Randomized Broadcast in Networks*

Uriel Feige[†], David Peleg[†‡], Prabhakar Raghavan[§], and Eli Upfal[¶]

ABSTRACT

In this paper we study the rate at which a rumor spreads through an undirected graph. This study has two important applications in distributed computation: in simple, robust and efficient broadcast protocols, and in the maintenance of replicated databases.

Key Words: rumor propagation, randomized algorithms, distributed computation, random graphs.

1. INTRODUCTION

Let $G = (V, E)$ be a connected, undirected graph on N vertices. One vertex initially knows of a "rumor" that has to be conveyed to every other vertex in V. The rumor is propagated as follows: at each step, every vertex that knows of the rumor chooses one of its neighbors in G uniformly at random, and informs it of the rumor. How many steps elapse before every vertex knows the rumor? The answer clearly depends on the nature of G; for instance, if G were the complete graph on N vertices, K_N, it is well known [9, 13] that $\Theta(\log N)$ steps suffice almost surely.

Consider the standard model of point to point communication networks, described by a connected, undirected graph: the vertices represent the processors

* A preliminary version of this paper will be presented at the SIGAL International Symposium on Algorithms, Tokyo, Japan, August 1990.

† The Weizmann Institute of Science, Rehovot 76100, Israel.

‡ Supported in part by an Allon Fellowship, a Bantrell Fellowship, and a Haas Career Development Award.

§ IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, USA. A portion of this work was done while the author was visiting the Weizmann Institute of Science.

¶ IBM Almaden Research Center, San Jose, CA, and Department of Applied Mathematics, The Weizmann Institute of Science, Rehovot, Israel. Work at the Weizmann Institute supported in part by a Bat-Sheva de Rothschild Award and a Revson Career Development Award.

Random Structures and Algorithms, Vol. 1, No. 4 (1990)
 CCC 1042-9832/90/040447-14$04.00

of the network and the edges represent bidirectional communication channels between the vertices. The study of rumor propagation has at least two applications to distributed computing in such networks. The first is in algorithms for *broadcasting*: a single processor wishes to broadcast a piece of information to all the other processors in the network (see [11] for a comprehensive survey). There are at least three advantages to a randomized broadcast algorithm such as the one we have described:

1. Simplicity—the entire algorithm is simple and local, with no need to know the current global topology and no need to remember whether the rumor has already been sent to a particular neighbor. Despite this simplicity, the algorithm achieves fast broadcast, as we shall see.
2. Scalability—the algorithm is independent of the size of the network; it does not grow more complex as the network grows.
3. Robustness—randomized broadcast is robust in that it works well even in the face of link/node failures in the network. Let us again consider broadcast on the complete graph K_N; it is possible to devise deterministic broadcast algorithms that achieve $O(\log N)$-step broadcast in the absence of faults, just as our random broadcast does (almost surely). Now let $f < cN$ be a positive integer for a constant $c < 1/3$; suppose that an adversary were allowed to choose f links in the network to "break." For any deterministic algorithm, an adversary can select the link failures such that the algorithm takes $\Omega(f + \log N)$ steps. To see this, observe that in any deterministic algorithm, the vertex that initially knows the rumor will attempt to transmit the rumor to its N-1 neighbors in a fixed sequence. The adversary breaks the links to the first f neighbors in this sequence. In contrast, we prove in Theorem 2.4 that randomized broadcast achieves $O(\log N)$ steps almost surely no matter how the adversary chooses the link failures.

A second application comes from the maintenance of replicated databases, for instance, in name servers in a large corporate network [6]. There are updates injected at various nodes, and these updates must propagate to all the nodes in the network. At each step, a processor and its neighbor check whether their copies of the database agree, and if not make the necessary reconciliation. The goal is that eventually all copies of the database converge to the same contents. The reader is referred to [6] for details.

The spread of rumors in a network bears a superficial resemblance to another stochastic process that has been studied much by computer scientists: the random walk [1]. In rumor propagation, every node that has already seen the rumor is broadcasting it at every subsequent step, whereas in the random walk propagation occurs from exactly one node at each step. The tools used in random walks, such as algebraic methods [2] or physical analogies [5] seem to be inappropriate to our study here. Also related to our work is the mathematical theory of epidemics [10, 12], although that theory deals with variants of rumor broadcast on the complete graph (rather than general graphs as we consider here).

In Section 2 we give some fundamental theorems governing the rate of propagation of rumors in graphs. While these theorems yield tight results for many classes of graphs, it is sometimes necessary to examine the structure of the

graph more closely to precisely determine the rate of rumor propagation. We illustrate this in Section 3, where we show that a rumor reaches all vertices of a hypercube in $\Theta(\log N)$ steps almost surely; elementary analysis only yields $O(\log^2 N)$ steps. The hypercube also illustrates the difficulties introduced in the analysis due to the statistical correlation of several copies of a rumor repeatedly running into each other. In Section 4 we study the rate of rumor propagation on random graphs with various densities.

