Remco van der Hofstad
Department of Mathematics and Computer Science
Eindhoven University of Technology, P.O. Box 513
5600 MB Eindhoven, The Netherlands.
rhofstad@win.tue.nl

Inaugural Lecture

Random networking: between

With the arrival of the Internet, a good understanding of networks has become important for everyone. Network theory, which originated in the eighteenth century with Euler, and in the nineteenth century with Markov, has until recently concentrated its attention mainly on regular types of graphs. In his inaugural lecture, Remco van der Hofstad shows us a shift towards highly irregular graphs having vertices with extremely high degrees. He argues that this irregularity is a main characteristic of real life networks such as the Internet, social networks and networks describing biophysical phenomena. On January 1, 2005, Remco van der Hofstad was appointed full professor in Probability at the University of Eindhoven.

Who has not, at some point, exclaimed: isn't it a small-world? Others may have heard about 'six degrees of separation'. These phrases are all about the nature of social relations. The notion of six degrees of separation originated as the title of a John Guare play. In it, one of the main characters says:

"Everybody on this planet is separated only by six other people. Six degrees of separation. Between us and everybody else on this planet. The president of the United States. A gondolier in Venice...It's not just the big names. It's anyone. A native in the rain forest. [...] An Eskimo. I am bound to everyone on this planet by a trail of six people. It is a profound thought."

Similarly, it is claimed that each Dutch citizen is at most five handshakes away from the Dutch queen. To put these facts in a more mathematical framework, we reformulate social relations in terms of graphs. Its elements correspond to the people under investigation, and two individuals are connected when they share a certain social relation, such as knowing each other on a first name basis or having shaken hands. The phrase six degrees of separation then means that any pair of individuals is connected by a chain of at most six intermediaries. This was first empirically observed by the sociologist Stanley Milgram [13].

Real Networks

In the last decade, real networks have drawn tremendous attention in scientific disciplines ranging from biology, telecommunications, economics, linguistics, to sociology. Examples are the World Wide Web (WWW), the Internet, networks describing e-mails in a population, but also protein interaction networks, food webs, and collaboration networks. The elements of the WWW are web pages, and there is a (directed) connection between two web pages when the first links to the second. In a collaboration network, the elements are scientists, and two scientists are connected when they have a joint publication. In the Internet, the elements are computers or routers, and two computers are connected when there is a physical cable linking them. Thus, while the WWW is virtual, the Internet is physical. As you can see, these examples are very different in nature!

The study of networks has boomed in the period between 1995 and now, and this *must* have a reason. Of course, computer speed has dramatically increased in the past decade and network data sets have become publicly available on the Internet. This, however, cannot be the whole explanation of the rapid growth of interest. Large data sets are just huge collections of numbers, which are only of interest when one can uncover some of the structure behind them!

Many networks turn out to share two properties: they are small worlds and scale-free, as I will explain in more detail below. Both properties tell us something about the structure of the networks under consideration. The structure of networks is particularly important for processes living on these networks. For example, the spread of a virus is affected not only by the characteristics of its elements, but also by the structure of social networks, such as the traveling habits of individuals, as the

Remco van der Hofstad

order and chaos

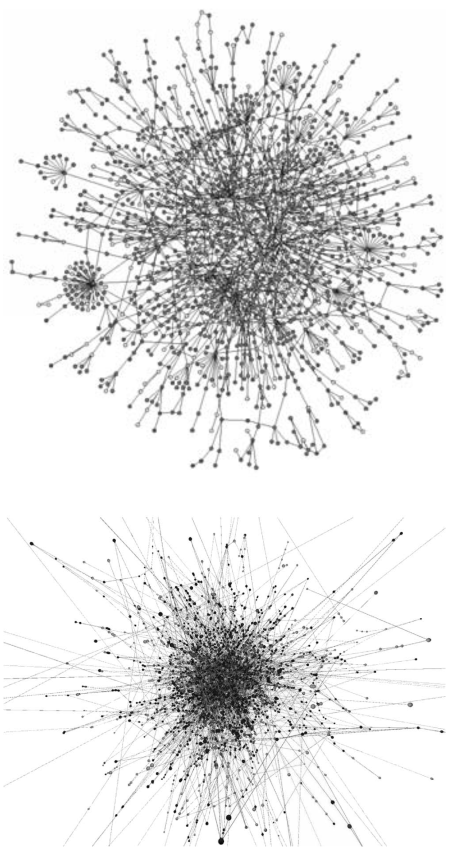

Figure 1 Yeast protein interaction network taken from [10], and the Internet topology in 2001 taken from www.fractalus.com/steve/stuff/ipmap, also reprinted on the front cover of this issue of the *Nieuw Archief voor Wiskunde*

recent SARS epidemic has shown. The spread of HIV depends sensitively upon properties of the network of sexual contacts. The 'success' of computer viruses and worms is influenced by how well they can be transmitted from one computer to the next, which, in turn, depends sensitively on the structure of the Internet and e-mail networks. We all know examples of successful computer viruses in the recent past. Their economical damage partly explains the interest in the study of the network structure, but there is more to it. Interestingly, networks turn out to be different from network models that have been studied over the past few decades, and the diverse networks above have many unexpected common features. For artist's impressions of some networks see Figure 1.

