THÈSE

dirigée par Clémence MAGNIEN

présentée pour obtenir le grade de

DOCTEUR EN SCIENCES
DE L'UNIVERSITÉ PIERRE ET MARIE CURIE

spécialité Informatique

École doctorale Informatique, télécommunications et électronique (Paris)

Analyse empirique et modélisation de la dynamique de la topologie de l'Internet

---

Sergey KIRGIZOV

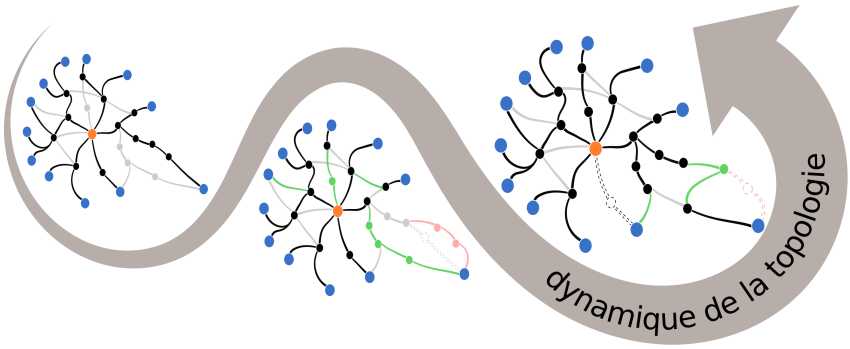

soutenue publiquement le 12 décembre 2014 devant le jury composé de

Rapporteurs :	Paulo GONÇALVES	Chargé de recherche, ENS Lyon, INRIA
	André-Luc BEYLOT	Professeur, IRIT/ENSEEIHT
Examinateurs :	Jeremie LEGUAY	Docteur, Thales Communications & Security
	Stefano SECCI	Maître de Conférences, UPMC
	Benoit DONNET	Professeur, Université de Liège
Directrice :	Clémence MAGNIEN	Directrice de recherche, UPMC, CNRS

for Liza

Τὰ πάντα ῥεῖ καὶ οὐδὲν μένει.[1]
— attributed to Heraclitus

[1] Everything flows, nothing stands still.

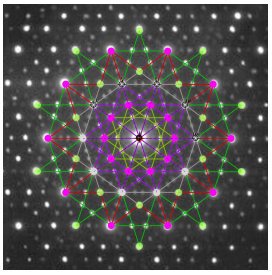

Preface

Networks... networks are everywere.

Food webs, lymphatic systems, brains, electrical circuits, crystals and quasicrystals, **the Internet**—many systems contain a lot of elements connected (or related) to each other in some sence. In mathematics we call these things "graphs", in other sciences we may use another words. But regardless of chosen language, the essence remains the same. Many real-world networks are, in fact, not so different. They share some common structural properties, and when we study a particular network, we actually study all of them in some way.

During my primary and secondary school, I was a great fan of the emerging Internet. I remember well the time when we used 56K modems in order to get something from the net. It was cool. Just after ten years, we begin to use another technology that is about a hundred times faster. This progress is really awesome. But not only the hardware part has been affected. Software has also been changed. Moreover, the Internet has been evolved structurally, it has been grown up by thousands of companies and millions of users. Now, when its prodigious structure is well-formed, we start to consider the Internet as a mature thing that deserves true scientific attention. Furthermore, the Internet is not stopped in its development, it continues to change, showing its own dynamics.

That's why I'm very happy to study the dynamics of the Internet here, in this thesis.

*Electron diffraction pattern of an icosahedral Zn-Mg-Ho quasicrystal, a kind of network (http://en.wikipedia.org/wiki/File:Ho-Mg-Zn_E8-5Cube.jpg license: CC BY-SA 3.0, author: jgmoxness).

Acknowledgements ❀

I thank my adviser Clémence Magnien for her continuous support. She showed me the right way through a labyrinth of scientific research.

I am grateful to Fabien Tarissan for his knowledge and magic of thought expressions. He is one of the co-authors of the key papers that formed the core of my thesis.

I would like to thank all my colleagues from the "Complex Networks" team, especially Maximilien Danisch, Lionel Tabourier, Jean-Loup Guillaume, Daniel Bernardes, François Queyroi, Amélie Medem, Raphaël Fournier, Jordan Viard, Matthieu Latapy, Noé Gaumont, Damien Nogues, Alice Albano, Qinna Wang, Elie Rotenberg, Bénédicte Le Grand and Sébastien Heymann. I thank them for numerous philosophical, scientific and ordinary conversations that allowed me to better understand life, science, and French language.

I also thank my colleagues and friends from other research teams: Julien Bensmail and Romaric Duvignau from LaBRI, Natalya Rozhnova from NPA (LIP6) and Olga Melekhova from MoVe (LIP6).

I sincerely thank Paulo GONÇALVES and André-Luc BEYLOT for reviewing my dissertation and for Jeremie LEGUAY, Stefano SECCI and Benoit DONNET for accepting the invitation to integrate the jury of my thesis defence.

I am very grateful to my previous advisers Jacques Malenfant from MoVe and Sergey Zhilin from Altai State University. Without their help it will be hardly possible for me to make my first steps in science.

I would also like to thank the administrative assistants of UPMC, especially Véronique Varenne, for their help and advices.

I gladly acknowledge the financial help I received from the *Agence Nationale de la Recherche* (grant ANR-10-JCJC-0202), from FP7 FIRE project EULER (Grant No.258307) and from EIT-ICT Labs action line Future Networking Solutions (FNS).

I express my gratitude to my wife Yulia, my parents, my best friends Daniil Churikov and Vladimir Tjut, for their constant inspiration and support.

I also thank my friends from Moscow (Ludmila Garkavaya, Olga Inyakina, Yesenia Shamonina, Alexander Malaev and Sergey Levashev) where I have stayed during my visa preparations.

Finally, I thank all Russian poets, especially Velimir Khlebnikov.

❀sorry if I forgot someone

Contents

Introduction **11**

1 Brief history and state of the art **13**
1.1 Graphs and complex networks . . . 13
1.2 Network metrology . . . 17
1.3 Dynamic networks and metrology . . . 18
1.4 Models of networks . . . 22

2 Model **25**
2.1 Motivation: `tracetree` measurements . . . 26
2.2 Causes of the observed dynamics . . . 29
2.3 Model description . . . 30
2.4 Preliminary results of simulation . . . 32
2.4.1 Evolution of the number of distinct nodes. . . . 32
2.4.2 Observation number vs. block number. . . . 34

List of definitions **37**

3 Characterising the dynamics **39**
3.1 Linearity of the evolution . . . 40
3.1.1 Correlation coefficient . . . 41
3.1.2 Coefficients of segmented linear regression . . . 42
3.2 Characterisation of the slope . . . 42
3.2.1 Random graphs . . . 45
3.2.2 Power-law graphs . . . 47
3.3 In search of unified laws . . . 49
3.3.1 Impact of the size of the shortest path subgraph . . . 52
3.3.2 Probability of shortest path subgraph modifications . . . 55
3.4 Real-world measurements . . . 55
3.4.1 Frequency of measurements . . . 55
3.4.2 Size of the shortest path tree vs slope . . . 57

4 Size of shortest path subgraphs **59**
4.1 Definitions . . . 62
4.2 Complete and quasi-complete graphs . . . 63
4.3 Dense random graphs (p is fixed, $n \to \infty$) . . . 64
4.4 Sparse random graphs with unbounded mean degree ($p \to 0$ and $np \to \infty$ as $n \to \infty$) . . . 68
4.4.1 Approximated expectation of the size . . . 69
4.4.2 Classification of sparse graphs according to the distance distribution . . . 70

5 Real and observed dynamics **75**
5.1 Impact of the measurement frequency . . . 76
5.2 Inferring the evolution speed . . . 78
5.2.1 Poisson process . . . 80
5.2.2 *sps*-process . . . 83
5.2.3 *spt*-process . . . 86
5.3 Nonuniform dynamics . . . 88

Conclusion **91**

Bibliography **102**

Introduction

The Internet is a complex structure that connects approximately three billions people[4]. No official map being available, researchers have to conduct costly measurement campaigns, and deal with the fact that the obtained data can be biased [40, 65]. Studying the dynamics of this topology is therefore an equally hard, if not harder, problem.

Many studies have already studied its topology and focused on the nodes' degree distribution, with the question whether or not it follows a power law (see for example [24, 27, 36]). However few of these studies have asked how this topology evolves in time. Indeed, it is well known that the Internet is also a living organism with many nodes and links added and removed every day. Characterising the Internet dynamics can provide interesting insights that may help in the design of future routing protocols or the conception of new types of applications that make use of its evolving topology.

Instead of trying to obtain a complete view of the Internet topology dynamics, it is possible to use an orthogonal approach to obtain insight on the dynamics of the routing topology observed at the IP-level. This approach has been introduced in [72]. We follow it in this manuscript and study *ego-centred views* of this topology. Given a monitor and a fixed set of destinations, one such view is obtained by measuring the routes from a monitor to a set of destinations. This can be performed quickly and with low network load with the `tracetree` tool [72]. Repeating this measurement periodically therefore allows us to study the dynamics of this view. Figure 1 shows five consecutive ego-centred views.

Previous work has shown that ego-centred views exhibit strong dynamics, and in particular that the set of observed nodes evolves much more quickly than what was previously expected [64]. Our goal in this thesis is to go further in this direction and effectively understand the dynamics of the Internet at the IP-level. To this end we use a combination of analysis and simulation. We use real data to extract key properties that characterise these dynamics. Based on those observations, we then propose a model for the underlying mechanisms of the Internet topology dynamics and we study both this model and the real data. Our main goal is to understand the impact of the model's parameters (and the impact of parameters of real measurements) on the observed dynamics. We will look for the most important parameters and general laws that govern these dynamics.

This manuscript is organised as follows. In Chapter 1 we present a brief historical context and the current state of network science. Next, in Chapter 2 we

[4] 2,802,478,934 people are connected according to "Internet World Stats" http://www.internetworldstats.com/stats.htm, December 31, 2013.

identify underlying mechanisms of the Internet topology dynamics and propose a model for the dynamic topology and ego-centred measurements. In Chapter 3 we study in details how the observed characteristics of dynamics change when we vary the parameters of our model, and compare it to what happens when we vary the parameters of the real-world measurement process (a part of the this Chapter was published as [80]). We find in particular that the number of routes between two computers, which corresponds to shortest paths in the model, plays an important role in the observed network dynamics. We therefore investigate the size of the shortest path subgraph between two nodes in Chapter 4 (a part of this Chapter was submitted to a journal [82]). Finally, in Chapter 5 we study how the frequency of measurements affects the observed network dynamics and how can we infer the real dynamics using only partial measurements (the first part of this Chapter was presented at a conference [79], and the second part exists in a preprint form [83]). In the Conclusion we discuss our main achievements and possible future works.

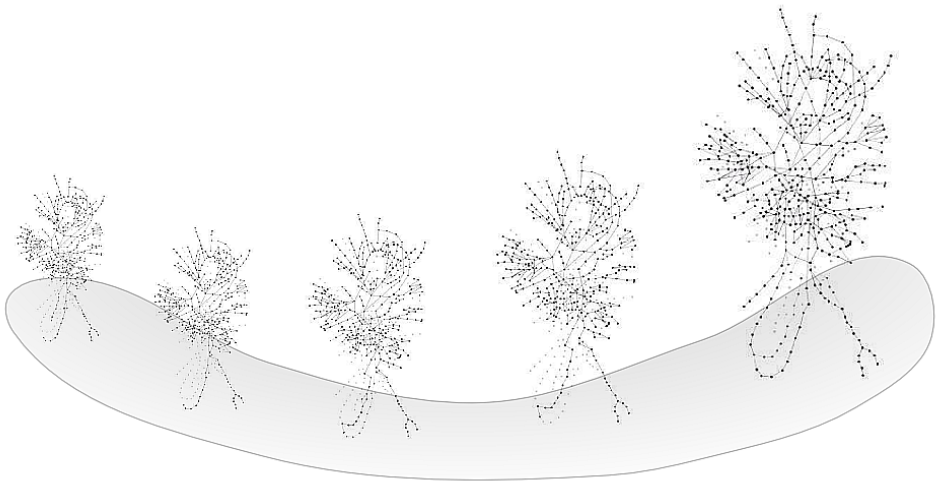

Figure 1 – *Ego-journey.* Visual representation of the Internet topology dynamics. Using ego-centred measurements [72] we obtain a sequence of topology snapshots. Time goes from left to right. Every little figure corresponds to one measurement round, i.e. one snapshot. All links observed at least once are shown: in black we display nodes and links seen at the current round, and in grey we display nodes and links seen in other rounds. Changes are most visible in the bottom part of the Figure. The original video, called "Dynamics of the internet topology around me", was created by Assia Hamzaoui and Matthieu Latapy.[5]

[5]See original video here: http://www.complexnetworks.fr/dynamics-of-the-internet-topology-around-me/

Chapter 1

Brief history and state of the art

Most real-world networks are continuously evolving structures that cannot be fully observed, in particular due to their tremendous sizes and dynamics. Scientists from various fields have studied real-world networks sometimes together, sometimes independently, for a long time. However, only in the last decades has this study become systematic. In this chapter we present a brief historical context and the current state of network science; both theoretical and empirical results are considered.

This chapter is organised as follows. In Section 1.1 we recall and discuss the notions and mathematical formulations of networks and related structures. Practical and theoretical aspects of network measurements are considered in Section 1.2. In Section 1.3 we review the definitions, examples and measurement methods in the case of dynamic networks. Finally, in Section 1.4 we describe several important models of static and dynamic networks.

1.1 Graphs and complex networks

The notion of graph goes back to the paper "Seven Bridges of Königsberg" written by Euler in 1736. However, the term "graph" was introduced by Sylvester in his paper "Chemistry and Algebra" [3] only 152 years after Euler's work. The first textbook about graphs *"Theorie der endlichen und unendlichen Graphen"* has been written by Dénes Kőnig and published in 1936 [91]. Two other notable books—*"Théorie des Graphes et ses applications"* (1958) by Claude Berge [85], and "Graph

Theory" (1969) by Frank Harary [89]—finally enabled scientists and engineers from different fields to communicate using the same language.

At the present time, researchers use words "graph" and "network" when they talk about an idealised representations of a real-world network. The formal definitions may vary, but the intuitive idea is the following:

> *A graph (in a broad sense) is a collection of nodes connected by links.*

Some nodes (or even all) can be isolated. Sometimes, nodes are called "vertexes" or even "points", and links are called "edges".

Using different mathematical definitions of "links" and "collection of nodes", researchers obtain different types of graphs: weighted graphs, directed graphs, multigraphs, coloured graphs, signed graphs, hypergraphs, infinite graphs, graphons, graphins, etc. Among these types, we are mainly interested in the simplest one. Following a convention from modern textbooks in graph theory, by a graph (in a strict sense) we mean undirected finite simple graph, that is:

> *An undirected finite simple graph G is a graph with finite sets of nodes V and links E. Every link connects two different nodes.*

Again some nodes (or even all) can be isolated.

We present here essential definitions that allow us to speak about a graph's topology:

- When a link connects two nodes v and u, we say that u and v are *adjacent*. One can go from v to u, by traversing this link. A *path* from a node v to another node w is a sequence of several nodes and links, passing along which we can go from v to w.

- We say that a graph is *connected* when there are paths between any two nodes.

- Two graphs are called *isomorphic* when there is a bijection f between their sets of nodes, that preserves adjacency. In other words there is a link between nodes u and v if and only if there is a link between their images $f(u)$ and $f(v)$.

- The *degree* of a node v is its number of direct neighbours, or equivalently the number of links that connects the node v to other nodes.

Other definitions will be given when its needed.

Researchers tend to consider larger and larger graphs. Four papers written by Rapoport [5], Solomonoff and Rapoport [6], Gilbert [8], and Erdős and Rényi [7], published between 1940 and 1960, initiated the study of some properties of random graphs when their sizes tends to infinity. In particular they studied whether random graphs are connected or not. The next step consisted in graphs with countable number of nodes. Surprisingly, it turns out that almost all random graphs with countable number of nodes are isomorphic. In other words, there is only one countable random graph (see Cameron's work [30]). This graph is called Rado graph [10]. However, one can go beyond countable set of nodes. For example Lovász in his recent book [92] considers two kinds of uncountable graphs: graphons and graphins. These structures can be regarded as the limits of sequences of random graphs when their size tends to infinity. Lovász showed that these notions are much more appropriate than the Rado graph. Developing of practical applications of these techniques is the very important direction of the future works. We hope it will be done in the coming decades.

Mathematicians are not the only scientists studying very large graphs. Many researchers from various areas have studied networks for a long time. In fact, there is an emerging interdisciplinary academic field devoted to studying very large and complex real-world networks: logistical networks, gene regulatory networks, computer networks, chemical reaction networks, social and linguistic networks, etc. Researchers use empirical, theoretical and computer-aided methods in order to understand the structures of these networks and their functions. The interested reader is referred to Newman's review [41].

Imagine that we have two graphs G_1 and G_2 perhaps of different size. Currently we do not know how to directly compare their structures. However, we can compare some characteristics of these structures. Below we present two most commonly used characteristics.

- The first characteristic is the *density*. The density of a graph is the ratio between the number of existing links and the maximum number of links that a graph can contain. In the case of simple graphs the density is equal to $\frac{2m}{n(n-1)}$, where m is the number of links, and n is the number of nodes in the graph.

- The second, and perhaps the most popular characteristic of a graph's topology, is the *degree distribution*. The degree distribution shows how many there nodes have degree one, two, etc. Random graphs have a binomial degree distribution. But, according to Barabási and Albert [23], many networks in

the real world have degree distributions that are very different from binomial one and approximately follow a power law.

There exist other, finer characteristics of the topology: clustering coefficient [20], assortativity coefficient [31], distance distributions, graph genus, graph spectra [67], subgraph densities and neighbourhood distributions [92].

Now let us consider the object of our study, the Internet. The Internet is a notable example of a complex network. It should be noted that the Internet, as any other complex network, attracts enormous attention from scientists, and we cannot review all papers and books about it in this thesis. We can construct different graphs that will reflect various levels of the Internet structure. Below we briefly describe important levels and cite some works as examples.

- Researchers study graphs of **physical** and radio connections between computers. For example, Akyildiz *et al.* published a survey on wireless mesh networks [50], and Narula-Tam considered the design of physical topology that will guarantee some level of protection against link failures in certain physical networks [44].

- At the **IP-level** the topology of the Internet is represented by a graph with IP-addresses of routers as nodes. A link exists between two nodes if the corresponding routers are connected at the IP level. Works that consider the IP-level topology of the Internet include (but are not limited to) a paper about connectivity of the IP-level graph written by Broido and kc claffy [29], and a paper written by F. de Montgolfier *et al.* about the treewidth and hyperbolicity of the Internet [74].

- The Internet topology can also be represented by a graph with Autonomous Systems (ASes)[1] as nodes and AS connections as links. Many studies have already characterised the Internet topology at the **AS-level** (see for example papers of Faloutsos *et al.* [24], Medina *et al.* [27], and Chen *et al.* [36]). As suggested by Wang *et al.* [75] the structure of IP-level graph differs from the structure of AS-level graph.

- In the **webgraph** vertexes represent WWW pages, and an edge connects two pages if there exists a hyperlink between these pages. Many researchers study the webgraph, see for example the works of Bontato [49] and Kleinberg *et al.* [25].

[1] An AS is a connected group of routers run by one or more network operators which has a single and clearly defined routing policy [17].

In this thesis we focus on the IP-level graph and its dynamics. The first question that arises is *how can we measure the actual topology of the Internet at the IP-level?*. In fact, the measurement problem is one of the most important problems in the study of real-world networks. The next section reviews the current situation in this topic.

1.2 Network metrology

The study of large graphs and networks is very important from practical and theoretical points of views, because many real and theoretical objects show a certain structure that can be well described using the language of graphs. But often such graphs are so large that we cannot study them using classical graph theoretical methods and algorithms. For example, if some graph has billions of nodes, algorithms with quadratic complexity become practically intractable, and even linear algorithms may take a very long time. Moreover, some real-world networks cannot be fully observed due to their tremendous sizes, while others cannot be fully observed in principle, because their structures continuously changes over time.

Perhaps the most rapidly developing area of mathematics and theoretical computer science that deals with these problems, is *property testing* [2]. Generally speaking, a property testing algorithm is a probabilistic algorithm with sublinear time complexity than tries to decide whether a given graph has some global property using only partial information about this graph (see for example Alon's paper [37]). There is even some connections between property testing and continuous graphs (see Lovász and Szegedy [69]).

In the case of real-world networks, it is not immediately clear which properties of the network structure we want to test. We just want to collect as much information about the structure as we can, using real-world measurement tools. And only then can we start to study the network's properties. Maybe this is why the majority of network scientists have not yet adopted the property testing approach. We need maybe a few decades to fully understand and develop this idea in order to apply it in practice. Guillaume and Latapy [46] argue for the necessity of developing a science of metrology of complex networks. Here we insist that this science of metrology should be developed using the force of contemporary mathematics.

Real-world networks are very different from one another, and we should use different measurement tools in order to obtain some information about their

[2] The interested reader may consider the following online reviews about current achievements of property testing: http://ptreview.sublinear.info/ and http://www.wisdom.weizmann.ac.il/~oded/test.html.

structure. Methods used for measuring brain networks [66] are clearly different from methods used to collect data from social networks [95]. Often, the available measurement methods are intrinsically biased.

Lakhina *et al.* [40] showed that IP-level topologies collected via traceroute-like measurements [93, 52] may be highly biased. In addition to the seminal paper of Lakhina *et al.* there are other papers in this area, for instance [72, 65, 75]. The AS-level topology is better understood mainly because there are better measurement tools [45]. The interested reader is also referred to a survey about the Internet topology discovery [57] written by Donnet and Friedman, and to a CAIDA technical report [76], wherein Huffaker *et al.* discussed the results of their "systematic comparison of Internet topologies derived from different data sources"[3].

Once the data is collected and transformed into a graph G, we can study some properties of G, or we can try to estimate the most likely underling real-world network topology that leads to G. If measurements are biased, there are several possible real-world topologies that can lead to G. How can we infer the most likely one? Some researchers, for example Coates *et al.* [34], start to use statistical methods in order to answer this question.

Finally, other works of interest concerning the metrology of the Internet are worth mentioning. Crespelle and Tarissan [70] evaluate a new method for measuring the Internet degree distribution. Latapy and Magnien [60] discussed the relevance of observed properties. Augustin *et al.* [61] discussed detection and prevention of anomalies that arises when we measure the Internet graph.

The structure of many real-world networks changes over time. This complicates both measurements and analysis. In the next section we discuss the notion of dynamic network and the problems that arise in this context.

1.3 Dynamic networks and metrology

Most real-world phenomena, including networks, are not static, they change over time, for example:

- The information flows in our brain. Electrical properties of neurons and neural connections may change (see for example Mozzachiodi and Byrne [68]). Moreover, new neurons may grow even in an adult human's brain, as indicated by Eriksson *et al.* [22]. Graph theoretical analysis of structural and functional brain networks are discussed in the review by Bullmore and Sporns [63].

- Computers are being connected to the Internet and start to communicate. The changes of network topology reflect, for example, the technical progress

[3]This is a direct quotation from CAIDA's page http://www.caida.org/publications/papers/2012/topocompare-tr/.

and the economic interests of users. Many researchers have observed the Internet dynamics at different levels, see for example papers by Magnien *et al.* [64] and Oliveira *et al.* [59].

