$

Therefore we have,

$$P_n(\Delta_{m,n} \geq \lfloor \delta L_n \rfloor) \leq \lceil L_n \rceil P'_n(\Delta'_{m,n} \geq \lfloor \delta L_n \rfloor).$$

Let $m = 1$, we need to estimate,

$$P_n\left(\frac{\Delta_{1,n} + c_2 + B_{L_n}}{B_{L_n}} > 1 + \epsilon\right) = P_n\left(\frac{\Delta_{1,n}}{B_{L_n}} + \frac{c_2}{B_{L_n}} + 1 > 1 + \epsilon\right).$$

Since $\frac{B_{L_n}}{L_n} \xrightarrow{P} A > 0$ it is sufficient to estimate $P_n(\Delta_{1,n} \geq \lfloor \delta L_n \rfloor)$. We have shown that,

$$P_n(\Delta_{1,n} \geq \lfloor \delta L_n \rfloor) \leq \lceil L_n \rceil P'_n(\Delta'_{1,n} \geq \lfloor \delta L_n \rfloor)$$

From theorem 3 in [7] we have,

$$P'_n(\Delta'_{1,n} \geq \lfloor \delta L_n \rfloor) < e^{-C\lfloor L_n \rfloor}.$$

Therefore,

$$\lim_{n \longrightarrow \infty} P_n(\Delta_{1,n} \geq \lfloor \delta L_n \rfloor) = 0$$

This proves that, $\forall \epsilon > 0$,

$$P_n\left(\frac{\Delta_{1,n} + c_2 + B_{L_n}}{B_{L_n}} > 1 + \epsilon\right) \longrightarrow 0$$

□

6. Conclusions

We have studied the problem of gossiping in a system where nodes, equipped with radio transmitters, are placed on a ring. The gossiping algorithm presented is asymptotically optimal. This work builds upon earlier results for the case of broadcasting and gossiping on a line presented in [6–8]. The next step, we believe, is to study the much harder problem of gossiping in two dimensions. We believe the results reported in this paper will prove very useful in solving some models in 2D (e.g. gossiping on a lattice) where information can travel between two nodes via multiple paths (note that in the line, information can travel along one path only, in the ring there are two paths).

References

1. L. Breiman, *Probability* (Addison-Wesley, Reading, Massachusetts, 1968).
2. J. D. Deuschel and D. W. Stroock, *Large Deviations* (Academic Press, San Diego, CA, 1989), pp. 4–5.
3. A. Farley and A. Proskurowski, Gossiping in Grid Graphs, *J. Combin. Inform. Systems Sci.* **5** (1980) 161–172.
4. S. T. Hedetniemi, S. T. Hedetniemi and A. L. Liestman, A Survey of Gossiping and Broadcasting in Communication Networks, *Networks* **18** (1988) 319–349.
5. D. J. Kleitman and J. B. Shearer, Further Gossip Problems, *Discr. Math.* **30** (1980) 151–156.
6. K. Ravishankar and S. Singh, "Broadcasting on $[0, L]$", *Discrete Applied Math.* **53** (1994) 299–319.
7. K. Ravishankar and S. Singh, "Asymptotically optimal gossiping in radio networks", *Discrete Applied Math.* **61** (1995) 61–82.
8. K. Ravishankar and S. Singh, "Central Limit Theorem for Time to Broadcast in Radio Networks", *Probability in the Engineering and Informational Sciences*, (to appear).
9. Á. Seress, Quick Gossiping Without Duplicate Transmissions, *Graphs and Combinatorics* **2** (1986) 363–383.
10. S. Singh and M. A. Sridhar, "Gossiping on Interval Graphs", *7th Int. Parallel Processing Symp.*, April 13–15, Newport Beach, CA, (1993), pp. 338–343.
11. P. J. Slater, E. Cockayne and S. T. Hedetniemi, Information Dissemination in Trees, *SIAM J. Computing* **10** (1981) 692–701.
12. D. B. West, Gossiping without Duplicate Transmissions, *SIAM J. Alg. Disc. Meth.* **3** (1982) 418–419.
13. D. B. West, A class of solutions to the gossip problem part I, *Discrete Math.* **39** (1982) 307–326.