Cambridge Series in Statistical and Probabilistic Mathematics

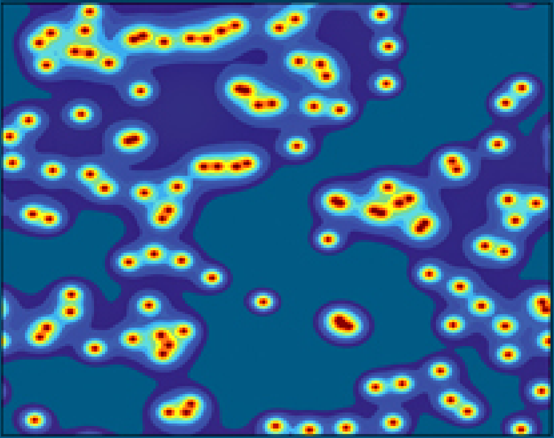

Random Networks for Communication

From Statistical Physics to Information Systems

Massimo Franceschetti, Ronald Meester

CAMBRIDGE www.cambridge.org/9780521854429

This page intentionally left blank

Random Networks for Communication

When is a network (almost) connected? How much information can it carry? How can you find a particular destination within the network? And how do you approach these questions – and others – when the network is random?

The analysis of communication networks requires a fascinating synthesis of random graph theory, stochastic geometry and percolation theory to provide models for both structure and information flow. This book is the first comprehensive introduction for graduate students and scientists to techniques and problems in the field of spatial random networks. The selection of material is driven by applications arising in engineering, and the treatment is both readable and mathematically rigorous. Though mainly concerned with information-flow-related questions motivated by wireless data networks, the models developed are also of interest in a broader context, ranging from engineering to social networks, biology, and physics.

MASSIMO FRANCESCHETTI is assistant professor of electrical and computer engineering at the University of California, San Diego. His work in communication system theory sits at the interface between networks, information theory, control, and electromagnetics.

RONALD MEESTER is professor of mathematics at the Vrije Universiteit Amsterdam. He has published broadly in percolation theory, spatial random processes, self-organised criticality, ergodic theory, and forensic statistics and is the author of *Continuum Percolation* (with Rahul Roy) and *A Natural Introduction to Probability Theory*.

CAMBRIDGE SERIES IN STATISTICAL AND PROBABILISTIC MATHEMATICS

This series of high-quality upper-division textbooks and expository monographs covers all aspects of stochastic applicable mathematics. The topics range from pure and applied statistics to probability theory, operations research, optimization, and mathematical programming. The books contain clear presentations of new developments in the field and also of the state of the art in classical methods. While emphasizing rigorous treatment of theoretical methods, the books also contain applications and discussions of new techniques made possible by advances in computational practice.

Already published

1. *Bootstrap Methods and Their Application*, by A. C. Davison and D. V. Hinkley
2. *Markov Chains*, by J. Norris
3. *Asymptotic Statistics*, by A. W. van der Vaart
4. *Wavelet Methods for Time Series Analysis*, by Donald B. Percival and Andrew T. Walden
5. *Bayesian Methods*, by Thomas Leonard and John S. J. Hsu
6. *Empirical Processes in M-Estimation*, by Sara van de Geer
7. *Numerical Methods of Statistics*, by John F. Monahan
8. *A User's Guide to Measure Theoretic Probability*, by David Pollard
9. *The Estimation and Tracking of Frequency*, by B. G. Quinn and E. J. Hannan
10. *Data Analysis and Graphics using R*, by John Maindonald and W. John Braun
11. *Statistical Models*, by A. C. Davison
12. *Semiparametric Regression*, by D. Ruppert, M. P. Wand, R. J. Carroll
13. *Exercises in Probability*, by Loic Chaumont and Marc Yor
14. *Statistical Analysis of Stochastic Processes in Time*, by J. K. Lindsey
15. *Measure Theory and Filtering*, by Lakhdar Aggoun and Robert Elliott
16. *Essentials of Statistical Inference*, by G. A. Young and R. L. Smith
17. *Elements of Distribution Theory*, by Thomas A. Severini
18. *Statistical Mechanics of Disordered Systems*, by Anton Bovier
19. *The Coordinate-Free Approach to Linear Models*, by Michael J. Wichura
20. *Random Graph Dynamics*, by Rick Durrett
21. *Networks*, by Peter Whittle
22. *Saddlepoint Approximations with Applications*, by Ronald W. Butler
23. *Applied Asymptotics*, by A. R. Brazzale, A. C. Davison and N. Reid

