LA-UR- 09-04231

Approved for public release;
distribution is unlimited.

Title: ORIENTATIONAL ORDER IN DENSE BACTERIAL COLONIES EMERGES THROUGH DIRECTIONAL REVERSALS

Author(s): YI JIANG
YILIN WU
DALE KAISER
MARK ALBER

Intended for: JOURNAL:
NATURE-PHYSICS

Los Alamos
NATIONAL LABORATORY
EST. 1943

Efficient Broadcast on Random Geometric Graphs

Milan Bradonjić[1] Robert Elsässer[2] Tobias Friedrich[3] Thomas Sauerwald[3] Alexandre Stauffer[4]

[1] Los Alamos National Laboratory, Los Alamos, NM 87545, USA, milan@lanl.gov
[2] Department of Computer Science, University of Paderborn, 33102 Germany, elsa@upb.de
[3] International Computer Science Institute, Berkeley, CA 94707, USA, {tobias,sauerwal}@icsi.berkeley.edu
[4] Computer Science Division, University of California, Berkeley, CA 94720, USA, stauffer@cs.berkeley.edu

Abstract

A *Random Geometric Graph* (RGG) is constructed by distributing n nodes uniformly at random in the unit square and connecting two nodes if their Euclidean distance is at most r, for some prescribed r. We analyze the following randomized broadcast algorithm on RGGs. At the beginning, there is only one informed node. Then in each round, each informed node chooses a neighbor uniformly at random and informs it.

We prove that this algorithm informs every node in the largest component of a RGG in $\mathcal{O}(\sqrt{n}/r)$ rounds with high probability. This holds for any value of r larger than the critical value for the emergence of a giant component. In particular, our result implies that the diameter of the giant component is $\Theta(\sqrt{n}/r)$.

1 Introduction

The study of information spreading in large networks has various fields of applications in distributed computing. One important example is the maintenance of replicated databases on name servers in a large network [4, 9]. There are updates injected at various nodes, and these updates must be propagated to all the nodes in the network. In each step, two neighboring nodes check whether their copies of the database agree and perform the updates, if necessary. In order to be able to let all copies of the database converge to the same content, efficient broadcasting algorithms have to be developed. Typically, these broadcast algorithms should be simple, resilient against failures and should work locally, i.e., the nodes do not have any knowledge of the global topology. One simple algorithm of this kind is randomized broadcast (a.k.a. push algorithm) we study here. In this algorithm, in each round each informed node chooses a neighbor uniformly at random.

One of the first random graph models, Erdős-Rényi random-graph model [7, 8], exhibits the independence property among edges of the graph. In modeling large networks, as well as in statistical testing, the graphs possess the triangular property, that is, different edges in the graph are not necessarily independent. One of the models that preserves this triangular property and is suitable for many real applications is the model of random geometric graphs [15]. Furthermore, since there are sharp transitions in the structure of a random geometric graph, it is of theoretical interest to study the behavior of algorithms on random geometric graphs, at, below, and above these thresholds. More precisely, in this paper we study the behavior of the broadcast algorithm on random geometric graphs, above the threshold for which the giant component appears (see Section 10 in [15]), as well as above the connectivity threshold (see Section 13 in [15]).

1.1 Related Work

The classic random broadcast has been first analyzed on complete graphs by Frieze and Grimmett [10] who proved that with probability $1 - o(1)$, the runtime is $\log_2 n + \ln n + o(\log n)$. This result was tightened by Pittel [16] showing that with the same probability, $\log_2 n + \ln n + \mathcal{O}(1)$ steps are necessary and sufficient. Feige et al. [9] proved that on any graph, the runtime is at most $\mathcal{O}(n \log n)$ *whp*[1], and that for any bounded

[1] *whp* means with probability tending to 1, as the number of nodes n tends to infinity.

degree graph, $\mathcal{O}(\text{diam}(G))$ steps are sufficient. Furthermore, they established a runtime of $\mathcal{O}(\log n)$ on hypercubes and sufficiently dense random graphs *whp*. In [6], two of the authors extended this result to other graphs by proving an upper bound of $\mathcal{O}(\log n + \text{diam}(G))$ for different Cayley graphs.

Random walks on (RGG) have been studied in [1, 3]. In [1], Avin and Ercal considered random geometric graphs for dimension 2. They proved for a radius that is a constant factor larger than the connectivity threshold, a random geometric graph has a cover time of $\Theta(n \log n)$ with high probability, which is asymptotically optimal. Cooper and Frieze in [3] gave a more precise estimate of the cover time that works for all dimensions larger than 2.

In this work we are specifically interested in the problem of broadcasting in random geometric graph model in two dimensions. We precisely state the definition of RGG in Section 2, and for further study on RGGs we refer the reader to the monograph by Penrose [15]. First, one of the transitions from the random graph models given by Erdős and Rényi [7, 8] to the models that may describe real processes in a more realistic manner is the model of RGGs. As the name of the model suggests, there is a notation of geometry introduced in that model. Inherently, the random geometric graph model can be used in many disciplines, such as, in modeling of sensor networks, cluster analysis, statistical physics, hypothesis testing, as well as in other related disciplines. As one example of random geometric graphs, we see the deployment of sensor networks when the devices are thrown over the field from the air. The devices may be treated as nodes of a random geometric graph, given with their positions. Then, two devices can communicate if and only if they are within some given distance. A further possibility of the usage of RGGs is that data in a higher dimensional space can be seen as vertices of an RGG, the coordinates of which represent the nodes' attributes. The metric imposed on RGG depicts the similarity between data elements in the higher dimensional space.

Although, we do not intend to comprehensively study RGGs in this section, for completeness, we state some of the properties satisfied on RGGs, which will be used in the further sections. In RGG model every single vertex has the same "coverage radius" r_n, and regarding percolation properties, there is the unique giant component, and the appearance of the giant component occurs sharply at a threshold radius $r_n = \sqrt{\lambda_c/n}$ [15]. Theoretically, the exact value of the constant λ_c is not known, but it is known experimentally from the simulations by Rintoul and Torquato [17] that for the dimension $d = 2$ the constant $\lambda_c \approx 1.44$. For $d = 2$, the rigorous bounds on λ_c are given in [13], that is $\lambda_c \in [0.696, 3.32C_2]$, while improvement With respect to connectivity property, the RGG model possesses the sharp threshold, at $r_n = \sqrt{\ln n/(\pi n)}$ [12, 14]. Furthermore, Goel, Rai, and Krishnamachari [11] have shown that every monotone property in RGG (including existence of the giant component and connectivity) possesses a sharp threshold.

1.2 Our Results and Techniques

We analyze the randomized broadcast algorithm on RGG for all values r being larger than the critical value for the existence of a giant component. We prove that this algorithm distributes a piece of information, initially known to one vertex in the graph, within $\mathcal{O}(\sqrt{n}/r)$ to all other vertices in the largest component (with high probability). In particular, if the graph is connected, then all vertices get informed after $\mathcal{O}(\sqrt{n}/r)$ rounds.

To prove this result, we first show that the diameter of the giant component is $\Theta(\sqrt{n}/r)$, whenever r is chosen such that a giant component exists. To the best of our knowledge, only for the connected case the diameter was known before [5]. Our techniques are inspired by percolation theory and we believe them to be useful for other problems, e.g., for considering the cover time of the giant component of RGG.

1.3 Organization

The structure of this paper is as follows. In Section 2 we define the basic notation and state some basic results that are used later on. In Section 3 we derive some graph-theoretic results (e.g., on the diameter) of random geometric graphs. In Section 4 we perform the runtime analysis of the push algorithm. We close with the conclusions in the last section.