Power laws and degrees

While many real networks appear to be highly chaotic, there is some hidden order in them: they tend to be small worlds and are scale-free. A network is a *small world* when distances in it are small. The phrase 'scale-free networks' cannot be understood so easily, and I will spend some time to explain it.

For an element of a network, its *degree* is the number of other elements it is connected to. Most networks have large fluctuations in the degrees of their elements. While most elements have small degrees, there also exist elements with high, and even extremely high, degrees. For example, one cannot be surprised that in the sexual network of the eighteenth century, Casanova would appear as an element with an unprecedented degree!

The role of Casanova in the collaboration network in mathematics is played by Paul Erdős, who has the staggering amount of over five hundred co-authors. When studying graphs, one is bound to run into Erdős, and he will feature more prominently later on. Erdős was quite a special mathematician, and his deep love of mathematics has been recorded in popular books as well as a movie.

There are also other networks with such special elements. The protein p53 is believed to play a special role in protein interaction networks, and is even related to the occurrence of cancer. In social networks, there are these rare individuals who seem to know everybody. In the WWW, Google now attracts zillions of links!

The common fact about these unusual elements is that they are highly interconnected, that is, they have a high *degree*. These highly connected elements are sometimes called *hubs*. The hubs in networks are the Erdőses, the Googles, the Casanovas, or the other special elements that are present. While the hubs play a special role in networks, the fact that networks have them turns out not to be special at all. Most networks do!

One can imagine that these elements are quite important in applications. If one wishes to attack the Internet, the impact of attacking a router with a gigantic degree is much larger than that of attacking my laptop. To spread a rumor quickly, one should tell it to a person who happens to know everyone! It is well-documented that the large number of AIDS victims in the west is to a large extent due to the high promiscuity of one of the first Westerners that was infected (see e.g. [2]).

In summary, many networks have highly connected elements or hubs playing a central role in their functionality. This points us to the study of degrees in networks. In many networks, the degrees are scale-free, meaning that they have no typical size. Not all properties are such. For example, for Dutch men, the average height is about 1.83, and while differences in height exist, almost all Dutch men are in between 1.40 and 2.10. Thus, we can think of 1.83 as being the *typical height*. Sometimes such a typical size does not exist. For example, when studying the wealth in populations, already Pareto observed a huge variability [12]. Most individuals do not earn so much, but there are these rare individuals that earn a substantial part of the total income. Pareto's Law was best known under the name '80/20 rule', indicating, for example, that 20 percent of the people earn 80 percent of the total income. This law appears to be true much more generally. For example, 20 percent of the people own 80 percent of the land, 20 percent of the employees earn 80 percent of the profit of large companies, and maybe even 20 percent of the scientists write 80 percent of the papers. In each of these cases, no typical size exists due to the high variability present, which explains why these properties are called scale-free.

A network is called *scale-free* when its degrees have no typical size. While many elements only have few connections, the hubs have extremely many connections. The above is a more informal description, but there is a beautiful mathematical description as well: scale-free networks have power-law degree sequences. This needs further explanation!

The *degree sequence of a network* is the vector consisting of the number of elements in the graph with some specified degree, for all of the possible degrees. Thus, the degree sequence $\{f(k)\}_{k=1}^{\infty}$ is such that $f(k)$ is the number of elements with degree equal to k. This explains the meaning of one part of the phrase power-law degree sequences. We now turn to the other part, the power laws.

Mathematicians are always drawn to simple, yet fascinating, relations, as we believe they explain the rules that gave rise to them. A *power-law relation* between two variables means that one is proportional to a power of the other. In a formula, the variable k and the characteristic $f(k)$ are in a power-law relation when $f(k)$ is proportional to a power of k, that is, for some τ and C,

$$f(k) = Ck^{-\tau}. \qquad (1)$$

Intuitively, when an 80/20 rule holds, a power law is hidden in the background! Power-law relations are one-step extensions of linear relations. Conveniently, even when one does not understand the mathematical definition of a power law too well, one can still observe them in a simple way: in a log-log plot, power laws are turned into straight lines! Indeed, taking the log yields

$$\log f(k) = \log C - \tau \log k,$$

so that $\log f(k)$ is in a linear relationship with $\log k$, with slope equal to $-\tau$.