Here we give a classification of dynamic networks from different points of view and briefly discuss theoretical and practical approaches used by researchers to study these networks.

Conceptually there are two main kinds of dynamic networks:

- when the network **topology changes**, for example new nodes (or links) appear (or disappear);
- when the topology remains the same, but there is some dynamic process that happens **on the network** (for example random walk or information transmission).

Sometimes, we have networks with **mixed dynamics**, i.e. there is a dynamic process on the network and the topology also changes. Eventually one may develop a formalism under which these two kinds of dynamic networks can be studied using the same language. But traditionally, researchers tend to make a clear difference between them. In this thesis by dynamics we mainly mean topological dynamics.

From the epidemiological point of view there are two kinds of dynamics:

- The dynamics **defined** by a model. This kind of dynamics is usually studied in the context of *discrete* dynamical systems wherein an object changes at each time step due to some fixed rule, and *continuous* dynamical systems in which time is continuous. Prisner in his monograph [94] considered a special case of discrete dynamical systems where the objects are graphs and changes are modelled by certain graph operators.
- The real-world dynamics **observed** by scientists. This kind of dynamics is very important in empirical research. The main problem here is that the real dynamics should be observed in some way. Thus, we have to answer at least to the following questions:
  - Which tools we can use to measure a dynamical real-world system?
  - Which tools will provide the best results?
  - How do we represent the results of our measurements?

The choice of the measuring tool intrinsically depends on the nature of the real system. We have many systems, and there are a lot of different measuring tools.

Performing measurements of dynamic networks, we are faced with the same problems as in the case of measurements of static networks: the large size of

networks and the measurement bias. Moreover, the dynamics itself complicates the measurement process. The measurements of dynamic networks can be even more biased and more incomplete than the measurements of static network.

There are several works that considered these problems in P2P networks. The create-based method, considered by Roselli *et al.* [26] and Saroiu *et al.* [39], is based on the observation that being able to only capture accurately the length of sessions that begin and end within the measurement window creates a bias towards short sessions. To remove this bias, the measurement window of length T is divided into two halves, and only the sessions that begin during the first half and last less than $T/2$ are considered. This leads to an unbiased estimation of sessions with length less than $T/2$.

The problem of measurement bias in the dynamic P2P network was further examined by Benamara and Magnien [73]. They studied several properties in large P2P systems and introduced an empirical methodology for deciding when the bias induced by the finiteness of the observation window in dynamic systems becomes negligible.

Oliveira *et al.* [59] formulated the *Liveness Problem*, that arises when an observed change in the AS-level topology of the Internet does not necessarily reflect the real change. Using three processes: birth, death, and revelation, Oliveira *et al.* developed an empirical model of the observed changes at the AS-level topology. Based on the model, the authors are able to distinguish real topology changes from transient routing changes at the AS-level topology with a given confidence level.

Representations of observed dynamic can also vary from one system to another. In the context of real-world networks with dynamic topology, there are several representations of the observed dynamics, for example:

- **Graph snapshots.** We perform periodic measurements of a network. We represent the observed dynamics as a sequence $\{g_t\}$ of measurement results, where by g_t is the graph obtained at time t. This representation is particularly relevant when one measurement gives a whole network topology.

- **Link streams** that were recently introduced by Latapy and Viard in [81]. The verbatim definition is as follows: link streams are series of triplets (t, a, b) meaning that a and b interacted at time t. This approach a typically relevant when we observe instantaneous interactions between nodes. For example in the case of email conversations a link (t, a, b) corresponds to a message from a to b sent at time t.

- **Time-varying graphs** (TVGs) are graphs in which nodes, or edges may vary in time. Wehmuth *et al.* [78] proposed an unifying model for representing time-varying graphs by $H = (V, E, T)$, where V is the set of nodes, T is the

finite set of time instants, and $E \subset V \times T \times V \times T$ is the set of dynamics edges. A dynamic edge e is defined as a tuple $e = (u, t_a, v, t_b)$, where u and v are two nodes from V (not necessary different), while t_a and t_b are the origin and destination time instants from T (again not necessarily different).

Fundamentally, graph snapshots and link streams are special cases of the following representation schema:

> *At time t we perform a measurement. This measurement gives us some, maybe partial, information about a structure of the real-world network. We denote this information by r_t. The symbol r can represent one link, subgraph, or even the whole graph. In this language, the dynamical graph observations is represented by a sequence $\{r_t\}$.*

We note that this schema can be modified in order to capture the lifetime of r_t, if our measurements can provide this information.

There are several monographs that cover some aspects of dynamic networks. The physicists' point of view is presented in the book by Dorogovtsev and Mendes [87]. The reader interested in random graph dynamics may consider Durrett's book [88], and the reader interested in dynamics of small-world networks (wherein most nodes are not adjacent, but one can go from a node to another using only a logarithmically small number of links comparing to the total number of nodes) may read Watts' PhD thesis [20].

The first paper about the Internet topology dynamics has been written in 1993 by Chinoy [16]. Chinoy considered the dynamics of NSFNET backbone network at the AS-level. The first thesis in this context has been defended by Paxson in 1997 [18]. In particular, Paxson "found that most Internet paths are heavily dominated by a single dominant route". In 1997 it was indeed true, but in 2014 we know that there are several different routes between almost any pair of computers (see for example Cunha *et al.* [71]).

Earlier papers in this area abound with technical details and terms, while in modern works, including this thesis, authors try to abstract a bit from specific networks. This makes sense, since it turns out that many real-world networks have similar properties. For example Krioukov *et al.* showed "that the causal network representing the large-scale structure of space-time in our accelerating universe is a power-law graph with strong clustering, similar to many complex networks such as the Internet, social, or biological networks" [77].

Many researchers studied the evolution of some structural characteristics of the Internet topology. For example, Fabien de Montgolfier *et al.* considered

the treewidth and hyperbolicity of the Internet [74]; large-scale topological and dynamical properties of the Internet were considered by Vázquez *et al.* [32]; the dynamics of the multicast topology was studied by Pansiot [58]; Park *et al.* in [42] compared static and dynamic measurements and models of the Internet's AS-level topology; and Lad *et al.* discussed the visualisation of the Internet dynamics [55].

Almost all works in the area of the Internet topology are about large-scale dynamics of the Internet, and only few works considered fine-scale dynamics. The Internet is very big, and it is impossible to actually observe how the complete topology changes every few minutes. However, it is possible to use an orthogonal approach to obtain insight on the dynamics of the routing topology observed at the IP-level [72]. In this thesis, we follow this approach and study *ego-centred views* of this topology. Given a monitor and a fixed set of destinations, one such view is obtained by measuring the routes from a monitor to a set of destinations. This can be performed quickly and with low network load with the `tracetree` tool [72]. Repeating this measurement periodically therefore allows us study the dynamics of this view. Previous work has shown that ego-centred views exhibit strong dynamics, and in particular that the set of observed nodes evolves much more quickly than what was previously expected [64].

1.4 Models of networks

In order to understand better the structure and function of real-world networks, researchers model them. Perhaps the first model was introduced by Rapoport in 1948 [5]. Rapoport considered a random graph as a possible model of some biological networks. A related idea has been developed by Erdős and Rényi in their seminal paper [7].

However, in 1965, Derek de Solla Price showed that the topology of paper-citation networks is different from the typical topology of a random graph. In particular the degree distribution of this network follows a power law [11], while in random graphs degrees are distributed according to a binomial law. In 1999 Faloutsos *et al.* showed that the degree distribution of AS-level Internet graph also approximately follows a power law [24]. Later,' researchers have found many other examples of real-world networks that have power-law degree distributions. We conclude that classical Erdős-Rényi model does not reflect the reality very well. But the superrealism has never been the main purpose of this model, and researchers continue to study random graphs due to their elegance and simplicity.

In order to improve Erdős-Rényi model, scientists have proposed many other models, notable examples include:

- The configuration model, that allows us to uniformly generate a random graph

with a given degree distribution, developed by Bender and Canfield [14].

- The model of small-world networks by Watts and Strogatz [21]. In such networks most nodes are not adjacent, but the typical distance is small.
- The preferential attachment model considered by Barabási and Albert [23]. In this case the probability that a node is connected to another node u is proportional to the degree of u.

In the Internet context, specific models have also been proposed. In [19] Zegura *et al.* presented a quantitative comparison of different Internet topology models. Tangmunarunkit *et al.* compared degree-based and structural network topology generators [35]. Quoitin and Uhlig in [48] considered a model of routing between autonomous systems.

As we know from Section 1.3 almost all real-world networks are intrinsically dynamic. And models should also take this into account. Large-scale dynamic models are widespread in the scientific literature. The Barabási–Albert model [23] is a classical example of a growing network. Vázquez *et al.* [33] studied the Internet at the AS-level and compared "the properties of growing network models with the present real Internet data analysis".

In this thesis we introduce a model of fast topology dynamics [64, 72]. This model adds some dynamics to a static graph of any topology. We discuss this model in details in Chapter 2.

> Many studies have already characterised the static topology of the Internet. However few of these studies have asked how this topology evolves with time. Indeed, it is well known that the Internet is a living organism with many of nodes and links added and removed every day. Therefore, characterising the Internet dynamics can provide interesting insights that may help in the design of future routing protocols or the conception of new types of applications that make use of its evolving topology.

Chapter 2

Model

In this chapter we review previous studies of the dynamics of the Internet topology that consist in periodic measurements of routing trees from a single monitor to a fixed set of destinations. We then identify invariant properties of these dynamics. Based on those observations, we then propose a model for the underlying mechanisms of the topology dynamics. Our model remains simple as it only incorporates load-balancing phenomena and routing changes. By extensive simulations, we show that, despite its simplicity, this model effectively captures the observed behaviours, thus providing key insights on relevant mechanisms governing the Internet routing dynamics.

Previous work has shown that the Internet topology changes much faster than it was expected [64]. In Section 2.1, we review the analysis of these results. In Section 2.2 we identify two factors that play a key role in the observed dynamics: load-balancing routers, and the evolution of the routing topology. Based on these observations, we propose in Section 2.3 a baseline model for the routing dynamics in the Internet that incorporates these two ingredients, using simple choices for modelling these factors. Schematically this model is represented in Figure 2.1. Finally, we present preliminary results of our simulations in 2.4. The next chapters will be devoted to an in-depth study of the model behavior.

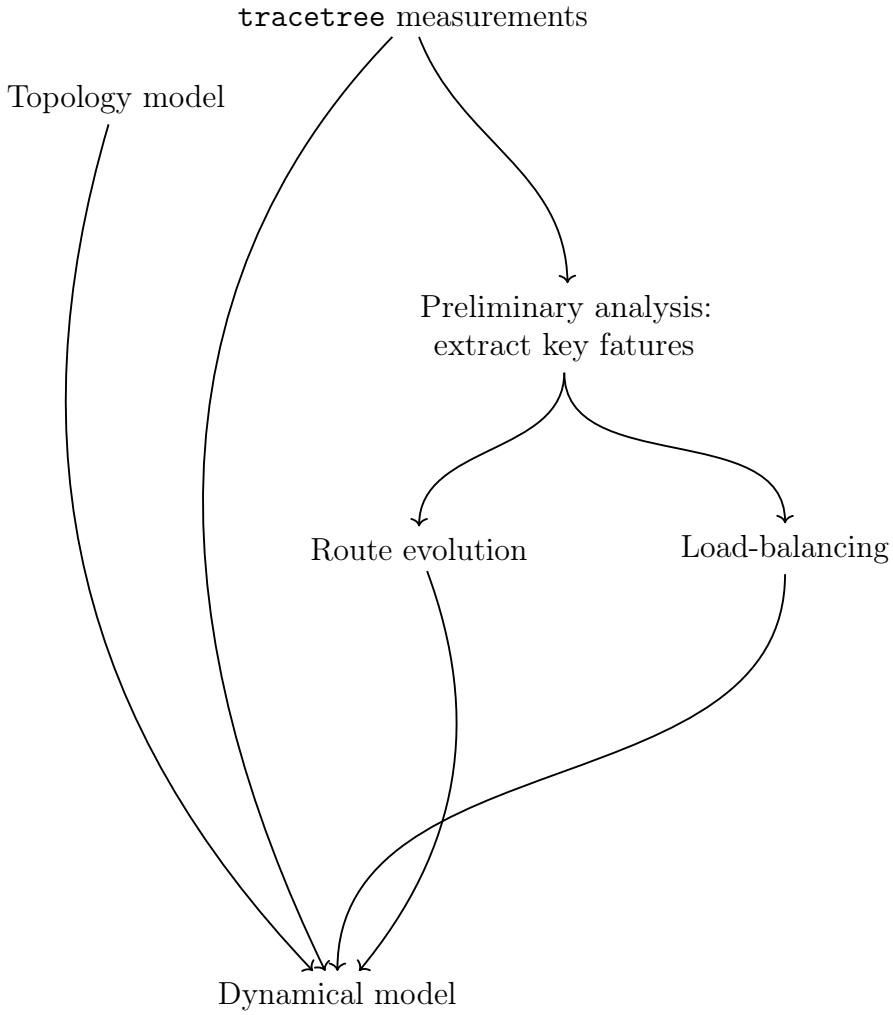

Figure 2.1 – Schematic representation of the model

2.1 Motivation: `tracetree` measurements

The `tracetree` tool [72] collects the *ego-centred view* from a given monitor to a given set of destinations by measuring the routes from this monitor to each destination. This corresponds to a subset of the routing topology, in which nodes are the IP-addresses of routers, and a link exists between two nodes if the corresponding routers are connected at the IP level. Note that the routing topology is different from the physical topology, as two routers may be physically connected by a link that is not used for routing. Note also that we only observe a subset of the whole routing topology, as measuring the routes from a single monitor to a limited set of destinations certainly does not allow the discovery of all nodes and links in this

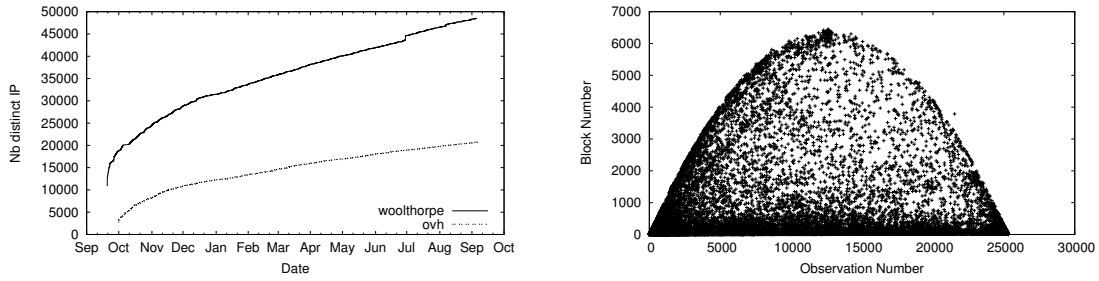

(a) Number of distinct IP addresses observed since measurement beginning.

(b) Observation number vs. block number.

Figure 2.2 – Properties of the observed dynamics.

topology. Moreover, this subset is not representative of the whole topology, see for instance [40, 65]. Keeping this in mind, we will see that we are still able to make interesting observations about the *dynamics* of this topology.

Running the `tracetree` tool periodically allows us to capture the dynamics of ego-centred views. We collected two datasets in this way. The first one, `woolthorpe`, was collected from a monitor in University Pierre and Marie Curie in Paris towards a set of 3,000 destinations. The collection frequency is of one measurement round every 15 minutes approximately. It started in September, 2011 and lasted approximately a year with some small interruptions due to power shortages. This represents a total of 32,018 rounds. The second one, `ovh`, was collected from a French server hosting company. Only 500 destinations were used in order to increase the measurement frequency, which is of one round every one and a half minute approximately. It started in October, 2010 and ended in September, 2011, which represents a total of 318,000 rounds. In both cases, the destinations were chosen by sampling random IP addresses that answered to a ping at the time of the list creation [1]. These datasets are publicly available [1]. Notice that previously measured datasets are also available, for different durations, at different times since 2008.

It is possible that, at a given time, several routes to a same destination co-exist, in particular because of load balancing. Therefore, two consecutive measurement rounds may capture different routes to a same destination even if no routing change has occurred. We study this in the next section, and present below the main characteristics of the *observed* dynamics. We performed our analysis on a representative set of the existing datasets, and made similar observations to the ones we present here. This shows the generality of our observations.

[1] Previous work has indeed shown that tracing routes to unused IP addresses can introduce measurement artefacts [47].

Discovery of new IP addresses. A previous study of the same type of data has shown that the set of observed IP addresses does not stabilize with time [64]. Instead, it was observed that these measurements continuously discover new IP addresses that had never been observed before, at a significant rate. These observations were made on two-months-long measurements. Fig. 2.2a shows that it is also true for very long measurements. It presents the number of IP addresses observed since the beginning of the measurement, for both datasets. A dot (x, y) in this figure means that y different addresses have been observed at least once before time x [2]. We see that, after an initial fast growth, the plot increases significantly for extended periods of time.

This plot presents the number of distinct IP addresses observed, and not the number of distinct routers, as in general several IP addresses, or interfaces, correspond to a same router. Detecting which interfaces correspond to which routers is a difficult task. Though several methods exist, none is 100% accurate. We used the MIDAR tool developed by CAIDA [2], and studied the number of discovered *routers* observed since measurement beginning. The corresponding plot, not presented here due to lack of space, clearly displays the same shape as those of Figure 2.2a. Moreover, previous work has studied the number of distinct ASes discovered by such measurements, and showed that it also increases significantly [64]. All in all, there is a good evidence that new routers are actually discovered at a significant rate, even if part of the observed growth may be caused by discovering new interfaces for already observed routers. As there is no method that allows us to know with certainty which interfaces correspond to a same router, we limit ourselves to the study of interfaces in the rest of the paper.

Stability of IP addresses. To analyse more in depth the dynamics of the ego-centred views, we compute two quantities for each IP address. Its *observation number* is simply the number of distinct rounds it was observed in. Moreover, an IP address is in general observed in blocks of several consecutive rounds, preceded and followed by one or more rounds during which it is not observed. More precisely, the *block number* of an IP address is the number of groups of consecutive rounds in which it is observed. For example, an IP address which was observed on rounds 1, 3, 4, 7, 8, 9, and 10 has an observation number of 7 and a block number of 3.

Fig. 2.2b presents the correlation between these quantities for the `woolthorpe` dataset [3]. Each dot corresponds to an IP address, and its coordinates are its observation number on the x-axis and its block number on the y-axis. The plot

[2] Since the `woolthorpe` dataset was collected after the end of the `ovh` dataset measurement, we shifted x-axis one year for the plot for the `woolthorpe` dataset, so that both plots appear in the same time span.

[3] We computed this plot for the longest uninterrupted part of the measurement, which represents 25322 rounds, i.e. approximately 264 days.

presents a clear parabola shape. This can be explained by load-balancing routers. If a load-balancing router randomly spreads traffic among k paths[4], each router belonging to any of these paths has a probability $p = 1/k$ of being observed at each round, leading to an observation number equal to rp approximately, where r is the total number of rounds performed. A given round is then the first of a consecutive block of observations for one of these routers with the probability p that this router was observed in this round, multiplied by the probability $1 - p$ that it was not observed in the previous round. Multiplying this probability by r gives the expected block number, which is then equal to $rp(1 - p)$ and is the equation of the parabola. This is a simplification of the real case in which a router may belong to paths used by several load balancers, themselves belonging to paths used by other load balancers. In practice, an IP address belonging to load-balanced paths can have any probability p, $0 < p < 1$, of being observed.

We can also observe many dots close to the $y = x/2$ line. They correspond to addresses that are observed only during a finite part of the measurement, and have during that time a probability $p = 1/2$ of being observed, due to load balancing. If such an IP address is observed with a probability $1/2$ during k rounds, its observation number will indeed be $x = k/2$, and its expected block number will be $y = k(1/2)^2 = x/2$.

Finally, many IP addresses are close to the x-axis. This means that, whether they are observed in a large or small number of rounds, they are mainly observed during blocks of consecutive rounds, with few interruptions.

2.2 Causes of the observed dynamics

It is acknowledged that load-balancing routers play a significant role in the observed dynamics of routes with `traceroute`-like measurements [71, 61]. Previous work also suggests that routing dynamics play a key role in the continuous discovery of new IP addresses in our measurements [64]. This section identifies the strong role played by these factors in our observations.

These two factors play different roles. Suppose first that there is no load balancing. In this case, a measurement will discover routing changes as they occur, and the longer a measurement lasts, the more IP addresses it will observe (because more changes will occur). If on the contrary there are no routing changes but load balancing is used, then performing more measurement rounds will lead to observe more IP addresses, independently of the time elapsed between consecutive rounds[5].

[4] It has been shown [61] that per-packet or per-flow load-balancing routers spread `traceroute` probes equally among all paths to the destination, which is roughly equivalent to randomly choosing a path.

[5] This is of course only true under certain conditions on the number of measurement rounds

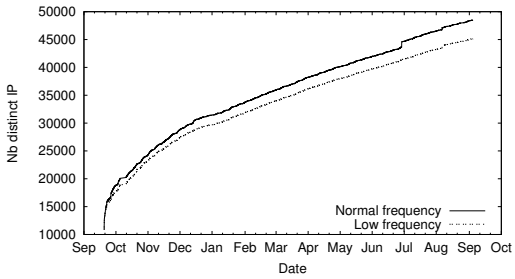

(a) Number of discovered IP addresses as a function of time.

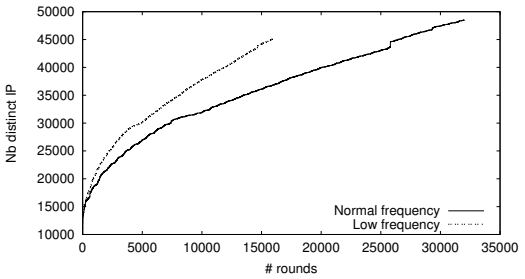

(b) Number of discovered IP addresses as a function of the number of measurement rounds.

Figure 2.3 – How the frequency impacts the number of discovered IP addresses.

The observed dynamics is a combination of these factors.

In order to study this rigorously, we use the `woolthorpe` data set and simulate slower measurements by considering only one out of every two rounds. Fig. 2.3 presents the number of distinct IP addresses observed with both these measurements, as a function of time elapsed since the beginning of the measurement, and the number of measurement rounds performed.

As expected, less IP addresses are observed over time with the slow measurements than with the faster ones. Fig. 2.3a shows that in a given time interval, performing more measurement rounds therefore allows us to discover more IP addresses. This confirms that several measurement rounds are needed to discover all existing routes. This is caused by factors such as load balancing. Conversely, Fig. 2.3b shows that the slow measurements discover more IP addresses *at each round* than the faster ones. Therefore if more time elapses between two consecutive rounds, then each round discovers more IP addresses. This indicates that routes evolve with time.

In both cases, the gap between the plots for the slow and faster measurements are significant, which shows that both factors play an important role in the observed dynamics. This is why we propose a model that incorporates load balancing and route dynamics.