2 Notation and Preliminaries

2.1 Random Broadcast

We consider the following randomized broadcast algorithm also known as push algorithm (cf. [9]). We are given an undirected, connected graph $G = (V, E)$. At the beginning round 0, a vertex $s \in V$ owns a piece of information (is informed). In each subsequent round $1, 2, \ldots$, every informed vertex chooses a neighbor uniformly at random and informs it. We denote the runtime of this algorithm by $\mathcal{R}(G)$, which is a random variable. Following the results from previous works, our aim is to prove bounds on $\mathcal{R}(G)$ that hold with high probability (w.h.p.), i.e., with probability at least $1 - n^{-1}$.

2.2 Random Geometric Graph Model

We first recall the most natural definition of random geometric graphs.

Definition 1 (cf. [15]). *For the d-dimensional space $\mathbb{R}^d$ provided with the distance norm $\|\cdot\|$, let $\mathcal{X}_n = \{X_1, X_2, \ldots, X_n\} \subset [0, \sqrt[d]{n}]^d$ be chosen independently and uniformly at random. The random geometric graph $G(\mathcal{X}_n; r)$ has the vertex set $V = \mathcal{X}_n$ and the edge set $E = \{(x, y) \colon x, y \in \mathcal{X}, \|x - y\| \leqslant r\}$.*

For our analysis however, the following definition is advantagous.

Definition 2 (cf. [15]). *Let N_n be a Poisson random variable with parameter n, independent of $\{X_1, X_2, \ldots\} \subset [0, \sqrt[d]{n}]^d$, and let $\mathcal{P}_n = \{X_1, X_2, \ldots, X_{N_n}\}$ be a Poisson process. The random geometric graph $G(\mathcal{P}_n; r)$ has the vertex set $V = \mathcal{P}_n$ and the edge set $E = \{(x, y) \colon x, y \in \mathcal{X}, \|x - y\| \leqslant r\}$.*

The following basic lemma says that any result that holds in the setting of Definition 2 with high probability, also holds in the setting of Definition 1 with high probability. Since the process $\mathcal{P}_n$ does not posses the spatial independence, it is usually easier to prove claims on $G(\mathcal{P}_n; r)$, and then to state the same results on $G(\mathcal{X}_n; r)$. This equivalence is allowed, since in the limit when $n \to \infty$, the processes, $\mathcal{P}_n$ and $\mathcal{X}_n$ are equivalent by *poissonization* method, pp.18, 19 [15].

Lemma 2.1. *Let $\mathcal{A}$ be any event that holds with probability at least $1 - \alpha$ in the $\mathcal{G}(\mathcal{P}_n; r)$ model. Then, $\mathcal{A}$ also holds in the $\mathcal{G}(\mathcal{X}_n; r)$ with probability at least $1 - \mathcal{O}(n^{1/2}\,\alpha)$.*

In what follows, we will always consider the $\mathcal{G}(\mathcal{P}_n; r)$ model.

For a given number of nodes n and radius r_n, we will denote a realization of a RGG as $\mathcal{G}_{n,r}$.

A list of Chernoff bounds we use can be found in Section B in the appendix.

3 The Diameter of the Giant Component

In this section we assume that $r > r_c$ and consider the random graph $G = \mathcal{G}(\mathcal{P}_n, r)$. For all $v_1, v_2 \in \mathcal{P}_n$, we define $d_G(v_1, v_2)$ as the distance between v_1 and v_2 on G, that is, $d_G(v_1, v_2)$ is the length of the shortest path from v_1 to v_2 in G. Also, we define $d(v_1, v_2)$ as the euclidean distance between the positions of v_1 and v_2.

In the remainder of this section we prove the following theorem, which gives an upper bound for $d_G(v_1, v_2)$. Note first that $d_G(v_1, v_2) \geqslant d(v_1, v_2)/r$ for all $v_1, v_2 \in \mathcal{P}_n$.

Theorem 3.1. *For any two nodes $v_1, v_2 \in \mathcal{P}_n$, $d_G(v_1, v_2) = \mathcal{O}(d(v_1, v_2)/r)$ with probability XXX.*

The theorem above yields the following corollary about the diameter of G.

Corollary 3.2. *If $r > r_c$, the diameter of G is $\mathcal{O}(\sqrt{n}/r)$.*

In order to prove the theorem, we first take two nodes v_1 and v_2 and show that $d_G(v_1, v_2) = \mathcal{O}(d(v_1, v_2)/r)$ with high probability. We assume that v_1 and v_2 belong to the largest component of G. We use Figure 1 as a reference to show how to find a path from v_1 to v_2. Take the line L that contains v_1 and v_2. We draw a sequence of adjacent rectangles with side length $r/2$ and r starting from v_1 until we draw a rectangle that contains v_2. The largest sides of the rectangles are parallel to L and are such that L splits each rectangles in two squares of side length $r/2$. We refer to the rectangles as $R_1, R_2, \ldots, R_\kappa$ and let S_{2k-1}, S_{2k} denote the squares contained in R_k. Our goal is to compute $d_G(v_1, v_2)$ by moving between adjacent squares. Note that a node in S_k is a neighbor of all nodes in S_{k-2} and S_{k+2}. Moreover, if there is a path crossing the region $\bigcup_{i=1}^{\kappa} R_i$, then because of the size of the rectangles, there must be a node from this path inside $\bigcup_{i=1}^{\kappa} R_i$.

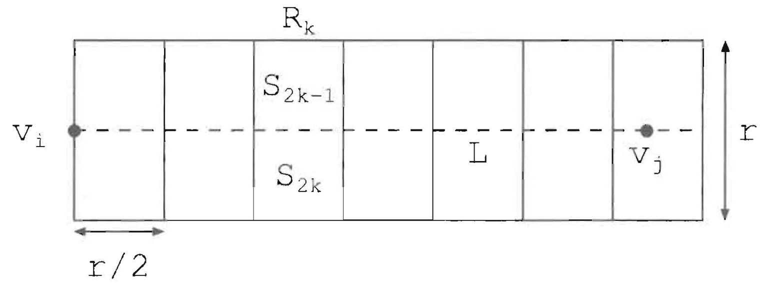

Figure 1: Illustration for the calculation of the diameter.

First we show that $d_G(v_1, v_2) \leqslant 2d(v_1, v_2)/r$ if $r \geqslant \sqrt{14 \log n}$ with probability $1 - O(1/n)$. The probability that there is a node inside S_k for all $1 \leqslant k \leqslant 2\kappa$ is $1 - e^{-r^2/4}$. By the union bound, the probability that there is a node inside each S_k is at least $1 - 2\kappa e^{-r^2/4}$. Clearly, $\kappa \leqslant 1 + 2\sqrt{2n}/r$ for any choice of v_1 and v_2. If there is a node inside each S_k, then $d_G(v_1, v_2) \leqslant 2d(v_1, v_2)/r$. Using the union bound over all choices of v_1, v_2, we obtain that if $r \geqslant \sqrt{14 \log n}$, we have one node in each rectangle for every pair of nodes v_1, v_2 with probability no smaller than $1 - O(n^{2.5})e^{-r^2/4} = 1 - O(1/n)$.

Henceforth we assume that $r < \sqrt{14 \log n}$ and we want to find a path that goes around the empty square. In this section, we refer to a square as empty if it does not contain a node from the largest connected component of G. We disregard the nodes from the other connected components of G. For any empty S_k, we follow the shortest path from a node in S_{k-1} to some $S_{k'}$ for $k' \geqslant k+1$. Note that there is always such a k' since $R_\kappa = S_{2\kappa-1} \cup S_{2\kappa}$ contains v_2.