Power laws play a crucial role in mathematics as well as in many applications, and have a long history. Zipf [14] was one of the first to find one in the study of the frequencies of occurrence of words in large pieces of text. He observed that the number of occurrences of words is roughly inversely proportional to its rank in the frequency table of all words. This is called *Zipf's Law*.

Already in the twenties, several other examples of power laws were found. Lotka [11] investigated references in the Chemical Abstracts in the period 1901-1916, and found that with $f(k)$ denoting the number of scientists appearing in k entries, (1) holds, where τ now is close to 2. This is *Lotka's Law*.

In both examples, no typical size exists. For Zipf's Law, in most texts, there exist words (often 'the') that occur very often, while most of the words are only used a couple of times. Similarly, most scientists were only referred to once in Lotka's Law, but there exist scientists of whom many papers were cited.

Scale-free networks

A *scale-free network* is such that its degrees are scale-free. As a result, a large variability in degrees exists: there exist hubs with very high degrees.

More formally, with $f(k)$ equal to the number of elements with degree equal to k, we have that (1) holds, and the graph has a power-law degree sequence. For example,

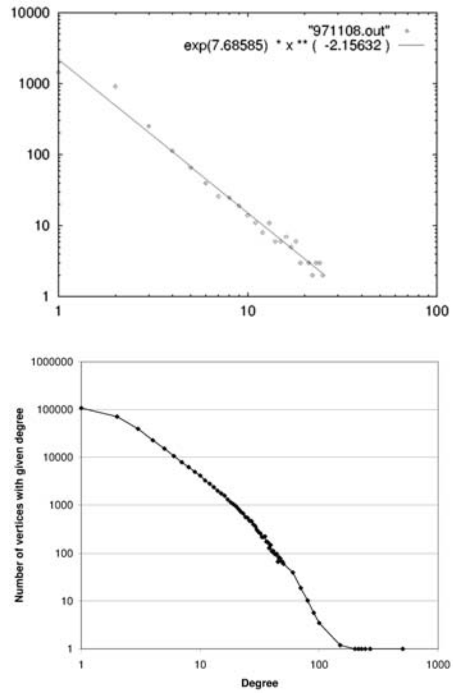

Figure 2 The degree sequence of the graph of Autonomous Systems in Internet in 1997 taken from [7] and of the collaboration graph among mathematicians taken from www.oakland.edu/enp

the degree sequences of the graph of autonomous systems in the Internet in 1997 and the collaboration graph of mathematicians are given in Figure 2, in log-log scale. While the log-log plots are not precisely equal to straight lines, they are quite close. For the collaboration graph, this means that the number of mathematicians with a given number of collaborators k is proportional to k to a certain power: we again find a power law in scientific productivity as Lotka did 80 years ago!

The first to realize that many networks are scale-free were Reká Albert and Albert-László Barabási, who dubbed such networks scale-free networks, igniting research in various disciplines, and, admittedly, causing a bit of a hype. For example, read the following quote about Barabási [5]: "What do the proteins in your body, the Internet, a cool collection of atoms and sexual networks have in common? One man thinks he has the answer, and it's going to transform the way we view the world . . ." This may be a bit too much, but the realization that similar network phenomena occur in social sciences, telecommunications and biology raises many modeling questions. The proposed models, in turn, lead to interesting mathematics. This is an appealing combination!

Applications do sometimes make a mathematician's life easier. Indeed, when asked for what I do professionally, the response "I am a mathematician" does not work so well. "Ah, I was always terrible at math in high school!" is the most frequent answer, and I keep on being surprised by the pride with which such a reply is uttered! Applications are a good way of explaining why we mathematicians do what we do. In truth, many mathematicians are guided by curiosity, by puzzles, yes, even by passion! Personally, I derive the most joy from a beautiful mathematical proof, an unexpected similarity between seemingly unrelated problems, or on finding a satisfactory explanation for a certain phenomenon. Such occasions are the icing on the cake!

Returning to the title of my speech, 'Random networking: between order and chaos', we have now seen many networks that are quite chaotic, yet share an intriguing order in being small worlds and scale-free. We now turn to randomness, which, as you may imagine, is inescapable for a probabilist such as me!

Modeling complexity by randomness

One cannot explicitly describe the structure of the network of all social relations on this planet with its size of about six billion people, or of the WWW, which has an estimated size of several billion web pages. Foremost, this is due to the fact that the data are not available, but also because the data are just too much to comprehend. This situation is well known in several disciplines, and most profoundly in particle systems. Most particle systems consist of around 10^{23} particles, and it is infeasible to describe the motions and locations of these particles explicitly. This fact has been realized a long time ago, and was resolved by introducing randomness, leading, in the case of interacting particles, to the theory of statistical mechanics. Boltzmann and Gibbs were the first to use this concept, which, at the time of Boltzmann, lead to serious opposition.

