Models of Sparse Random Graphs and Network Algorithms

Nicolas Broutin (Inria, France),
Luc Devroye (McGill, Canada),
Gábor Lugosi (Pompeu Fabra and ICREA, Spain)

Feb 5–10, 2012

1 Overview of the Field

There is no doubt now that the current trend that every electronic device should be connected in one way or another (usually many) implies a greater need for efficient networks. These networks should be connected, exhibit small-world behavior and their topology should be robust to local modifications like device movement; all this should be achieved using minimal and distributed processes. The workshop we held at BIRS between Feb 5–10 2012 was aimed at a deeper understanding of some of the major network models.

The topic of the workshop was mainly *sparse geometric graphs* and their use as models for wireless, bluetooth and ad-hoc networks, but some more general models of sparse networks have also been discussed. We focused in particular on quantitative indicators of the quality of the network, and of performance of the main communications algorithms. We first describe the important aspects of the subject area: the models, the quantities of interest and the relationships with neighboring fields.

The community working in the field aims at designing and understanding models of networks. A good analysis is crucial in that it permits to improve the design and to adapt it to the needs. The models of interest arise naturally either from concrete applications and the constraints imposed by physics (random geometric graphs) or as essential objects in more fundamental questions (Erdős–Rényi random graphs, Achlioptas processes). The parameters that are crucial to the quality of a model are

- CONNECTIVITY: in general, it should be possible to go from any point to any other; but it it is sometimes acceptable if a vast majority of the nodes are interconnected (there is a giant connected component).
- MAGNITUDE OF DISTANCES: the diameter should be rather short to ensure that it is (at least theoretically) possible to move quickly from any point to any other; again, constraining the diameter is sometimes too strong a requirement and one settles for short typical distances.
- SPARSITY: it is rather easy to design networks that are very connected, and with short distances, just take a complete graph. Such a topology is however unacceptable for reasons of cost and/or scalability, and one would like the graph to be as sparse as possible.
- NAVIGABILITY: for sparse connected graph with short distances, one is only certain that short paths exist, but they may be difficult to find using local information only. To be useful in practice, one would like that the network be navigable, i.e., that short paths are easily found by some distributed algorithm using local information.

- DIFFUSIVITY: finally, one would like broadcast algorithms to perform very well. The analysis here is usually important since one would like good information to spread quickly, but also to be able to stop efficiently the propagation of viruses.

We will finish this quick overview of the field by defining the two main classical models of interest, Erdős-Rényi random graphs, and random geometric graphs. Most of the interesting and more complex models build on either of these. We will then discuss the recent developments and open questions in Section 2.

ERDŐS–RÉNYI RANDOM GRAPHS. The model introduced by Erdős–Rényi [25, 26] is one of the most studied models of random networks. In spite of its simplicity it exhibits very interesting phenomena. It is usually the model one keeps in mind to understand more complex networks. A random graph $G(n, p)$ is a graph (V, E) on n labelled vertices $V = \{1, 2, \ldots, n\}$ where the edge set E consists of a random sample of the $\binom{n}{2}$ possible edges where each one is present with probability p, independently of the others. The structure of typical graphs varies greatly depending on the value of p with respect to n. With the parametrization $p = c/n$, where c is a fixed constant, one sees that the largest component of $G(n, p)$ has size at most $O(\log n)$ when $c < 1$, $\Theta(n^{2/3})$ when $c = 1$ or $\Theta(n)$ when $c > 1$. This last phase is most interesting since a large part of the graph is then interconnected. The graph only becomes completely connected when p is of the order of $\log(n)/n$ so that the average degree is of order $\log n$.