Our aim is to give a bound to the length of the detour around empty squares. The path starts at $v_1 \in R_1$. For $3 \leqslant k \leqslant 2\kappa$, if S_k is empty and S_{k-2} is not empty, let D_k be the length of the shortest path from S_{k-2} to some $S_{k'}$ for $k' \geqslant k+1$. If S_k is not empty or both S_k and S_{k-2} are empty, then we set $D_k = 0$ since the detour around S_{k-2} will also go around S_k. With these definitions we can write $d_G(u, v) \leqslant \kappa + \sum_{k=3}^{2\kappa} D_k$.

In order to calculate D_k, we exploit the idea of crossings for continuum percolation. First we need some definition. Given a point $x = (x_1, x_2) \in [0, \sqrt{n}]^2$ and two numbers $s > 0$ and $\gamma > 1$, let $A(x, s, \gamma)$ be the annulus with center at x and sides of length s and γs. More formally, if we denote $Q(x, L, s)$ as a square with center at x, sides parallel and perpendicular to L and side length s, then $A(x, L, s, \gamma) = Q(x, L, \gamma s) \setminus Q(x, L, s)$ (see Figure 2(a)). An annulus $A(x, L, s, \gamma)$ can be decomposed into two horizontal rectangles ($Z_1 Z_4 Z_5 Z_{12}$ and $Z_{11} Z_6 Z_7 Z_{10}$ in Figure 2(b)) and two vertical rectangles ($Z_1 Z_2 Z_9 Z_{10}$ and $Z_3 Z_4 Z_7 Z_8$ in Figure 2(b)). For a horizontal rectangle, we define a horizontal crossing as a path in G completely contained in the rectangle and that connects the left to the right side of the rectangle. Similarly for a vertical rectangle, we define a vertical crossing as a path in G that is completely contained in the rectangle and that connects the top to the bottom side of the rectangle. For an annulus $A(x, L, s, \gamma)$ we define $F(A(x, L, s, \gamma))$ as the event that both horizontal rectangles of $A(x, L, s, \gamma)$ have a horizontal crossing and that both vertical rectangles of $A(x, L, s, \gamma)$ have a vertical crossing. This event is illustrated in Figure 2(c).

We explain now how to use the annuli to find detours around an empty square S_k. Suppose that S_{k-2} is not empty but S_k is, and take the point x to be the middle point of the edge between S_{k-2} and S_k. Clearly $Q(x, L, r)$ contains S_{k-2} and S_k and does not intersect S_{k+2}. Also, for all $\gamma > 0$, $A(x, L, r, \gamma)$ contains neither S_{k-2} nor S_k but does intersect with at least two other squares $S_{k'}$ and $S_{k''}$ such that $k' \leqslant k-4$ and $k'' \geqslant k+2$. Then, if for some $\gamma > 0$, $F(A(x, L, r, \gamma))$ happens, then we can exploit the crossings of $A(x, L, r, \gamma)$ to conclude that there exists a path from the node inside S_{k-2} to some $S_{k'}$, $k' \geqslant k+1$, entirely

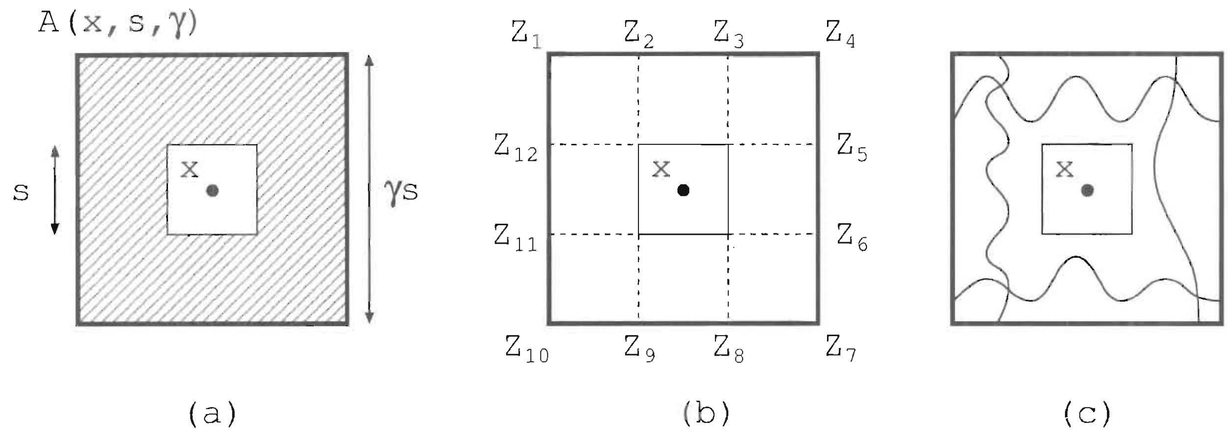

Figure 2: Annulus.

contained inside $Q(x, L, \gamma r)$. (Recall that a path cannot cross $\bigcup_{i=1}^{\kappa} R_i$ without having a node inside that region.) We rely on the fact that the node inside S_{k-1} belongs to the largest connected component, which guarantees that there is a path from S_{k-2} to the crossings of $A(x, L, r, \gamma)$. Furthermore, if $S_{2\kappa-2}$ is empty and the smallest annulus around $S_{2\kappa-2}$ containing the crossings is such that the crossings do not cross $S_{2\kappa-1}$ or $S_{2\kappa}$, then since v_2 belongs to the largest component, there is a path from v_2 to the crossings of this annulus. A similar observation can be drawn if $S_{2\kappa-3}$ is empty. Once we know that such an annulus exists arounds an empty square S_k, we can easily bound D_k by the following trivial geometric lemma.

Lemma 3.3. *Let x be a point in $[0, \sqrt{n}]^2$, s be a nonnegative number, L be a line, and u_1 and u_2 be two nodes of G inside $Q(x, L, s)$. If there exists a path between u_1 and u_2 entirely contained in $Q(x, L, s)$, then $d_G(u_1, u_2) = \mathcal{O}(s^2/r^2)$.*

Proof. The smallest path u_1 and u_2 that is contained inside $Q(x, L, s)$ has the property that for any two non-consecutive nodes w and w' in the path, their distance is larger than r. Otherwise, we can take the edge (w, w') and make the path smaller. This means that if we draw a ball of radius $r/2$ around each node of the path, then two balls overlap only if their respective nodes are consecutive in the path. Let m be the length of the path. There are $m/2$ non-overlapping balls of radius $r/2$. For each ball, at least $1/4$ of its area is contained inside $Q(x, L, s)$. Therefore, it must hold that

$$m \leqslant 2\frac{\text{Area}(Q(x, L, s))}{\pi(r/2)^2/4} = \frac{32s^2}{\pi r^2}.$$

□

The lemma below gives an upper bound for the probability that $A(x, L, r, \gamma)$ does not have a crossing.