The main advantage of statistical mechanics is that one only needs to describe the local interactions between molecules, and these local rules enforce global behavior. Perhaps surprisingly, randomness does not seem to play any role when observing a glass of water. Indeed, in huge systems, random fluctuations are washed out, and only averages remain. Yet, a statistical mechanical model can predict global behavior, such as the existence of a phase transition separating the phases between water and ice. Thus, statistical mechanics is all about relating microscopic properties of molecules to the macroscopic properties of the material in question. In summary, we model complexity using randomness. Since the world around is getting ever more complex, probability is bound to play an increasingly important role.

A similar approach can be taken for complex networks. Their size makes their complete description utterly impossible. Therefore, one could suffice by describing how many elements the network has, and by which probabilistic rules elements are connected to one another. This leads us to consider *random graphs* as models for real networks, and introduces *randomness* in network theory.

Erdős-Rényi random graphs

The field of random graphs was established in the late fifties and early sixties of the last century. While there were a few papers appearing around (and even before) that time, one paper is generally considered to have founded the field [6]. The authors Erdős and Rényi studied the simplest imaginable random graph, which is now named after them. Their graph has n elements, and each pair of elements is independently connected with a fixed probability. When we think of this graph as describing a social network, then the elements denote the individuals, while two individuals are connected when they know one another. The probability for elements to be connected is sometimes called the *edge probability*. In their seminal paper [6], Erdős and Rényi clearly showed their predictive powers: "It seems to us further that it would be worth while to consider besides graphs also more complex structures from the same point of view, i.e. to investigate the laws governing their evolution in a similar spirit. [...] In fact, the evolution of graphs can be seen as a rather simplified model of the evolution of certain communication nets ..." In [6], Erdős and Rényi investigate basic properties of their random graph when their size is large, for various values of the edge probability. It turns out that the Erdős-Rényi random graph exhibits a phase transition, similarly to the water-ice transition. Indeed, when the average number of neighbors per element is less than one, the connected components are tiny and scattered, even for large graphs. When it is larger than one, on the other hand, a giant component emerges containing a number of elements proportional to the size of the graph.

The phase transition for the random graph is clearly demonstrated in Figure 3. It shows realizations of an Erdős-Rényi random graph of size 100, with an average number of neighbors of $1/2$ and $3/2$, respectively below and above the critical value 1. The components are drawn in blue, ordered in size from dark to light. While there is little difference in the relative sizes of the small components, the largest components are quite different indeed! In many systems, at the value of the phase transition peculiar behavior can be observed that is quite different from the behavior in either the sub- or the supercritical regimes. For example, in the water-ice transition, precisely at $0°$C, water and ice can co-exist. On the critical Erdős-Rényi random graph, the largest component has approximate size $n^{2/3}$, i.e., the number of elements to the power $2/3$.

While the Erdős-Rényi random graph is a beautiful model displaying fascinating scaling behavior for large graphs and varying edge probabilities, its degrees are not scale-free, rendering it unrealistic as a network model. Indeed, its typical degree size is the average degree, and there is little variability in its degrees. In particular, no hubs exist. Therefore, to model networks more appropriately, we are on the hunt for scale-free random graph models! Remarkably, the fact that the Erdős-Rényi random graph is not a suitable network model

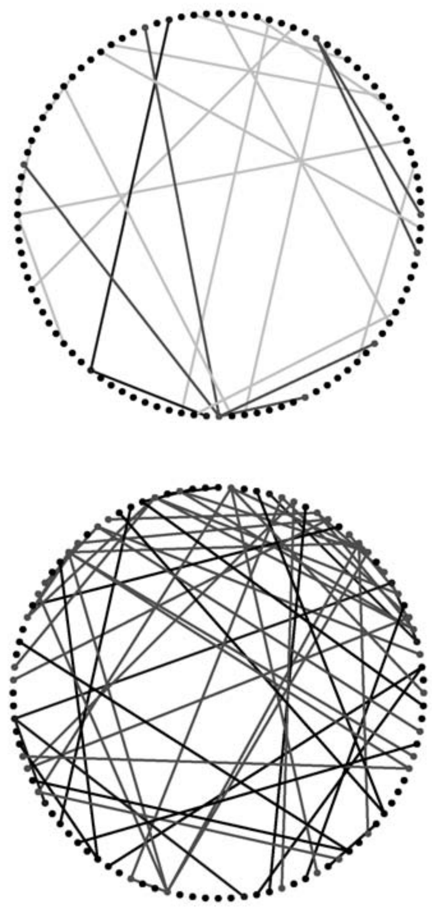

Figure 3 Two realizations of Erdős-Rényi random graphs with 100 elements and edge probabilities $1/200$, respectively $3/200$