RANDOM GEOMETRIC GRAPHS. The model dates back to Gilbert [28] who introduced an underlying spatial topology for the network. One is given a connected domain, $[0, 1]^d$ for simplicity, and draws a set of n uniformly random points $\mathcal{X} = \{X_1, \ldots, X_n\}$. The graph $G(n, r)$ is the graph with vertex set $\mathcal{X}$ where two nodes are tied by an edge if the Euclidean distance between them is at most r. Again, the structure of a typical graph depends highly on the respective values of r and n. When $r = (\lambda/n)^{1/d}$ for some fixed constant λ, the average degree is of order λ. As in the $G(n, p)$ model, the is a phase transition for the connectivity: there is a critical value λ_c such that the largest connected component has linear size if and only if $\lambda > \lambda_c$. A typical graph is connected only when r is of order $(\log(n)/n)^{1/d}$ so that, again, the average degree is about $\log n$.

2 Recent Developments and Open Problems

2.1 Sparse connected graphs

As we mentioned above, the typical graphs only get connected when the average degree is logarithmic in the size. This is usually not acceptable since this raises the question of scalability: a typical node of the network cannot have a number of links that grow with the size of the graph.

SPARSE RANDOM GEOMETRIC GRAPH MODELS. The natural model for a spatial network is the *random geometric graph* $G(n, r)$, where n points are randomly distributed in space, say the unit square $[0, 1]^2$, and edges are added between any pair of vertices whose Euclidean distance does not exceed some value r [33, 32]. Above the threshold $r^\star(n)$ for connectivity, the average degree is of order $\Omega(\log n)$ so that graph is too dense for pratical reasons. This is why a number of related models have been proposed, that built a *sparse connected overlay* of the geometric graph. An example is provided by the *irrigation graph* introduced by Dubhashi, Häggström, Johansson, Panconesi, and Sozio [23]. In this model, each node chooses independently a number $c(n)$ of its geometric neighbours (at distance at most $r(n)$) with which it establishes a connection. One would like to study how small one can pick $c(n)$ and still obtain a nice enough network. More precisely, the properties that are crucial for applications are related to the quantitative measures of the *connectivity* and the ability to design efficient *communications algorithms*, both highly dependent upon the network model.

2.2 The problem of distances and navigability

The good connectivity of the graph is only a minimal requirement for the network. There should obviously be a trade-off between connectivity and sparsity of the network. This trade-off must also take into account the fact that the network should be usable: it should be relatively fast and easy to find short routes between vertices.

DISTANCES AND SMALL-WORLD PHENOMENON. One important property concerns the magnitude of distances: indeed the communication time between two nodes is lower bounded by the shortest distance between them.Branching arguments show that the scale of distances is at least of order $1/r(n)$; so one would want the distances to have a magnitude approaching this lower bound. In other words, one would like the graph to exhibit *small-world* behavior. Many non-geometric models exhibit such properties, and for these precise asymptotics have been proved for quantities such as the typical distance or the diameter (maximum shortest pairwise distance) [36, 10, 34, 16]. However, for geometric models, the known results about distances mostly characterize the growth rate and more precise results are still lacking.

NAVIGABILITY, CONDUCTANCE AND BROADCASTING. For the navigability, one would like that the short paths, if they exists, are easily found in a random way. Two main parameters of interest here are *mixing time* and the *cover time*. The mixing time is the length of the transition period before the behavior stabilizes to the stationary distribution. It is related to the spectral properties of the graph and to its conductance, which measures the uniform expansion quality of the network. The cover time measures the number of steps required for the simple random walk to visit every node. In some sense, it is the simplest (of course inefficient in practice) toy model for broadcast in a graph where a *unique messenger* must inform every user in person; the practical version —in which every informed person becomes a messenger— is considered below. For random geometric graphs, the cover time has been addressed by [17] in dimension three and up; the most important case of dimension two is still unknown.