2.3 Model description

Our purpose here is to propose relevant and simple mechanisms that reproduce the observations made in Section 2.1. We do not aim at proposing a realistic model, but rather at providing a first and significant step towards understanding the impact

and the time elapsed between consecutive rounds.

of simple mechanisms on the observed dynamics. This model incorporates four ingredients: the routing topology, the routes from the monitor to the destinations in this topology, load balancing, and routing changes. For modelling each ingredient, we try to make the simplest choice possible, our goal being to obtain a baseline model which makes it possible to investigate the role of each component, and to which future and more realistic models should be compared.

Topology model. We represent the topology by a random graph. In order to strengthen the conclusions drawn from our study, we used two different models generating different topologies: the Erdös-Rényi model [7] which makes no hypothesis on the structure of the graph and is therefore the simplest model possible, and the configuration model [14] in order to generate graphs with power-law degree distributions. The random graph model has two parameters: the numbers n of nodes and m of links. The configuration model has two parameters: the number n of nodes and the exponent γ of the power-law. As we will see in the next section, the comparison between results obtained with both generation processes gives insights on the impact of the topology on the observed dynamics.

Measurements model. Given a graph representing the topology, we assume that the route between the monitor and a destination is a shortest path, which can be obtained by performing a *breadth-first search (BFS)*.

Load-balancing. In order to simulate load balancing, each node chooses at random the next node on a shortest path to the destination, and we therefore implement a *random BFS*. It generates a shortest-path tree from the monitor to the destinations by considering the neighbours of explored nodes in a random order. These routing trees will therefore be different from one random BFS to the next, even if the underlying graph does not change.

Routing evolution. Second, we need to model changes in the routing topology. We use a simple approach based on link rewiring, or *swap*. It consists in choosing uniformly at random two links (u, v) and (x, y)[6] and swap their extremities, *i.e.* replace them by (u, y) and (x, v).

Finally, our simulation setup consists in the following. First, we generate a graph G_1. From G_1, we randomly select one node as the monitor and d nodes as the destinations. We then simulate r measurement rounds by iterating the following steps:

[6]We choose them such that the four nodes are distinct.

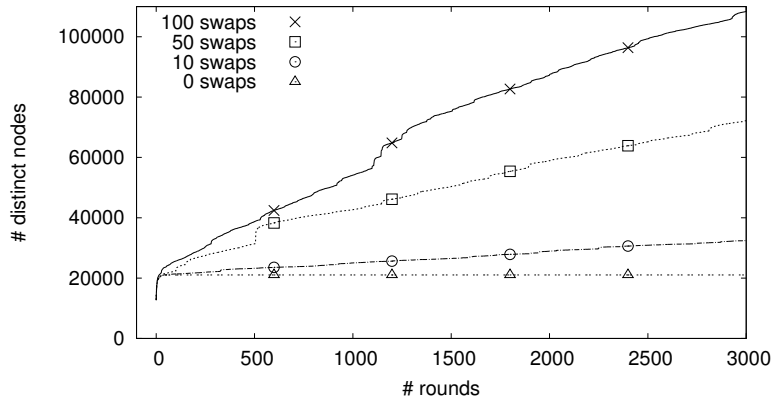

Figure 2.4 – Number of nodes observed since measurement beginning for various values of s (random graphs, $n = 500,000$, $m = 1,000,000$, $d = 3,000$).

1. extract a routing tree T_i from G_i ($i \in [1..r]$) by performing a random BFS from the monitor towards the destinations;

2. modify the graph G_i by performing s random swaps, which produces the graph G_{i+1}. s is a parameter of the model.

This process generates a series of r trees $T_1, T_2, \ldots, T_r$ which simulate periodic `tracetree` measurements, on which we can conduct similar analysis as those we performed on real data.

2.4 Preliminary results of simulation

In this section we show that this model is relevant to explain the dynamic properties presented in Section 2.1. To that purpose, we perform several simulations varying the parameters of the model: the numbers n of nodes, m of links, d of destinations, and s of swaps per round. Our goals are to find (1) whether the simulations reproduce the observations and (2) how the different parameters impact the results and what are the relations between them.

2.4.1 Evolution of the number of distinct nodes.

In order to answer the first question, we present in Figure 2.4 the evolution of the number of distinct nodes observed over time for Erdös-Rényi random graphs with $n = 500,000$, $m = 1,000,000$, $d = 3,000$ and various values of the number s of swaps. It shows a similar behaviour to the one we observed in real data (see

Fig. 2.2a). In particular all the curves present clearly a fast initial growth[7] and then a more or less linear progression. Moreover, the slopes of the curves increase with the number of swaps. This is due to the fact that with a higher number of swaps, the paths to the destinations change more quickly and thus more nodes are discovered at each step.

This figure also shows that when the underlying graph does not evolve ($s = 0$), there is only an initial growth in which all shortest paths are explored. Once all nodes on these paths have been discovered, the curve becomes constant. This confirms the intuition that the regular discovery of new IP addresses in real data may stem from route dynamics.

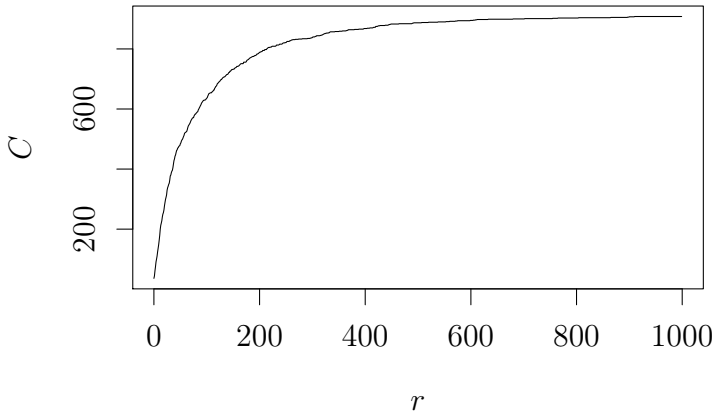

Figure 2.5 – Number of nodes observed since measurement beginning (small and fast evolving random graph, $n = 1000, m = 2000, d = 10, s = 200$). We rapidly arrive at the point when almost all nodes have been observed. We do not observe all 1000 nodes, because there are some nodes of degree 0. We cannot access them with a BFS and their degree does not change with swaps.

We should note that parameters of our model should be chosen carefully. While performing `tracetree` measurements we see only a small part of the Internet and the curve of the number of observed nodes is almost linear. However, this increase must slow down at some time because there are only a finite number of IP addresses. Imagine that we have observed all possible 2^{32} IP addresses, after that our curve will no longer grow, because we have already observed all IPv4 nodes. But we are very far from that point, which probably means that we observe a very small part

[7] this phase lasts more than 1 round, although it is difficult to visualise it on the plot.

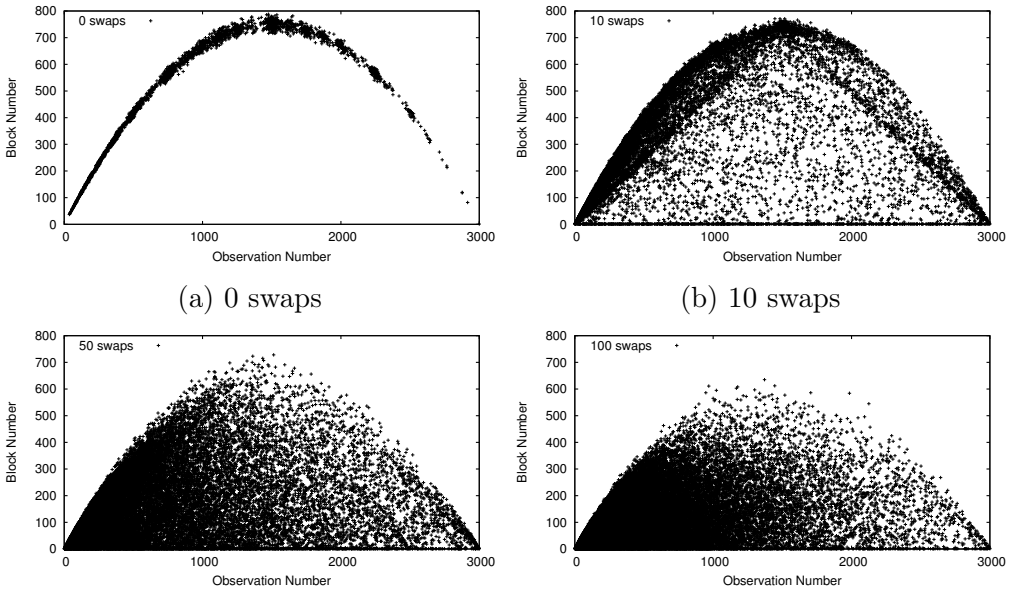

(a) 0 swaps (b) 10 swaps

(c) 50 swaps (d) 100 swaps

Figure 2.6 – Observation number vs. block number for various values of s (random graphs, $n = 500,000$, $m = 1,000,000$, $d = 3,000$).

of the whole Internet topology. On the other hand, in the model we may use very small and (or) fast evolving graphs. In this case we rapidly arrive at the point when almost all nodes (links) have been observed and there is no linear dependence (see for example Figure 2.5). We must therefore be careful to use large enough graphs and small enough numbers of swaps in order to reproduce the observations.

2.4.2 Observation number vs. block number.

We present in Figure 2.6 the correlations between the observation number and block number for the same simulations as in Figure 2.4.

For $s = 10$ (Fig. 2.6b), the main invariants we observed in Fig. 2.2b are reproduced: the parabola, the $y = x/2$ line and a dense strip close to the x-axis. As already explained in Section 2.1, the line $y = x/2$ corresponds to nodes that are observed with probability $p = 1/2$ for a given duration, and are not observed before or after. We also observe a high density of nodes on a line with equation $y = (r - x)/2$, r being the total number of rounds performed. This line has a similar explanation: it corresponds to nodes which are observed with probability $p = 1/2$ for a given duration, and are observed *in all rounds* before and after that. Although this line is not present in Fig. 2.2b, it sometimes can be observed in other datasets, although not as clearly as here.

When no route dynamics is simulated ($s = 0$, Fig. 2.6a), only the parabola is present, thus confirming that this phenomenon observed in real data is due to load balancing mechanisms which are well captured by the random BFS model. At the opposite, when the number of swaps is too high (Fig. 2.6c and 2.6d), route dynamics get the better of load balancing phenomena and the parabola tends to vanish.

In this chapter we conducted periodic measurements of ego-centred views of the Internet topology and studied their dynamics. We isolated invariant characteristics of these dynamics, and identified load balancing and evolution of the routing topology as key factors in the observed properties. Based on this observation, we proposed a model for the dynamics of the topology, which integrates both load balancing and routing changes. Simulations show that this model captures the main characteristics of the dynamics of the ego-centred views.

List of notations

Here you can find the most frequently used notations. Note that this list is not necessarily complete.

G	—	simple graph
V	—	set of nodes
E	—	set of links
n	—	number of nodes in graph (the size of graph)
m	—	number of links in graph
p	—	probability of an edge between two nodes in Erdős-Rényi model
d	—	number of destinations
d	—	number of destinations
s	—	number of swaps per round
SPT	—	shortest path tree
SPS	—	shortest path subgraph
spt	—	number of nodes in *shortest path tree*
sps	—	number of nodes in *shortest path subgraph*
spt_l	—	number of links in shortest path tree
sps_l	—	number of links in shortest path subgraph
S	—	random variable that is equal to the number of nodes in SPS
S_l	—	random variable that is equal to the number of links in SPS
r	—	measurement round number
Δ	—	delay between measurements
α	—	slope of the curve of the number of observed nodes/links
λ	—	mean rate of change (in the context of Poisson processes)
$\mathcal{Pr}_{change}$	—	probability of observing some new links at each measurement round

Chapter 3

Characterising the dynamics

The simulations presented in the previous chapter indicate that the model that we have introduced for simulating the internet topology and its dynamics succeeds in reproducing the main characteristics of the dynamics of the ego-centred views.

In this chapter we will study how the observed characteristics change when we vary the parameters of our model. Our goal here is to understand the influence of parameters on the observations. Ideally, we should be able to predict the behaviour of the model from the parameters without running expensive simulations.

We also examine the case of real-world measurements, keeping in mind that in this case we can control only a few parameters such as: the number of destinations and the frequency of measurements.

Recall that we study is a sequence graph snapshots. Such snapshots may have different shapes and sizes. It is not immediately clear how to study them and which proprieties are the most interesting and relevant. In Section 2.4 we explored some global characteristics of such dynamic graphs, and it turns out that the number of observed routers (or links) observed since measurement beginning captures an important part of the Internet dynamic. Figure 3.1 displays this number as a function of the measurement round. Considering this curve we observe a clear linear trend after an initial phase of faster increase. We are therefore interested in the slope of this curve, because it captures a large portion of the graph's dynamics

in a single number.

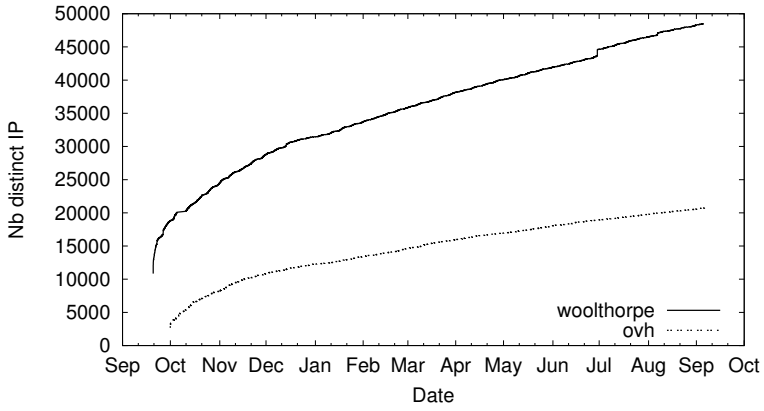

Figure 3.1 – Number of distinct IP addresses observed since measurement beginning: `woolthorpe` and `ovh` datasets.

The main question that arises is whether the curve is truly linear, i.e. whether the plot follows $\mathcal{C} = \alpha r + \beta$ (here $\mathcal{C}$ is the number of nodes or links, and r is the measurement round), or just seems so, both for real data and simulations. It is clear that, due to the random nature of our model and to the number of factors which influence our real-world measurements, we cannot expect an exact linear dependence between r and $\mathcal{C}$. Thus, we should use some tools in order to measure the degree of linearity. Next, it should be noted that our results vary from one simulation to another, due to the random nature of our simulations. However we expect these variations to be small fluctuations around an expected behaviour. Thus, we want to be able to compute *the typical slope* of the curve in the case of the model.

In Section 3.1 we will study whether the curve of the number of observed nodes (or links) is linear or not. In Section 3.2 we will show how we compute *the typical slope* of the curve and discuss how the slope depends on the model parameters. We experimentally determine some general laws of observed network dynamics in Section 3.3. And finally, we validate our findings by considering real-world `radar` measurements (Section 3.4).

3.1 Linearity of the evolution

The curve of the number of observed nodes (or links) presents clearly a fast initial growth and then a more or less linear progression. Stojmenovic, Nayak and Zunic in [54] discussed several measures of linearity that can be applied in order to test

whether a curve (or any shape in general) is (almost) linear. Moreover, they studied which methods provide the closest results to human perception. Inspired by this work we used two methods to test our hypothesis about the linear nature of the curve of the number of observed nodes (or links):

- Pearson's correlation coefficient;
- consistency of coefficients of segmented (piecewise) linear regression;

In the following we describe these methods in details.

3.1.1 Correlation coefficient

Pearson's product-moment coefficient or Pearson's correlation coefficient [4] is the most familiar measure of linear dependence between two variables X and Y. Usually it is called *the correlation coefficient*. Its absolute value is close to 1 when there is a strong linear dependence between two variables, and it is equal to 0 when there is no linear dependence. Pearson's correlation coefficient, denoted by ρ, is calculated using the following formula:

$$\rho_{X,Y} = \frac{\operatorname{cov}(X,Y)}{\sigma_X \sigma_Y} = \frac{\mathbb{E}[(X - \mathbb{E}[X])(Y - \mathbb{E}[Y])]}{\sigma_X \sigma_Y},$$

where cov is the covariance, σ is the standard deviation and $\mathbb{E}$ is the expected value. If we have samples drawn from the variables X_i and Y_i, we calculate a sample correlation coefficient as follows:

$$\hat{\rho}_{X,Y} = \frac{\sum_{i=1}^{n}(X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^{n}(X_i - \bar{X})^2}\sqrt{\sum_{i=1}^{n}(Y_i - \bar{Y})^2}},$$

where $\bar{X}$ and $\bar{Y}$ are sample means.

Note that the correlation, as a measure of linearity, makes sense not only for random variables, but can be applied to any finite sets of points in $\mathbb{R}^2$ (see [54] for example).

We manually identify a value r_0 such that the initial fast increase is over after r_0 rounds. We are not interested here in the smallest such value, so it is not difficult. Then we calculate the correlation coefficient between the measurement round r and the number of observed nodes (or links) $\mathcal{C}$ for the for the right part of our curve ($r > r_0$). We also compute this coefficient for the left part ($r < r_0$) that corresponds to the fast initial growth and should be less linear than the right part.

Perfectly linear data should have a correlation coefficient equal to 1. But our data is not ideal: there are small fluctuations. We say that our data is almost perfectly linear, if the correlation coefficient is greater than or equal to 0.995.

Figure 3.2 shows the correlation coefficients for real and simulated measurements. We use red background colour when the correlation coefficient is less than 0.995 and green otherwise. We see that in all cases, except small and (or) fast evolving graph (see Fig. 3.2a), the right half of the curve is almost perfectly linear and the left part is highly nonlinear. This agrees with visual analysis.

3.1.2 Coefficients of segmented linear regression

Using a simple linear regression we can fit a straight line $\mathcal{C} = \alpha r + \beta$ through the set $\{(r_i, \mathcal{C}_i)\}$ of data points by minimising the sum of squared residuals, i.e. squared vertical distances between the original points and the fitted line. For detailed explanations and formulæ see [9].

Let us divide our curve into several segments, compute the linear fit for each segment, and check that all obtained values for the slopes are consistent. The results are similar for real and simulated data. Let us discuss here only the `woolthorpe` dataset. Figure 3.3 presents the number of observed routers (black curve) and corresponding slopes of segmented linear regression (blue points). Each blue point corresponds to the slope of one segment of the black curve. Each segment has a length of 200 rounds, and starts at the abscissa of the corresponding blue point. We see clearly that after an initial fast decrease, the value of the slope stabilises around an average. Between round 25000 and 30000 we observe an outlier. This happens because our curve is not completely linear, and sometimes there are sharp increases at some points. Such increases may have a strong impact on the local slope but a very small impact on the overall slope.

3.2 Characterisation of the slope

In the previous section we showed that the curve of the number of observed nodes is almost linear after an initial fast increase. Since this curve is a representation of the network dynamics, we can characterise an important part of such dynamics by its slope. Informally, the slope corresponds to the observed speed of the network evolution. In this section we study the dependence between the slope and the model parameters (and the parameters of real-world measurements).

Our results naturally vary from one simulation to another and the plot may present some sharp increases at some points. In order to capture the notion of *typical* slope of a curve in the model, we run a large number of simulations with a given set of parameters and consider the average plot of the result over all simulations[1]. It should be noted that we do not apply this averaging procedure to

[1] Usually we compute an average over 200 simulations.

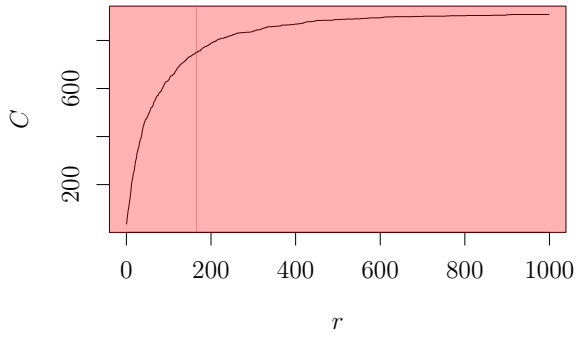

(a) Small and fast evolving random graph, $n = 1000, m = 2000, d = 10, s = 200$.

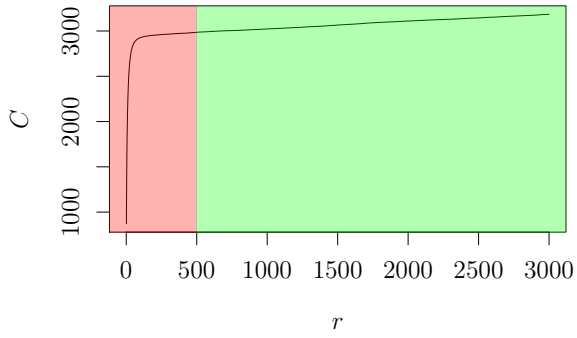

(b) Random graphs, $n = 100,000; m = 800,000, d = 300, s = 2$.

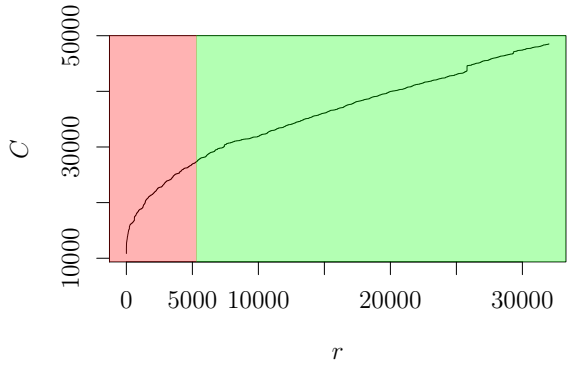

(c) `woolthorpe` dataset.

Figure 3.2 – Correlation coefficient for several curves of the number of nodes observed since measurement beginning. Red background means low linearity (corr. < 0.995), green background means high linearity.

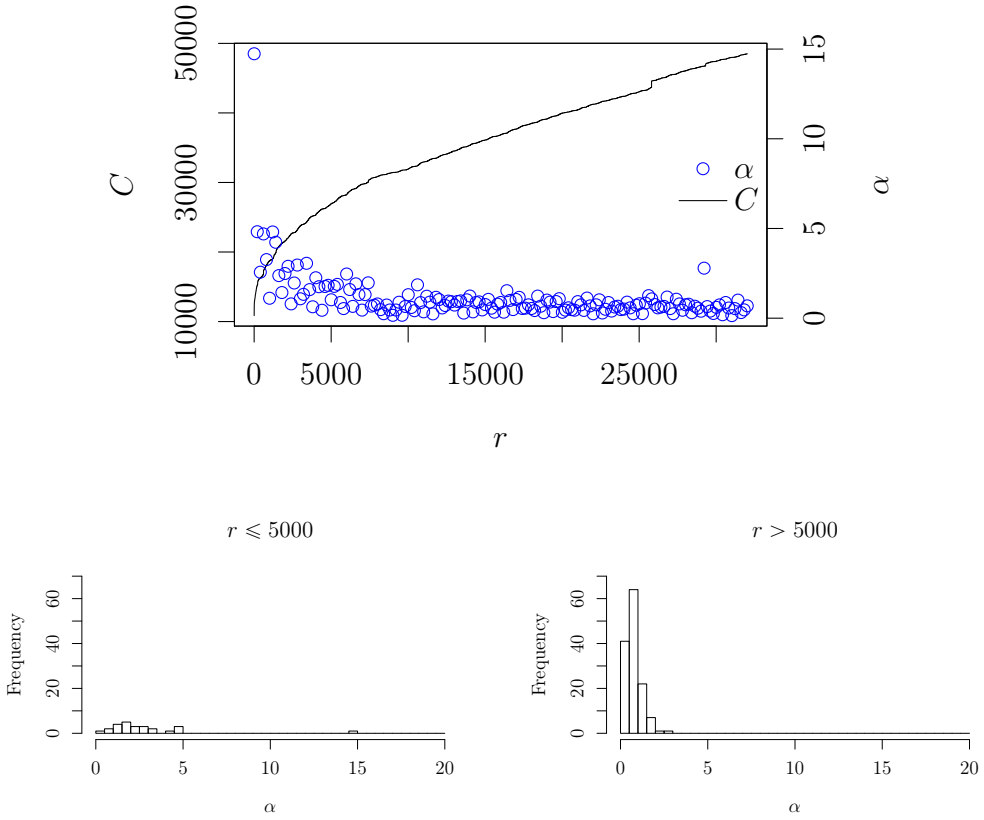

Figure 3.3 – Coefficients of segmented linear regression. The black curve represents the number of observed routers as a function of the number of measurement rounds performed (we use the `woolthorpe` dataset). Each blue point corresponds to the slope of one segment of the black curve. Each segment has a length of 200 rounds, and starts at the abscissa of the corresponding blue point. The histograms present the distribution of the slopes in the initial phase ($r \leqslant 5000$) and after.

the real data, because we have only one curve for each dataset and the averaging across different datasets is meaningless.