Lemma 3.4. *There exist constants c and $\gamma_0 > 1$ such that for all $\gamma > \gamma_0$, all points $x \in [0, \sqrt{n}]^2$, and any line L,*

$$1 - \mathbf{Pr}\left[F(A(x, L, r, \gamma))\right] \leqslant \exp\left(-c\gamma r\right).$$

Proof. We use the ideas from the proof [15, Lemma 10.5]. Recall the decomposition of $A(x, L, r, \gamma)$ into rectangles (refer to Figure 2(b)) and take the top rectangle $Z_1Z_4Z_5Z_{12}$. Its sides have length γr and $r(\gamma - 1)/2$. Therefore, the aspect ratio of the rectangle is $(\gamma - 1)/(2\gamma) \leqslant 1/2$. We want to calculate the probability that such a rectangle has a horizontal crossing as γ increases. This is slightly different from the calculation in [15, Lemma 10.5], since there the aspect ratio is fixed and the side of the rectangle is allowed to vary. But clearly, since $\gamma_0 \leqslant \gamma \leqslant 1/2$, for any rectangle with side lengths γr and $r(\gamma-1)/2$ we can stretch the largest sides (while keeping the smallest sides fixed) to make the aspect ratio be $(\gamma_0 - 1)/2$. Clearly, if there is a horizontal crossing in the enlarged rectangle, there must be a horizontal crossing in the original one. Following along the lines of the proof [15, Lemma 10.5] we can then conclude that there are constants γ_0 and c such that for all $\gamma \geqslant \gamma_0$ a rectangle of side lengths γr and $r(\gamma-1)/2$ has a horizontal crossing with probability larger than $1-e^{-c\gamma r}/4$. Applying the union bound over the 4 rectangles composing the $A(x, L, r, \gamma)$ concludes the proof. $\square$

Now we use this lemma to bound the length of a detour. Suppose that S_k is empty and partition S_{k-2} is not empty. Again, let x be the middle point of the edge between S_{k-1} and S_k. We want to obtain a upper bound to the value of $\gamma \geqslant \gamma_0$ such that the event $F(A(x, L, r, \gamma))$ happens. Note that if for a fixed γ, $F(A(x, L, r, \gamma))$ happens, then from Lemma 3.3 we have $D_k = O(\gamma^2)$. Since there are $O(n^2)$ choices for v_1 and v_2, and for each v_1 and v_2 we have at most $2\kappa = O(\sqrt{n})$ terms D_k, Lemma 3.4 gives that there is a constant c_1 such that with probability $1-O(1/n)$ we obtain $D_k \leqslant c_1 \log n$ for all k. Let $\mathcal{E}$ be this event. In order to apply Azuma's inequality to $\sum_{i=1}^{2\kappa} D_k$, we first derive $\mathbf{E}[D_k]$.

The probability that S_k is empty is $e^{-r^2/4}$. Note that if $F(A(x, L, r, \gamma))$ holds then $D_k \leqslant c_2\gamma^2$, for some constant c_2 and for all $\gamma \geqslant \gamma_0$. Therefore, $\mathbf{Pr}[D_k \geqslant \ell] \leqslant 1 - \mathbf{Pr}\left[F(A(x, L, r, \sqrt{\ell/c_2}))\right] \leqslant \exp(-c\sqrt{\ell/c_2}\,r)$. We can then write $\mathbf{E}[D_k] = e^{-r^2/4}\sum_{\ell=1}^{\infty}\mathbf{Pr}[D_k \geqslant \ell] \leqslant e^{-r^2/4}\int_0^{\infty}\mathbf{Pr}[D_k \geqslant \ell]\,d\ell$, where the last inequality follows from $\mathbf{Pr}[D_k \geqslant \ell]$ being a nonincreasing function of ℓ. We have an exponential upper bound to $\mathbf{Pr}[D_k \geqslant \ell]$ for any ℓ with $\ell \geqslant c_2\gamma_0^2$, therefore

$$
\begin{aligned}
\mathbf{E}[D_k] &\leqslant e^{-r^2/4}c_2\gamma_0^2 + e^{-r^2/4}\int_{\ell \geqslant c_2\gamma_0^2} e^{-c\sqrt{\ell/c_2}\,r}\,dk \\
&= e^{-r^2/4}c_2\gamma_0^2 + e^{-r^2/4}O(e^{-c\gamma_0 r}/r)) \\
&= O(1),
\end{aligned}
$$

for all $r > r_c$. Therefore, by linearity of expectation we have $\mathbf{E}[d_G(v_1, v_2)] = O(\kappa) = O(d(v_1, v_2)/r)$.

We have $\mathbf{E}[D_k] = O(1)$ for each k, and if the event $\mathcal{E}$ holds, we have $D_k \leqslant c_1 \log n$ for all k. Also, if $\mathcal{E}$ holds, then the size of the maximum annulus we need to consider is $c_1 \log n$, which implies that the annulus around S_k and $S_{k'}$ are disjoints if $|k-k'| \geqslant 4c_1 \log n/r$ and, consequently, the random variables D_k and $D_{k'}$ are independent. Let $\Delta = 1 + 4c_1 \log n/r$ and define the index set $I_j = \{k \colon 1 \leqslant k \leqslant 2\kappa, k \equiv j \pmod{\Delta}\}$. We can write $d_G(u, v) = \kappa + \sum_{j=1}^{\Delta}\sum_{k \in I_j} D_k$. But the second sum contains independent random variables and we can apply a version of Azuma's inequality to derive that for each j

$$
\mathbf{Pr}\left[\sum_{k \in I_j} D_k - \sum_{k \in I_j}\mathbf{E}[D_k] \geqslant c_3|I_j|\right] \leqslant 1 - \mathbf{Pr}[\mathcal{E}] + 2\exp\left(-\frac{c_3^2|I_j|^2}{2c_1^2\log^2 n}\right).
$$

Since $|I_j| \geqslant \kappa/\Delta = \Omega(d(v_1, v_2)/\log n)$, the probability above is smaller than $\exp\left(-\frac{c_4 d^2(v_1, v_2)}{\log^4 n}\right)$, for some constant c_4. We solve the first sum by the union bound, obtaining

$$
\mathbf{Pr}\left[\sum_{k=1}^{2\kappa} D_k - \sum_{k=1}^{2\kappa}\mathbf{E}[D_k] \geqslant c_5\kappa\right] \leqslant \Delta\exp\left(-\frac{c_4 d^2(v_1, v_2)}{\log^4 n}\right) \leqslant n^{-3},
$$

for any v_1 and v_2 such that $d(v_1, v_2) \geqslant \sqrt{3/c_4}\log^2 n$. Hence, for a fixed pair of nodes v_1, v_2 such that $d(v_1, v_2) \geqslant \sqrt{3/c_4}\log^2 n$, $d_G(v_1, v_2) = \mathcal{O}(d(v_1, v_2)/r)$ with probability $1-n^{-3}$. Applying Lemma 3.5 yields the claim.

Lemma 3.5. *Let $\mathcal{E}(u, v)$ be an event associated to a pair of vertices $u, v \in G(\mathcal{P}_n, r)$. Assume that for all pairs of vertices, $\mathbf{Pr}[\mathcal{E}(u, v)] \geqslant 1-p$ with $p > 0$. Then,*

$$
\mathbf{Pr}\left[\wedge_{u,v \in V(G(\mathcal{P}_n, r))}\mathcal{E}(u, v)\right] \geqslant 1 - 2n\log n \cdot p + n^{-2}.
$$

Proof. Let $\mathcal{F}$ be the event that the number of vertices in G is at most $2n\log n$. Using a Chernoff bound, it follows easily that $\mathbf{Pr}[\neg\mathcal{F}] \leqslant n^{-2}$. Using the definition of conditional probabilities and the union bound (over conditional probabilities), we have

$$
\begin{aligned}
&\mathbf{Pr}\left[\vee_{u,v\in V(G(\mathcal{P}_n,r))}\neg\mathcal{E}(u,v)\right] \\
&\leqslant \mathbf{Pr}\left[\vee_{u,v\in V(G(\mathcal{P}_n,r))}\neg\mathcal{E}(u,v) \mid \mathcal{F}\right]\cdot\mathbf{Pr}\left[\mathcal{F}\right]+\mathbf{Pr}\left[\neg\mathcal{F}\right] \\
&\leqslant 2n\log n\cdot\mathbf{Pr}\left[\neg\mathcal{E}(u,v) \mid \mathcal{F}\right]+n^{-2} \\
&= 2n\log n\cdot\frac{\mathbf{Pr}\left[\neg\mathcal{E}(u,v)\wedge\mathcal{F}\right]}{\mathbf{Pr}\left[\mathcal{F}\right]}+n^{-2} \\
&\leqslant 2n\log n\cdot\frac{\mathbf{Pr}\left[\neg\mathcal{E}(u,v)\right]}{1-n^{-2}}+n^{-2} \\
&\leqslant 2n\log n\cdot p+n^{-2}.
\end{aligned}
$$

□

4 Broadcast Time

In the previous section we showed how to find a path between two nodes v_1 and v_2 such that if $d(v_1,v_2)\geqslant\sqrt{3/c_4}\log^2 n$ then $d_G(v_1,v_2)=O(d(v_1,v_2)/r)$. Assume that only node v_1 is initially informed and let $\tau(v_1,v_2)$ be the number of rounds it takes for the random broadcast procedure to inform v_2. We show that $\tau(v_1,v_2)=O(d(v_1,v_2)/r)$ with high probability.