In practice, one of the broadcast strategies that has received a lot of attention is based on *rumor spreading* and relates to the propagation of epidemics in a population [21]. There are three versions, *push* and *pull* and *push–pull*. Some piece of information is originally in the hands of some user of the network. The algorithm consists in spreading the information in the network by pushing (each user randomly chooses a neighbor to transmit information to) or pulling (each user asks information from a randomly selected neighbor) or both. The performance of the algorithm has been studied under various assumptions on the underlying graph. Only recently [15] have used a more general approach that characterizes the broadcast time in terms of the graph conductance. In some sense, it relates the cover time of the branching random walk to the spectral properties of a single random walk and opens a very promising route towards more general characterizations of the performance of broadcast algorithms.

2.3 Diffusions in random networks

To of the main questions about propagations in networks are about diffusion of new technology (bootstrap percolation) and diffusion of a rumor.

BOOTSTRAP PERCOLATION Some processes on random networks are of great interest. The first one, bootstrap percolation, has been used to model the adoption of new technology by a population. The process takes place on any (connected) graph. One is given a set of vertices that is initially infected (has the new technology). The infection then propagates deterministically: any node that with at least k infected neighbors becomes infected. These new infected vertices can, in turn, contributed to the propagation of the infection. One says that we have percolation if the entire graph gets infected. The most important question consists in determining the proportion p_c of nodes that have to be infected at random in the first place to ensure percolation. In order to understand the influence of locality on the process, random geometric graphs is a natural model to study bootstrap percolation. Recently, some first bounds on the critical threshold p_c have been obtained by Bradonjić and Saniee [14].

RUMOR SPREADING USING CONDUCTANCE. A first step towards more general results about broadcasting algorithms consists in characterizing the time for rumor spreading in terms of an important parameter of the connectivity of the graph that would be computed for specific examples. The first result in this direction is due to Chierichetti, Lattanzi, and Panconesi [15] who provide almost tight bounds for the broadcast time using a *push–pull* strategy in terms of the graph conductance. We will discuss ways to nail down the correct asymptotics and discuss what other important graph invariant might be more suitable to express the asymptotic bounds in a useful way.

2.4 Fundamental questions of universality

Aside from the parameters of practical interest, one would like to understand better the models themselves and the common behaviors that they exhibit.

UNIVERSALITY OF RESCALED COMPONENT SIZES. A first step towards the understanding of the metric structure consists in understanding the component sizes. Aldous [3] has shown that the component sizes of critical Erdős–Rényi random graphs were asymptotically following a very specific coagulation process, the multiplicative coalescent. It seems that this process also describes the evolution of the asymptotic sequence of component sizes in many other random graphs processes, like inhomogeneous random graphs [8], the Bohman–Frieze process [7] as well as other Achlioptas processes (where one is given rules to choose an edge from a list of random available ones). The question of the universality of the multiplicative coalescent for natural models of random graphs is one of the fundamental and difficult questions.

TOWARDS UNIVERSALITY FOR GRAPH PARAMETERS. One of the more challenging questions concerns the universality of the behavior of graph models. Rather than analyzing the models one by one, one would gain a lot of insight by adopting a more abstract point of view. The main obstacle here is to define a suitable measure of similarity between graphs. The question of metrics on the space of graphs and continuity of the parameters in the induced topology is crucial here. The question has been very successfully answered for dense graphs with the notion of graph limits by [30] (see also [20]). For sparse graphs, the question is more subtle and many natural metrics (like the Gromov–Hausdorff metric) do not yield good topologies on the spaces of sparse graph that are not "critical" (only supercritical (branching) graphs are crucial in practice for their expansion properties.) The recent survey by [12] provides possible approaches. The question of convergence of geometric networks is also the main theme of Aldous' talk at the ICM [2].

3 Presentation Highlights

The talks covered the spectrum of topics we intended to discuss. In particular, there were talks about properties of random geometric graphs, the structure and distance in random sparse graphs, some applications of sparse graphs, as well as some presentations about related more fundamental questions.