We manually identify a value r_0 such that the initial fast increase is over after r_0 rounds, i.e. the curve is linear for all $r > r_0$. We recall that we are not interested here in the smallest such value, so it is not difficult. Then we calculate the slope of the remaining stable part, using a linear regression.

We apply this methodology to explore the simulations over Erdös-Rényi graphs, which we call *random graphs* (Subsection 3.2.1) and we then we compare the

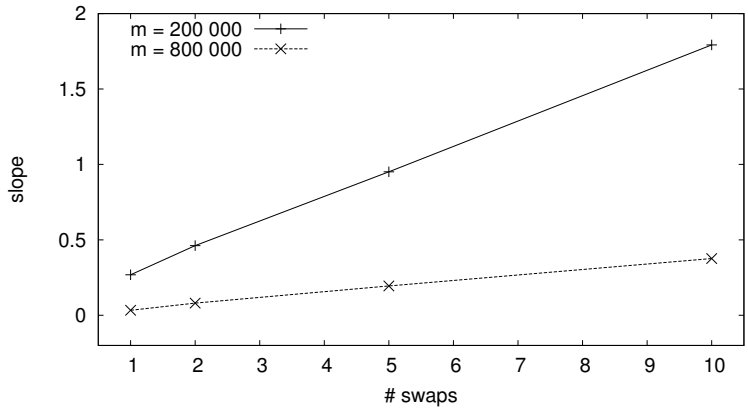

Figure 3.4 – Impact of the number of swaps on the slope (random graphs, $n = 100,000$, $d = 300$).

obtained results with simulations performed on random graphs with a power-law degree distribution, which we call *power-law graphs* (Subsection 3.2.2).

3.2.1 Random graphs

We applied the methodology presented above on simulations for random graphs with $n = 100,000$ nodes and $d = 300$ destinations for various numbers of links and swaps. It makes it possible to study the impact of the number of swaps s on the slope α. Results are presented Figure 3.4. We observe that the slope increases almost linearly with the number of swaps. This indicates a strong correlation between the observation of new nodes and the underlying dynamics. With a higher number of swaps, the topology changes more frequently and, consequently, more paths are affected at each round.

Besides, the plot also shows that the relation between the two quantities is affected by the total number of links in the graph. Intuitively, the swaps are less likely to impact the paths from the monitor to the destinations if the graph is more dense. This is confirmed on the plot: for a given number of swaps, the slope is higher for graphs with $200,000$ links than for graphs with $800,000$ links.

Swaps vs. links. We studied more deeply the relation between the number of swaps s and the number of links m by varying the two parameters at the same time. We set $n = 100,000$ and $d = 300$ and made simulations with several values of s and m. The results are presented on Figure 3.5. We first observe that, for a given number of swaps, the larger the number of links, the smaller the slope of the corresponding curve, which confirms that when the number of links increases, a smaller fraction of them is affected by the swaps. Notice however that, for

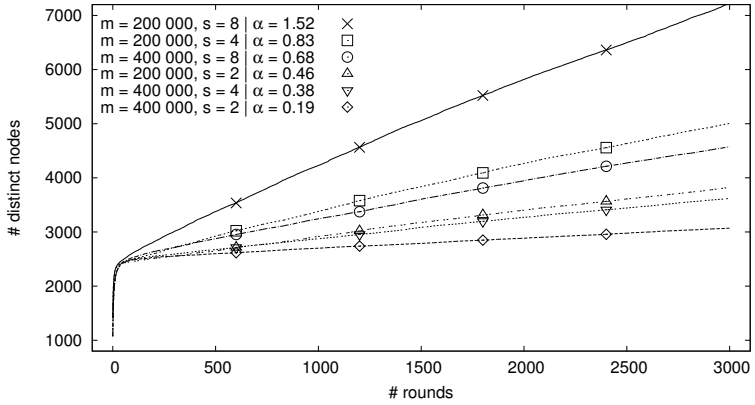

Figure 3.5 – Relation between links and swaps (random graphs, $n = 100,000$, $d = 300$).

different simulations with a same ratio s/m, the corresponding slopes are not equal. For instance, we can observe that for $s/m = 10^{-5}$ (the two curves marked with triangles), the slopes are equal to 0.46 and 0.38, respectively.

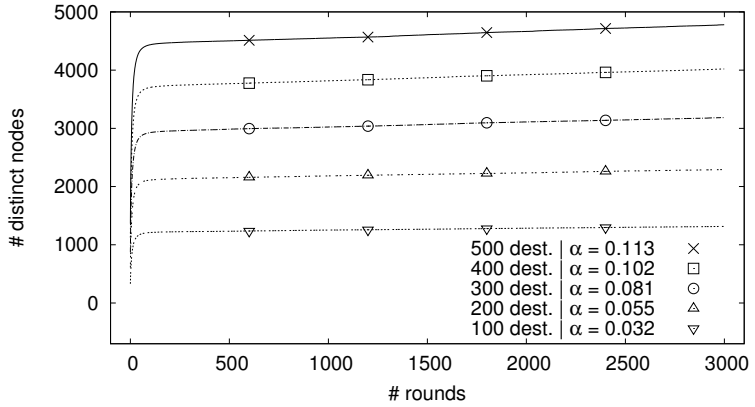

Figure 3.6 – Impact of the number of destinations (random graphs, $n = 100,000$, $m = 800,000$, $s = 2$).

Number of destinations. Finally, we also studied the impact of the number of destinations d for graphs with $n = 100,000$, $m = 800,000$ and $s = 2$ (Figure 3.6). Intuitively, increasing the number of destinations causes the number of nodes on the shortest paths to the destinations to increase. Indeed, we observe that the initial growth phase, which corresponds to the discovery of all nodes on all shortest paths to the destinations, reaches a higher value when the number of destinations

increases. As before, this phase is followed by a linear progression. Notice that increasing the number of destinations also increases the slope. This is clearly due to the fact that since the size of routing trees from the monitor to the destinations increases with the number of destinations, the probability for a swap to affect such a routing tree increases likewise.

3.2.2 Power-law graphs

In order to study the impact of the underlying structure, we compared the behaviours observed above for random graphs to those obtained for power-law graphs.

In summary, we use the following procedure [28, 14]: (1) given an exponent γ, we randomly generate a list of degrees that respects the following power law:

$$f(d) \sim d^{-\gamma}, \gamma > 0,$$

where $f(d)$ is the fraction of nodes with degree d; (2) for each node, create as many half links as the value of its degree; (3) randomly choose pairs of half-links and connect them to form links.

The observations made for power-law graphs are the same, qualitatively, as the ones made for random graphs. The number of nodes observed since measurement beginning displays a linear progression after a fast initial growth.

We do however observe a *quantitative* difference: the slope of the curve of the number of observed nodes for power-law graphs is smaller than the slope corresponding to random graphs. This can be observed in Figure 3.7, which presents the number of nodes observed since measurement beginning for a power-law graph with exponent 2.3, which corresponds to approximately 200,000 links, and two random graphs. The slope of the curve for the power-law graph is indeed much smaller than the one for a random graph with the same number of nodes and links.

The average distance may play a role in this. It has indeed been proven that the average distance is smaller for power-law graphs with exponents in the range $2 < \gamma < 3$ (for which it is in the order of $\log \log n$ [38]), than for random graphs with bounded mean degree (for which it is in the order of $\log n$ [86]). This implies that shortest path trees from the monitor to the destinations will have fewer nodes in power-law graphs than in random graphs, naturally inducing the observation of fewer new nodes. However, though this certainly plays a role, this is not enough to explain the observed difference. Figure 3.7 shows that the slope of the curve for the power-law graph is also smaller than the one for a random graph with the same average distance (and hence with $m = 1,200,000$ links).

As we can observe from the three graphs used above as an example, the structural differences between random and power-law graphs lead to important differences in the average distance and/or in the number of links. As observations made in

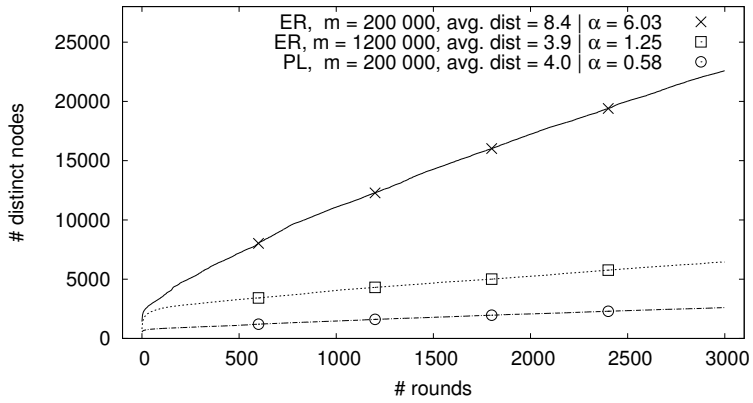

Figure 3.7 – Comparison between random (ER) and power-law (PL) graphs ($n = 100,000$, $d = 300$, $s = 50$).

	Avg. dist.	spt_l
power-law, $m = 2 \cdot 10^5$	4.0	545
random, $m = 12 \cdot 10^5$	3.9	792
random, $m = 2 \cdot 10^5$	8.4	1420

Table 3.1 – Average distance and size of shortest path trees for different graphs. $n = 100,000$, $d = 300$.

the previous section suggest, the slope of the curve for the number of link should intuitively be proportional to the probability that a given swap will change the shortest path tree from the monitor to the destinations. Let us call spt_l the *typical*[2] number of *links* in a shortest path tree from the monitor to the destinations.

Table 3.1 compares the value of spt_l for a power-law graph and two random graphs: one with the same average distance and one with the same number of links. We observe that it highlights additional structural differences between these two types of graphs: though the average distance is approximately equal in a power-law graph with exponent 2.3 and a random graph with $12 \cdot 10^5$ links, the numbers of links in shortest path trees from the monitor to the destinations are significantly different.

[2]In the same way that we performed several simulations and averaged the results in order to obtain the typical behaviour for given model parameters, we generate several graphs with the same size and number of destinations to compute the typical size of a shortest path tree.

3.3 In search of unified laws

After analysing the results of our simulations, we have seen that the relations between the slope of the curve of the number of observed nodes and the model parameters are qualitatively similar in the cases of random and power-law graphs. The slope increases with the number of swaps and the number of destinations. In this context the following question naturally arises: *can we find some common patterns that will unify the model's behaviour for random and power-law dynamic graphs?* We propose in this subsection some general laws that approximately hold in both cases.

Since swaps affect links, the probability that a shortest path tree is affected by a swap depends on its size in terms of links, and it is relevant to study the number of *links* observed since measurement beginning, which has the same behaviour as the number of observed nodes. Therefore, we start by investigating the impact of the typical number of links in a shortest path tree, spt_l, on the observed value of the slope. Figure 3.8 plots the slope of the curve of the observed number of links since measurement beginning vs spt_l for different types of graphs. For each type of graph, we perform simulations using different numbers of destinations, which induces different slopes and different values of spt_l. We observe that the slope increases linearly with spt_l, but at different rates for different types of underlying graphs.

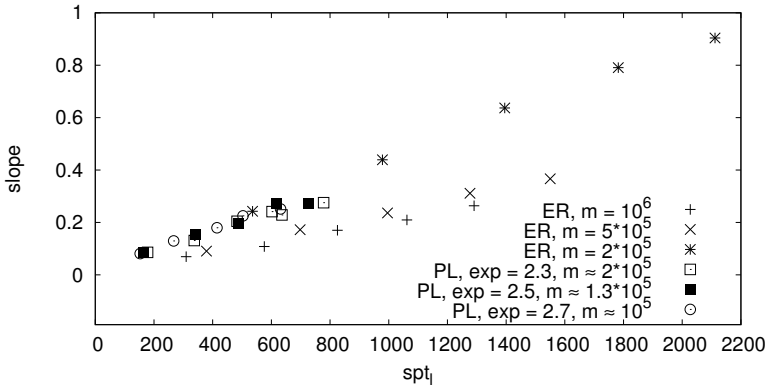

Figure 3.8 – Slope of the curve of the observed number of links vs spt_l for different graphs and different destination numbers. $s = 2$. For all graphs we used several values of d: $d = 100, 200, 300, 400, 500$.

Our intuition suggests that the slope, should be proportional to the ratio between the typical size of the shortest path subgraph and the total number of links when. Thus, trying to obtain a general law, we divide the size of the shortest

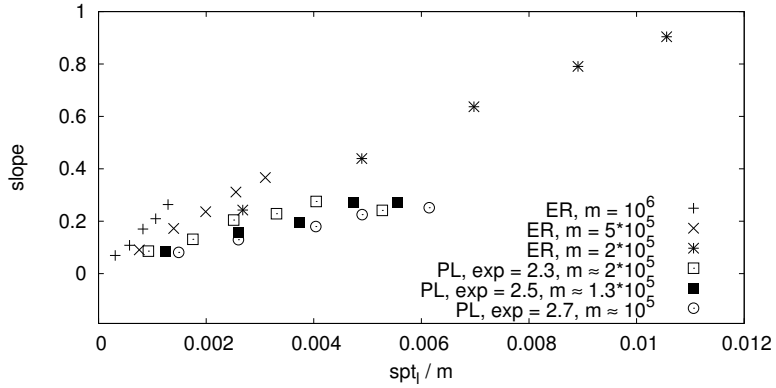

Figure 3.9 – Slope of the curve of the observed number of links vs spt_l/m for different graphs and different destination numbers. $s = 2$. For all graphs we used several values of d: $d = 100, 200, 300, 400, 500$

path subgraph by the total number of links. Figure 3.9 plots the slope vs $\frac{spt_l}{m}$. We observe a strong correlation between these two quantities, meaning that $\frac{spt_l}{m}$ plays an important role in the observed behaviour, for all types of graph. Though this does not fully allow us to understand the model's behaviour, this shows that the number of links and the size of a shortest path routing tree are key parameters for understanding the quantitative difference between random and power-law graphs.

Besides the relation between $\frac{spt_l}{m}$ and the slope of the curve of observed links, there are others that affect the model's behaviour. Here we present the evolution of the slope when we simultaneously vary the number of links and the number of destinations. We use 10^5 nodes, 2 swaps per round, and we compute the typical slope by averaging over 460 simulations with the same parameters, we proceed as follows:

- For random graphs we vary the number of links $m \in \{2 \cdot 10^5, 3 \cdot 10^5, 4 \cdot 10^5, 5 \cdot 10^5, \ldots 9 \cdot 10^5, 10^6\}$ and the number of destinations $d \in \{100, 200, 500\}$.
- For power-law graphs we vary the exponent of power-law distribution $\gamma \in \{2.1, 2.3, 2.4, 2.5, 2.6, 2.7, 2.8, 2.9\}$ and the number of destinations $d \in \{100, 200, 500\}$.

Figure 3.10 presents the interpolated[3] surface $(m, d) \mapsto \alpha$ for random graphs. Figure 3.11 shows the surface of the same type, but for power-law graphs. Clearly, both surfaces do not look like planes. Even more, they are not monotonic and we observe some oscillations. Particularly, the slope evolution is not monotonous

[3]We use `gnuplot`'s splines, which is an interpolation based on "thin plate splines" [13].

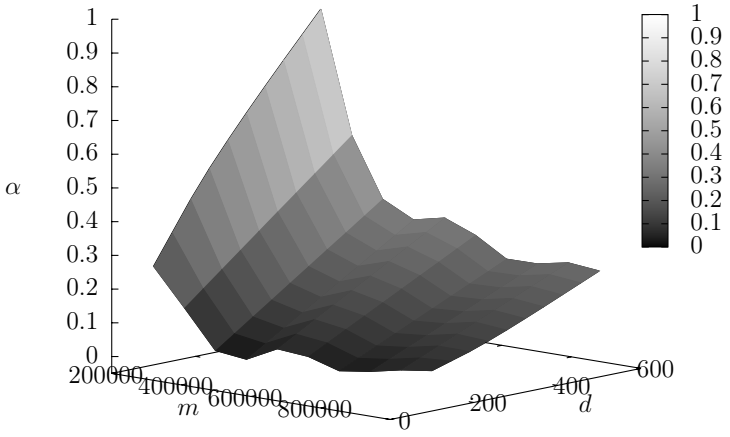

(a) Interpolated surface.

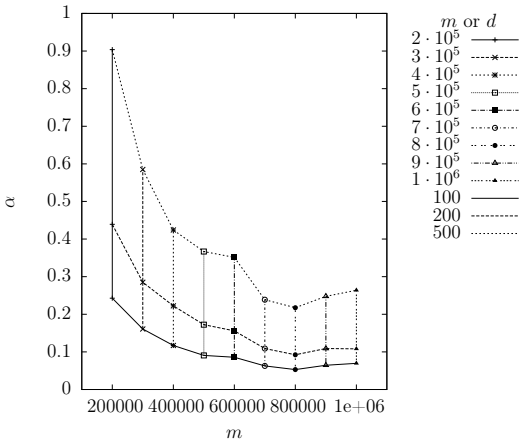

(b) Projection $m \mapsto \alpha$.

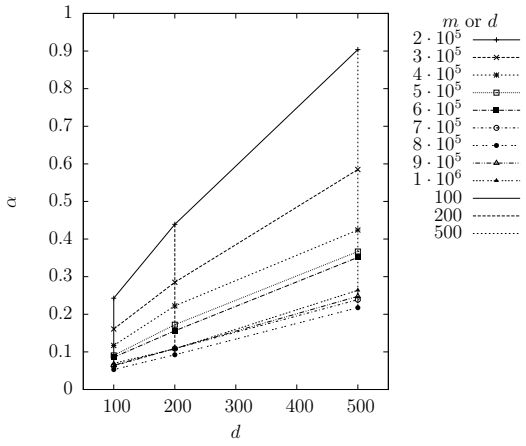

(c) Projection $d \mapsto \alpha$.

Figure 3.10 – Surface $(m, d) \mapsto \alpha$ and its projections. Random graphs ($n = 100000$, $m \in \{200000, 300000, \ldots 1000000\}$, $d \in \{100, 200, 500\}$, $s = 2$).

when we increase the density of a graph by increasing the number of links (see Figures 3.10b and 3.11b).

3.3.1 Impact of the size of the shortest path subgraph

As Figures 3.10b and 3.11b show the slope is not monotonous as the number of links increases. The state of the art has shown that similar oscillations arises when we consider the average proportion of links that lie on all shortest paths from a given node to all other nodes [46].

Here we denote by SPS the subgraph of all shortest paths between the monitor and destinations. Because of load-balancing, one measurement gives us only one shortest path tree (SPT), but the shortest path subgraph may contain several trees. Thus, the size of the shortest path subgraph should play an important role in the observed dynamics.

Let sps_l be the size of the shortest path subgraph between the monitor and destinations, i.e. the number of links that are on the union of all shortest paths between the monitor and destinations. Consider a notion of a unit slope (or slope per unit of the network structure), i.e. the ratio $\frac{\alpha}{sps_l}$. We expect that the unit slope should decrease inversely proportional to the number of links in the graph. Figure 3.12 shows that this is almost true for different random and power-law graphs.

Using this relation between α, spt_l and m, we are able to predict approximately the slope without performing costly simlations. For example, if we have Erdős-Rényi graph with $m = 250000$ links [4], $n = 100000$ nodes and we perform $s = 2$ swaps per round, the slope will be approximately equal to $0.000175 \cdot spt_l$ according to Figure 3.12a. However, we need to know the value of spt_l. For the moment, we can estimate spt_l only experimentally. In Chapter 4 we will make a first step towards analytical expression for spt_l.

We also observe that that there is a non-linear component that causes oscillations of the slope. On the one hand, the presence of non-linear dependence complicates the characterisation of the observed dynamics. But on the other hand it naturally raises another interesting question: how can we explain them? For the moment, we don't know the exact answer. In the next Chapter we will give an explanation of a similar phenomenon: "sps_l oscillates when the density of random graphs grows". We believe that the underlying mechanisms of these oscillations and the mechanisms of the oscillations discussed in this section are the same. Clearly, in order to understand better the dynamics of the network we should study in details the size of the shortest path subgraph. This will be the focus of Chapter 4.

[4] Note that we didn't performe simulations with this value of m.

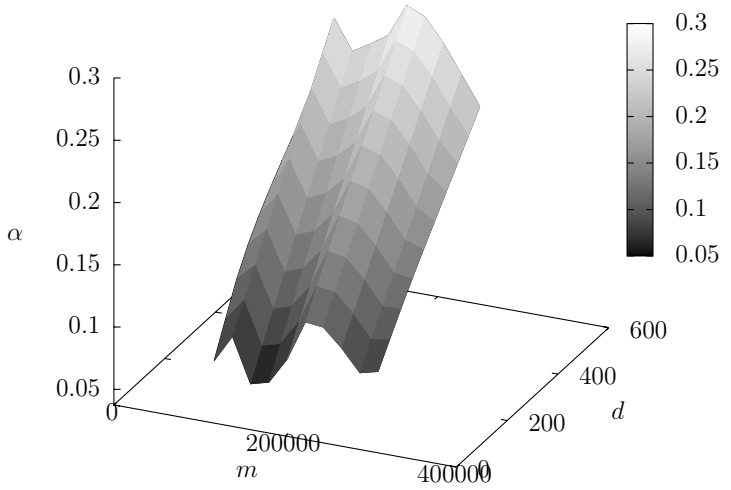

(a) Interpolated surface.

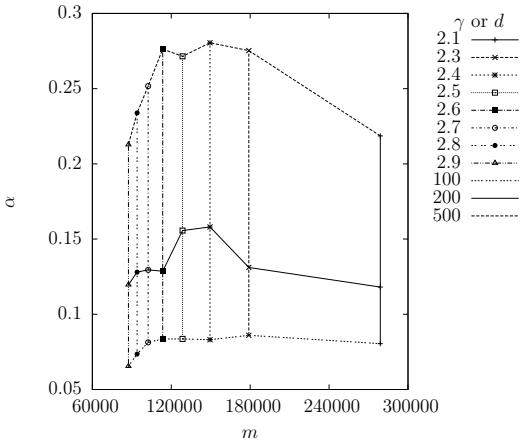

(b) Projection $m \mapsto \alpha$.