Random Networks for Communication

From Statistical Physics to Information Systems

Massimo Franceschetti
and
Ronald Meester

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo

Cambridge University Press
The Edinburgh Building, Cambridge CB2 8RU, UK
Published in the United States of America by Cambridge University Press, New York

www.cambridge.org
Information on this title: www.cambridge.org/9780521854429

First published in print format 2008

ISBN-13 978-0-511-37246-9 eBook (Adobe Reader)

ISBN-13 978-0-521-85442-9 hardback

Contents

Preface

What is this book about, and who is it written for? To start with the first question, this book introduces a subject placed at the interface between mathematics, physics, and information theory of systems. In doing so, it is not intended to be a comprehensive monograph and collect all the mathematical results available in the literature, but rather pursues the more ambitious goal of laying the foundations. We have tried to give emphasis to the relevant mathematical techniques that are the essential ingredients for anybody interested in the field of random networks. Dynamic coupling, renormalisation, ergodicity and deviations from the mean, correlation inequalities, Poisson approximation, as well as some other tricks and constructions that often arise in the proofs are not only applied, but also discussed with the objective of clarifying the philosophy behind their arguments. We have also tried to make available to a larger community the main mathematical results on random networks, and to place them into a new communication theory framework, trying not to sacrifice mathematical rigour. As a result, the choice of the topics was influenced by personal taste, by the willingness to keep the flow consistent, and by the desire to present a modern, communication-theoretic view of a topic that originated some fifty years ago and that has had an incredible impact in mathematics and statistical physics since then. Sometimes this has come at the price of sacrificing the presentation of results that either did not fit well in what we thought was the ideal flow of the book, or that could be obtained using the same basic ideas, but at the expense of highly technical complications. One important topic that the reader will find missing, for example, is a complete treatment of the classic Erdös–Rényi model of random graphs and of its more recent extensions, including preferential attachment models used to describe properties of the Internet. Indeed, we felt that these models, lacking a geometric component, did not fit well in our framework and the reader is referred to the recent account of Durrett (2007) for a rigorous treatment of preferential attachment models. Other omissions are certainly present, and hopefully similarly justified. We also refer to the monographs by Bollobás (2001), Bollobás and Riordan (2006), Grimmett (1999), Meester and Roy (1996), and Penrose (2003), for a compendium of additional mathematical results.

Let us now turn to the second question: what is our intended readership? In the first place, we hope to inspire people in electrical engineering, computer science, and physics to learn more about very relevant mathematics. It is worthwhile to learn these mathematics, as it provides valuable intuition and structure. We have noticed that there is a tendency to re-invent the wheel when it comes to the use of mathematics, and we thought it would be very useful to have a standard reference text. But also, we want to inspire mathematicians to learn more about the communication setting. It raises specific questions that are mathematically interesting, and deep. Such questions would be hard to think about without the context of communication networks.

In summary: the mathematics is not too abstract for engineers, and the applications are certainly not too mechanical for mathematicians. The authors being from both communities – engineering and mathematics – have enjoyed over the years an interesting and fruitful collaboration, and we are convinced that both communities can profit from this book. In a way, our main concern is the interaction between people at either side of the interface, who desire to *break on through to the other side*.

A final word about the prerequisites. We assume that the reader is familiar with basic probability theory, with the basic notions of graph theory and with basic calculus. When we need concepts that go beyond these basics, we will introduce and explain them. We believe the book is suitable, and we have used it, for a first-year graduate course in mathematics or electrical engineering.

We thank Patrick Thiran and the School of Computer and Communication Sciences of the École Politechnique Fédérale de Lausanne for hosting us during the Summer of 2005, while working on this book. Massimo Franceschetti is also grateful to the Department of Mathematics of the Vrije Universiteit Amsterdam for hosting him several times. We thank Misja Nuyens who read the entire manuscript and provided many useful comments. We are also grateful to Nikhil Karamchandani, Young-Han Kim, and Olivier Lévêque, who have also provided useful feedback on different portions of the manuscript. Massimo Franceschetti also thanks Olivier Dousse, a close research collaborator of several years.

List of notation

In the following, we collect some of the notation used throughout the book. Definitions are repeated within the text, in the specific context where they are used. Occasionally, in some local contexts, we introduce new notation and redefine terms to mean something different.