3.1 Random geometric graphs

Nicolas Fraiman — *Connectivity of Bluetooth graphs*

We study the connectivity of random Bluetooth graphs, these are obtained as irrigation subgraphs of the well-known random geometric graph model. There are two parameters that control the model: the radius r that determines the visible neighbors of each node and the number of edges c that each node is allowed to have. The randomness comes from the distribution of nodes in space and the choices of each vertex. We characterize the connectivity threshold (in c) for values of r close the critical value for connectivity in the underlying random geometric graph. This is joint work with Nicolas Broutin, Luc Devroye and Gabor Lugosi [13].

Tobias Muller — *Colouring random geometric graphs*

If we pick points $X_1, \dots, X_n$ at random from d-dimensional space (i.i.d. according to some probability measure) and fix a $r > 0$, then we obtain a random geometric graph by joining points by an edge whenever their distance is $< r$. I will talk about some results on the chromatic number and the clique number of this model [31].

Matthew Penrose — *Connectivity of G(n,r,p)*

Consider a graph on n vertices placed uniformly independently at random in the unit square, in which any two vertices distant at most r apart are connected by an edge with probability p. This generalizes both the classical random graph and the random geometric graph. We discuss the chances of its being disconnected without having any isolated vertices, when n is large, for various choices of the other parameters.

Joseph Yukich — *Probabilistic Analysis of Some Geometric Networks*

We survey some techniques for establishing general limit theorems in stochastic geometry (laws of large numbers, variance asymptotics, and central limit theorems). We show how the general theorems may be applied to deduce the limit theory for various functionals of random geometric graphs, including, for example, network connectivity functionals, clique count, total edge length, and component count. The talk is based on joint work with M. Penrose, T. Schreiber, and Y. Baryshnikov.

3.2 Structure and distances in sparse graphs

Shankar Bhamidi — *Limited choice and randomness in evolution of networks*

The last few years have seen an explosion in network models describing the evolution of real world networks. In the context of math probability, one aspect which has seen an intense focus is the interplay between randomness and limited choice in the evolution of networks, ranging from the description of the emergence of the giant component, the new phenomenon of "explosive percolation" and power of two choices. I will describe on going work in understanding such dynamic network models, their connections to classical constructs such as the standard multiplicative coalescent and local weak convergence of random trees.

Justin Salez — *Joint distribution of distances in large random regular networks.*

We study the array of point-to-point distances in large random regular graphs equipped with exponential edge-weights. The asymptotic marginal distribution of a single entry is now well-understood, thanks to the work of Bhamidi, van der Hofstad and Hooghiemstra (2010). In this talk, we will show that the whole array, suitably re-centered, converges in the weak sense to a rather simple infinite random array. Our proof consists in analyzing the invasion of the network by several mutually exclusive flows emanating from different sources and propagating simultaneously at unit rate along the edges. The result applies to both the random regular multi-graph produced by the configuration model and the uniform regular simple graph.

3.3 Applications of sparse graphs

David Aldous — *Some thoughts on data compression and entropy for sparse graphs with vertex-names*

After an informal review of classic Shannon theory of entropy and data compression for random sequences, I will speculate on analogs for sparse graphs with vertex-names.

Marc Lelarge — *A new approach to the orientation of random hypergraphs*

A h-uniform hypergraph $H = (V, E)$ is called (ℓ, k)-orientable if there exists an assignment of each hyperedge $e \in E$ to exactly ℓ of its vertices $v \in e$ such that no vertex is assigned more than k hyperedges. Let $H_{n,m,h}$ be a hypergraph, drawn uniformly at random from the set of all h-uniform hypergraphs with n vertices and m edges. In this paper, we determine the threshold of the existence of a (ℓ, k)-orientation of $H_{n,m,h}$ for $k \geq 1$ and $h > \ell \geq 1$, extending recent results motivated by applications such as cuckoo hashing or load balancing with guaranteed maximum load. Our proof combines the local weak convergence of sparse graphs and a careful analysis of a Gibbs measure on spanning subgraphs with degree constraints. It allows us to deal with a much broader class than the uniform hypergraphs.