We need to work with a subgraph of G. Let $G'=(\mathcal{P}_n,r')$ such that $r_c<r'=(1-\varepsilon)r$, where $\varepsilon<1$ is a constant. Note that since $r'>r_c$ we know that $d_{G'}(v_1,v_2)=O(d(v_1,v_2)/r')$.

Before proceeding to the proof of the random broadcast time, we need first to introduce some definition. Let u_1 and u_2 be two nodes such that u_1 and u_2 have at least one neighbor in common in G'. For any point $x\in[0,\sqrt{n}]^2$, let $B(x,s)$ be the ball with center at x and radius s. For any node w, let $B(w,s)$ be defined as $B(x,s)$ where x is the location of w. Let $H(u_1,u_2)=B(u_1,r)\cap B(u_2,r')$ (refer to Figure 3(a)). Assume that u_1 is the only informed node and let $T(u_1,u_2)$ be the number of rounds taken by the random broadcast procedure over G until the first node inside $H(u_1,u_2)$ is informed (note that $H(u_1,u_2)$ is not empty by the definition of u_1 and u_2). In a similar way, consider a node u and a region $X\subseteq[0,\sqrt{n}]^2$ such that X contains at least one node (besides u if $u\in X$) and for each node in X, u and this node have at least one neighbor in common in G'. Then, assume that a node w chosen uniformly at random from X is informed and let $T(X,u)$ be the number of rounds taken by the random broadcast procedure over G until the first node in $H(w,u)$ is informed.

Note that $T(u_1,u_2)$ is distributed as between the area of $H(u_1,u_2)$ and $B(u_1,r)$. The following lemma gives a upper bound for $\mathbf{E}\left[T(u_1,u_2)\right]$.

Lemma 4.1. *For any pair of nodes u_1 and u_2, $\mathbf{E}\left[T(u_1,u_2)\right]=O(1)$.*

Proof. Disregarding the nodes u_1 and u_2, let Y_1 be the number of nodes in $H(u_1,u_2)$ and Y_2 be the number of nodes inside $B(u_1,r)$. Therefore, $\mathbf{E}\left[T(u_1,u_2)\right]=\mathbf{E}\left[Y_2/Y_1\right]$. We know that there is a node inside $H(u_1,u_2)$ therefore, Y_1-1 and Y_2-1 are Poisson random variables with mean Area$(H(u_1,u_2))$ and πr^2, respectively. Conditional on $Y_2-1=k$, the value of Y_1-1 is given by a Binomial distribution with mean kp, where p is given by the ratio of the area of $H(u_1,u_2)$ to the area of πr^2. Let this ratio be denoted by

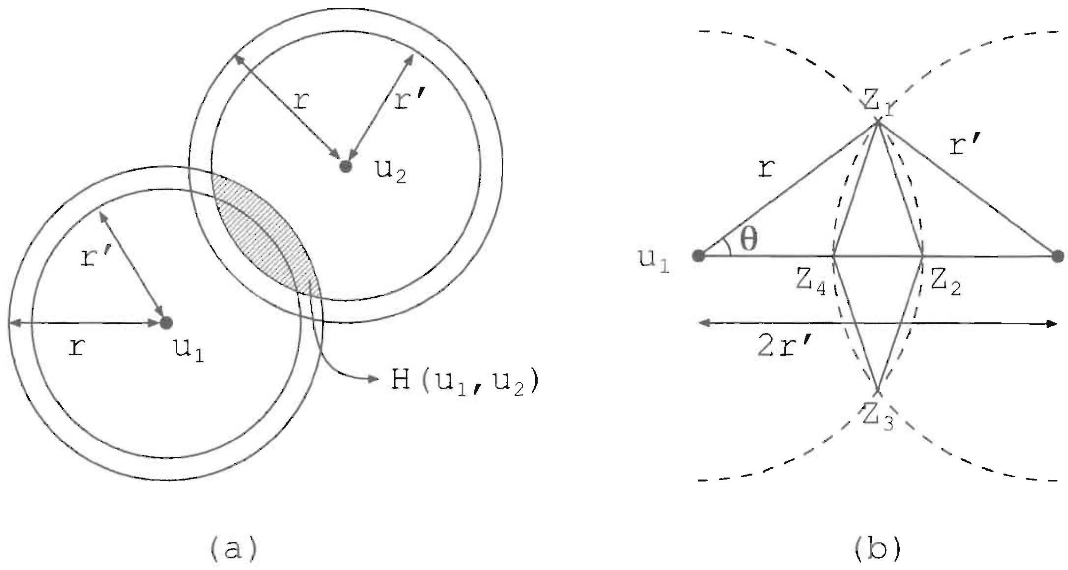

Figure 3: Broadcast

β, we obtain

$$
\begin{aligned}
\mathbf{E}\left[T(u_1, u_2)\right] &= \sum_{k=0}^{\infty}\sum_{i=0}^{k} \frac{k+1}{i+1} e^{-\pi r^2} \frac{(\pi r^2)^k}{k!} \binom{k}{i} \beta^i (1-\beta)^{k-i} \\
&= \frac{1}{\beta} \sum_{k=0}^{\infty} e^{-\pi r^2} \frac{(\pi r^2)^k}{k!} \sum_{i=0}^{k} \binom{k+1}{i+1} \beta^{i+1} (1-\beta)^{k-i} \\
&= \frac{1}{\beta} \sum_{k=0}^{\infty} e^{-\pi r^2} \frac{(\pi r^2)^k}{k!} \left(\sum_{i=0}^{k+1} \binom{k+1}{i} \beta^i (1-\beta)^{k+1-i} - (1-\beta)^{k+1} \right) \\
&= \frac{1}{\beta} \sum_{k=0}^{\infty} e^{-\pi r^2} \frac{(\pi r^2)^k}{k!} \left(1 - (1-\beta)^{k+1} \right) \\
&\leqslant \frac{1}{\beta} e^{-\pi r^2 \beta}.
\end{aligned}
$$

Therefore $\mathbf{E}\left[T(u_1, u_2)\right] = O(1)$ if $\beta = \Omega(1)$. In order to show that $\beta = \Omega(1)$, note that $d(u_1, u_2) \leqslant 2r'$ and the case that minimizes the area of $H(u_1, u_2)$ is when $d(u_1, u_2) = 2r'$ (as illustrated in Figure 3(b)). We show that the quadrilateral $Z_1Z_2Z_3Z_4$ in the figure, which is contained inside $H(u_1, u_2)$, has area $\Omega(r^2)$. It is easy to check that the line Z_2Z_4 has length $r - r' = \varepsilon r$. Now we need to show that the distance from Z_1 to the line Z_2Z_4 has length $\Omega(r)$. Applying the law of cosines to the triangle $Z_1u_1u_2$, we obtain $\cos(\theta) = \frac{1+3(1-\varepsilon)^2}{4(1-\varepsilon)}$ and, consequently, for a constant ε, there exists a constant δ such that $\sin(\theta) \geqslant \delta$. Then, the area of the quadrilateral $Z_1Z_2Z_3Z_4$ is larger than $\varepsilon\delta r^2$ and, consequently, $\beta \geqslant \varepsilon\delta/\pi = \Omega(1)$. □

For any region $X \subseteq [0, \sqrt{n}]$ containing at least one node and a node u so that every node in X has a neighbor in common with u, let W be a node chosen uniformly at random from X. Then, $T(X, u)$ is also distributed according to a geometric random variable and $\mathbf{E}\left[T(X, u)\right] = \mathbf{E}\left[\mathbf{E}\left[T(w, u) \mid W = w\right]\right] = O(1)$.