$\lvert\cdot\rvert$	Lebesgue measure
	Euclidean distance
	L_1 distance
	cardinality
$\lfloor\cdot\rfloor$	floor function, the argument is rounded down to the previous integer
$\lceil\cdot\rceil$	ceiling function, the argument is rounded up to the next integer
$\mathcal{A}$	an algorithm
	a region of the plane
a.a.s.	asymptotic almost surely
a.s.	almost surely
β	mean square constraint on the codeword symbols
B_n	box of side length $\sqrt{n}$
	box of side length n
$B_n^{\leftrightarrow}$	the event that there is a crossing path connecting the left side of B_n with its right side
$C(x)$	connected component containing the point x
C	connected component containing the origin
	channel capacity
$C(x,y)$	channel capacity between points x and y
	chemical distance between points x and y
C_n	sum of the information rates across a cut
$\partial(\cdot)$	inner boundary
$D(G)$	diameter of the graph G
$D(\mathcal{A})$	navigation length of the algorithm $\mathcal{A}$
d_{TV}	total variation distance
$E(\cdot)$	expectation
$g(x)$	connection function in a random connection model

$\overline{g}(\|x\|)$	connection function depending only on the Euclidian distance, i.e., $\overline{g} : \mathbb{R}^+ \to [0, 1]$ such that $\overline{g}(\|x\|) = g(x)$
G	a graph
G_X	generating function of random variable X
γ	interference reduction factor in the SNIR model
$I(z)$	shot-noise process
$\tilde{I}(z)$	shifted shot-noise process
I	indicator random variable
i.i.d.	independent, identically distributed
k_c	critical value in nearest neighbour model
λ	density of a Poisson process, or parameter of a Poisson distribution
λ_c	critical density for boolean or random connection model
$\Lambda(x)$	density function of an inhomogeneous Poisson process
$\ell(x, y)$	attenuation function between points x and y
$l(\|x-y\|)$	attenuation function depending only on the Euclidian distance, i.e., $l : \mathbb{R}^+ \to \mathbb{R}^+$ such that $l(\|x-y\|) = \ell(x, y)$
N	environmental noise
$N(A)$	number of pivotal edges for the event A
$N_\infty(B_n)$	number of Poisson points in the box B_n that are also part of the unbounded component on the whole plane
$N(n)$	number of paths of length n in the random grid starting at the origin
O	origin point on the plane
P	power of a signal, or just a probability measure
$Po(\lambda)$	Poisson random variable of parameter λ
p_c	critical probability for undirected percolation
$\vec{p}_c$	critical probability for directed percolation
p_c^{site}	critical probability for site percolation
p_c^{bond}	critical probability for bond percolation
p_α	critical probability for α-almost connectivity
$\psi(\cdot)$	probability that there exists an unbounded connected component
Q	the event that there exists at most one unbounded connected component
r_α	critical radius for α-almost connectivity in the boolean model
r_c	critical radius for the boolean model
R	rate of the information flow
$R(x, y)$	achievable information rate between x and y
$R(n)$	simultaneous achievable per-node rate in a box of area n
SNR	signal to noise ratio
SNIR	signal to noise plus interference ratio
T	a tree
	a threshold value
$\theta(\cdot)$	percolation function, i.e., the probability that there exists an unbounded connected component at the origin
U	the event that there exists an unbounded connected component

U_0	the event that there exists an unbounded connected component at the origin, when there is a Poisson point at the origin
W	channel bandwidth
	sum of indicator random variables
w.h.p.	with high probability
X	Poisson process
	a random variable
X_n	a sequence of random variables
X^m	a codeword of length m
$X(A)$	number of points of the Poisson process X falling in the set A
$X(e)$	uniform random variable in $[0, 1]$,
	where e is a random edge coupled with the outcome of X
$x \leftrightarrow y$	the event that there is a path connecting point x with point y
Z_n	nth generation in a branching process

1

Introduction

Random networks arise when nodes are randomly deployed on the plane and randomly connected to each other. Depending on the specific rules used to construct them, they create structures that can resemble what is observed in real natural, as well as in artificial, complex systems. Thus, they provide simple models that allow us to use probability theory as a tool to explain the observable behaviour of real systems and to formally study and predict phenomena that are not amenable to analysis with a deterministic approach. This often leads to useful design guidelines for the development and optimal operation of real systems.