2. DEFINITIONS AND GENERAL RESULTS

Let $deg(v)$ denote the degree of a vertex v, and let $deg(G)$ be $\max_{v \in G} deg(v)$. The length of the shortest path connecting two vertices v and u is denoted by $dist(v, u)$, and the length of the shortest path connecting a vertex v to some member of a set of vertices C is denoted by $dist(u, C)$. The diameter of a network G is denoted by $diam(G)$. All logarithms in the paper are to the base 2.

In this work we analyze the following random broadcast procedure; initially, at least one vertex in V starts off knowing the rumor.

Procedure $\mathcal{RB}$
repeat
  for all $v \in V$, **in parallel do**
    if v has already received the rumor **then**
      v randomly chooses a neighbor uniformly at random, and sends the rumor to it;
end

Given a network G, the number of iterations of the procedure $\mathcal{RB}$ until the rumor reaches all the vertices of G is a random variable that depends on the topology of G.

Definition 1. *We say that $\mathcal{T}(G)$ is the almost sure rumor coverage time of a network G, if after $\mathcal{T}(G)$ iterations of the procedure all vertices in G receive the rumor with probability $1 - 1/N$.*

Theorem 2.1. *For a general network G, with N vertices,*

1. $\log N \leq \mathcal{T}(G) \leq 12N \log N$,
2. *There are networks G for which $\mathcal{T}(G) = \Omega(N \log N)$,*
3. *There are networks G for which $\mathcal{T}(G) = O(\log N)$.*

Proof. Each iteration of the procedure $\mathcal{RB}$ can at most double the number of vertices that received the rumor; thus $\mathcal{T}(G) \geq \log N$.

Let $v = x_0, x_1, \ldots, x_l = u$ denote a shortest path connecting v to u in G. A vertex w is connected to two vertices x_i and x_j on the path only if they are at most one apart on the path (else there is a shorter path). Thus, $\sum_{j=0}^{l-1} deg(x_j) \leq 3N$. The expected number of iterations between the time x_j receive a rumor until the time x_j sends the rumor to x_{j+1} is bounded by $deg(x_j)$. Thus, the expected number of

iterations until u receives the rumor is $3N$, and after $6N$ iterations the probability that u does not receive the rumor is bounded by $1/2$. After $12N \log N$ iterations of the procedure, the probability that there exists a vertex that did not receive the rumor is bounded by $N2^{-2 \log N} \leq 1/N$.

The bounds are tight within a constant factor since $O(\log N)$ iterations suffice for the complete graph [9], and a star graph clearly requires $\Omega(N \log N)$ iterations. ∎

Theorem 2.2. *For every network G*

$$\mathcal{T}(G) = O(deg(G)(diam(G) + \log N)) .$$

Proof. Given a shortest path between two vertices in G, the probability that the rumor fails to traverse that path in $3deg(G)(diam(G) + 2 \log N)$ iterations is bounded (by the Chernoff bound [4]) by $1/N^2$. Thus, the probability that a rumor does not reach all N vertices of the graph within this number of steps is bounded by $1/N$. ∎

Corollary 2.3. *For a bounded degree graph G, $\mathcal{T}(G) \leq O(diam(G))$.*

We now prove the result on the rate of propagation in the complete graph with f link failures.

Theorem 2.4. *Let $f < cN$ be a positive integer for a constant $c < 1/3$. Let G be a graph derived from the complete graph K_N by deleting any f of its edges. The $\mathcal{T}(G)$ is $O(\text{Log } N)$.*

Proof. Let X_t denote the set of nodes of G that have received the rumor before step t; thus $X_1 = 1$. Let $x_t = |X_t|$, and $y_t = N - x_t$. We will show that (1) if $x_t \leq N/3$, then $E[x_{t+1}] - x_t \geq c_1 x_t$, and (2) if $x_t > N/3$, then $E[y_{t+1}] \leq c_2 y_t$, for some constants c_1, c_2 in $(0, 1)$.

We will refer to a node that sends the rumor to another as making a *proposal* to the recipient. We say that a node in X_t is *successful* in step t if it proposes to a node v not in X_t *and* no other node in X_t proposes to v in step t.

Case 1 ($x_t \leq N/3$): Let f_i be the number of faulty edges between $v_i \in X_t$ and nodes not in X_t; thus $\Sigma f_i \leq f$. The probability that a node $v_i \in X_t$ fails to propose to some node in $G - X_t$ using a nonfaulty edge is at most $(x_t + f_i)/N$. The probability that some other node in X_t also proposes to the same node as v_i is at most x_t/N. Thus the probability that v_i is unsuccessful is at most $(2x_t + f_i)/N$; summing over all the $v_i \in X_t$, the expected number of unsuccessful nodes in X_t is (since $x_t \leq N/3$) at most $2x_t/3 + f/N$. Thus the expected number of successful nodes $E[x_{t+1}] - x_t \geq c_1 x_t$ for some constant c_1.