Now we are in position to get a upper bound for $\tau(v_1, v_2)$. Take a path $u_1, u_2, \ldots, u_m$ in G' such that $u_1 = v_1$ and $u_m = v_2$. Let X_1 be the point where u_1 is located and define the region X_i recursively for $2 \leqslant i \leqslant m-1$ as $X_i = H(W_{i-1}, u_{i+1})$, where W_{i-1} is a node chosen uniformly at random from X_{i-1} (hence X_i depends on X_{i-1}).Note that $u_i \in X_i$ and for any node $w \in X_i$, the region $H(w, u_{i+2})$ is not empty (in particular, $u_{i+1} \in H(w, u_{i+2})$). We derive the time it takes for the random broadcast procedure to inform v_2 by following the information along the path $u_1, H(X_1, u_3), H(X_2, u_4), \ldots, H(X_{m-2}, u_m), u_m$.

Note that when u_1 informs a node inside the region $H(X_1, u_3)$, the node that gets informed is chosen uniformly at random from $H(X_1, u_3)$. The same happens for all $H(X_i, u_{i+2})$. Let $\theta(X_{m-1}, u_m)$ be the time it takes for a node chosen uniformly at random from X_{m-1} to inform u_m. Hence, $\tau(v_1, v_2) \leqslant \theta(X_{m-1}, v_2) + \sum_{i=1}^{m-2} T(X_i, u_{i+2})$.

For a fix pair of nodes v_1 and v_2 we take $u_1, u_2, \ldots, u_m$ to be the shortest path obtained from the procedure used to derive a upper bound for $d_G(v_1, v_2)$ in Section 3. That path has a nice property. If $\mathcal{E}$ holds, then the maximum annulus considered in a detour around an empty square has length $c_1 \log n$. Let S_i and S_j be two empty squares, let x_i be the middle point of the edge between S_{i-1} and S_i, and x_j be the middle point of the edge between S_{j-1} and S_j. Let L be the line containing v_1 and v_2. Then, for any two nodes $w_i \in A(x_i, L, r, c_1 \log n/r)$ and $w_j \in A(x_j, L, r, c_1 \log n/r)$, the region $B(w_i, r) \cap B(w_j, r)$ is empty if $d(x_i, x_j) \geqslant c_1 \log n + 2r$. Since each annulus have $O(\log^2 n/r)$ nodes in the path, we obtain that two nodes u_i and u_j in the path are such that $B(u_i, r) \cap B(u_j, r)$ is empty if $|i-j| \geqslant c_5 \log^3 n/r$, for some constant c_5. Let $\Delta' = 3 + c_5 \log^3 n/r$, for any i, the random variables $T(X_i, u_{i+2})$ and $T(X_j, u_{j+2})$ are independent if $|i-j| \geqslant \Delta'$.

Let the index set $J_j = \{1 \leqslant i \leqslant m : i \equiv j \pmod{\Delta'}\}$. We can then write $\tau(v_1, v_2) = \theta(X_{m-1}, v_2) + \sum_{j=1}^{\Delta'} \sum_{i \in J_j} T(X_i, u_{i+2})$, where for all j, the term $\sum_{i \in J_j} T(X_i, u_{i+2})$ is given by the sum of independent geometric random variables. We apply the following Chernoff bound for this term.

Lemma 4.2. *Let $0 < \delta < 1$ be arbitrary. Suppose that $X_1, \ldots, X_n$ are independent random variables on $\mathbb{N}$ with $\mathbf{Pr}[X_i = k] = (1-\delta)^{k-1}\delta$ for every $k \in \mathbb{N}$. Let $X = \sum_{i=1}^{n} X_i$. Then it holds for every $\varepsilon > 0$ that*

$$\mathbf{Pr}[X \geq (1+\varepsilon)n/\delta] \leq e^{-\varepsilon^2 n/2(1+\varepsilon)}.$$

Proof. Consider transforming every $X_i = k$ into a binary string $B_k = (000...01)$ with $(k-1)$ zeroes. Then the series of $X_1 = k_1, X_2 = k_2, X_3 = k_3, \ldots$ can be represented by a string B of the form of $\{B_{k_1} B_{k_2} B_{k_3} \ldots\} = \{00...100...100...1...00...1\}$. Note that B contains n many 1's and the total number of positions (0 or 1) in B is $K = \sum_i k_i$.

Now, consider instead an infinite set of binary random variables $Y_1, Y_2, Y_3, \ldots$ with $\mathbf{Pr}[Y_i = 1] = \delta$. Viewing Y_i as representing the i-th position in B, it is not difficult to check that

$$\mathbf{Pr}\left[\sum_{i=1}^{n} X_i \geq k\right] = \mathbf{Pr}\left[\sum_{j=1}^{k} Y_j \leq n\right]$$

Let $k = (1+\varepsilon)\mu$ with $\mu = \mathbf{E}[X]$, $Y = \sum_{j=1}^{(1+\varepsilon)\mu} Y_j$, and $\mu' = \mathbf{E}[Y]$. It follows from the Chernoff bounds for binomial random variables (cf. Lemma B.1) that for any $\varepsilon' \in [0,1)$,

$$\mathbf{Pr}\left[Y \leq (1-\varepsilon')\mu'\right] \leq e^{-(\varepsilon')^2\mu'/2}$$

Since $\mathbf{E}[X_i] = 1/\delta$ and therefore $\mu = n/\delta$, it follows that $\mu' = (1+\varepsilon)\mu \cdot \delta = (1+\varepsilon)n$. Hence, if we set $n = (1-\varepsilon')\mu'$, we get that $(1-\varepsilon')(1+\varepsilon) = 1$ and therefore $\varepsilon' = \varepsilon/(1+\varepsilon)$. Plugging this into the equations above, we obtain

$$\mathbf{Pr}\left[\sum_{i=1}^{n} X_i \geq (1+\varepsilon)n/\delta\right] \leq e^{-\varepsilon^2 n/2(1+\varepsilon)}$$

which proves the lemma. □

The lemma above gives that for each j,

$$\mathbf{Pr}\left[\sum_{i \in J_j} T(X_i, u_{i+2}) \geqslant (1+\varepsilon) \sum_{i \in J_j} \mathbf{E}[T(X_i, u_{i+2})]\right] \leqslant \exp\left(-\frac{\varepsilon^2 m}{2(1+\varepsilon)\Delta'}\right) \leqslant \exp\left(-c_6 \frac{d(v_1, v_2)}{\log^3 n}\right).$$

We know that there exists a constant c_6 such that $\mathbf{E}[T(X_i, u_{i+2})] \leqslant c_6$ for all i. We can then take the union bound over all j and conclude that

$$\mathbf{Pr}\left[\sum_{j=1}^{\Delta'}\sum_{i\in J_j} T(X_i, u_{i+2}) \geqslant (1+\varepsilon)c_6(m-2)\right] \leqslant \Delta' \exp\left(-c_7 \frac{d(v_1, v_2)}{\log^3 n}\right),$$

for some constant c_7.