Pat Morin — *Maximum interference in the highway and related models.*

Given a set D of n disks, the interference of a point p is defined as the number of disks of D that contain p. The interference of D is the maximum interference over all centers of disks in D. In this talk, we discuss upper and lower bounds on maximum interference in 1 dimension, 2 dimensions, in the worst case, and in probabilistic settings.

3.4 Fundamental questions

Charles Bordenave — *How does a uniformly sampled Markov chain behave ?*

This is joint work with P. Caputo and D. Chafai. In this talk, we will consider various probability distributions on the set of stochastic matrices with n states and on the set of Laplacian/Kirchhoff matrices on n states. They will arise naturally from the conductance model on n states with i.i.d conductances. With the help of random matrix theory, we will study the spectrum of these processes. An emphasis will be put on the case of the simple random walk on a sparse directed Erdős-Rényi graph.

Csaba Toth — *Convex partitions.*

A convex partition is a planar straight-line graph where every bounded face is convex and the complement of the outer face is also convex. Two results are presented in this talk: (1) For every n points in the plane, there is a convex partition G such that the total edge length of G is at most $O(\log n/\log\log n)$ times that of a Euclidean minimum spanning tree (EMST) for the n points, and this bound is the best possible. (2) If G is a convex partition and the outer face has $O(1)$ edges, then G contains a monotone path of at least $\Omega(\log n/\log\log n)$ edges, and this bound is the best possible. (Joint work with Adrian Dumitrescu.)

4 Some open questions

One of the main objectives of the workshop was to foster new collaborations between the participants. To this aim, we organized an open problem session the day of arrival and scheduled informal working sessions every afternoon. Some of the proposed questions are famous and notoriously difficult, and were brought up to get the opinions of participants with a fresh eye and a different background, others have spawn from the participants current interests. In any case, the set of open questions which have been proposed has great scientific value, and we reproduce them here:

LENGTH OF THE GREEDY TOUR (Presented by C. Bordenave). Thow n uniform points on the unit square, and consider the total length L_n of the greedily constructed traveling salesman path. More precisely, start from a random point, then at any stage move to the closest yet non-visited point. Let L_n denote the length of this path. It is possible to show that L_n is of order $\sqrt{n}$ [6]. Prove that $L_n/\sqrt{n}$ converges to a constant. Many problems of this kind, when the path constructed is in some sense optimal, have been solved using sub-additive arguments (see, e.g., [35]). Here the main argument collapse since the path is constructed in a greedy way.

GREEDY TOUR IN THE PLANE (Presented by C. Bordenave). In an other version of the previous problem, one starts from a homogeneous Poisson point process $\mathcal{P}$ in the plane, conditioned to have a point at the origin. Starting from the origin, proceed to a walk which always visits the nearest point of $\mathcal{P}$ not yet visited. Does this walk eventually visit every point of $\mathcal{P}$?

GREEDY ALGORITHM UNDER POISSON RAIN. Consider a homogeneous Poisson point process on $\mathbb{R}^2 \times [0,\infty)$. The first two coordinates are interpreted as space, and the third one as the time at which the point defined by the first two coordinates arrives. Suppose that an agent always goes towards the closest non-visited point. In particular, if at some time t a point appears that is closer than all the others, the agent changes direction to aim at the newly arrived point. Let V_t be the number of points visited before time t. Is it true that $\liminf_{t\to\infty} V_t/t > 0$?

ASYMPTOTICS FOR GENERALIZED U-STATISTICS (presented by J. Yukich). U-statistics are a way to bootstrap an r-sample estimator into an n-sample estimator. Given a real-valued function f of r variables, the U-statistic $f_n : \mathbb{R}^n \to \mathbb{R}$ is the average over distinct ordered r-subsamples. Generalized U-statistics of the form $\sum_{i,j} f_n(X_i, X_j, D_n)$ where the X_i are independent and identically distributed in $\mathbb{R}^d$, $D_n = (X_1, ..., X_n)$, and f_n is some translation invariant function are of great interest. Can they be understood using a variation of the stabilization method?