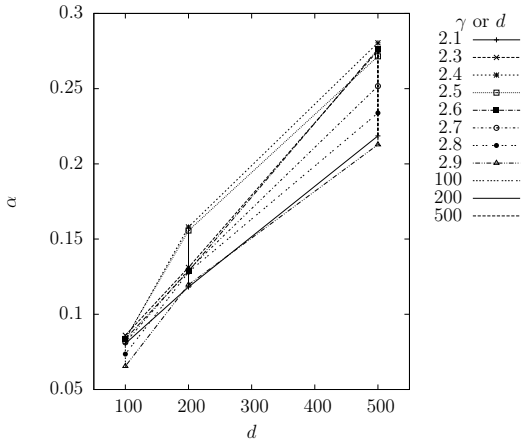

(c) Projection $d \mapsto \alpha$.

Figure 3.11 – Surface $(m, d) \mapsto \alpha$ and its projections. Power-law graphs ($n = 100000$, $\gamma \in \{2.1, 2.3, 2.4, 2.5, 2.6, 2.7, 2.8, 2.9\}$, $d \in \{100, 200, 500\}$, $s = 2$).

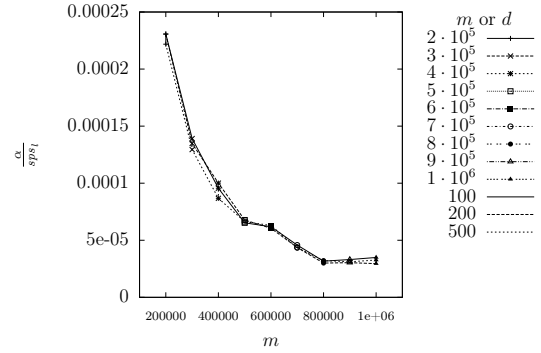

(a) Random graphs.

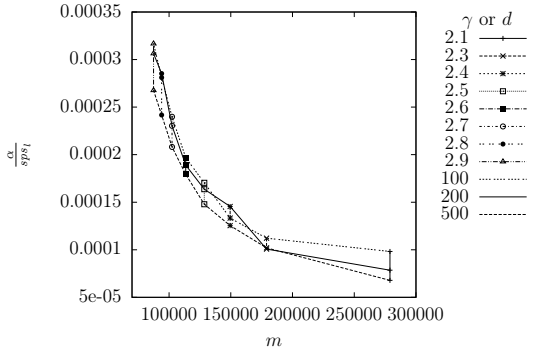

(b) Power-law graphs.

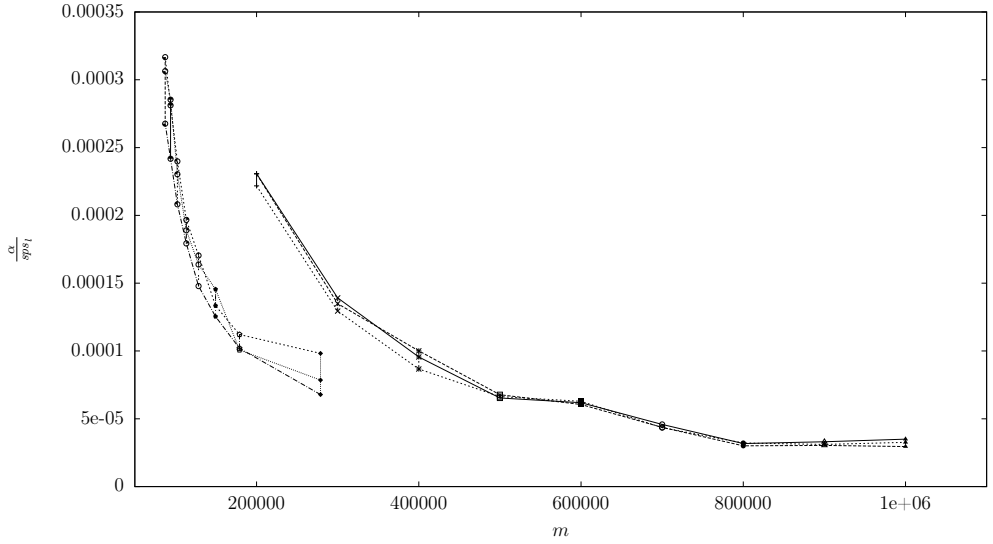

(c) Random and power-law graphs.

Figure 3.12 – The number of links m vs. $\frac{\alpha}{sps_l}$.

3.3.2 Probability of shortest path subgraph modifications

Let $\mathcal{P}r_{change}$ be the probability of observing some new links at each measurement round. We estimate empirically $\mathcal{P}r_{change}$ by computing the number of rounds in which we see a new link divided by total number of rounds. $\mathcal{P}r_{change}$ is another quantity that describes the curve of the number of observed links. Our intuition suggests that $\mathcal{P}r_{change}$, like the slope, should be proportional to $\frac{sps_l}{m}$:

$$\mathcal{P}r_{change} \propto \frac{sps_l}{m} .$$

Figure 3.13a confirms our intuition: we see the same linear law for random graphs with different densities. The case of power-law graphs is slightly different (see Fig. 3.13b), but a strong correlation is present.

Experimenting with various relations between different model parameters and different characteristics of observed dynamic, we have found one very interesting relation:

> *The probability that a shortest path subgraph will change before the next measurement occurs depends almost linearly on the ratio between the size of shortest path subgraph and the total number of links in our graph.*

3.4 Real-world measurements

Studying the model is an important task, but also we should also study the real data in order to confirm our findings. Here we explore how the frequency of measurements and the number of destinations change the observed slope of the curve of the number of observed nodes. Since the results for the `woolthorpe` and `ovh` datasets are qualitatively similar, we present here only the results for the `ovh` dataset.

3.4.1 Frequency of measurements

In the model, the number of swaps corresponds to the frequency of measurements. For example, when we perform twice as many swaps, it means that we measure two times less frequently. Therefore, while in real measurements we cannot control the number of swaps, we can change the frequency, which is equivalent. Figure 3.14 shows how the slope increases when we increase the delay between measurements. These results are similar to the model's results presented in Figure 3.4. Note that here the slope is computed for the curve $r \mapsto \mathcal{C}$, that is the number of links observed since measurement beginning *as a function of the number of measurement*

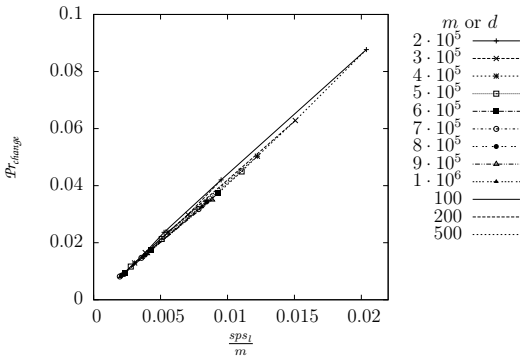

(a) Random graphs.

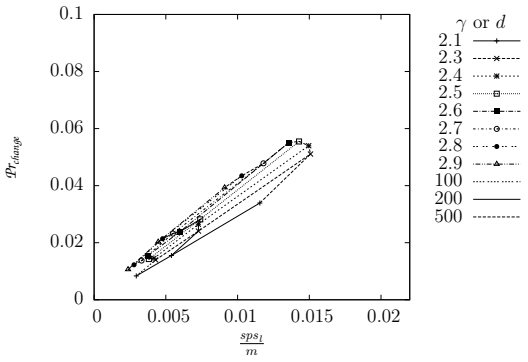

(b) Power-law graphs.

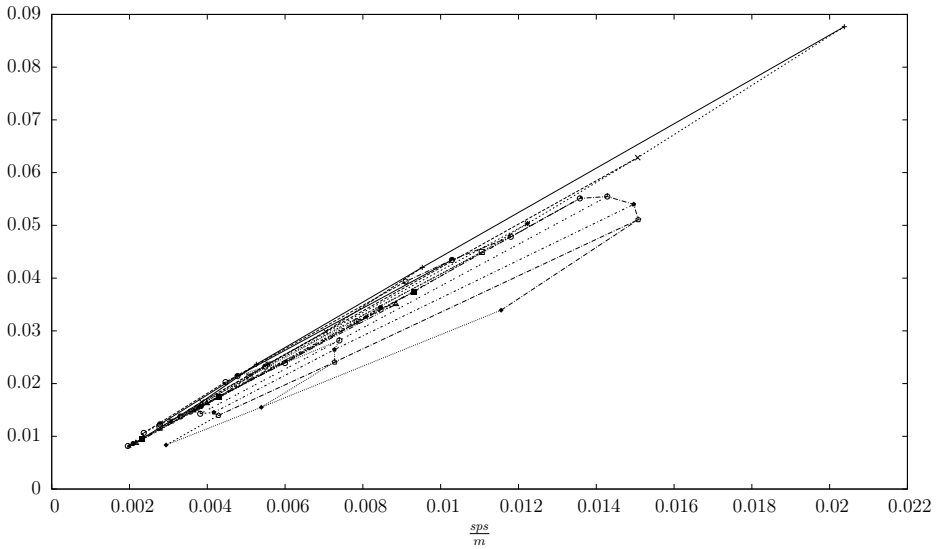

(c) Random and power-law graphs.

Figure 3.13 – Probability of change vs. $\frac{sps_l}{m}$.

rounds performed. The case, where we study the number of observed links *as a function of time* ($t \mapsto \mathcal{C}$, where t is timestamp), will be considered in the first section of Chapter 5. Actually these two cases are almost equivalent, because we typically perform our measurements at constant rate, every Δ seconds, and there is a transformation $t = t_0 + r\Delta$. But the first case is more illustrative in the context of this chapter.

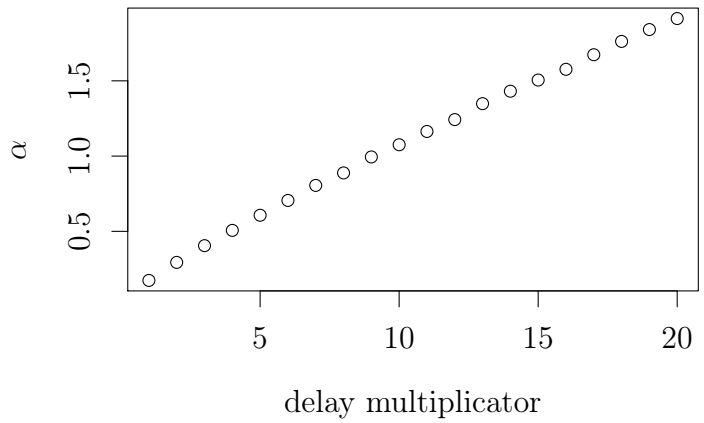

Figure 3.14 – Impact of the frequency for `ovh` dataset: how the slope of the curve $r \mapsto \mathcal{C}$ changes when we multiply the delay between measurements.

3.4.2 Size of the shortest path tree vs slope

As Figure 3.8 shows, in the model the slope is proportional to the number of links on the shortest path tree if the other parameters of the model remain the same. Intuitively, real-world data should exhibit the same behaviour. In order to vary spt_l we change the number of destinations. The original `ovh` dataset contains a sequence of shortest path trees from a monitor to 500 destinations. We obtain measurements with smaller number of destinations by applying the following procedure:

1. We choose x destinations out of 500 available ones.
2. From each shortest path tree we remove nodes and links that do not lead to any of the previously chosen x destinations.

In Figure 3.15 we plot the slope against spt_l. The dependence is approximately linear like in the case of the model.

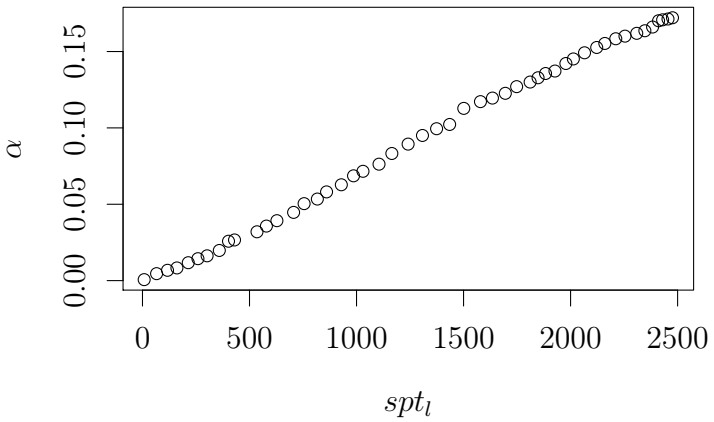

Figure 3.15 – Slope of the curve of the observed number of links vs spt_l for `ovh` dataset.

However, we cannot directly study the dependence between the size of *sps*-sequence and $\mathcal{P}r_{change}$ in the case of real-world measurements, because in this case have only a sequence of *spt*-s. We will reconsider this problem in slightly different setting in Chapter 5.

As a conclusion, by exploring the impact of the parameters on the observed dynamics, we experimentally identified several relations or laws. Some of them are almost exact, but others give only an idea of how the observed dynamics will change, if we consider, for instance a different type of underlying graph. We validated our findings using real world data. Now, we are able to predict some characteristics of the model from the parameters without running expensive simulations.

However, along with successfully identified relations, the model shows some non-linear and non-monotonic behaviours, that complicate analysis. A search for the explanation of these phenomena gives a direction for future works. For example in the next Chapter we will show that the typical size of shortest path subgraph, which play a key role in the observed dynamics, also increases non-monotonously when then graph's density grows.

Chapter 4

Size of shortest path subgraphs

In the previous chapter we have seen that the size of the shortest path subgraph between the source and destinations plays an important role in the observed network dynamics.

In this chapter we study in details the size of shortest path subgraph in the case where there is a single destination. We consider a random graph $G(n,p)$, and we denote by $SPS(u,v)$ the subgraph of all shortest paths between two nodes u and v. We show that the size of $SPS(u,v)$ follows a nontrivial probability law with several local maximum values. In particular, in some cases the average size is a combination of these maxima and not a value that can be reached in practice. Also we approximate the expected number of nodes in $SPS(u,v)$, when the distance between u and v is known (we give the exact distribution, when the distance is equal to 2).

We have seen in Chapter 3 that the size of the shortest path subgraph between the source and destinations plays an important role in the observed dynamics. In this chapter we start to develop a theoretical approach in order to describe this size. We use a random graph [5, 6, 7, 8] as a model of the Internet topology. We recall that if the Internet is modelled by a random graph and if we assume that information follows a shortest paths, it is only possible to observe links that are on the shortest paths between the measurement node and other nodes. Previous works studied through massive numerical simulations the particular case where all nodes are destinations, i.e. the shortest paths from a given node to all other

nodes [46]. It has been showed that the average ratio between the number of links in these shortest paths and the total number of links existing in the graph oscillates when the density grows. These oscillations have bean analysed in [56]. We can see that the average number of links (not the proportion) also oscillates when p grows (see Figure 4.1).

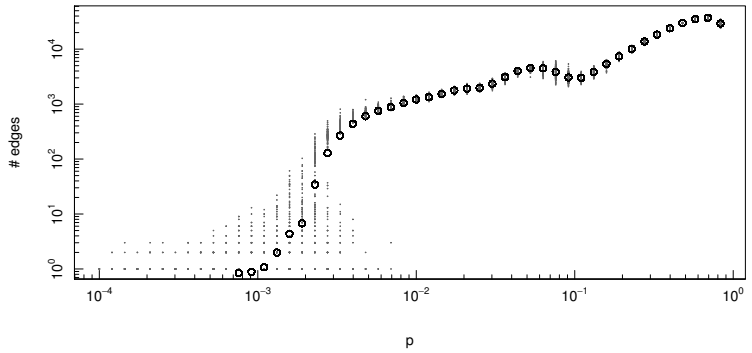

Figure 4.1 – Evolution with p (log scale) of the number of links (log scale) that belong to a shortest path from a node to all other nodes in $G(n, p)$ with $n = 500$ nodes. For each value of p we generated 200 graphs. Each grey point corresponds to an observed value. Black circles represent the average over these graphs. A fast increase around $p = 2 * 10^{-3}$ corresponds to the connectivity threshold.

A measurement that uses all IP addresses as destinations takes a very long time, therefore such measurements cannot be performed with sufficient frequency. For this reason, in reality we use only a small portion of available IP addresses as destinations. Thus, in our model we also take only a small portion of all nodes as the destinations. For example, if we have 500 nodes in total, we take only 3 of them. Figure 4.2a and 4.2b show the size of shortest path subgraph from a node to 3 other nodes in random graphs of different density. We see that: (i) the average size oscillates when the density grows; (ii) in contrast to Figure 4.1, points in Figures 4.2a and 4.2b are less homogeneous, i.e. there are a few dense areas separated by empty space.

In this chapter we refine these analysis by studying the number of nodes (or links) which belong to a all shortest path between two nodes. In this case also, we observe that the averages fluctuates when p varies. The average number of links (Figure 4.2c) oscillates in similar way to the average number of nodes (Figure 4.2d). A sharp increase of the average around $p = 2 * 10^{-3}$ corresponds to well-known phenomenon in Erdős–Rényi graphs: the phase transition that corresponds to the appearance of a giant component (see [86] for detailed explanations). For smaller

values of p, with high probability, u and v do not belong to the same connected component and therefore the average number of nodes or links existing on the shortest paths between them is very small.

Surprisingly, we observe that for many values of p the distribution is non-trivial, with several local maxima; correspondingly the average is a combination of these maxima and not a value that can be observed in practice. Figure 4.3 shows such distributions for different values of p.

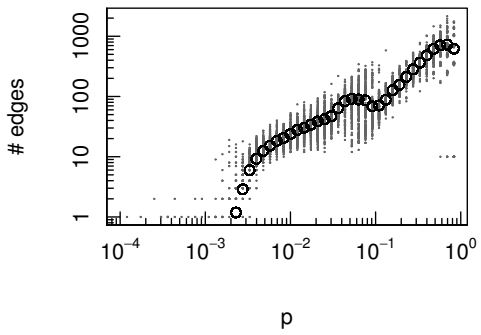

(a) 3 destinations, links

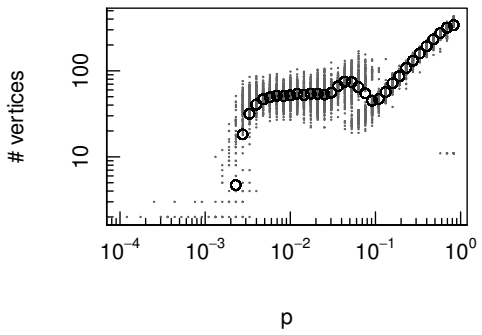

(b) 3 destinations, nodes

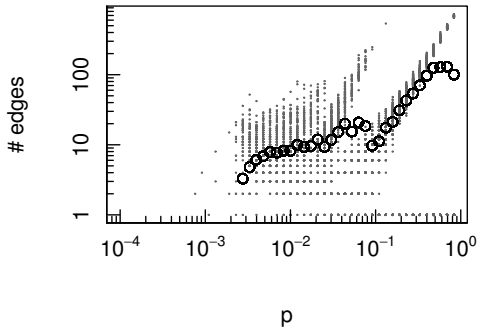

(c) 1 destination, links

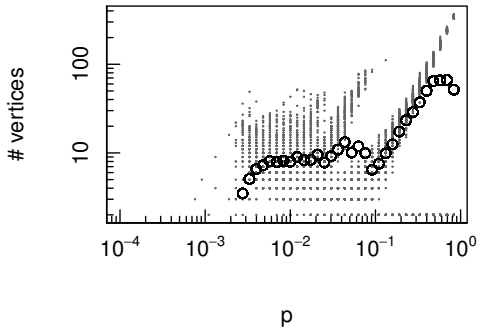

(d) 1 destination, nodes

Figure 4.2 – Evolution with p (log scale) of the number of links and nodes (log scale) that lie on a shortest path from a node to 1 or 3 other nodes from $G(n = 500, p)$. For each value of p we generated 200 graphs. Each grey point corresponds to an observed value. Black circles represent the average over these graphs. A fast growth around $p = 2 * 10^{-3}$ corresponds to the connectivity threshold.

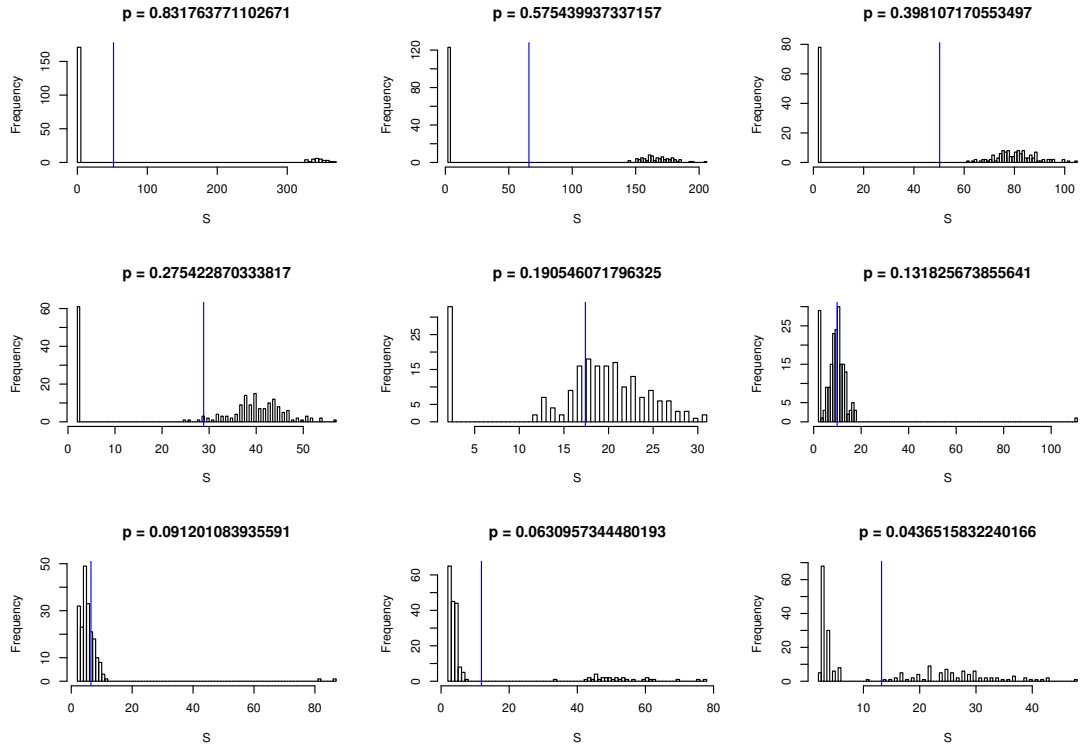

Figure 4.3 – Histograms of empirically measured sizes (e.g. number of nodes) of shortest path subgraphs between two nodes for different random graphs with 500 nodes. For each value of p we generated 200 graphs. The size of shortest path subgraph is denoted by S. The blue line represents the average S.

The rest of this chapter is organised as follows. In Section 4.1 we give the necessary definitions. Next, in Section 4.2 we study complete and quasi-complete graphs (i.e. graphs obtained from complete graphs by removing an link), in order to form an intuition about the results that we will show next: in Section 4.3 we consider dense random graphs and in Section 4.4 we study sparse random graphs with unbounded mean degree. In all cases (except complete graphs) we find a non-trivial multimodal distribution of the number of nodes which belong to a shortest path between two nodes and we explain this phenomenon. Finally, we summarise our main results and discuss possible future works.