Historically, random networks has been a field of study in mathematics and statistical physics, although many models were inspired by practical questions of engineering interest. One of the early mathematical models appeared in a series of papers starting in 1959 by the two Hungarian mathematicians Paul Erdös and Alfréd Rényi. They investigated what a 'typical' graph of n vertices and m edges looks like, by connecting nodes at random. They showed that many properties of these graphs are almost always predictable, as they suddenly arise with very high probability when the model parameters are chosen appropriately. This peculiar property generated much interest among mathematicians, and their papers marked the starting point of the field of random graph theory. The graphs they considered, however, were abstract mathematical objects and there was no notion of geometric position of vertices and edges.

Mathematical models inspired by more practical questions appeared around the same time and relied on some notion of geometric locality of the random network connections. In 1957, British engineer Simon Broadbent and mathematician John Hammersley published a paper introducing a simple discrete mathematical model of a random grid in which vertices are arranged on a square lattice, and edges between neighbouring vertices are added at random, by flipping a coin to decide on the presence of each edge. This simple model revealed extreme mathematical depth, and became one of the most studied mathematical objects in statistical physics.

Broadbent and Hammersley were inspired by the work they had done during World War II and their paper's motivation was the optimal design of filters in gas masks. The gas masks of the time used granules of activated charcoal, and the authors realised that proper functioning of the mask required careful operation between two extremes. At one extreme, the charcoal was highly permeable, air flowed easily through the cannister, but the wearer of the mask breathed insufficiently filtered air. At the other extreme, the charcoal pack was nearly impermeable, and while no poisonous gases got through, neither did sufficient air. The optimum was to have high charcoal surface area and tortuous paths for air flow, ensuring sufficient time and contact to absorb the toxin. They realised that this condition would be met in a critical operating regime, which would occur with very high probability just like Erdös and Rényi showed later for random graph properties, and they named the mathematical framework that they developed *percolation theory*, because the meandering paths reminded them of water trickling through a coffee percolator.

A few years later, in 1961, American communication engineer Edgar Gilbert, working at Bell Laboratories, generalised Broadbent and Hammersley's theory introducing a model of random planar networks in continuum space. He considered nodes randomly located in the plane and formed a random network by connecting pairs of nodes that are sufficiently close to each other. He was inspired by the possibility of providing long-range radio connection using a large number of short-range radio transmitters, and marked the birth of continuum percolation theory. Using this model, he formally proved the existence of a critical transmission range for the nodes, beyond which an infinite chain of connected transmitters forms and so long-distance communication is possible by successive relaying of messages along the chain. By contrast, below critical transmission range, any connected component of transmitters is bounded and it is impossible to communicate over large distances. Gilbert's ingenious proof, as we shall see, was based on the work of Broadbent and Hammersley, and on the theory of branching processes, which dated back to the nineteenth-century work of Sir Francis Galton and Reverend Henry William Watson on the survival of surnames in the British peerage.

Additional pioneering work on random networks appears to be the product of communication engineers. In 1956, American computer scientist Edward Moore and information theory's father Claude Shannon wrote two papers concerned with random electrical networks, which became classics in reliability theory and established some key inequalities, presented later in this book, which are important steps towards the celebrated threshold behaviours arising in percolation theory and random graphs.

As these early visionary works have been generalised by mathematicians, and statistical physicists have used these simple models to explain the behaviour of more complex natural systems, the field of random networks has flourished; its application to communication, however, has lagged behind. Today, however, there is great renewed interest in random networks for communication. Technological advances have made it plausible to envisage the development of massively large communication systems composed of small and relatively simple devices that can be randomly deployed and 'ad hoc' organise into a complex communication network using radio links. These networks can be used for human communication, as well as for sensing the environment and collecting and exchanging data for a variety of applications, such as environmental and habitat monitoring, industrial process control, security and surveillance, and structural health monitoring. The behaviour of these systems resembles that of disordered particle systems studied in statistical physics, and their large scale deployment allows us to appreciate in a real setting the phenomena predicted by the random models.

Various questions are of interest in this renewed context. The first and most basic one deals with connectivity, which expresses a global property of the system as a whole: can information be transferred through the network? In other words, does the network allow at least a large fraction of the nodes to be connected by paths of adjacent edges, or is it composed of a multitude of disconnected clusters? The second question naturally follows the first one: what is the network capacity in terms of sustainable information flow under different connectivity regimes? Finally, there are questions of more algorithmic flavour, asking about the form of the paths followed by the information flow and how these can be traversed in an efficient way. All of these issues are strongly related to each other and to the original 'classic' results on random networks, and we attempt here to give a unifying view.