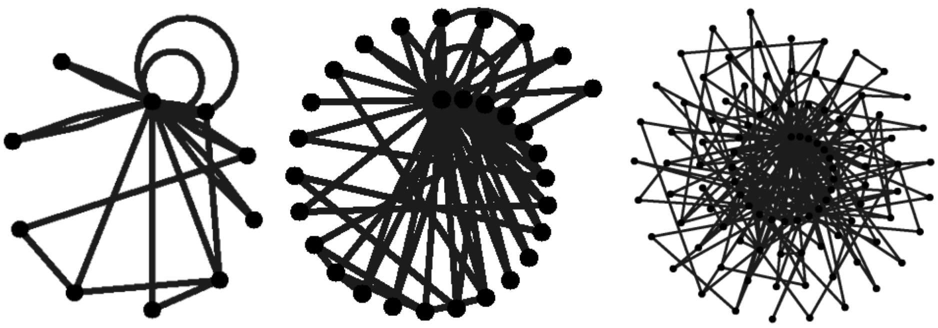

Figure 4 Preferential attachment random graphs of sizes 10, 30 and 100

was already foreseen by the masters themselves [16]: “*Of course, if one aims at describing such a real situation, one should replace the hypothesis of equiprobability of all connections by some more realistic hypothesis.*”

How do power laws arise then in networks, and what can we learn from that?

Preferential attachment

When it was realized that the Erdős-Rényi random graph is not an appropriate network model, several models were invented that adapt it in such a way that power-law degree sequences do arise. These models describe networks, in some sense, quite satisfactorily. However, they do not explain how the networks came to be as they are. A possible explanation for the occurrence of scale-free behavior was given by Albert-Barabási [3].

Most real networks grow. For example, the WWW has increased from a few web pages in 1990 to an estimated size of a few billion now. *Growth* is an aspect that is not taken into account in Erdős-Rényi random graphs, but it would not be hard to reformulate them as a growth process where elements are successively added, and connections are added and removed. Thus, growth by itself is not enough to explain the occurrence of power laws. However, viewing real networks as evolving in time does give us the possibility to investigate just how they grow.

So, how do real networks grow? Think of a social network describing a certain population in which a newcomer arrives, increasing it by one element. He or she will start to socialize with people in the population, and this process is responsible for the connections to the newly arrived person. In an Erdős-Rényi random graph, the connections to the newcomer will be spread uniformly over the population. Is this realistic? Is the newcomer not more likely to get to know people who are socially active, and, therefore, already have a larger degree? Probably so! We do not live in a perfectly egalitarian world. Rather, we live in a self-reinforcing world, where people who are successful are more likely to become even more successful! Therefore, rather than equal probabilities for our newcomer to acquaint him- or herself to other individuals in the population, there is a bias towards individuals who already know many people. When we think of the degree of elements as describing the wealth of the individuals in the population, we live in a world where the rich get richer!

This phenomenon is now mostly called *preferential attachment*, and was first described informally by Albert and Barabási [3]. While the above explanation is for social networks, also in other examples some form of preferential attachment is likely to be present. For example, in the WWW, when a new web page is created, it is more likely to link to an already popular site, such as Google, than to my personal web page. For the Internet, it may be profitable for new routers to be connected to highly connected routers, since these give rise to short distances.

Phrased in a more mathematical way, preferential attachment models are such that new elements are more likely to attach to elements with high degree compared to elements with small degree.

For example, we can successively add new elements with a fixed amount of connections to the older elements equal to m, each connection to an older element with a probability which is proportional to the degree of the older element plus a constant δ. Interestingly, preferential attachment gives rise to scale-free random graphs, where the power-law degree exponent equals $\tau = 3 + \frac{\delta}{m}$ (see e.g., [4] for the case where $\delta = 0$).

Preferential attachment offers a convincing explanation as to why power-law degree sequences occur. As Barabási puts it [2]

“...the scale-free topology is evidence of organizing principles acting at each stage of the network formation. [...] No matter how large and complex a network becomes, as long as preferential attachment and growth are present it will maintain its hub-dominated scale-free topology.”

On the other hand, preferential attachment cannot be the only mechanism describing network growth. For example, two people living in the same city are more likely to know each other than people living in a different continent. Thus, there is bound to be some dependence of real networks on geometry, which is not taken into account so far. A random network in which geometry plays a crucial role is percolation.

Percolation

Percolation can best be described as a random maze or labyrinth. Imagine a square in which an ant is trying to go from the left part where his nest is, to the right part where its food supply is. The labyrinth is made out of a collection of walls or obstacles, which the ant cannot pass through. Thus, its task is to walk past the walls from its nest to the food (and then back, naturally).