ABOUT INTERFERENCE (presented by P. Morin) Given a set D of n disks, the interference of a point p is defined as the number of disks of D that contain p. The interference of D is the maximum interference over all centers of disks in D.

For a point set V, and a graph G, the set D consists of the disks centered at the points v of V with radius equal to the length of the longest edge adjacent to v. The interference $I(G)$ of the graph G is the interference of that set D.

- Given a point set $V \subset \mathbb{R}$, can we compute at graph G that approximately minimizes $I(G)$?
- Is the following statement true: "for any $V \subset \mathbb{R}^d$, there exists a graph $G = (V, E)$ with $I(G) = O(\sqrt{n})$"
- Is there an algorithm for finding a graph G that approximately minimizes $I(G)$? A $5/4$-approximation is best possible.
- If V consists of n i.i.d. points uniform on $[0, 1]$ the minimum spanning tree has interference $\Theta(\sqrt{n})$. Is there a better graph that gives $o(\sqrt{n})$? What is the interference of the minimum spanning tree of n i.i.d. uniform points in $[0, 1]^d$?
- Let $G^\star$ minimize the interference $I(G)$ for a n i.i.d. points uniform on $[0, 1]^d$. What is $\mathbb{E}[I(G^\star)]$? Previous construction show an upper bound of $O((\log n)^{1/2})$.

DIAMETER OF THE EUCLIDEAN MST (presented by L. Addario-Berry). The minimum spanning tree (MST) is one of the most important sparse graphs. A lot is known about the local properties of minimum spanning trees; much less is known about the typical distances, or about the diameter. In a geometric setting, take n i.i.d. uniform points in the square and use Euclidean distance to weight the edges of a complete graph on n vertices. What can we say about the expected diameter D_n of the corresponding random Euclidean minimum spanning tree? Only the trivial bounds $D_n = O(n)$ and $D_n = \Omega(\sqrt{n})$ are known.

THE EXTRA-COST FOR GUARDING SCULPTURES. (presented by L. Addario-Berry). An art gallery is a simple polygon. One is asked to place a minimal set of guards to that every point of the interior perimeter of the gallery (where the paintings lie) is seen by at least one guard. A related question concern the extra cost needed to also ensure that the interior of the polygon (where the sculptures lie) is also guarded. Addario, Amini, Séréni and Thomassé proved that if n guards are needed for the walls, then the extra cost for sculptures is at most $4n - 6$ [1]. It seems that $n - 2$ extra guards should suffice.

RESILIENT SPANNERS (presented by V. Dujmovic). Highly connected and yet sparse graphs (such as expanders or graphs of high tree-width) are fundamental, widely applicable and extensively studied combinatorial objects. Can we find such graphs that are robust to failures of some of the nodes in the following sense: Given a point set V, $|V| = n$, is it possible to construct a graph $G = (V, E)$ such that for any $V' \subseteq V$, with $|V'| \leq f(n)$ the subgraph G induced on $V \setminus V'$ has a connected component of size $n - o(n)$ that is a sparse (at most $g(n)$ edges) t-spanner? For instance, it is possible with $g(n) = O(n)$ and $f(n) = \sqrt{n}$?

5 Scientific Progress Made

The schedule was made with only a limited number of presentations in order to leave most of the time for discussions between the participants. In particular, we organized an open problem session on the first day, and working sessions every afternoon. We believe that the workshop was very successful with respect the exchanges it fostered. On top of the informal discussions whose long term impact is difficult to estimate, a number of open problems have been solved, some of which have also already been submitted to a conference or a journal [9, 19]. Most concrete outcomes that have been produced or will be published in the very near future concern groups of people with similar backgrounds, and this is inevitable. However some of the most promising work involve participants with different backgrounds; such collaborations take of course more time to reach a ripe state.