Case 2 ($x_t > N/3$): For a node v_i not in X_t, let f_i be the number of edges between v_i and nodes in X_t; thus $\Sigma f_i \leq f$. The probability that v_i receives no proposal from any of the nodes in X_t is

$$= (1 - 1/n)^{x_t - f_i}$$

$$\leq 1 - \frac{x_t - f_i}{n} + \frac{(x_t - f_i)^2}{2n^2}$$

$$\leq 1 + \frac{f_i}{n} - \frac{x_t}{n} + \frac{x_t^2}{2n^2} .$$

Summing over the y_t nodes v_i in $G - X_t$, we have that

$$E[y_{t+1}] \leq y_t\left(1 - \frac{x_t}{N} + \frac{x_t^2}{2N^2}\right) + \frac{f}{N}$$

A simple calculation now shows that since $x_t/N \in (1/3, 1]$, $E[y_{t+1}] \leq c_2 y_t$ for some c_2. ∎

Note that the proof in fact holds even in the face of an *adaptive adversary*: one who can, at time step t, decide anew which f edges to render faulty (rather than fix them a priori at the beginning of the broadcast process). This demonstrates the resilience of the random broadcast procedure: it will work as claimed in Theorem 2.4 even in the face of dynamically changing faults, provided the number of faults meets the prescribed bound.

3. RANDOM BROADCAST ON THE HYPERCUBE

We now focus on the hypercube, an important network for parallel computation for which the general theorems do not give the correct upper bound. For instance, Theorem 2.2 only gives an upper bound of $O(\log^2 N)$ on the rumor coverage time. Here we determine the rumor coverage time of the hypercube, illustrating some of the difficulties and techniques in proving such a bound. Let $H_n = (V_n, E_n)$ denote the n-dimensional hypercube, where $V_n = \{0, 1\}^n$ and

$$E_n = \{(x, y) \mid x, y \in V_n, x \text{ and } y \text{ differ in exactly one bit}\} .$$

The network has $|V_n| = 2^n = N$ nodes, $|E_n| = n \cdot N/2$ edges, and diameter $n = \log N$.

For the analysis that follows we need the following two estimates, proved in the Appendix.

Lemma 3.1. *For any constants* $0 < \alpha, \beta < 1$ *such that* $2\beta < \alpha$,

$$\left(e\left(\frac{\alpha}{\beta} - 1\right)\right)^{\beta n} < \binom{\alpha n}{\beta n} < \left(4\left(\frac{\alpha}{\beta} - 1\right)\right)^{\beta n}$$

for sufficiently large n.

Lemma 3.2. *For every* $0 < \alpha < 1/2$,

$$\prod_{i=1}^{(1/2+\alpha)n} \left(1 - \left(1 - \frac{i}{n}\right)^k\right) > 2^{-\theta n/k},$$

where $\theta = (\pi^2/6) \log e \cong 2.37$.

Theorem 3.3. *For the n-dimensional hypercube* H_n, $\mathcal{T}(H_n) = \Theta(n) = \Theta(\log N)$.

Proof. The diameter of the hypercube is n, but the expected number of steps to traverse any single path of length $O(n)$ on the hypercube is $\Theta(n^2)$. To prove the theorem we need to analyze the progress of the rumor along many paths in parallel. The main difficulty in this analysis is that the paths are not disjoint, and thus introduce dependencies to the analysis. We overcome this difficulty by analyzing the progress of the procedure in two phases. The first phase brings the rumor to within a distance of αn from all vertices of the hypercube. The second phase completes the distribution.

Definition 2. *A set of vertices* C, α-approximates *a vertex* t *if* $dist(t, C) \leq \alpha n$. C *is an* α-cover *of the hypercube if it* α-*approximates all vertices of* V_n.

Lemma 3.4. *For any* $0 < \alpha < 1$, *after* $3n/\alpha$ *iterations of the procedure* $\mathcal{RB}$, *with probability* $1 - 2^{-2n}$, *the rumor has reached every vertex of some* α-*cover* C.

Proof. In analyzing the process of α-approximating a vertex t, we concentrate on one path $s = x_0, x_1, \ldots$ *generated by the procedure* $\mathcal{RB}$ in the following way:

1. s is the vertex that initiated the rumor.
2. x_{i+1} is the first vertex satisfying $dist(x_{i+1}, t) < dist(x_i, t)$ that received the rumor from x_i.