We now want to spend a few words on the organisation of the book. It starts by introducing random network models on the infinite plane. This is useful to reveal *phase transitions* that can be best observed over an infinite domain. A phase transition occurs when a small variation of the local parameters of the model triggers a macroscopic change that is observed over large scales. Obviously, one also expects the behaviour that can be observed at the infinite scale to be a good indication of what happens when we consider finite models that grow larger and larger in size, and we shall see that this is indeed the case when considering scaling properties of finite networks. Hence, after discussing in Chapter 2 phase transitions in infinite networks, we spend some words in Chapter 3 on connectivity of finite networks, treating full connectivity and almost connectivity in various models. In order to deal with the information capacity questions in Chapter 5, we need more background on random networks on the infinite plane, and Chapter 4 provides all the necessary ingredients for this. Finally, Chapter 5 is devoted to studying the information capacity of a random network, applying the scaling limit approach of statistical physics in an information-theoretic setting, and Chapter 6 presents certain algorithmic aspects that arise in trying to find the best way to navigate through a random network.

The remainder of this chapter introduces different models of random networks and briefly discusses their applications. In the course of the book, results for more complex models often rely on similar ones that hold for simpler models, so the theory is built incrementally from the bottom up.

1.1 Discrete network models

1.1.1 The random tree

We start with the simplest structure. Let us consider a *tree* T composed of an infinite number of vertices, where each vertex has exactly $k > 0$ children, and draw each edge of the tree with probability $p > 0$, or delete it otherwise, independently of all other edges. We are then left with a random infinite subgraph of T, a finite realisation of which is depicted in Figure 1.1. If we fix a vertex $x_0 \in T$, we can ask how long is the line of descent rooted at x_0 in the resulting random network. Of course, we expect this to be on average longer as p approaches one. This question can also be phrased in more general

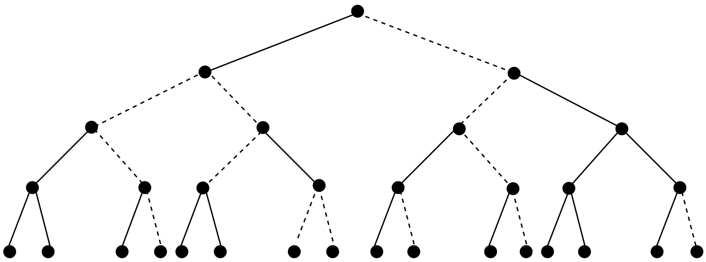

Fig. 1.1 A random tree $T(k, p)$, with $k = 2$, $p = 1/2$; deleted edges are represented by dashed lines.

terms. The distribution of the number of children at each node of the tree is called the *offspring distribution*, and in our example it has a Bernoulli distribution with parameters k and p. A natural way to obtain a random tree with arbitrary offspring distribution is by a so-called *branching process*. This has often been used to model the evolution of a population from generation to generation and it is described as follows.

Let Z_n be the number of members of the nth generation. Each member i of the nth generation gives birth to a random number of children, X_i, which are the members of the $(n+1)$th generation. Assuming $Z_0 = 1$, the evolution of the Z_i can be represented by a random tree structure rooted at Z_0 and where

$$Z_{n+1} = X_1 + X_2 + \cdots + X_{Z_n}, \tag{1.1}$$

see Figure 1.2. Note that the X_i are random variables and we make the following assumptions,

(i) the X_i are independent of each other,
(ii) the X_i all have the same offspring distribution.

The process described above could in principle evolve forever, generating an infinite tree. One expects that if the offspring distribution guarantees that individuals have a sufficiently large number of children, then the population will grow indefinitely, with positive probability at least. We shall see that there is a critical value for the expected

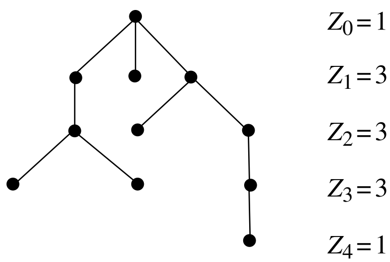

Fig. 1.2 A random tree obtained by a branching process.

offspring that makes this possible and make a precise statement of this in the next chapter. Finally, note that the branching process reduces to our original example if we take the offspring distribution to be Bernoulli of parameters k and p.