INTERFERENCE GRAPHS. Pat Morin and Luc Devroye have worked on interference in random geometric graphs. They improved the bounds on the interference of graphs on a set V of n uniformly random points in $[0, 1]^d$. In particular, they showed that there exists a connected graph on V that has interference of order $(\log n)^{1/3}$ and that no connected graph on V has interference $o((\log n)^{1/4})$. They also showed that the minimum spanning tree on V has interference $\Theta((\log n)^{1/2})$. Their results have been collected in a manuscript that has been submitted [19].

ROBUST GEOMETRIC SPANNERS. Jit Bose, Pat Morin and Vida Dujmovic have worked on the robustness of highly connected yet sparse graphs. In particular, they initiated the study of such graphs that are in addition geometric spanners. Following the suggestions the open problem proposed by V. Dujmovic, they defined a robustness property and proved that robust spanners must have a superlinear number of edges, even in one dimension. On the positive side, they also give constructions, for any dimension, of robust spanners with a near-linear number edges. These results have already been submitted [9].

IRRIGATION GRAPHS WITH CONSTANT OUT-DEGREE. Full connectivity of a network is of little importance in practice, there will always be some remote and isolated points, and this does not affect the quality of the network from the provider's point of view. What matters is the size of the largest connected component. S. Boucheron, N. Broutin, L. Devroye and G. Lugosi have worked on this question on the model that was the subject of N. Fraiman's lecture, the connectivity of bluetooth graphs. They have proved that even when the range of reach r of a point is only slightly above the connectivity threshold of the random geometric graph, keeping only two random neighbors suffices to ensure that the largest connected component has size $n-o(n)$.

References

[1] L. Addario-Berry, O. Amini, J.S. Sereni, and S. Thomassé, , Guarding art galleries: The extra cost for sculptures is linear. In *Proceedings of the 11th Scandinavian workshop on Algorithm Theory*, Lecture Notes in Computer Science vol. 5124, pp. 41–52, Springer, 2008.

[2] D.J. Aldous. Exchangeability and continuum limits of discrete random structures. In *Proceedings of the International Congress of Mathematicians*, 2010.

[3] D. Aldous. Brownian excursions, critical random graphs and the multiplicative coalescent. *The Annals of Probability*, vol. 25, pp. 812–854, 1997.

[4] D. Aldous and V. Limic. The entrance boundary of the multiplicative coalescent. *Electronic Journal of Probability*, vol. 3, pp. 1–59, 1998.

[5] D.J. Aldous and J. Shun. Connected spatial networks over random points and a route-length statistic. arXiv:1003.3700 [math.PR], 2010.

[6] J.L. Bentley and J. Saxe. An analysis of two heuristics for the Euclidean traveling salesman problem. In *Proceedings of the 18th Annual Allerton Conference on Communication, Control, and Computing*, pp. 41–49, 1980.

[7] S. Bhamidi, A. Budhiraja, and X. Wang. Bohman-Frieze processes at criticality and emergence of the giant component. Submitted, 2011. arXiv:1106.1022 [math.PR]

[8] S. Bhamidi, R. van der Hofstad, and J.S.H. van Leeuwaarden. Novel scaling limits for critical inhomogeneous random graphs. *The Annals of Probability*, to appear, 2012. arXiv:0909.1472 [math.PR].

[9] P. Bose, V. Dujmovic, P. Morin and M. Smid Robust geometric spanners. Submitted, arXiv:1204.4679 [CS.CG], 2012.

[10] B. Bollobás and O. Riordan. The diameter of a scale-free random graph. *Combinatorica*, 24:5–34, 2004.

[11] B. Bollobás and O. Riordan. *Percolation*. Cambridge University Press, 2006.