The probability that in $3n/\alpha$ iterations of the procedure $\mathcal{RB}$ the path does not α-approximate the vertex t is bounded by the probability that there is a set of αn coordinates that were not hit in $3n/\alpha$ successive trials. This probability is bounded by

$$\binom{n}{\alpha n}(1-\alpha)^{3n/\alpha} < 2^n e^{-3n} \leq 2^{-2n-1}.$$

Thus, the probability that there exists a vertex that is not α-approximated is bounded by $1 - 2^{-n-1}$. ∎

Definition 3. *For any vertex* v *and integer* $1 \leq h \leq n$, *the band* $B(v, h)$ *is the set of vertices at distance exactly* h *from* v, *i.e.*,

$$B(v, h) = \{u | dist(u, v) = h\}.$$

Lemma 3.5. *Given a 1/64-cover C, for any vertex t there exist an integer h and a set of vertices $S_h(t)$ such that the following properties hold.*

1. $31n/64 < h < 33n/64$,
2. $S_h(t) \subseteq B(t, h) \cap C$,
3. $|S_h(t)| \geq 2^{n/4}$, *and*
4. *the distance between every two vertices of $S_h(t)$ is at least $n/8$.*

Proof. Let $S' \subseteq C$ be some minimal collection of vertices of C that 1/64-approximate the vertices of the band $B(t, n/2)$. The number of vertices in this band is

$$\left|B\left(t, \frac{n}{2}\right)\right| = \binom{n}{n/2} \geq \frac{2^n}{\sqrt{2\pi n}}.$$

A vertex may serve as an α-approximation only for vertices at distance αn or less from it. The number of vertices that can have the same 1/64-approximation is bounded, using Lemma 3.1, by

$$Y = \sum_{i=1}^{n/64} \binom{n}{i} \leq \frac{n}{64} \binom{n}{n/64} \leq n(4(64-1))^{n/64} \leq n \cdot 252^{n/64},$$

and

$$|S'| \geq \frac{|B(t, n/2)|}{|Y|}.$$

By the definition of 1/64-approximation, S' contains vertices at distances $n/2 \pm n/64$ from t. Therefore, there is a band $B(t, h)$ for some h in this range such that the set $S'' = S' \cap B(t, h)$ satisfies

$$|S''| = |S' \cap B(t, h)| \geq \frac{64|B(t, n/2)|}{n|Y|}.$$

Using the greedy method, we now pick vertices from S'' into our final set $S_h(t)$, each at distance $n/8$ or more from all previously selected vertices. The number of vertices that each new pick rules out is bounded, using Lemma 3.1 again, by

$$Z = \sum_{i=1}^{n/8} \binom{n}{i} \leq \frac{n}{8} \binom{n}{n/8} \leq n(4(8-1))^{n/8} \leq n \cdot 28^{n/8}.$$

Thus the total number of vertices we can pick into $S_h(t)$ is at least

$$\frac{|S''|}{|Z|} \geq \frac{64|B(t, n/2)|}{n|Y||Z|} \geq \frac{2^{n+6}}{252^{n/64} 28^{n/8} n^3 \sqrt{2\pi n}} \geq 2^{n/4}$$

for sufficiently large n. ∎

In analyzing the process of reaching a vertex t in the second phase of distributing the rumor, we concentrate only on paths that start at vertices in $S_h(t)$. In order to define a sequence of independent events, we further restrict our discussion to paths that were built at a predefined set of time intervals. Formally, consider a set of paths $L(t, \tau, h, k)$, generated by a subprocess of the procedure $\mathcal{RB}$ in the following way:

1. The process starts at time τ after all vertices in $S_h(t)$ have already received the rumor.
2. Each $s \in S_h(t)$ is a start-point of a path in $L(t, \tau, h, k)$.
3. Let $x_0, x_1, \ldots$ denote a path in $L(t, \tau, h, k)$. For every i, $x_i \in B(t, h-i)$, and x_{i+1} is the first vertex in $B(t, h-i-1)$ that received the rumor from x_i in the time interval $[\tau + ik, \ldots \tau + (i+1)k - 1]$. If no vertex in $B(t, h-i-1)$ received the rumor from x_i in that time interval, x_i is the last vertex in this path.

Note that the set of paths $L(t, \tau, h, k)$ are independent from the set of paths $L(t, \tau + k, h, k)$, since for every i, the i-vertices of the paths in the first set, and the i-vertices of paths in the second set, were chosen at different time intervals.

Lemma 3.6. *Assume that all vertices in $S_h(t)$ have received the rumor at time τ. The probability that at least one of the paths in $L(t, \tau, h, 11)$ reaches the vertex t is at least* 2/3.

Proof. Let $m = |S_h(t)|$, and let $s_1, \ldots, s_m$ denote the vertices of $S_h(t)$. Denote by p_d the probability that a path from $s \in S_h(t)$ reaches the band $B(t, d)$. By symmetry, this probability is equal for all $s \in S_h(t)$. Let x_i be the random variable defined by

$$x_i = \begin{cases} 1 & \text{if the path from } s_i \text{ reaches } t, \\ 0 & \text{otherwise.} \end{cases}$$

Let $p = p_0 = Prob(x_i = 1)$, and let $X = \Sigma_{i=1}^{m} x_i$.