Depending on the locations of the walls, the nest and the food, the ant may be able to reach the food or not. Of course, there are instances where there are walls all around the nest or the food, and then the poor ants will die of starvation. This will particularly happen when there are many walls. On the other hand, it is possible that there is a passage way, and then the clever ant will always find it, see Figure 6.

We randomly position the walls, each possible location having a wall with some fixed probability p, and no walls with probability $1 - p$. Then, intuitively, when there are few walls, i.e., when p is small, the ant will be able to go to the food, while, when there are many walls, i.e., when p is large, the ant will likely not be able to. Now we make the task for the ant even more daunting by making the labyrinth very large. Then it turns out that there again is a phase transition. Indeed, when $p < 1/2$, with high probability, the ant will get to its food, while for $p > 1/2$, it will

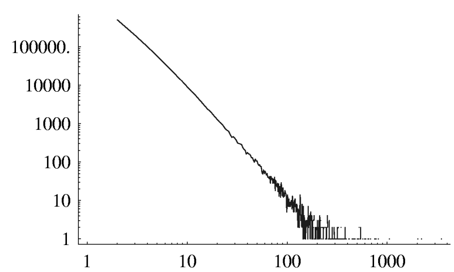

Figure 5 The degree sequences of a preferential attachment random graph of size 300,000 in log-log scale

not be able to. For $p > 1/2$, percolation occurs, in the sense that the collection of walls contains one dominant connected piece, which basically covers the entire square. For $p < 1/2$, on the other hand, the connected components of walls remain small, and do not obstruct the ant too much.

The most interesting phenomena occur close to this critical value. For example, for p slightly smaller than $1/2$ and a huge labyrinth, we can make the walls very small in such a way that the extent of the labyrinth remains fixed. In this case, the poor ant is having a harder and harder time to arrive at the food, even though we know that, with very high probability, it should be able to get there.

The above random labyrinth is an instance of percolation, a subject in probability theory and theoretical physics that has drawn enormous attention in the past decades due to the fact that the model is simple to state, while it is mathematically extremely rich. It is a model for porous media, where the materials consist of holes and substance. Think of a porous stone. When we pour water on it, will its center become wet?

While the above example of the ant in the labyrinth is intrinsically two-dimensional, the problem of wetting the core of a porous stone is clearly three-dimensional. There is no problem of posing the percolation problem in three dimensions. In this case, the walls are small square pieces of material, and ants cannot pass through these walls, while they can pass when a wall is absent (luckily, ants can climb walls). Again, we can think about an ant trying to work its way through this three-dimensional porous structure, such as a termite in a trunk. Many of the above properties, such as the existence of a critical percolation value for p, remain valid, even though, interestingly, it is not known exactly what this critical value is.

Realistic porous media are clearly restricted to two or three dimensions, but it is also possible to envision percolation in higher dimensions. Mathematically, considering higher dimensions does not cause any complications, and percolation can be, and has been, introduced and studied in general dimensions. It is believed that the qualitative behavior of percolation, particularly close to criticality, is quite different in different dimensions. For example, in high dimensions, the critical behavior of percolation on finite tori is similar to the one on the Erdős-Rényi random graph (see [9] and the references therein).

While percolation is a thoroughly investigated subject, many aspects of it are still ill understood, and it is likely to remain a source of inspiration for years to come. An excellent mathematical source is the book by Grimmett [8], who quite successfully conveys his fascination for percolation, for example in his quote: “Percolation is sexy!”

The road ahead

While percolation is a beautiful example of a random network that naturally incorporates geometry, it is not a good network model. Indeed, percolation violates the two key phenomena that networks display, the small-world and the scale-free properties. Indeed, distances between far away elements are large, and the degrees are bounded by 4 in dimension 2 and 6 in dimension 3. How to deal with geometry in truthful scale-free network models remains a mystery to this day!

There are more riddles to be resolved though …For example, networks tend to be clustered, which is just a difficult way of expressing that my friends tend to know one another. Clustering creates triangles in networks, and, even stronger, communities of individuals all connected to each other. How can we model these appropriately, and what is their relation to the scale-free nature of networks?

The preferential attachment models that are currently used predict that the elements with the largest degrees are the first ones present. While this may be true for some networks, in reality this is not always the case. Google is a late arrival on the WWW, but has become one of the most successful web site. How can we take the fitness of elements into account in the growth of networks?