1.1.2 The random grid

Another basic structure is the random *grid*. This is typically used in physics to model flows in porous media (referred to as percolation processes). Consider an infinite square lattice $\mathbb{Z}^2$ and draw each edge between nearest neigbours with probability p, or delete it otherwise, independently of all other edges. We are then left with a random infinite subgraph of $\mathbb{Z}^2$, see Figure 1.3 for a realisation of this on a finite domain. It is reasonable to expect that larger values of p will lead to the existence of larger connected components in such subgraphs, in some well-defined sense. There could in principle even be one or more infinite connected subgraphs when p is large enough, and we note that this is trivially the case when $p = 1$.

What we have described is usually referred to as a *bond percolation model* on the square lattice. Another similar random grid model is obtained by considering a *site percolation model*. In this case each box of the square lattice is occupied with probability p and empty otherwise, independently of all other boxes. The resulting random structure, depicted in Figure 1.4, also induces a random subgraph of $\mathbb{Z}^2$. This is obtained by calling boxes that share a side neighbours, and considering connected neighbouring boxes that are occupied. It is also interesting to note that if we take a tree instead of a grid as the underlying structure, then bond and site percolation can be viewed as the same process, since each bond can be uniquely identified with a site and vice versa.

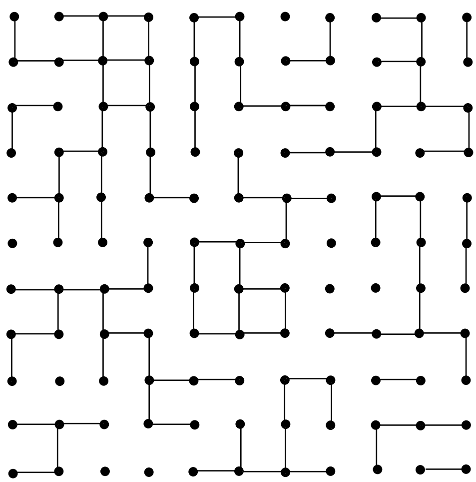

Fig. 1.3 The grid (bond percolation).

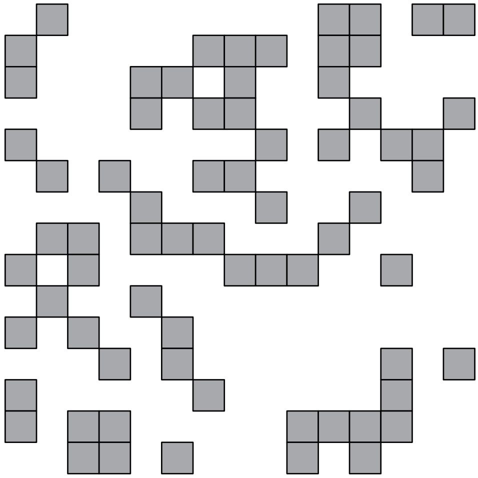

Fig. 1.4 The grid (site percolation).

1.2 Continuum network models

1.2.1 Poisson processes

Although stochastic, the models described above are developed from a predefined deterministic structure (tree and grid respectively). In continuum models this is no longer the case as the positions of the nodes of the network themselves are random and are formed by the realisation of a *point process* on the plane. This allows us to consider more complex random structures that often more closely resemble real systems.

For our purposes, we can think of a point process as a random set of points on the plane. Of course, one could think of a more formal mathematical definition, and we refer to the book by Daley and Vere-Jones (1988) for this. We make use of two kinds of point processes. The first one describes occurrences of unpredictable events, like the placement of a node in the random network at a given point in space, which exhibit a certain amount of statistical regularity. The second one accounts for more irregular network deployments, while maintaining some of the most natural properties.

We start by motivating our first definition listing the following desirable features of a somehow regular, random network deployment.

(i) *Stationarity*. We would like the distribution of the nodes in a given region of the plane to be invariant under any translation of the region to another location of the plane.

(ii) *Independence*. We would like the number of nodes deployed in disjoint regions of the plane to be independent.

(iii) *Absence of accumulation.* We would like only finitely many nodes in every bounded region of the plane and this number to be on average proportional to the area of that region.