4.1 Definitions

Here we are interested in the number of nodes or links that belong to the shortest paths between two nodes. We recall that in the general case there are more than

one shortest path between two nodes and therefore the number of links (or nodes) can be very different compared to the distance between them. We will use the following definitions:

Definition 4.1.1: Given a graph and two of its nodes u and v, let $SPS(u,v)$ be the subgraph of all shortest paths between u and v, i.e. $SPS(u,v)$ contains all links and nodes that belong to a shortest path between u and v. When u and v do not belong to the same connected component $SPS(u,v)$ does not exist.

Definition 4.1.2: Given a graph and two of its nodes u and v, let $S(u,v)$ be the number of nodes in $SPS(u,v)$. Similarly, $S_E(u,v)$ denotes the number of links in $SPS(u,v)$. When u and v do not belong to the same connected component $S(u,v)$ and $S_E(u,v)$ are equal to 0.

Definition 4.1.3: For a given graph, we denote by $\mathcal{D}(u,v)$ the distance between nodes u and v, i.e. the length of a shortest path from u to v. If there is no path between u and v, we define $\mathcal{D}(u,v) = \infty$. For brevity we often write $\overset{x}{\widetilde{uv}}$ instead of $\mathcal{D}(u,v) = x$.

Definition 4.1.4: A random graph $G(n,p)$ is a graph with n nodes such that each link belongs to the graph with probability p.

We fix two distinct nodes u and v from a given set V of n nodes, and consider all realisations of $G(n,p)$ over V as a probability space.

Definition 4.1.5: Let $f_d = \Pr\left[\mathcal{D}(u,v) = d\right]$ and $f_{>d} = \Pr\left[\mathcal{D}(u,v) > d\right]$.

Abusing the notation we will denote by S the random variable for the number of nodes in $SPS(u,v)$ when there is no ambiguity. Analogously, we use S_E for the random variable for the number of links in $SPS(u,v)$.

4.2 Complete and quasi-complete graphs

Let us consider complete and quasi-complete graphs in order to form an intuition about the distribution of S. In the case of the complete graph K_n the structure of $SPS(u,v)$ is trivial, because $SPS(u,v)$ contains only the nodes u, v and the link uv. Consider now a quasi-complete graph $K_n - ab$, i.e. a graph obtained from K_n by removing an link ab.

Proposition 4.2.1: *For any distinct nodes u and v from the quasi-complete graph $K_n - ab$ we have:*

$$S(u,v) = \begin{cases} n & \text{if } \{u,v\} = \{a,b\}\,, \\ 2 & \text{otherwise}\,. \end{cases}$$

Proof. Suppose $u = a$ and $v = b$ (or conversely). Nodes a and b are not adjacent, but they are connected by $n-2$ paths of length 2 (see for example Figure 4.4). The union of these paths contains all nodes of our graph, so $S(a, b) = n$. Otherwise, u and v are adjacent, and $S(u, v) = 2$. □

The size distribution of shortest path subgraphs of the quasi-complete graph contains two peaks: 2, n. The average size is equal to $2 + \frac{2(n-2)}{n(n-1)}$, but there is no subgraphs with this size.

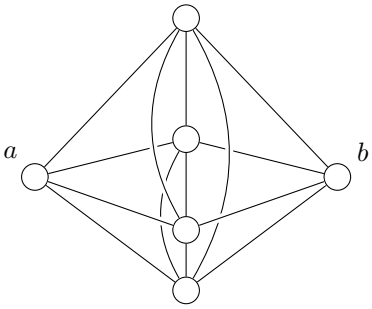

Figure 4.4 – $K_6 - ab$. There are 4 different shortest paths between nodes a and b, while there is only 1 shortest path between any other two nodes.

4.3 Dense random graphs (p is fixed, $n \to \infty$)

In this section we study dense random graphs, i.e. graphs with constant density p. First, we recall that such graphs almost surely have diameter 2. This allows us to consider only two cases: $\overset{1}{\widetilde{uv}}$ and $\overset{2}{\widetilde{uv}}$. We study these cases, and we show that the size distribution of $SPS(u, v)$ looks similar to the two-peak distribution from Proposition 4.2.1.

It is well known that almost all random graphs have diameter 2 (see for example [12]). We use the following formulation for this result:

Theorem 4.3.1 (Random graphs with constant density have diameter 2)**:** *For any given $p > 0$ we have $f_1 = p$, $f_2 \to 1 - p$, and $f_{>2} \to 0$ as $n \to \infty$.*

Proof. By definition of $G(n, p)$ we have $f_1 = p$ and $f_{>1} = 1 - p$. Next, the distance between two nonadjacent nodes u and v is larger than 2 if and only if they have no common neighbours. There are $n - 2$ nodes distinct from u and v, and each of them is not their common neighbour with probability $1 - p^2$. So, we have:

$$f_{>2} = f_{>1}(1 - p^2)^{n-2}\,.$$

Since $p > 0$ is fixed, $\lim_{n\to\infty} f_{>2} = 0$ and $\lim_{n\to\infty} f_2 = 1 - p$. □

Theorem 4.3.1 shows that in dense random graphs there is almost surely only two cases: $\overset{1}{\widetilde{uv}}$ and $\overset{2}{\widetilde{uv}}$. The structure of $SPS(u, v)$ in the former case is trivial, because $SPS(u, v)$ contains only the nodes u, v and the link uv.

In the rest of this section we study the latter case. Let Y denote the number of nodes that are directly connected to both u and v. Note that Y is equal to $S - 2$.

Lemma 4.3.2: *Y is a binomial random variable with parameters $n-2$ and success probability p^2:*

$$Y \sim B(n-2, p^2)\,.$$

Proof. The probability that any node c is directly connected to both u and v is equal to p^2. We have $n-2$ nodes which are independently susceptible to lie between u and v. □

Theorem 4.3.3: *When the distance between u and v is equal to 2, the probability function* $\Pr\left[Y = k | \overset{2}{\widetilde{uv}}\right]$ *is equal to*

$$\begin{cases} 0 & \text{if } k = 0\,, \\ \frac{\Pr[Y=k]}{1-(1-p^2)^{n-2}} & \text{if } k \geqslant 1\,. \end{cases}$$

Proof. From the definition of conditioned probability we have

$$\Pr\left[Y = k | \overset{2}{\widetilde{uv}}\right] = \frac{\Pr\left[Y = k \text{ and } \overset{2}{\widetilde{uv}}\right]}{\Pr\left[\overset{2}{\widetilde{uv}}\right]}\,.$$

Let A be an event "there is no link between u and v". Observe that

$$\overset{2}{\widetilde{uv}} \iff A \text{ and } Y \geqslant 1\,.$$

A is independent from Y and $\Pr[A] = 1 - p$, so

$$\Pr\left[Y = k | \overset{2}{\widetilde{uv}}\right] = \frac{(1-p)\Pr\left[Y = k \text{ and } Y \geqslant 1\right]}{(1-p)(1-(1-p^2)^{n-2})}\,.$$

Note that

$$\Pr\left[Y = k \text{ and } Y \geqslant 1\right] = \begin{cases} 0 & \text{if } k = 0\,, \\ \Pr[Y = k] & \text{if } k \geqslant 1\,. \end{cases}$$

The claimed formula easily follows. □

Corollary 4.3.4: *When the distance between u and v is equal to 2, we have the following expressions for the expectation and the variance of S:*

$$\mathbb{E}\left[S|\,\widetilde{uv}^{2}\right] = 2 + \frac{(n-2)p^2}{1-(1-p^2)^{n-2}}\,,$$
$$\mathrm{Var}\left[S|\,\widetilde{uv}^{2}\right] = \frac{(n-2)\,p^2\,(1-p^2+(n-2)p^2)}{1-(1-p^2)^{n-2}} - \left(\frac{(n-2)p^2}{1-(1-p^2)^{n-2}}\right)^2.$$

Proof. When the distance between u and v is equals to 2, $SPS(u,v)$ contains $2+Y$ nodes. Thus, we have $\mathbb{E}\left[S|\,\widetilde{uv}^{2}\right] = 2 + \mathbb{E}\left[Y|\,\widetilde{uv}^{2}\right]$. From Lemma 4.3.2 we know that Y is a binomial random variable with parameters $n-2$ and success probability p^2. From Theorem 4.3.3 we know the probability function $\Pr\left[Y=k|\,\widetilde{uv}^{2}\right]$. Next, we write

$$\begin{aligned}\mathbb{E}\left[Y|\,\widetilde{uv}^{2}\right] &= \sum_{k=0}^{\infty} k\Pr\left[Y=k|\,\widetilde{uv}^{2}\right]\\ &= \sum_{k=1}^{\infty}\frac{k\Pr[Y=k]}{1-(1-p^2)^{n-2}}\\ &= \frac{\mathbb{E}[Y]}{1-(1-p^2)^{n-2}}\,.\end{aligned}$$

Now, the claimed formula for the expectation can be easily obtained.

Concerning the variance, we have:

$$\begin{aligned}\mathrm{Var}\left[S|\,\widetilde{uv}^{2}\right] &= \mathrm{Var}\left[Y|\,\widetilde{uv}^{2}\right]\\ &= \mathbb{E}\left[Y^2|\,\widetilde{uv}^{2}\right] - \left(\mathbb{E}\left[Y^2|\,\widetilde{uv}^{2}\right]\right)^2.\end{aligned}$$

Note that

$$\begin{aligned}\mathbb{E}\left[Y^2|\,\widetilde{uv}^{2}\right] &= \sum_{k=1}^{\infty}\frac{k^2\Pr[Y=k]}{1-(1-p^2)^{n-2}}\\ &= \frac{\mathbb{E}[Y^2]}{1-(1-p^2)^{n-2}}\\ &= \frac{(n-2)\,p^2\,(1-p^2+(n-2)p^2)}{1-(1-p^2)^{n-2}}\,.\end{aligned}$$

And finally

$$\mathrm{Var}\left[Y|\overset{2}{\widetilde{uv}}\right] = \frac{(n-2)\,p^2\,(1-p^2+(n-2)p^2)}{1-(1-p^2)^{n-2}} - \left(\frac{(n-2)p^2}{1-(1-p^2)^{n-2}}\right)^2.$$

□

In order to illustrate Corollary 4.3.4 we performed some numerical simulations. Figure 4.5a shows the values of S in the case when n is fixed and $p \in [0,1]$. Different colours correspond to different distances between u and v. The black line represents $\mathbb{E}[S|\overset{2}{\widetilde{uv}}]$. From Chebyshev's inequality we know that at least 88% of the distribution's values are within 3 standard deviations of the mean, we use red lines in order to delimit these bounds. Green points correspond to the distance 2, and we see that Corollary 4.3.4 perfectly describes their behaviour when the density grows. Figure 4.5b shows what happens when p is fixed but n grows: (i) the probability that we have a pair of nodes at a distance greater than 2 tends to zero; (ii) the average size (repsented by black circles) of shortest path subgraph lies between $\mathbb{E}[S|\overset{2}{\widetilde{uv}}]$ and $\mathbb{E}[S|\overset{1}{\widetilde{uv}}]$.

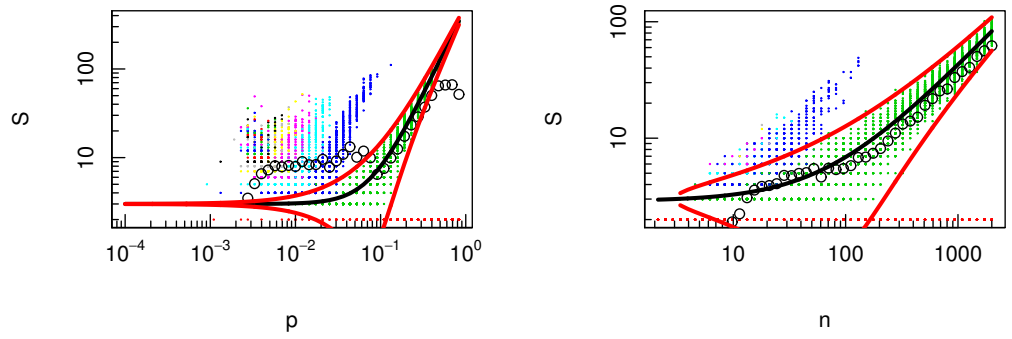

(a) $n = 500, p \in [0,1]$ (b) $n \in [2,2000], p = 0.2$

Figure 4.5 – Empirically measured values of S for different random graphs. For each value of p we generated 200 graphs. Each point corresponds to an observed value of S. Black circles represent the average size of SPS. Different colours correspond to different distances between u and v (red: $\overset{1}{\widetilde{uv}}$, green: $\overset{2}{\widetilde{uv}}$, blue: $\overset{3}{\widetilde{uv}}$, cyan: $\overset{4}{\widetilde{uv}}$, magenta: $\overset{5}{\widetilde{uv}}$). The black line represents $\mathbb{E}[S|\overset{2}{\widetilde{uv}}]$. At least 88% of the distribution's values are within 3 standard deviations of the mean, we use red lines in order to delimit these bounds. A log scale is used for both axes.

Corollary 4.3.5: *When the distance between u and v is equal to 2, we have:*

$$\Pr\left[S_E = 2k \middle| \overset{2}{\widetilde{uv}}\right] = \Pr\left[Y = k \middle| \overset{2}{\widetilde{uv}}\right].$$

Proof. It is sufficient to realise that for each node $c \in SPS(u,v), c \notin \{u,v\}$ there are exactly two distinct links in $SPS(u,v)$, i.e. (c,u) and (c,v). □

Finally, the probability mass function of S is a mixture of two functions: the first corresponds to the case $\overset{1}{\widetilde{uv}}$ and the second to $\overset{2}{\widetilde{uv}}$ (see Figure 4.6):

$$S = \begin{cases} 2 & \text{if } \overset{1}{\widetilde{uv}}, \\ 2+Y & \text{if } \overset{2}{\widetilde{uv}}, \text{ where } \Pr\left[Y = k \middle| \overset{2}{\widetilde{uv}}\right] = \begin{cases} 0 & \text{if } k = 0, \\ \frac{\Pr[Y=k]}{1-(1-p^2)^{n-2}} & \text{if } k \geqslant 1. \end{cases} \end{cases}.$$

The distribution of S has two local maxima, and the average size of $SPS(u,v)$ lies between these maxima. This finally explains why the values of $S(u,v)$ and $S_E(u,v)$ may be very different from the average.

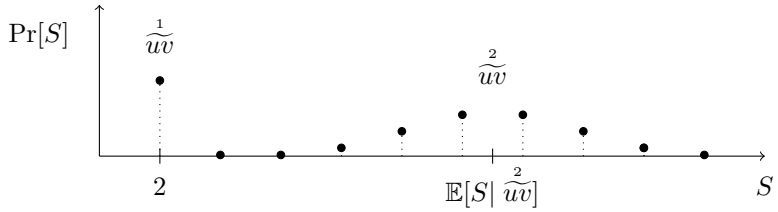

Figure 4.6 – Schematic representation of the probability mass function of S.

4.4 Sparse random graphs with unbounded mean degree ($p \to 0$ and $np \to \infty$ as $n \to \infty$)

We say that a random graph is sparse when its density p tends to zero as n goes to infinity. There are two classes of sparse graphs: (i) mean degree is bounded by some constant that does not depend on the size of the graph ($np \leqslant c$), (ii) mean degree is unbounded ($np \to \infty$). Here we study the second case.

As Figure 4.5a suggests, when p decreases, the number of peaks in the distribution of S grows. For example, a third peak (dark-blue points on the Figure 4.5a) appears when the probability that $\mathcal{D}(u,v) = 3$ becomes non-negligible, it happens when $p \approx 0.1318$, and Figure 4.3 shows a histogram for this case. Let us give an

intuitive explanation for the fact that there are several peaks in the size distribution of $SPS(u, v)$. We observe that $S(u, v) \geqslant \mathcal{D}(u, v) + 1$, so, when $\mathcal{D}(u, v)$ grows, $S(u, v)$ also grows. Intuitively, when our graphs are not similar to trees, $S(u, v)$ grows much faster than $\mathcal{D}(u, v)$:

$$\mathcal{D}(u, v) > \mathcal{D}(u', v') \Rightarrow S(u, v) \gg S(u', v') \,.$$

Therefore, each observed value for $\mathcal{D}(u, v)$ will correspond to a peak in the distribution of S.

In subsection 4.4.1 we give an approximation of the expected number of nodes in $SPS(u, v)$ in the case when we know the distance between u and v. In subsection 4.4.2 we give a classification of sparse graphs with unbounded mean degree, and we study the size distribution of $SPS(u, v)$ according to this classification.

4.4.1 Approximated expectation of the size

Recall that we denote by f_d (resp. $f_{>d}$) the probability that the distance between two nodes is equal to d (resp. greater than d). Authors in [56] showed that f_d can be approximated by a recurrent formula:

$$\begin{aligned} f_{>0} &= 1 - \frac{1}{n} \,, \\ f_{>d} &= \big(1 - \frac{1}{n}\big)(1-p)^{(1-f_{>d-1})n} \,, \\ f_d &= f_{>d-1} - f_{>d} \,. \end{aligned}$$

We refer interested reader to [56] for details about $f_{>d}$. Note however that the authors considered the case where the two chosen nodes are not necessarily distinct. Therefore, their definition of f_d is a bit different from ours, but asymptotically they coincide.

Note that the expectation of S, conditioned on the distance between u and v being equal to x, is equal to

$$2 + (n-2) * \Pr\left[c \in SPS(u, v) \middle| \overset{x}{\widetilde{uv}}\right] .$$

But it seems difficult to calculate this probability, so we present here only an approximation of the expected value.

Approximation 4.4.1: *When the distance between u and v is equal to x, we have the following approximation for the expectation of S.*

$$\mathbb{E}\left[S \middle| \overset{x}{\widetilde{uv}}\right] \approx x + 1 + (n - x - 1) \sum_{y=1}^{x-1} f_y f_{x-y} \,.$$

Idea. When $\overset{x}{\widetilde{uv}}$, we know that $S \geqslant x+1$. There are $n-x-1$ possible nodes which also can lie on a shortest path between u and v, so we have

$$\mathbb{E}\left[S \middle| \overset{x}{\widetilde{uv}}\right] \approx x+1+\Pr\left[c \in SPS(u,v) \middle| \overset{x}{\widetilde{uv}}\right](n-x-1)\,. \tag{4.1}$$

Note that a node c is on a shortest path between u and v if and only if $\mathcal{D}(u,v) = \mathcal{D}(u,c)+\mathcal{D}(c,v)$. Therefore we have:

$$\Pr\left[c \in SPS(u,v) \middle| \overset{x}{\widetilde{uv}}\right] = \frac{\sum_{y=1}^{x-1} \Pr\left[\overset{y}{\widetilde{uc}} \text{ and } \overset{x-y}{\widetilde{cv}} \text{ and } \overset{x}{\widetilde{uv}}\right]}{\Pr\left[\overset{x}{\widetilde{uv}}\right]}$$

Assuming that events $\overset{y}{\widetilde{uc}}$, $\overset{x-y}{\widetilde{cv}}$ and $\overset{x}{\widetilde{uv}}$ are mutually independent and identically distributed (actually it is not true, because there is triangular inequality that creates some dependencies), we approximate:

$$\Pr\left[\overset{y}{\widetilde{uc}} \text{ and } \overset{x-y}{\widetilde{cv}}\right] \approx \Pr\left[\overset{y}{\widetilde{uc}}\right]\Pr\left[\overset{x-y}{\widetilde{cv}}\right] \approx f_y f_{x-y}\,,$$

and finally

$$\Pr\left[c \in SPS(u,v) \middle| \overset{x}{\widetilde{uv}}\right] \approx \sum_{y=1}^{x-1} f_y f_{x-y}\,. \tag{4.2}$$

Using Blondel *et al.* relation and our formulæ (4.1) and (4.2), we are able to calculate this.

□

Figure 4.7 illustrates our approximation when $\mathcal{D}(u,v) \in \{3,4\}$. Clearly we see several peaks, i.e. typical values of S. Our approximation corresponds to the centres of these peaks. There are valleys between peaks. But these valleys vanish when p is very small, due to the variance of S.

Typically our approximation gives very good estimations compared to experimental data. However, when $\mathcal{D}(u,v)$ is very large (compared to the average distance), the result of approximation is slightly inadequate. This happens due to the following reasons: (i) Blondel *et al.* expression for f_d is not exact; (ii) we neglect the dependencies between some events.

4.4.2 Classification of sparse graphs according to the distance distribution

In appendix of [56] the following was shown:

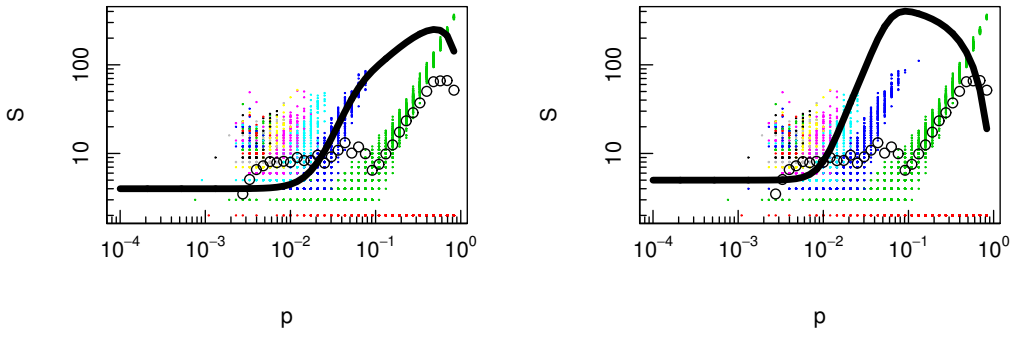

(a) $\overset{3}{\widetilde{uv}}$ (b) $\overset{4}{\widetilde{uv}}$

Figure 4.7 – Empirically measured values of S in random graph with 500 nodes. For each value of p we generated 200 graphs. Each point corresponds to an observed value of S. Black circles represent the average size of *SPS*. Different colours correspond to different distances between u and v (red: $\overset{1}{\widetilde{uv}}$, green: $\overset{2}{\widetilde{uv}}$, blue: $\overset{3}{\widetilde{uv}}$, cyan: $\overset{4}{\widetilde{uv}}$, magenta: $\overset{5}{\widetilde{uv}}$). Approximated $\mathbb{E}\left[S|\ \overset{x}{\widetilde{uv}}\right]$ is represented by black lines. A log scale is used for both axes.

Theorem 4.4.2: *For any given $d \geqslant 2$ and any given $\lambda \in (0, \infty)$, if $n^{d-1}p^d = \lambda$ we have:*

$$\begin{aligned}
\lim_{n\to\infty} f_{<d} &= 0\,,\\
\lim_{n\to\infty} f_{d} &= 1 - e^{-\lambda}\,,\\
\lim_{n\to\infty} f_{d+1} &= e^{-\lambda}\,,\\
\lim_{n\to\infty} f_{>d+1} &= 0\,.
\end{aligned}$$

Informally, when $p = \sqrt[d]{\frac{\lambda}{n^{d-1}}}$ and n is sufficiently large, there are only two possibilities: $\overset{d}{\widetilde{uv}}$ and $\overset{d+1}{\widetilde{uv}}$. It can be shown that $\lim_{n\to\infty} \mathbb{E}\big[S|\ \overset{d}{\widetilde{uv}}\ \big]$ exists and depends only on λ, while $\mathbb{E}\big[S|\ \overset{d+1}{\widetilde{uv}}\ \big]$ grows with n. Correspondingly, we have seen that for random graphs with constant density the following is true: $\mathbb{E}[S|\ \overset{1}{\widetilde{uv}}] = 2$ and $\mathbb{E}\left[S|\ \overset{2}{\widetilde{uv}}\right] \approx np^2$ (see Proposition 4.3.3 and Corollary 4.3.4). Therefore, the

distributions of S in the case of dense and sparse (with unbounded mean degree) random graphs are quite similar.