The above research questions can be modeled in various ways. In fact, already for scale-free graphs there are numerous different models available, all of which are naturally caricatures of reality. How can we be sure that the conclusions drawn from these models have any predictive power for real networks? That is, if we change the model slightly, will the conclusions not change dramatically? It is here that the philosophy of statistical physics proves useful again. In this philosophy, universality plays a fundamental role. Universality states that independently of how we model a phenomenon locally, the global picture agrees, and is a well-established notion in theoretical physics, even though its mathematical foundations are often less clear. Models for real networks should be universal, being caricatures of reality in the first place! For example, for different network models having identical degrees, are the distances between elements similarly behaved? The above are just a few examples of challenges that are currently being taken up, or will be so in the near future. One of the truly beautiful aspects of networking is its *interdisciplinary nature.* Duncan Watts, with Barabási one of the pioneers of networking, says about this [15]: “*But in the world of networks, sociologists, economists, mathematicians, computer scientists, engineers and physicists all have something to offer each other, and much to learn.*” Different scientific disciplines often have a hard time talking to one another, so this clearly poses a challenge!

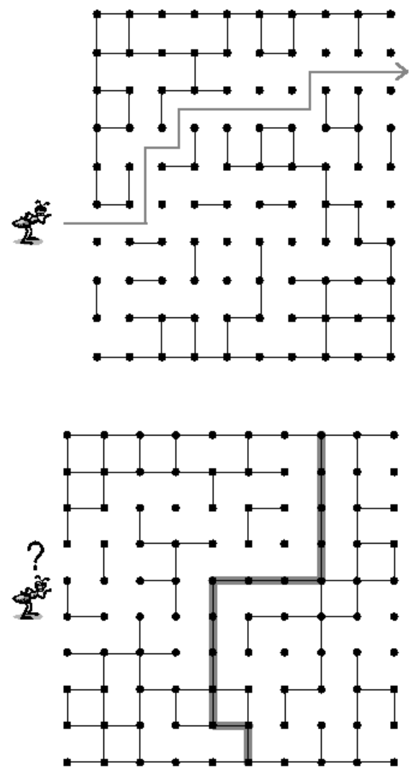

Figure 6 Two random labyrinths where an ant is trying to cross from the left to the right. It succeeds at the top, but not at the bottom.

Foremost, I hope and expect that the study of networks will lead to interesting mathematics. Mathematics is a beautiful discipline leading to extremely challenging riddles or puzzles, to be resolved using logic only. Solving these puzzles requires considerable creativity. While mathematics has established itself as a useful science, many mathematicians also have a very clear idea of beauty. Some proofs, which Erdős used to call ‘The proofs from the Book’, are astonishingly pretty, and the prettiest have been collected in [1]. In this sense, math is all about creativity and beauty, and is thus not so different from the arts. Many mathematicians are truly passionate about their work, and strive for the prettiest proof imaginable. Unfortunately, sometimes these are hard to find. In the past, the trend in science has to a large extent been towards applications, and recently even strongly towards valorization. While the economics of science are clearly important, I feel that we should not lose track of the more fundamental aspect of research, bearing in mind that what may seem to be fundamental now, may have important applications in the (near?) future.

From the recent success of Sudoku, I think that many among us share a passion for puzzles and logic! I fear that the image of mathematics of the general audience is rather different from the one described in this article, which at least partly explains the lack of interest of students in this beautiful field. I hope that, with a combined effort, the image of math can be improved and the interest in it revived!

Acknowledgement This work was supported in part by Netherlands Organisation for Scientific Research (NWO). I thank Marko Boon for creating the figures, and Mark Holmes for useful comments on an earlier version.

References

1 M. Aigner and G.M. Ziegler, *Proofs from The Book*, Springer-Verlag, Berlin, third edition, 2004.

2 A.-L. Barabási, *Linked: The New Science of Networks*, Perseus Publishing, Cambridge, Massachusetts, 2002.

3 A.-L. Barabási and R. Albert, 'Emergence of scaling in random networks', *Science* **286**(5439) (1999), pp. 509–512.

4 B. Bollobás, O. Riordan, J. Spencer, and G. Tusnády, 'The degree sequence of a scale-free random graph process', *Random Structures Algorithms* **18**(3) (2001), pp. 279–290.

5 D. Cohen, 'All the world's a net', *New Scientist* **2338** (2002).

6 P. Erdős and A. Rényi, 'On the evolution of random graphs', *Magyar Tud. Akad. Mat. Kutató Int. Közl.*, **5** (1960), pp. 17–61.

7 C. Faloutsos, P. Faloutsos, and M. Faloutsos, 'On power-law relationships of the internet topology', *Computer Communications Rev.* **29** (1999), pp. 251–262.

8 G. Grimmett, *Percolation*, Springer, Berlin, 2nd edition, 1999.

9 M. Heydenreich and R. van der Hofstad, 'Random graph asymptotics on high-dimensional tori', To appear in *Commun. Math. Phys.*

10 H. Jeong, S.P. Mason, A.-L. Barabási, and Z.N. Oltvai, 'Lethality and centrality in protein networks', *Nature* **411**(6833) (2001), pp. 41–42.

11 A.J. Lotka, 'The frequency distribution of scientific productivity', *Journal of the Washington Academy of Sciences* **16**(12) (1926), pp 317–323.