Assuming that the path from s_i successfully reaches distance d from t, the probability that it reaches distance $d-1$ is $(1 - (1 - d/n)^{11})$. In other words,

$$p_d = p_{d+1} \cdot \left(1 - \left(1 - \frac{d+1}{n}\right)^{11}\right).$$

Thus, by Lemma 3.2, for sufficiently large n,

$$p \leq \prod_{d=33n/64}^{1} \left(1 - \left(1 - \frac{d}{n}\right)^{11}\right) \geq 2^{-\theta n/11} > 2^{-n/4+4}.$$

By Lemma 3.5, $m > 2^{n/4}$, hence the expected number of successful paths that reach t satisfies $E[X] = mp > 12$.

To bound the probability that at least one of the m paths reaches t we need a bound on $Var[X] = \Sigma_{i,j} E[x_i x_j] - (E[X])^2$.

Let $\mathcal{E}_{i,j}^{1}$ denote the event that s_j reaches t without intersecting s_i's path, and let

$\mathcal{E}^2_{i,j}$ denote the event that s_j intersects s_i's path (and from then on the two paths merge). Then

$$Prob(x_i x_j = 1) = Prob(x_i = 1) \cdot (Prob(\mathcal{E}^1_{i,j}) + Prob(\mathcal{E}^2_{i,j})) .$$

Clearly, $Prob(\mathcal{E}^1_{i,j}) \le p$. It remains to bound $Prob(\mathcal{E}^2_{i,j})$. Let $\mathcal{E}^2_{i,j}(d)$ denote the event: s_j's path intersects s_i's path in band $B(t, d)$ and they did not intersect before. Clearly, $\mathcal{E}^2_{i,j} = \cup_{d>0} \mathcal{E}^2_{i,j}(d)$.

The path from s_j can meet the path from s_i only in vertices v satisfying

$$v \in B(t, d) \cap B(s_j, h - d) \cap B(s_i, h - d) .$$

Furthermore, there is only one set of $h - d$ coordinates that the path for s_2 can cross in its first $h - d$ transitions in order to meet the other path in $B(t, d)$.

The two origins s_i and s_j are at a distance of at least $n/8$ apart, and thus there exists some $d' \le h - n/16$ such that

$$Prob(\mathcal{E}^2_{i,j}(d)) \le \begin{cases} 0 & d > d' , \\ p_d / \binom{h}{d} & d \le d' . \end{cases}$$

Thus, using the bound of Lemma 3.1,

$$\sum_{d=n/16}^{h} Prob(\mathcal{E}^2_{i,j}(d)) \le \frac{33n}{64} Prob\left(\mathcal{E}^2_{i,j}\left(h - \frac{n}{16}\right)\right)$$
$$\le \frac{n}{\binom{31n/64}{n/16}} \le n \cdot 21^{-n/16} \le \frac{p}{100}$$

for sufficiently large n.

For the last segment of the path we need a different analysis. Let R_d denote the ratio between the probabilities of intersection at band d and at band $d + 1$. Using (1),

$$R_d = \frac{Prob(\mathcal{E}^2_{i,j}(d))}{Prob(\mathcal{E}^2_{i,j}(d+1))} = \frac{p_d \binom{h}{d+1}}{p_{d+1} \binom{h}{d}}$$
$$= \left(1 - \left(1 - \frac{d+1}{n}\right)^{11}\right) \frac{h - d}{d + 1} .$$

We are interested in R_d only for $0 \le d \le n/16$. In this range $R_d > 2$. Furthermore, for constant d and sufficiently large n,

$$R_d \ge \frac{11(d+1)}{n} \cdot \frac{h}{d+1} \ge \frac{11(d+1)}{n} \cdot \frac{31n}{64(d+1)} > 5 .$$

Recalling that the probability of intersection at $d = 0$ is at most p,

$$Prob(\mathcal{E}^2_{i,j}) < p \sum_{d=1}^{8} 5^{-d} + \frac{p}{5^8} \sum_{d=1}^{n/16-8} 2^{-d} + \frac{p}{100} < \frac{p}{4} .$$

Thus, $Ex_ix_j \leq p^2(1+1/4)$ when $i \neq j$ and p otherwise.

$$Var[X] = \sum_{i,j} Ex_ix_j - (E[X])^2 \leq m^2p^2\left(1+\frac{1}{4}\right) + mp - (mp)^2 \leq \frac{p^2m^2}{3}$$

for sufficiently large n (recalling that $m \geq 2^{n/4}$ by Lemma 3.5).