We performed numerical simulations using two families of sparse random graphs: (a) $np^2 = 1$ and (b) $n^2p^3 = 1$. Figure 4.8 presents the results. For each value of p we generated 200 random graphs $G(n, p)$ and we measured S (each point in Figure 4.8 corresponds to a measured value of S, different colours correspond to different distances between nodes). We see that $E\big[S|\, \overset{d}{\widetilde{uv}}\,\big]$ stabilises around some value, but $E\big[S|\, \overset{d+1}{\widetilde{uv}}\,\big]$ grows unboundently. The distance between these two typical values of S grows also with n.

We see that the average S (represented by black circles) almost always lies in the valley between $E\big[S|\, \overset{d}{\widetilde{uv}}\,\big]$ and $E\big[S|\, \overset{d+1}{\widetilde{uv}}\,\big]$. We see some shortest path subgraphs with sizes that are close to the average, but it happens very rarely. In general, sizes are concentrated around two peaks.

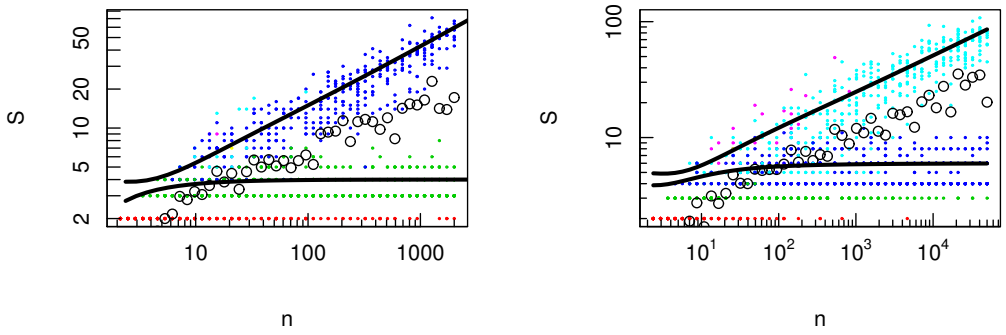

(a) $np^2 = 1$ (b) $n^2p^3 = 1$

Figure 4.8 – Empirically measured values of S. For each value of p we generated 30 graphs. Each point corresponds to an observed value of S. Black circles represent the average size of SPS. Different colours correspond to different distances between u and v. Our approximation for $\mathbb{E}\big[S|\, \overset{x}{\widetilde{uv}}\,\big]$ is represented by black lines (on the left: $x \in \{2, 3\}$, on the right: $x \in \{3, 4\}$). A log scale is used for both axes.

The size of the shortest path subgraph between the source and destinations plays an important role in the observed network dynamics, as we know from the previous chapter. In this chapter we have made the first step towards a rigorous characterisation of the size of shortest path subgraphs.

Particularly, we studied the size of the shortest path subgraph between two nodes. We denote by S the number of nodes in that subgraph. Our study results in a characterisation of S for dense random graphs and sparse random graphs with unbounded mean degree.

The probability mass function of S has several local maxima (peaks). Each peak corresponds to a possible distance between u and v. Between such peaks we have valleys of "improbable" sizes of *SPS*, in other words the distribution of S is multimodal.

For dense random graphs we have asymptotically exact results, and for sparse random graphs with unbounded mean degree we have an approximated representation of the expected number of nodes in $SPS(u, v)$.

Better approximations (or even exact distributions) are parts of a future work. One may also investigate the important class of random graphs with constant mean degree, using the methodology described in [62]. Another direction consists in studying real-world networks or other models of random graphs (e.g. power-law graphs). In this chapter, we considered only the shortest paths subgraph between two nodes (one monitor and one destination). The case of several destinations should also be considered in the future.

Moreover, the notion of $SPS(u, v)$ can be considered as a similarity measure between nodes u and v. In the domain of community detection methods there are a lot of related notions, for example: *connection subgraph* [43] and *proximity graphs* [53]. Thus, it may be very useful to develop a general theory about *SPS*-like objects as an analogue of mathematical theory of metric spaces, because the notion of *SPS* generalises, in some sense, the notion of a distance.

Chapter 5

Real and observed dynamics

In this chapter we study how the frequency of measurements affects the observed network dynamics. We are interested in the underlying processes which cause the observed dynamics. We introduce a non-classical method of stochastic process parameter estimation and we apply this method to real-world and modelled measurements in order to attempt to characterise the rate of the evolution of the topology. We also show that dynamic of the network is nonuniform, in the sense that different parts of the network can have different rates of evolution.

This chapter is organised as follows. In Section 5.1 we will show that the slope of the curve of the number of observed links depends intrinsically on the measurement frequency, and that it is very difficult to actually observe the real rate of the topology's evolution. Next, in Section 5.2 we give a skeleton of a theory of a non-classical sampling where the goal is to infer the real underlying process that causes such observations. We applied this theory to real-world and modelled measurements in order to approximate the true frequency of modifications of ego-centred views. Finally, in Section 5.3 we will study whether the dynamic of the network is uniform or not, in other words whether different parts of the network have different dynamical properties.

5.1 Impact of the measurement frequency

In order to show how the measurement frequency affects the observed behaviour we analyse both the real data and the behaviour of the model of the topology dynamics that we introduced in Chapter 2. In this section we focus on the slope of the curve of the number of observed links, which was introduced and discussed in Chapter 3. Intuitively, the observed slope depends both on the rate of routing changes and load balancing: the more changes happen, the more new links will be observed over time; and the more routes exist between two nodes, the more consecutive measurements will observe previously existing but unobseerved routes. Figure 5.1 illustrates the following: *when we decrease the frequency of measurements the observed slope also decreases.*

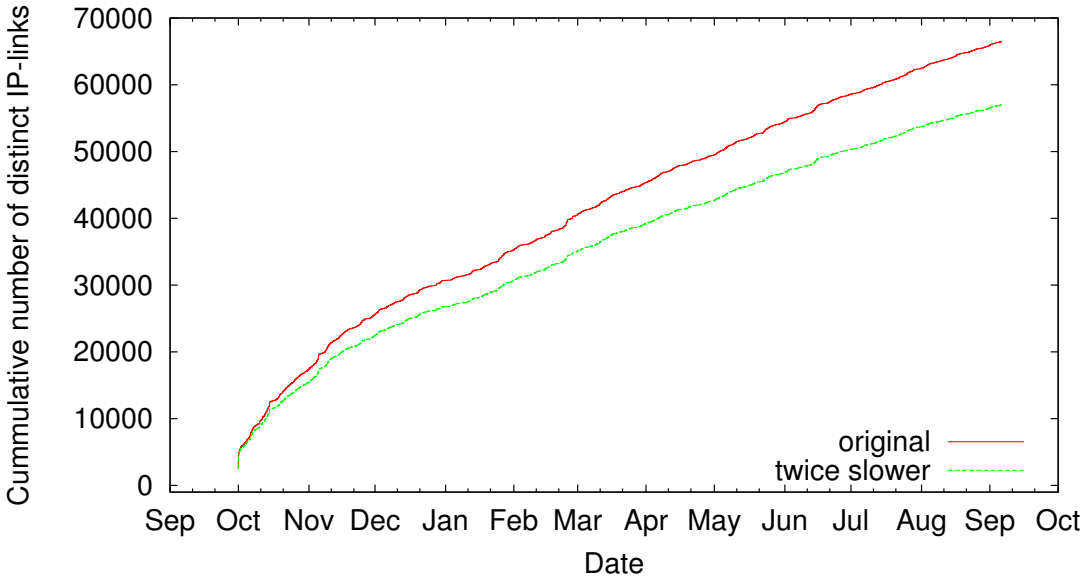

Figure 5.1 – Number of links observed since measurement beginning (`ovh` dataset). We observe a difference between the slopes for the original measurements (red) and measurements performed twice more slowly (green).

If the measurement frequency is low, we will fail to observe some links, because they will have disappeared before the corresponding route is explored. On the contrary, if measurements are performed fast enough, all routes from the monitor to the destinations should be observed. More formally, let Δ be the time interval between two consecutive rounds and let α be the corresponding slope. We expect that there exists a true slope α_m and corresponding ideal delay Δ_m such that for all $\Delta \leqslant \Delta_m$ we should observe the same slope:

$$\Delta \leqslant \Delta_m \Rightarrow \alpha = \alpha_m \,.$$

Thus, Δ_m is the optimal interval between measurements.

Figure 5.2a illustrates this.

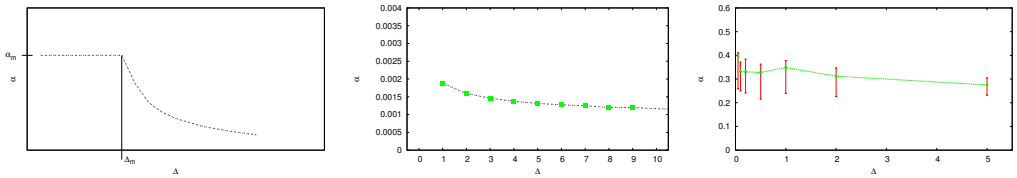

(a) Expected behaviour (b) `ovh` dataset (c) Model (Erdős–Rényi graph, $n = 10^5$, $m = 2{\cdot}10^5$, $d = 300$).

Figure 5.2 – Impact of the measurement frequency. Each figure represents α (the slope of the plot of the number of observed links as a function of time) as a function of the time interval between consecutive rounds Δ. For the model simulations, we present results for 100 experiments. Each point is the average of the slope, and the errorbars represent the 25- and 75-percentiles.

In order to test whether we are measuring with high enough frequency or not, we compute the slopes for different values of Δ. Ideally, for performing a rigorous analysis with real data, one should perform several measurements at different frequencies. For these measurements to be comparable, they should be performed from the same monitor towards the same destination set, and at the same time. As this is not feasible in practice, we use real measurements, performed with $\Delta_{original}$, to simulate other measurements with different, lower, frequencies. We do so by taking into account only every n-th measurement round, so that the simulated interval Δ will be equal to $n \cdot \Delta_{original}$. We are then able to compute the corresponding slope α_n. We use the `ovh` dataset that consists of measurements performed at a high frequency ($\Delta_{original}$ is equal to 1m 25s) and for a long time, so that we are able to simulate measurements with a wide frequency interval.

The model has no such parameter as Δ. However we can assume that topology changes, i.e. swaps, happen at a constant rate on average. Let $\Delta_{original}$ stand for the elapsed period of time between two consecutive rounds if only one swap is performed at each round. We can assume that $\Delta_{original}$ is equal to 1. Then we can simulate a lower frequency in the model by "performing n swaps at each round, which will represent a period that is equal to $n \cdot \Delta_{original} = n$. Consequently, we can simulate a higher frequency by performing only one swap every n' rounds, which will represent a period that is equal to $\Delta_{original}/n' = 1/n'$.

Figure 5.2 shows the observed results. Concerning simulations, we found that the variability between different experiments performed with the same parameters is quite high (the plot shows the average observed values, as well as the 25- and 75-percentiles). This makes it difficult to draw a rigorous conclusion, but the observations are compatible with the presence of a plateau for $\Delta \leqslant 1$. The median of the observed values, not presented here, also strongly suggests the presence of a plateau. Then for $\Delta \geqslant 1$, the slope decreases with Δ, showing that for these parameters, a frequency smaller than one round each swap is too slow to observe all changes.

In the case of empirical data, we again observe that, the longer the interval between measurements, the smaller the slope α. This confirms our expectations in the case of $\Delta \geqslant \Delta_m$. However, in this case, there is no plateau at the beginning of the curve. It might be the case that the plateau begins at $\Delta = 1$ but we have no data to confirm or refute this. Therefore, we don't know whether $\Delta = 1$ is optimal or not.

5.2 Inferring the evolution speed

In the previous section we demonstrated that the observed dynamics of ego-centred views depends intrinsically on the measurement frequency and that it is very difficult to reveal the actual speed of the Internet topology dynamics.

The goal of this section is to find a method that will allow us to estimate the actual speed of the Internet topology dynamics. However, it is not easy to directly study the real-world measurements. We need to introduce some abstractions: we should forget some details in order to obtain clear comprehension of the underlying processes. After this, we will put all the removed details back. We will use several models, from very detailed ones to very simple and analytically solvable ones. Figure 5.3 represents a hierarchy of processes that can be used to describe the dynamics of ego-centred views of the Internet at different levels of abstraction. Below we give give a detailed description of all levels shown in Figure 5.3.

***spt*-sequence.** Inspired by the model introduced in Section 2, we denote by *spt*, i.e. "shortest path tree", a result of one ego-centred measurement that corresponds to the result of one round in the model. Typically we have several different *spt*-s between the monitor and any destination.

***spt*-sequence $\xmapsto{\text{forget load-balancing}}$ *sps*-sequence.** In order to simplify our problem, we can decide to study not a sequence of shortest path trees, but rather a sequence of shortest path subgraphs. Recall that a shortest path subgraph, denoted

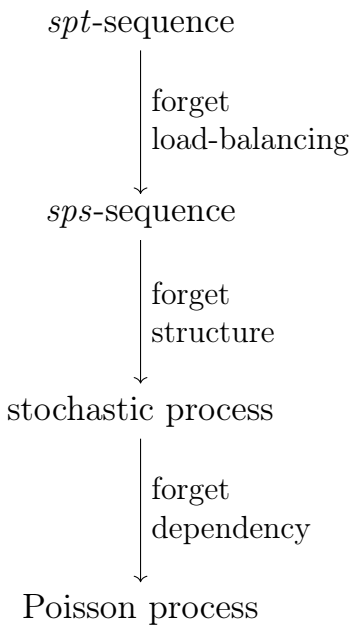

Figure 5.3 – Hierarchy of dynamical processes: from real-world to analytically solvable.

by *sps*, contains all shortest paths between two nodes. One *sps* is formed by the union of several *spt*-s.

We cannot directly study the *sps*-sequence in the case of real-world measurements, because in the real-world we have only a sequence of *spt*-s. Currently we don't know how to infer the desired *sps*-sequence from the real-world *spt*-sequence, and it seems to be a very difficult task. In the model, however, we can easily obtain a sequence of *sps*-s.

***sps*-sequence $\xmapsto{\text{forget structure}}$ stochastic process.** Having a *sps*-sequence, we can forget about its structure, i.e. we can consider each subgraph in the sequence as one solid object. The object may change at some times. We cannot observe changes directly as they happen, instead we have to perform measurements every Δ units of time.

Stochastic process $\xmapsto{\text{forget dependency}}$ Poisson process If now we assume that changes occur due to some Poisson process parameterised by λ, we can estimate the λ from the sequence of object's states measured every Δ units of time.

Subsection 5.2.1 describes the case of partially observed Poisson process in details. Next, in Subsection 5.2.2 we show whether changes in *sps*-sequence can

be modelled by Poisson process. And finally, in Subsection 5.2.3 we come back to *spt*-sequences.

5.2.1 Poisson process

The goal of this section is to introduce a method of stochastic process parameter estimation from partial observations. In order to keep this introduction as general as possible, we forget, for the moment, about dynamic networks, instead we consider just an generic object.

The object may change at some times, but it never reverts to a previous state. We also suppose that changes (or modifications) occur due to some Poisson process parameterised by the mean number of changes per unit of time. We denote this mean number by λ. We cannot observe changes of the object directly, but instead we should rely on periodic observations: we can then detect whether the object has changed from one observation to the next, but we cannot know how many changes have occurred.

Let $\langle X \rangle$ be the realisation of a Poisson process. We will call this process *underlying process*. Let $\langle Y \rangle$ be the sequence of our observations. Figure 5.4 schematically shows the differences between the underlying and observed processes.

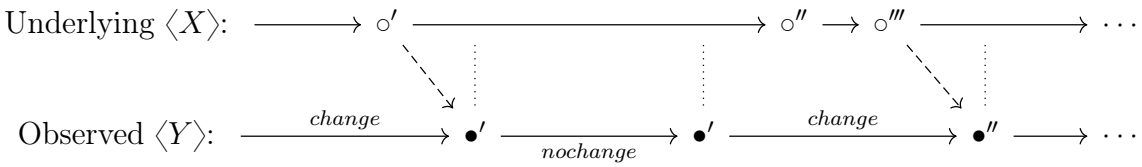

Figure 5.4 – Schematic representation of underlying and observed process. The circle represents our object. At the time of a change we draw a new circle. We use primes in order to show that the object was changed. Discs represent our observations.

Every Δ units of time we observe the object. If we compare two successive measurements there are only two possibilities:

(no change) $\bullet \to \bullet$;

(a change) $\bullet \to \bullet'$.

When there is no change ($\bullet \to \bullet$), we know that the object has not been modified between these two measurements. In the case of ($\bullet \to \bullet'$) we know only that the object has been modified at least once, but we *don't know the exact number of changes*. Figure 5.5 is a graphical form of our main question: How can we infer the most likely λ from a sequence of measurements?

In the rest of this subsection we propose an estimator $\hat{\lambda}$ and calculate its bias.

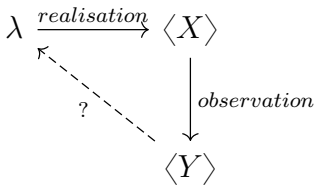

Figure 5.5 – Inference diagram. How can we infer λ of $\langle X \rangle$ using only a partial information represented by a sequence of observations $\langle Y \rangle$?

Estimator

Consider a sequence of $N+1$ observations $\langle Y \rangle$: $\bullet_0, \bullet_\Delta, \bullet_{2\Delta}, \ldots, \bullet_{N\Delta}$. Between any successive pair of observations there is an interval of Δ units of time. Therefore, we have a sequence of N intervals. Denote by W the number of intervals without changes ($\bullet \to \bullet$). We first present our estimator and show why this estimator is relevant.

Estimator

$$\hat{\lambda} = -\Delta^{-1} \log \frac{W}{N} \,,$$

where:

Δ – *size of the time interval between observations*;
W – *number of intervals without changes*;
N – *total number of intervals*.

For any Poisson process we have the following expression for the probability of exactly k changes happening in an interval of Δ units of time:

$$\Pr[k] = \frac{e^{-\lambda\Delta}(\lambda\Delta)^k}{k!} \qquad k = 0, 1, \ldots,$$

Note that due to the memorylessness of Poisson processes this probability does not depend of what happens in other Δ intervals. We write:

$$\Pr[\bullet \to \bullet] = \Pr[0] = e^{-\lambda\Delta} \,. \tag{5.1}$$

We have W observed intervals without changes, therefore we can approximate:

$$\Pr[\bullet \to \bullet] \approx \frac{W}{N} \,. \tag{5.2}$$

Combining (5.1) and (5.2) we obtain the estimator:

$$\begin{aligned} e^{-\lambda\Delta} &\approx \frac{W}{N}\,, \\ \lambda\Delta &\approx -\log\frac{W}{N}\,, \\ \lambda &\approx -\Delta^{-1}\log\frac{W}{N}\,. \end{aligned}$$

Bias of the estimator

Below we present the approximated bias of our estimator $\hat{\lambda}$ and a method by which this approximation was obtained.

Bias

$$\mathbb{E}\left[\hat{\lambda}-\lambda\right] \approx \frac{1}{2N\Delta}\left(e^{\lambda\Delta}-1\right)$$

In order to calculate the bias of $\hat{\lambda}$ we proceed as follows:

$$\mathbb{E}\left[\hat{\lambda}\right] = \mathbb{E}\left[-\Delta^{-1}\log W'\right] = -\Delta^{-1}\mathbb{E}\left[\log W'\right]\,, \text{with } W' = \frac{W}{N}\,.$$

Using a second order Taylor expansion [51] we approximate:

$$\mathbb{E}\big[\log W'\big] \approx \log\mathbb{E}[W'] - \frac{\mathrm{Var}[W']}{2\big(\mathbb{E}[W']\big)^2}\,. \tag{5.3}$$

It is easy to see that W is distributed binomially:

$$W \sim B(N, e^{-\lambda\Delta})\,,$$

where N is the number of trials, and $e^{-\lambda\Delta}$ is the probability of success. Now, we are able to calculate the expected value of W':

$$\mathbb{E}[W'] = \frac{\mathbb{E}[W]}{N} = \frac{Ne^{-\lambda\Delta}}{N} = e^{-\lambda\Delta}\,,$$

and the variance of W':

$$\mathrm{Var}[W'] = \frac{\mathrm{Var}[W]}{N^2} = \frac{Ne^{-\lambda\Delta}(1-e^{-\lambda\Delta})}{N^2} = \frac{e^{-\lambda\Delta}(1-e^{-\lambda\Delta})}{N}\,.$$

Now, we return to (5.3):

$$\begin{aligned}\mathbb{E}\big[\log W'\big] &\approx \log\left(e^{-\lambda\Delta}\right) - \frac{1}{N}\frac{e^{-\lambda\Delta}(1-e^{-\lambda\Delta})}{2e^{-2\lambda\Delta}} \\ &\approx -\lambda\Delta - \frac{1}{2N}\left(e^{\lambda\Delta}-1\right)\,.\end{aligned}$$

And finally:

$$\mathbb{E}\left[\hat{\lambda}\right] \approx \lambda + \frac{1}{2N\Delta}\left(e^{\lambda\Delta}-1\right)\,,$$

$$\mathbb{E}\left[\hat{\lambda}-\lambda\right] \approx \frac{1}{2N\Delta}\left(e^{\lambda\Delta}-1\right)\,.$$

The estimator $\hat{\lambda}$ is asymptotically unbiased in the following sense:

$$\lim_{N\to\infty}\mathbb{E}\left[\hat{\lambda}\right] = \lambda\,.$$

Experimental results

In order to illustrate how our estimator works, we perform some experiments in the following way:

1. We generate a time-sequence of poissonian events with parameter λ.
2. Every Δ-units of time we perform a measurement and we say that the object changes when at least one event happened since the previous measurement.

The results are presented in Figure 5.6a. Note that $\hat{\lambda}$ remains almost constant when we increase Δ.

5.2.2 *sps*-process

Let us consider *sps*-process obtained from our model with an Erdős–Rényi underlying graph with $n = 100000$ nodes and $m = 200000$ links). We use 1 source and 500 destinations; we perform one swap at each round. Our goal here is to check whether the *sps*-process is poissonian or not. Suppose that changes in *sps*-process occur by some Poisson process. What happens if we try to estimate the rate of that process?