12 V. Pareto, *Cours d'Economie Politique*, Droz, Geneva, Switserland, 1896.

13 J. Travers and S. Milgram, 'An experimental study of the small world problem', *Sociometry* **32** (1969), pp 425–443.

14 G.K. Zipf, 'Relative frequency as a determinant of phonetic change', *Harvard Studies in Classical Philology* **15** (1929), pp 1–95.

15 D. Watts, *Six degrees. The science of a connected age*, W.W. Norton & Co. Inc., New York, 2003.

16 D. Watts, *Six degrees. The science of a connected age*, W.W. Norton & Co. Inc., New York, 2003.

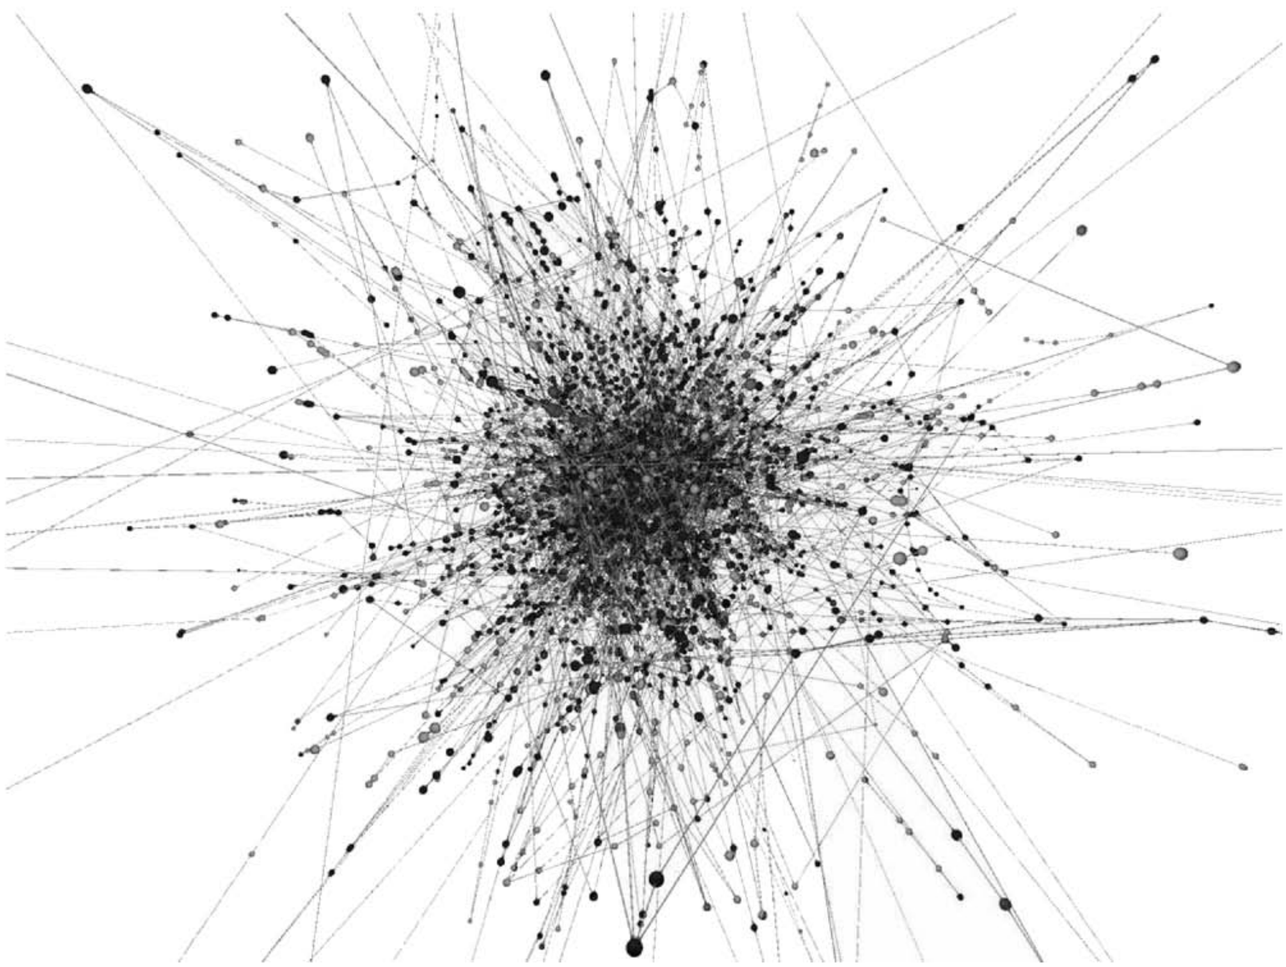

www.fractalus.com/steve/stuff/jpmap