Using Chebyshev's inequality,

$$Prob\{X=0\} \leq \frac{Var[X]}{(E[X])^2} \leq 1/3 . \quad \blacksquare$$

Lemma 3.7. *Let C be a $1/64$-cover, and assume that at time τ_0 all the vertices of C already received the rumor. There exists a constant $c>0$ such that after cn additional iterations of the Procedure $\mathcal{RB}$ all vertices receive the rumor with probability $1-2^{-2n}$.*

Proof. Let $\mathcal{L}(t, \tau, h, k)$ denote the event: at least one of the paths in $L(t, \tau, h, k)$ reached the vertex t. Let $k = 11$. Consider the following sequence of l events: $\mathcal{L}(t, \tau_0, h, k)$, $\mathcal{L}(t, \tau_0+k, h, k)$, $\mathcal{L}(t, \tau_0+2k, h, k), \ldots, \mathcal{L}(t, \tau_0+(2n-1)k, h, k)$.

The l events are independent, since each event considers the performance of vertices in each band at a different interval of k iterations.

By Lemma 3.6 each event holds with probability greater than $2/3$, thus, at time $\tau_0 + 2nk + kh$ the probability that t did not receive the rumor is bounded by $(1/3)^{2n} \leq 2^{-2n-1}$, and the probability that any vertex did not receive the rumor is bounded by 2^{-n-1}. ∎

We are now ready to complete the proof of the main theorem. By Lemma 3.4 after $\tau_0 = 3n/\alpha$ iterations, with probability $1-2^{-n-1}$ there exists an α-cover, for $\alpha = 1/64$, in which all vertices received the rumor. By Lemma 3.6 an additional cn iterations guarantee that the rumor reaches every vertex with probability $1-2^{-n-1}$, thus $O(n) = O(\log N)$ iterations are sufficient for distributing the rumor among all vertices of the hypercube with probability $1-1/N$. ∎

4. RANDOM BROADCAST ON RANDOM GRAPHS

In this section we show that, in a certain sense, "almost all" connected graphs have fast rumor coverage time. For $p \in [0, 1]$, we say that G is drawn from $\mathcal{G}_{N,p}$ if G is an N-vertex graph each of whose edges is present independently with probability p. (Thus $\mathcal{G}_{N,1/2}$ is the space of all N-vertex graphs chosen equiprobably.) Similarly, $\mathcal{G}_{N,M}$ denotes the space of graphs on N vertices with M randomly placed edges.

The following theorem characterizes the rumor coverage time on random graphs.

Theorem 4.1. *For almost all $G \in \mathcal{G}_{N,p}$ ($G \in \mathcal{G}_{N,M}$),*

1. *If* $p \le (\log N - \omega(N))/N$ *(resp.* $M \le (\log N - \omega(N))N/2$*), where* $\omega(N) \to \infty$*, then* $\mathcal{T}(G) = \infty$.
2. *If* $p = (\log N + \omega(N))/N$ *(resp.* $M = (\log N + \omega(N))N/2$*), where* $\omega(N) \to \infty$*,* $\omega(N) = O(\log \log N)$*, then* $\mathcal{T}(G) = \Theta(\log^2 N)$.
3. *If* $p \ge (1+\epsilon)(\log N)/N$ *(resp.* $M \ge ((1+\epsilon)N \log N)N/2$*), for some fixed* $\epsilon > 0$*, then* $\mathcal{T}(G) = \Theta(\log N)$.

Proof. 1. If $p \le (\log N - \omega(N))/N$ (resp. $M \le (\log N - \omega(N))N/2$), where $\omega(N) \to \infty$, then almost all graphs in $\mathcal{G}_{N,p}$ (resp. $\mathcal{G}_{N,M}$) are not connected [7], thus $\mathcal{T}(G) = \infty$.

2. If $p = (\log N + \omega(N))/N$ (resp. $M = (\log N + \omega(N))N/2$), where $\omega(N) = (k-1)\log\log N$ ($k \ge 1$ constant), then the minimum degree of almost all graphs in $\mathcal{G}_{N,p}$ (resp. $\mathcal{G}_{N,M}$) is not larger than k [8]. Let X_{min} be a vertex with minimum degree in G. The probability that X_{min} has a neighbor with degree smaller than $l = (\log N)/4$ is bounded in $\mathcal{G}_{N,p}$ by

$$N\binom{N}{k}p^k(1-p)^{N-k-1}k\binom{N}{l}p^l(1-p)^{N-l-1} \le \frac{1}{N^{1/4}},$$

similar estimate holds in $\mathcal{G}_{N,M}$. If all the neighbors of X_{min} have degree at least l, then even if all of them have the rumor at time 0, the probability that the rumor does not reach X_{min} in $(\log^2 N)/8k$ steps is larger than $1/N$, thus for almost all graphs in this probability space, $\mathcal{T}(G) = \Omega(\log^2 N)$. Since with high probability the maximum degree of G is $O(\log N)$ and its diameter is $O(\log N/\log\log N)$ [3], by Theorem 2.2, $\mathcal{T}(G) = O(\log^2 N)$.