It only remains to prove a upper bound for $\theta(X_{m-1}, v_2)$. We show that for any pair of nodes w_1 and w_2, we obtain $\theta(w_1, w_2) \leqslant$

Lemma 4.3. *For all pair of nodes w_1 and w_2 such that $d(w_1, w_2) \leqslant r$, the following holds with probability $1 - O(1/n)$,*

$$\theta(w_1, w_2) \leqslant c_9 \frac{\log^2 n}{\log\log n},$$

for some constant c_9.

Proof. Note that if the degree of w_1 in G is λ, then the number of rounds until w_1 sends the information to w_2 is given by a geometric random variable with mean λ. We first show that the maximum degree of G is smaller than $c_8 \log n/\log\log n$ with high probability. We first tesselate the region $[0, \sqrt{n}]$ in squares of size r. For each square, from Lemma 4.4, the probability that there are more than $(c_8/9)\log n/\log\log n$ nodes is smaller than $1/n^2$, for some constant c_8. Since we have n/r^2 squares, the probability that each square has less than $(c_8/9)\log n/\log\log n$ nodes is $1 - O(1/n)$. Since for each node, its neighbors belong to at most 9 squares we have that $\lambda \leqslant c_8 \log n/\log\log n$. Therefore,

$$\mathbf{Pr}[\theta(w_1, w_2) \geqslant t] \leqslant \exp\left(-c_8 t \frac{\log\log n}{\log n}\right).$$

If we set $t = 3\log^2 n/(c_8 \log\log n) = c_9 \log^2 n/\log\log n$, we obtain that $\mathbf{Pr}\left[\theta(w_1, w_2) \geqslant c_9 \log^2 n/\log\log n\right] \leqslant 1/n^3$ and, by the union bound over w_1, w_2, the probability that $\theta(w_1, w_2) \leqslant c_9 \log^2 n/\log\log n$ for all w_1, w_2 is larger than $1 - O(1/n)$. □

Lemma 4.4. *If X is a Poisson random variable with mean μ, then $\mathbf{Pr}[X \geqslant t] \leqslant e^{-t\log(t/\mu)+t-\mu}$.*

Proof. We have $\mathbf{E}\left[e^{\theta X}\right] = e^{\mu(e^{\theta}-1)}$ and by Markov's inequality $\mathbf{Pr}[X \geqslant t] \leqslant \exp(\mu(e^{\theta} - 1) - \theta t)$, setting $\theta = \log(\mu/t)$ concludes the proof. □

Therefore, we get that for any two nodes v_1 and v_2 in the largest component of G' such that $d(v_1, v_2) \geqslant (3/c_7)\log^3 n$, $\tau(v_1, v_2) = O(d_{G'}(v_1, v_2)) = O(d(v_1, v_2)/r)$ with probability $1 - O(1/n)$. We now need to show that after all the largest component of G' is informed, the nodes that belong to the largest component of G are informed after a small amount of rounds.

5 Conclusion

We analyzed random broadcast in random geometric graphs. We proved that the algorithm completes in $\mathcal{O}(\sqrt{n}/r)$ steps, where r can be an arbitrary value above the critical value for the emergence of the giant component. In particular, we also showed that the diameter of the largest component is $\mathcal{O}(\sqrt{n}/r)$.

An open problem is to extend our results to higher dimensions, which seems to be challenging, as certain ideas from percolation theory we use are restricted to two dimensions. A second possibility would be to try to apply our techniques to the analysis of the cover time of the giant component of RGGs. This would nicely complement recent results by Cooper and Frieze [2, 3] for the giant component of Erdos-Renyi-Random-Graphs and for connected RGGs.

Acknowledgments

Milan Bradonjić would thank Los Alamos National Laboratory. Part of this work was funded through the Laboratory Directed Research and Development Program, and Center for Nonlinear Studies. Alexandre Stauffer acknowledges partial support from CAPES/Fulbright.

References

[1] C. Avin and G. Ercal. On the Cover Time and Mixing Time of Random Geometric Graphs. *Theoretical Computer Science*, 380(1-2):2–22, 2007.

[2] C. Cooper and A. Frieze. The cover time of the giant component of a random graph. *Random Structures & Algorithms*, 32(4):401–439, 2008.

[3] C. Cooper and A. Frieze. The cover time of random geometric graphs. In *19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'09)*, pages 48–57, 2009.

[4] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic algorithms for replicated database maintenance. In *6th Annual ACM Principles of Distributed Computing (PODC'87)*, pages 1–12, 1987.

[5] R. B. Ellis, J. L. Martin, and C. Yan. Random geometric graph diameter in the unit ball. *Algorithmica*, 47(4):421–438, 2007.

[6] R. Elsässer and T. Sauerwald. Broadcasting vs. Mixing and Information Dissemination on Cayley Graphs. In *24th International Symposium on Theoretical Aspects of Computer Science (STACS'07)*, pages 163–174, 2007.

[7] P. Erdős and A. Rényi. On random graphs. *Publ. Math. Inst. Hungar. Acad. Sci.*, 1959.

[8] P. Erdős and A. Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hungar. Acad. Sci.*, 1960.

[9] U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized Broadcast in Networks. *Random Structures and Algorithms*, 1(4):447–460, 1990.

[10] A. Frieze and G. Grimmett. The shortest-path problem for graphs with random-arc-lengths. *Discrete Applied Mathematics*, 10:57–77, 1985.

[11] A. Goel, S. Rai, and B. Krishnamachari. Sharp thresholds for monotone properties in random geometric graphs. In *36th Annual ACM Symposium on Theory of Computing (STOC'04)*, pages 580–586, 2004.

[12] P. Gupta and P. R. Kumar. Critical power for asymptotic connectivity. In *Proceedings of the 37th IEEE Conference on Decision and Control*, volume 1, pages 1106–1110, 1998.

[13] R. Meester and R. Roy. *Continuum percolation*. Cambridge University Press, 1996.

[14] M. D. Penrose. The longest edge of the random minimal spanning tree. *The Annals of Applied Probability*, 7(2):340–361, 1997.

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

[16] B. Pittel. On spreading rumor. *SIAM Journal on Applied Mathematics*, 47(1):213–223, 1987.

[17] M. D. Rintoul and S. Torquato. Precise determination of the critical threshold and exponents in a three-dimensional continuum percolation model. *Journal of Physics A. Mathematical and General*, 30 (16):L585–L592, 1997.