[12] B. Bollobás and O. Riordan. Metrics for sparse graphs. In *Surveys in Combinatorics*, volume 365 of *LMS Lecture Notes Series*, pages 211–287, Cambridge, UK, 2009. Cambridge University Press.

[13] N. Broutin, L. Devroye, N. Fraiman, and G. Lugosi. Connectivity threshold for Bluetooth graphs. Submitted, 2011. arXiv:1103.0351 [math.PR]

[14] M. Bradonjić and I. Saniee. Bootstrap percolation on random geometric graphs. arXiv:1201.2953 [cs.DM], 2012.

[15] F. Chierichetti, S. Lattanzi, and A. Panconesi. Almost tight bounds for rumour spreading with conductance. In *Proceedings of STOC*, 2010.

[16] F. Chung and L. Lu. The diameter of random sparse graphs. *Advances in Applied Mathematics*, 26: 257–279, 2001.

[17] C. Cooper and A. Frieze. The cover time of random geometric graphs. *Random Structures and Algorithms*, 2010.

[18] P. Crescenzi, C. Nocentini, A. Pietracaprina, and G. Pucci. On the Connectivity of Bluetooth-Based Ad Hoc Networks. *Concurrenty and Computation: Practice and Experience*, 21:875–887, 2009.

[19] L. Devroye and P. Morin. A note on interference in random networks. ArXiv:1202.5945 [CS.DS]. Submitted, 2012.

[20] P. Diaconis and S. Janson. Graph limits and exchangeable random graphs. *Rendiconti di Matematica e delle sue Applicazioni*, 28:33–61, 2008.

[21] M. Draief and L. Massoulié. *Epidemics and Rumours in Complex Networks*, volume 369 of *London Mathematical Society Lecture Notes*. Cambridge University Press, Cambridge, 2010.

[22] D. Dubhashi and A. Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, Cambridge, UK, 2009.

[23] D. Dubhashi, O. Häggström, C. Johansson, A. Panconesi, and M. Sozio. Irrigating ad hoc networks in constant time. In *Proceedings of SPAA'05*, pages 106–115, 2005.

[24] R. Durrett. *Random Graph Dynamics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2006.

[25] P. Erdős aand A. Rényi. On random graphs I *Publ. Math. Debrecen*, vol. 6, pp. 290–297, 1959.

[26] P. Erdős and A. Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hungar. Acad. Sci.*, vol. 5, pp. 17–61, 1960.

[27] M. Franceschetti and R. Meester. *Random Networks for Communication: From Statistical Physics to Information Systems*. Cambridge University Press, 2008.

[28] E.N. Gilbert. Random plane networks. *Journal of the Society for Industrial and Applied Mathematics*, vol 9, pp. 533–543, 1961.

[29] H. Kesten. The critical probability for bond percolation on the square lattice equals $\frac{1}{2}$. *Communications in Mathematical Physics*, 74:41–59, 1980.

[30] L. Lovász and B. Szegedy. Limits of dense graph sequences. *Journal of Combinatorial Theory, Series B*, 96:933–957, 2006.

[31] C. McDiarmid and T. Müller. On the chromatic number of random geometric graphs. *Combinatorica*, vol. 31, pp. 1–66, 2011.

[32] R. Meester and R. Roy. *Continuum Percolation*. Cambridge Tracts in Mathematics. Cambridge University Press, 1996.

[33] M. Penrose. *Random Geometric Graphs*. Oxford Studies in Probability. Oxford University Press, 2003.

[34] O. Riordan and N.C. Wormald. The diameter of sparse random graphs. arXiv:0808.4067 [math.PR], 2008.

[35] J.M. Steele. *Probability Theory and Combinatorial Optimization*. CBMS-NSF Regional Conference Series in Applied Mathematics, SIAM, Philadelphia, 1997.

[36] D.J. Watts and S.H. Strogatz. Collective dynamics of ‘small-world’networks. *Nature*, 393:440–442, 1998.