3. If $p \ge (1+\epsilon)(\log N)/N$ (resp. $M \ge ((1+\epsilon)N \log N)N/2$), for some fixed $\epsilon > 0$, then there exist two constants $\alpha < 1$ and $\beta > 1$, such that for almost all graphs in $G \in \mathcal{G}_{N,p}$ (resp. $G \in \mathcal{G}_{N,M}$) the degrees in G are bounded between $\alpha D(N)$ and $\beta D(N)$, $D(N) = p(N-1)$ (resp. $D(N) = 2M/N$). Let v and u be two vertices in G. We show that the probability that a rumor that was initiated in v fails to reach u in $O(\log N)$ steps is bounded by $1/N^3$, which implies that for any start-point in G, $\mathcal{T}(G) = O(\log N)$.

We analyze the progress of the broadcasting process in three stages. In stage A we consider only the first four vertices used by each vertex and ignore transmission through any other vertex in the graph. Stage A can be analyzed as a discrete branching process in which each element has a chance to generate 5 offsprings. The probability that an offspring survives is the probability that a vertex chooses to send the rumor to a vertex that has not received it yet. Let X_i denote the number of elements in the ith generation of this branching process. Then

$$Prob\{X_{i+1} < 2X_i \mid X_i \le N/4\} \le \binom{N}{X_i}\left(\frac{2X_i}{N}\right)^{5X_i} \le \frac{\log N}{N^4},$$

thus $Prob\{X_{\log N} < N/\} \le (\log^2 N)/N^4$. Stage A takes no more than $5 \log N$ steps.

Stage B lasts $6 \log N$ steps. In analyzing this stage we consider only messages sent from vertices that received the rumor in the first stage, and only messages transmitted through edges that were not used in the first stage. The probability that vertex with degree d, $\alpha D(N) \le d \le \beta D(N)$ that received the rumor in stage

A fails to transmit the rumor through at least $(\alpha \log N)/2$ distinct edges in stage B is bounded by

$$\binom{d}{(\alpha \log N)/2}\left(\frac{(\alpha \log N)/2}{d}\right)^{6 \log N} \leq \frac{1}{N^5}.$$

Thus, with probability $1 - 1/N^4$, in stage B the rumor traversed at least $N(\alpha - 5)/4$ distinct edges that were not used in stage A.

Let $deg(u)$ denote the degree of u. The probability that fewer than half of the neighbors of u received the rumor in stage B is bounded by

$$\binom{deg(u)}{deg(u)/2}\left(1 - \frac{deg(u)}{2N}\right)^{N(\alpha \log N - 5)/4} \leq \frac{1}{N^4}.$$

In stage C we consider only messages from neighbors of u. At least half of the neighbors have already received the rumor before stage C. The degree of each neighbor of u is bounded by $\beta D(N)$, and there are at least $\alpha D(N)$ neighbors. Thus, the probability that u does not receive the rumor in $(8\beta \log N/)\alpha$ steps of stage C is bounded by

$$\left(1 - \frac{1}{\beta D(N)}\right)^{(\alpha D(N) 8\beta \log N)/2\alpha} \leq \frac{1}{N^4}.$$

Thus, with probability greater than $1 - 1/N^3$ the rumor that started in v reaches u in no more than $5 \log N + 6 \log N + (8\beta \log N)/\alpha = O(\log N)$ steps. ∎

ACKNOWLEDGMENTS

The authors would like to thank Rafi Heiman, Muli Safra, and Moshe Tennenholtz for helpful discussions.

APPENDIX

Proof of Lemma 3.1. By Stirling's formula

$$\binom{\alpha n}{\beta n} \cong \frac{(\alpha n/e)^{\alpha n}\sqrt{2\pi\alpha n}}{(\beta n/e)^{\beta n}\sqrt{2\pi\beta n}((\alpha - \beta)n/e)^{(\alpha n - \beta n)}\sqrt{2\pi(\alpha - \beta)n}}$$

$$= O\left(\frac{1}{\sqrt{n}}\right)\frac{\alpha^{\alpha n}(\alpha - \beta)^{\beta n}}{\beta^{\beta n}(\alpha - \beta)^{\alpha n}}$$

Rearranging and neglecting $O(1/\sqrt{n})$,

$$\binom{\alpha n}{\beta n} \cong \left(\frac{\left(\frac{\alpha}{\beta} - 1\right)}{\left(1 - \frac{\beta}{\alpha}\right)^{\alpha/\beta}}\right)^{\beta n}$$

Noting that $1/4 < (1-\beta/\alpha)^{\alpha/\beta} < 1/e$ for $2\beta < \alpha$, we conclude that

$$\left(e\left(\frac{\alpha}{\beta}-1\right)\right)^{\beta n} < \binom{\alpha n}{\beta n} < \left(4\left(\frac{\alpha}{\beta}-1\right)\right)^{\beta n}. \qquad \blacksquare$$