A Omitted Proofs

Lemma 2.1 (from page 3). *Let $\mathcal{A}$ be any event that holds with probability at least $1-\alpha$ in the $\mathcal{G}(\mathcal{P}_n; r)$ model. Then, $\mathcal{A}$ also holds in the $\mathcal{G}(\mathcal{X}_n; r)$ with probability at least $1-\mathcal{O}(n^{1/2}\,\alpha)$.*

Proof. In this proof, we shall use subscripts to indicate the space over which the probabilites are taken. Let $\mathrm{vol}(G_{n,r})$ denote the number of vertices in a realization of $\mathcal{G}(\mathcal{P}_n; r)$. Then it follows that by Stirling's formula that

$$\mathbf{Pr}_{\mathcal{G}(\mathcal{X}_n;r)}\left[\mathrm{vol}(G_{n,r}) = n\right] = e^{-n}\cdot\frac{n^n}{n!} = e^{-n}\cdot\frac{n^n}{\Theta((\frac{n}{e})^n\cdot\sqrt{2\pi n})} = \Theta\left(\frac{1}{\sqrt{n}}\right).$$

Note however, that conditioned on $\mathrm{vol}(G_{n,r}) = n$ for a relization in $\mathcal{G}(\mathcal{P}_n; r)$, $G_{n,r}$ is also a realization of $\mathcal{G}(\mathcal{X}_n; r)$. Therefore

$$\begin{aligned}\mathbf{Pr}_{\mathcal{G}(\mathcal{X}_n;r)}\left[\neg\mathcal{A}\right] &= \mathbf{Pr}_{\mathcal{G}(\mathcal{P}_n;r)}\left[\neg\mathcal{A}\mid \mathrm{vol}(G_{n,r}) = n\right] \\ &\leqslant \frac{\mathbf{Pr}_{\mathcal{G}(\mathcal{P}_n;r)}\left[\neg\mathcal{A}\right]}{\mathbf{Pr}_{\mathcal{G}(\mathcal{P}_n;r)}\left[\mathrm{vol}(G_{n,r}) = n\right]} \\ &= \mathcal{O}\left(n^{1/2}\,\alpha\right).\end{aligned}$$

□

B Chernoff Bounds

Lemma B.1 (Chernoff Bound for Sums of Binary Variables). *Let $X_1, X_2, \ldots, X_n$ be independent binary random variables. Let $X = \sum_{i=1}^{n} X_i$ and $\mu := \mathbf{E}[X]$. Then it holds or all $\delta > 0$ that*

$$\mathbf{Pr}\left[X \geqslant (1+\delta)\mu\right] \leqslant \exp\left(-\min\{\delta^2, \delta\}\mu/3\right).$$

Moreover it holds for all $0 < \delta < 1$ that

$$\mathbf{Pr}\left[X \leqslant (1-\delta)\mu\right] \leqslant \exp\left(-\delta\mu/2\right).$$

siam® acm Association for Computing Machinery

ACM-SIAM Symposium on Discrete Algorithms

January 17-19, 2010
Hyatt Regency Austin
Austin, Texas

ACM-SIAM Symposium on Discrete Algorithms (SODA10)

January 17-19, 2010
Hyatt Regency Austin
Austin, Texas

The twelfth Workshop on Algorithm Engineering and Experiments (ALENEX10) and the seventh Workshop on Analytic Algorithmics and Combinatorics (ANALCO10) will be held immediately preceding the conference, on January 16, at the same location.

SODA is jointly sponsored by the *ACM Special Interest Group on Algorithms and Computation Theory* and the *SIAM Activity Group on Discrete Mathematics*.

Announcements

NSF Student Travel Grants are available.

Program Committee

Gagan Aggarwal, Google Inc.
Matthew Andrews, Bell Laboratories
Lars Arge, Aarhus University, Denmark
Moses Charikar (chair), Princeton University
Esther Ezra, Duke University
Lisa K. Fleischer, Dartmouth College
Parikshit Gopalan, Microsoft Research, Silicon Valley
Venkatesan Guruswami, University of Washington and Carnegie Mellon University
Thore Husfeldt, ITU Copenhagen, Denmark, and Lund University, Sweden

In This Section

- Home
- Best Paper Awards
- Child Care
- Exhibits
- General Information
- Hotel Information
- Invited Presentations
- Other SODA Meetings
- Program
- Registration Information
- Related Links
- Sponsors
- Students
- Submissions
- Titles on Display
- Travel Support
- Workshops

Also See

- Workshop on Algorithm Engineering and Experiments (ALENEX10)
- Workshop on Analytic Algorithms and Combinatorics (ANALCO10)

Ravi Kannan, Microsoft Research, India
Sanjeev Khanna, University of Pennsylvania
Robert Krauthgamer, Weizmann Institute of Science, Israel
Eyal Lubetzky, Microsoft Research, Redmond
Andrew McGregor, University of Massachusetts, Amherst
Cris Moore, University of New Mexico and Santa Fe Institute
Rina Panigrahy, Microsoft Research, Silicon Valley
Kirk Pruhs, University of Pittsburgh
Rajmohan Rajaraman, Northeastern University
Sofya Rashkhodnikova, Pennsylvania State University
Asaf Shapira, Georgia Institute of Technology
Mohit Singh, Microsoft Research, New England
Daniel Stefankovic, University of Rochester
Maxim Sviridenko, IBM T.J. Watson Research Center
Wojciech Szpankowski, Purdue University
Kavitha Telikepalli, Indian Institute of Science, India
Kasturi Varadarajan, The University of Iowa
Suresh Venkatasubramanian, University of Utah
Jan Vondrak, IBM Almaden Research Center
Peter Winkler, Dartmouth College

Description

This symposium focuses on research topics related to efficient algorithms and data structures for discrete problems. In addition to the design of such methods and structures, the scope also includes their use, performance analysis, and the mathematical problems related to their development or limitations. Performance analyses may be analytical or experimental and may address worst-case or expected-case performance. Studies can be theoretical or based on data sets that have arisen in practice and may address methodological issues involved in performance analysis.

Funding Agency

Funding agencies will be listed here when available.

Themes

Themes and application areas include, but are not limited to, the following topics:

Aspects of Combinatorics and Discrete Mathematics, such as:

- Algebra
- Combinatorial Structures
- Discrete Optimization
- Discrete Probability
- Finite Metric Spaces
- Graph Theory
- Mathematical Programming

- Number Theory
- Random Structures
- Topological Problems

Aspects of Computer Science, such as:

- Algorithm Analysis and Complexity
- Algorithmic Game Theory
- Algorithmic Mechanism Design
- Combinatorial Scientific Computing
- Communication Networks
- Computational Geometry
- Computer Graphics and Computer Vision
- Computer Systems
- Cryptography and Computer Security
- Data Compression
- Data Structures
- Databases and Information Retrieval
- Distributed and Parallel Computing
- Experimental Algorithmics
- Graph Algorithms
- Internet and Network Algorithms
- Machine Learning
- On-line Problems
- Quantum Computing
- Pattern Matching
- Robotics
- Scheduling and Resource Allocation Problems
- Symbolic Computation

Applications in the Sciences and Business such as:

- Bioinformatics
- Economics
- Manufacturing
- Finance
- Sociology

Important Deadlines

SUBMISSION DEADLINES
June 29, 2009, 4:59 PM EDT - Paper Registration Deadline
July 6, 2009, 4:59 PM EDT - Final Submission Deadline

PRE-REGISTRATION DEADLINE
December 14, 2009

HOTEL RESERVATION DEADLINE
December 14, 2009

3600 Market Street, 6th Floor | Philadelphia, PA 19104-2688 USA
Phone: +1-215-382-9800 | FAX: +1-215-386-7999

Subject: charge codes: Re: LAUR
From: Milan Bradonjic <bradonjic@gmail.com>
Date: Mon, 29 Jun 2009 11:38:28 -0600
To: Nicole Romero <niromero@lanl.gov>

Hi Nicole: Here are the codes:

70% on X9HN
30% on X9M1-3000-BRAD

Best, Milan

On Mon, Jun 29, 2009 at 9:27 AM, Nicole Romero<niromero@lanl.gov> wrote:

> You do not need another LAUR number unless something was changed in the paper.
>
> Nicole
>
> Milan Bradonjic wrote:
>
>> Hi Nicole,
>>
>> I got the LAUR number for a paper to submit to a conference. That paper has been rejected at that conference.
>>
>> Now, I want to resubmit the paper to the other conference, Do I need a new LAUR number or not?
>>
>> Best, Milan