Proof of Lemma 3.2. We evaluate an expression A which is certainly smaller than that in the lemma, namely,

$$A = \prod_{i=2}^{i}\left(1-\left(1-\frac{i}{n}\right)^k\right) < \prod_{i=1}^{(1/2+\alpha)n}\left(1-\left(1-\frac{i}{n}\right)^k\right).$$

Substituting $j = n - i$, we get

$$A = \prod_{j=0}^{n-2}\left(1-\left(\frac{j}{n}\right)^k\right).$$

Taking natural logarithms,

$$\ln A = \sum_{j=0}^{n-2}\ln\left(1-\left(\frac{j}{n}\right)^k\right).$$

Changing sums to integrals and adjusting the range of integration,

$$\ln A > \int_{j=0}^{n-1}\ln\left(1-\left(\frac{j}{n}\right)^k\right)dj.$$

Substituting $x = j/n$, $dj = n\,dx$,

$$\ln A < n\int_{x=0}^{(n-1)/n}\ln(1-x^k)\,dx.$$

For $0 < y < 1$, $\ln(1-y) = -\Sigma_{i=1}^{\infty} y^i/i$, and so

$$B = \int_{x=0}^{(n-1)/n}\ln(1-x^k)\,dx = -\int_{x=0}^{(n-1)/n}\sum_{i=1}^{\infty}\frac{(x^k)^i}{i}\,dx.$$

Exchanging the order between summation and integration, and integrating,

$$B = -\sum_{i=1}^{\infty}\int_{x=0}^{(n-1)/n}\frac{(x^k)^i}{i}\,dx = -\left[x\sum_{i=1}^{\infty}\frac{(x^k)^i}{i(ki+1)}\right]_0^{(n-1)/n}$$

$$> -\frac{n-1}{n}\sum_{i=1}^{\infty}\frac{\left(\left(\frac{n-1}{n}\right)^k\right)^i}{ki^2}.$$

But $\Sigma_{i=1}^{\infty} 1/i^2 = \pi^2/6$. Substituting into the last expression we get

$$\frac{1}{n}\ln A > B > -\frac{n-1}{n}\cdot\frac{\pi^2}{6}\cdot\frac{((n-1)/n)^k}{k} > -\frac{\pi^2}{6k},$$

and finally

$$A > e^{-\pi^2 n/6k} > 2^{-\theta n/k}$$

for θ as in the lemma. ∎

REFERENCES

[1] R. Aleliunas, R. M. Karp, R. J. Lipton, L. Lovász, and C. Rackoff, Random walks, universal traversal sequences, and the complexity of maze problems, in *20th Annual Symposium on Foundations of Computer Science*, San Juan, Puerto Rico, October 1979, pp. 218–223.

[2] A. Z. Broder and A. R. Karlin, Bounds on covering times, in *29th Annual Symposium on Foundations of Computer Science*, White Plains, NY, October 1988, pp. 479–487.

[3] B. Bollobás, The diameter of random graphs, *Proc. Am. Math. Soc.*, **83**, 433–436 (1981).

[4] H. Chernoff, A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations, *Ann. Math. Stat.*, **23**, 493–509 (1952).

[5] A. K. Chandra, P. Raghavan, W. L. Ruzzo, R. Smolensky, and P. Tiwari, The electrical resistance of a graph captures its commute and cover times, in *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, Seattle, WA, May 1989, pp. 574–586.

[6] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry, Epidemic algorithms for replicated database management, in *6th ACM Symposium on Principles of Distributed Computing*, 1987, pp. 1–12.

[7] P. Erdös and A. Rényi, On random graphs I, *Publ. Math. Debrecen*, **6**, 290–297 (1959).

[8] P. Erdös and A. Rényi, On the strength of connectedness of a random graph, *Acta Math. Acad. Sci. Hung.*, **12**, 261–267 (1961).

[9] A. M. Frieze and G. R. Grimmett, The shortest-path problem for graphs with random arc-lengths, *Discrete Appl. Math.*, **10**, 57–77 (1985).

[10] W. Goffman and V. A. Newill, Generalization of epidemic theory—An application to the transmission of ideas, *Nature* **204**, 225–228 (1964).

[11] S. M. Hedetniemi, S. T. Hedetniemi, and A. L. Liestman, A survey of gossiping and broadcasting in communication networks, *Networks*, **18**, 319–349 (1988).

[12] H. G. Landau and A. Rapoport, Contribution to the mathematical theory of contagion and spread of information: I. spread through a thoroughly mixed population, *Bull. Math. Biophys.*, **15**, 173–183 (1953).

[13] B. M. Pittel, On spreading a rumour, *SIAM J. Appl. Math.*, **47**, 213–223 (1987).

Received May 10, 1990