On Figure 5.6b we plotted the estimated λ as a function of Δ for several *sps*-processes obtained from our model with an Erdős–Rényi underlying graph. It is clear that the estimator $\hat{\lambda}$ is very close to the λ of the underlying process when $\Delta \approx 1$, because in this case we observe all changes of the *sps*. We notice that the estimator remains almost constant when we increase Δ. We have already seen

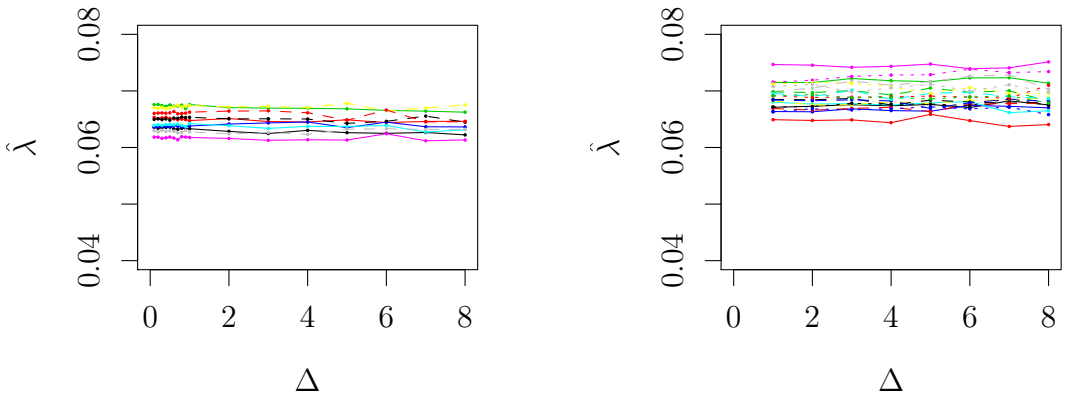

(a) Poisson process

(b) Model with an Erdős–Rényi underlying graph, *sps* is measured at each round.

Figure 5.6 – Estimated λ as a function of Δ. For Figure (a) we simulated 20 Poisson processes with $\lambda = 0.065$. For Figure (b) We used 20 realisations of the model, and we measure *sps* at each round. (Erdős–Rényi underlying graphs, $n = 100000, m = 200000, d = 500$, 1 swap per unit of time). Each process was observed using different delays Δ between measurements. Different lines corresponds to different realisations of the model (or the process).

a similar behaviour in the case of Poisson process on Figure 5.6a. This gives an intuition that the *sps*-process is indeed a Poisson one.

The main property of a Poisson process is its memorylessness. Moreover, the only memoryless continuous stochastic process is Poissonian. In the rest of this subsection we consider $\Delta = 3$ and we use another method to confirm whether the *sps*-process is likely to be memoryless or not[1]. In order to do this we use a test for serial independence [90]. One also can test whether the distribution of delays between changes follows an exponential law. In the case of Poisson process both these tests should give positive results. The application of the second test is a bit complicated in the case of the model, because the model is not completely continuous. We describe here only the first test.

[1]Similar results can be obtained for any other value of Δ.

Serial independence

We introduce a transformation of the sequence of object's states into a binary string. Consider Figure 5.7: the top level represents a sequence of measurements (different colours and primes mean different object's states), and the middle level is a result of our transformation. We have a 0 when two consecutive observations are the same, and 1 otherwise.

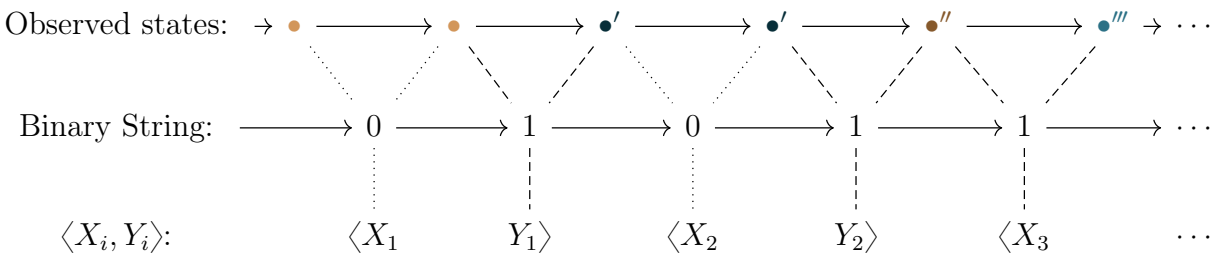

Figure 5.7 – Two transformations: (i) from the sequence of object's states to a binary string; (ii) from the binary string to a sequence of pairs $\langle X_i, Y_i \rangle$.

Next, from our binary string we form a sequence of pairs $\langle X_i, Y_i \rangle$ as illustrated on the bottom level of Fig. 5.7. Suppose that our process is memoryless, this should imply that $\langle X_i, Y_i \rangle$ is a sequence of independent observations of independent random variables X and Y. In order to test this we use the classical χ^2-test of independence [84, p. 36]. The obtained p-values along with expected and actual counts of patterns are presented in Tab. 5.1. In both cases p-values are good.

Finally, we conclude that our *sps*-process raised from the model is not far from memoryless and can be modelled by a Poisson process in the case when the underlying graph is random.

Pattern	**00**	**01**	**10**	**11**
Observed	2438	512	542	127
Expected	*2429*	*521*	*551*	*118*
χ^2	$= 0.88$,			
p-value	$= 0.35$			

(a) Poisson process

Pattern	**00**	**01**	**10**	**11**
Observed	2232	513	481	107
Expected	*2234*	*511*	*479*	*109*
χ^2	$= 0.5$,			
p-value	$= 0.82$			

(b) Model with Erdős–Rényi underlying graph

Table 5.1 – Count of patterns and p-values for χ^2-test of independence.

5.2.3 *spt*-process

Because of load-balancing, two consecutive *spt*-s can be different even when the underlying *sps*did not change. Our goal here is to estimate the λ of the underlying *sps*-process using only a sequence of *spt*-s. It is no longer sufficient to compare only two consecutive measurements as we did in Subsection 5.2.2. In Subsection 5.2.2, we considered each measurement of *sps* as an object and we applied the estimation method from 5.2.1. But now, in order to reduce the effect of load-balancing we will use "cumulative measurements": our object will be equal to the union of all observations performed since measurement beginning. These measurements corresponds to the curves of the number of nodes (or links) observed since measurement beginning (see Chapter 3).

Note that usual and cumulative measurements of *sps*-process have almost identical rates of evolution, because *sps* practically never reverts to the previous state.

Cumulative *spt*-sequences are very different. In particular our test of memorylessness fails. All together, this means that changes in cumulative *spt*-sequence do not occur due to a Poisson process. The estimator $\hat{\lambda}$, being applied, may give strange results. Indeed, in Figure 5.8a we see that $\hat{\lambda}$ is not constant when we increase Δ. However, we observe two facts:

1. $\hat{\lambda}$ is always larger for *spt*-sequences than for the case of *sps*-sequences;
2. $\hat{\lambda}$ calculated using *spt*-sequences decreases as Δ increases, while the $\hat{\lambda}$ corresponding to *sps*-sequences is constant.

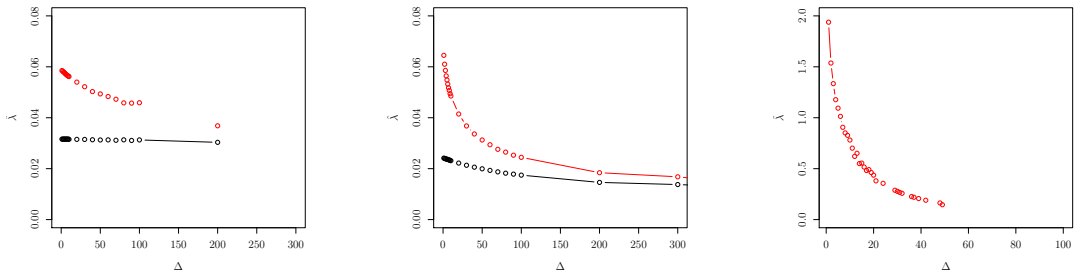

(a) Model (Erdős–Rényi graph, $n = 10^5, m = 2 \cdot 10^5, d = 300$)

(b) Model (Power-law graph, $n = 10^5, \gamma = 2.3, d = 300$)

(c) `woolthorpe` dataset

Figure 5.8 – Estimated values of λ calculated using *spt*- (red) and *sps*- (black) sequences.

These facts suggest that $\hat{\lambda}$ calculated for *spt*-sequence may approach the underlying rate of *sps*-process as the delay between measurements tends to infinity, or mathematically speaking:

$$\lim_{\Delta\to\infty} \left|\hat{\lambda}_{sps} - \hat{\lambda}_{cspt}\right| = 0\,,$$

where $\hat{\lambda}_{sps}$ denotes the estimated λ from *sps*-sequence, and $\hat{\lambda}_{cspt}$ is equal to estimated λ for the case of cumulative *spt*-measurements.

Actually, we can prove this limit, if we suppose there are $\epsilon > 0$ and $M < \infty$ such that $\epsilon \leqslant \frac{W_{cspt}}{W_{sps}} \leqslant M$ independently of Δ and N[2]:

$$\begin{aligned}
&\lim_{\Delta\to\infty} \left|\hat{\lambda}_{sps} - \hat{\lambda}_{cspt}\right| = \\
&\lim_{\Delta\to\infty} \Big| -\frac{1}{\Delta}\log\frac{W_{sps}}{N} + \frac{1}{\Delta}\log\frac{W_{cspt}}{N}\Big| = \\
&\lim_{\Delta\to\infty} \frac{1}{\Delta}\Big|\log\frac{W_{cspt}}{N} - \log\frac{W_{sps}}{N}\Big| = \\
&\lim_{\Delta\to\infty} \frac{1}{\Delta}\Big|\log\frac{W_{cspt}}{W_{sps}}\Big| = 0\,.
\end{aligned}$$

The hypothesis $\epsilon \leqslant \frac{W_{cspt}}{W_{sps}} \leqslant M$ simply means that numbers of observed changes in cumulative *spt*-sequence and usual *sps*-sequence differ only by a finite factor. It seems to be reasonable, because one *sps* typically contains a finite number of *spt*-s.

These results can be represented in the form of the following diagram:

$$\begin{array}{ccc}
\hat{\lambda}_{sps} & \xrightarrow{\;N\to\infty\;} & \lambda_{sps} \\
\| & & \\
\hat{\lambda}_{csps} & \xleftarrow[\;\Delta\to\infty\;]{} & \hat{\lambda}_{cspt}
\end{array}$$

As Figure 5.8a suggest this diagram is valid when we use Erdős–Rényi graph as the underlying network structure. Our estimator stops working when Δ becomes too large, because we start to observe new links (nodes) at every measurement. In this situation the number of intervals without changes W is equal to 0, and our estimator $\hat{\lambda}$ is equal to infinity. The bias of our estimator is approximately equal to $\frac{1}{2N\Delta}\left(e^{\lambda\Delta} - 1\right)$. Thus, when we increase the delay between measurements, we should increase *exponentially* the total number of measurements in order to insure that our estimator gives satisfactory results.

[2]Recall, by N we denote the number of measurements, and W is equal to the number of intervals without changes (see also Subsection 5.2.1).

Consider now the model with underlying power-law graphs (see Fig. 5.8b). Contrary to the case of Poisson processes (see Fig. 5.6), we see that the estimated value of λ for *sps*-sequence is not stable. Thus, we conclude that *sps*-process is not poissonian in the case of power-law topology. This can be further confirmed by the serial independence test discussed above.

In the case of real-world measurements we cannot observe any *sps*-es, and we cannot check whether the *sps*-sequence is memoryless or not. We can however apply our estimator to the *spt*-sequence. The result for the `woolthorpe` dataset is presented in Figure 5.8c. The estimated value of λ decreases as Δ grows, but we cannot say to what value it converges.

5.3 Nonuniform dynamics

In Section 5.1 we have introduced a methodology for testing whether measuremens are performed fast enough to observe all changes that happen in the underlying topology. Unfortunately it seems that it is impossible to obtain this in real measurements. However, we can characterise these unobserved changes. More specifically, we will answer the question: "Which nodes do we fail to observe when performing our nonideal measurements?" We will use the following notation:

o – the number of measurements that contains an arbitrary node;
Δ – interval between measurements;
t_{first} – the first time when this node was observed;
t_{last} – the last time when this node was observed.

Here we consider only the nodes as the parts the topology, but without loss of generality we can speak about links, paths of length 2, stars, etc.

We take an arbitrary node and we denote by δ the lifetime of this node, i.e. a period in which the node can be observed in principle. During the lifetime of the node we perform about $n = \frac{\delta}{\Delta}$ measurements. In practice, because of load-balancing, we do not observe this nodet at every measurement.

Using our sequence of measurements, we approximate the lifetime δ and the probability p of being observed in one measurement for a node:

$$\delta = t_{\text{last}} - t_{\text{first}}, \; p = \frac{o}{n}. \tag{5.4}$$

Assume for simplicity that the probability p of being observed in one measurement is constant over the lifetime of the node. Now we can write the probability that a particular node was never observed:

$$\Pr_{\text{never obs.}} = (1-p)^n = (1-p)^{\frac{\delta}{\Delta}}.$$

From this we conclude that *unobserved parts of topology have a short lifetime or have a very small chance of being observed.*

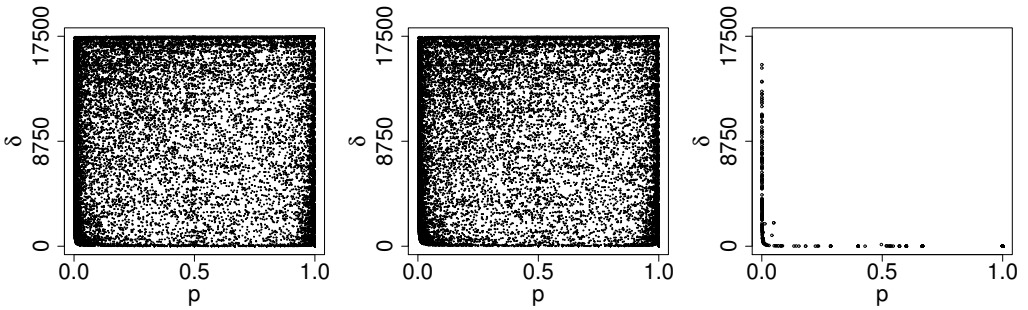

(a) Original measurements (b) Two times slower (c) "Original" minus "two times slower measurements"

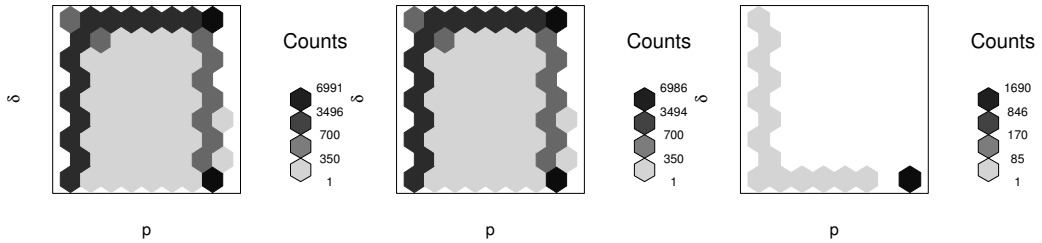

Hexbin version of (a) Hexbin version of (b) Hexbin version of (c)

Figure 5.9 – Probability p and lifetime δ of every observed node from `woolthorpe` dataset. Bottom figures are hexbin [15] versions of figures (a), (b) and (c). The brightness of each hexagon is proportional to the number of points in the corresponding region: "dark" means "there are many points", "white" means "there is almost no points here".

In order to justify this claim using real measurements, we use the `woolthorpe` dataset and proceed as follows:

(a) For each observed node we estimate p and δ using formula (5.4). Figure 5.9a presents δ vs. p for all nodes from `woolthorpe` dataset. As we can see there are a lot of points in some regions. This complicates the visual analysis of the figure. In order to clarify this, we present a hexbin version [15] of the same figure.

(b) Next, we simulate a frequency twice slower by considering only every second measurement. We estimate p and δ for each node observed by every second measurement, and we plot the results in Figure 5.9b.

(c) Finally, we "subtract' (b) from (a), in other words we plot on Figure 5.9c only the points presented in Fig. 5.9a but not in Fig. 5.9b. This allows us to see which nodes have not been observed when we perform measurements twice less frequently. Indeed, on Figure 5.9c we see that all unobserved nodes are concentrated near the axes. Thus, we conclude that nodes that are not observed have a short lifetime or have a very small chance of being observed. Also, at hexbin version of 5.9c we see that there are many nodes with $p = 1$, and these nodes were observed only once.

We studied the impact of the measurement frequency on the observed network dynamics. We showed that it is very difficult to quantify the actual speed of the internet topology evolution. Indeed, the *observed* dynamics depends intrinsically on the measurement frequency. We have introduced a methodology for testing whether our measurements are performed fast enough to observe all changes that happen in the underlying topology. Unfortunately it seems that it is impossible to obtain a complete picture of the changes in the case of real-world network (even when we perform measurements with a very high frequency).

When the changes occur according to some Poisson processes, we can very accurately estimate the mean number of changes per unit of time. With sufficient number of observations, we can guarantee any accuracy that we want, regardless of the delay between measurements. We showed that the changes in the sequence of shortest path subgraphs in our model with underlying Erdős–Rényi graph can be modelled by a Poisson processes. We also showed that the dynamic power-law graph cannot be properly modelled by a Poisson process. Future works should continue to study this case.

However, we showed that changes that are not observed are not very important in the following sense: network's parts that are not observed have a short lifetime or they have a very small chance of being observed.

Conclusion

In this work we considered periodic measurements of ego-centred views of the Internet topology and studied their dynamics. We isolated invariant characteristics of these dynamics, and identified load balancing and evolution of the routing topology as key factors in the observed properties.

Based on this observation, we proposed a model for the dynamics of the topology and ego-centred measurements, which integrates both load balancing and routing changes. Simulations show that this model captures the main characteristics of the dynamics of the ego-centred views. We performed extensive simulations with the model. We showed that the underlying structure and the measurement frequency play a strong role in the observed behaviour.

By exploring in the impact of the parameters on observed dynamics, we experimentally identified several relations between measurements parameters and the characteristics of dynamics. Using these relations, we are able to approximately predict some characteristics of dynamics from the parameters without running expensive simulations. We validated our findings using real-world data. Some relations are almost exact, but others give only a general idea of how the observed dynamics depends, for instance, on the type of underlying graph. The model shows some non-linear and non-monotonic behaviours, that complicate the analysis. We have shown that the number of routes between two computers—in the model routes correspond to shortest paths—play a key role in the observed dynamics and we studied in detail the size of the shortest path subgraph between two nodes in random graphs.

We showed that it is very difficult to quantify the actual speed of the Internet topology evolution. In order to study this, we considered the impact of the measurement frequency on the observed network dynamics. Under the hypothesis that topology changes occur according to some Poisson processes, we can very accurately estimate the mean number of changes per unit of time. With sufficient number of observations, we can guarantee any accuracy that we want, regardless of the delay between measurements. Unfortunately, we showed that real-world topology changes are not Poisson, and a precise characterisation of the topology evolution speed is still missed.

We also showed that changes that are not observed by our periodical sampling are also not very important in the following sense: network parts that are not observed have a short lifetime or they have a very small chance of being observed, i.e. they are a small part of a large set of routes between two computers.

Future work lies in several directions. We strongly believe that our model can be used to estimate some properties of the actual IP-level routing topology that are not directly available through measurements. Applying our knowledge obtained from the study of the model to real-world data would allow us to estimate the real-world values corresponding to these parameters, such as for instance the frequency of link changes in the whole topology. Since our model is based on random graphs and simple mechanics for load balancing and routing dynamics, it lends itself well to formal analysis. Thus, we should continue the formal investigations started in Chapters 4 and 5.

The field of Internet topology modelling is very active, and models far more realistic than random graphs are available. One should explore the combination of our routing mechanisms principles with these topology models, to investigate the role played by the topology structure on the observed dynamics. In particular, our model does not take into account the long term topology evolution, since it does not model node birth or death. Coupling the ingredients of our routing dynamics with, e.g. a growing model for the Internet topology which would reflect its long term dynamics would surely lead to insightful results. Another direction consists in developing of more realistic model of topology changes (currently we use random link swaps) and more realistic model of routing (we use shortest paths as a model of routes, but in real-world not all routes are shortest paths).

One may adapt contemporary mathematical theory of large networks, especially *property testing*[3] and *continuous graphs*[4] described in Chapter 1, in order to use them in practice. We also believe that it is possible to eventually develop a general theory about shortest path subgraphs, considering these objects as an analogue (or generalisation) of the notion of a distance. There may be several vertexes at the same distance but with different shortest path subgraphs between them. The distance is just *a number* that measures a shortest path between two elements, while the shortest path subgraph is a *whole structure* that lies between them.

Finally, one may study other complex networks using the methods described in this manuscript. For example, we may consider the dynamics of transmission paths of the information in a social network. In such a case we also have two factors:

[3]Generally speaking, a property testing algorithm is a probabilistic algorithm with sublinear time complexity than tries to decide whether a given graph have some global property using only a partial information about this graph.

[4]A graph with more than countable number of nodes. These structures can be regarded as the limits of sequences of random graphs when their size increases.

the fact that there are several paths that a rumour can use to spread is somewhat similar to load-balancing, and the topology also evolves. Therefore our work could be applied to the characterisation of social network dynamics from information of rumour spreading only.

Bibliography

[1] *A Radar for the Internet – Publicly available datasets.* http://data.complexnetworks.fr/Radar/.

[2] CAIDA, *MIDAR antialiasing tool.* http://www.caida.org/tools/measurement/midar/.

[3] J. J. Sylvester, *Chemistry and algebra*, Nature, 17 (1878), p. 284.

[4] K. Pearson, *Notes on regression and inheritance in the case of two parents*, Proceedings of the Royal Society of London, 58 (1895), pp. 240–242.

[5] A. Rapoport, *Cycle distributions in random nets*, The bulletin of mathematical biophysics, 10 (1948), pp. 145–157.

[6] R. Solomonoff and A. Rapoport, *Connectivity of random nets*, Bulletin of Mathematical Biology, 13 (1951), pp. 107–117.

[7] P. Erdős and A. Rényi, *On random graphs*, Publicationes Mathematicae Debrecen, 6 (1959), p. 290.

[8] E. N. Gilbert, *Random graphs*, Annals of Mathematical Statistics, 30 (1959), pp. 1141–1144.

[9] J. F. Kenney and E. S. Keeping, *Linear regression and correlation*, in Mathematics of Statistics, Van Nostrand, Princeton, NJ, 3 ed., 1962, ch. 15, pp. 252–285.

[10] R. Rado, *Universal graphs and universal functions*, Acta Arithmetica, 9 (1964), pp. 331–340.

[11] D. J. de Solla Price, *Networks of scientific papers*, Science, 149 (1965), pp. 510–515.

[12] J. W. Moon and L. Moser, *Almost all (0,1) matrices are primitive*, Studia Scientiarum Mathematicarum Hungarica, 1 (1966), pp. 153–156.

[13] J. Duchon, *Splines minimizing rotation-invariant semi-norms in sobolev space*, in Constructive Theory of Functions of Several Variables, vol. 571 of Lecture Notes in Mathematics, Springer-Verlag, Berlin, 1976.

[14] E. A. Bender and E. R. Canfield, *The asymptotic number of labeled graphs with given degree sequences*, Journal of Combinatorial Theory (A), 24 (1978), pp. 357–367.

[15] D. B. Carr, R. J. Littlefield, and W. L. Nichloson, *Scatterplot matrix techniques for large n*, in Proceedings of the Seventeenth Symposium on the Interface of Computer Sciences and Statistics on Computer Science and Statistics, New York, NY, USA, 1986, Elsevier North-Holland, Inc., pp. 297–306.

[16] B. Chinoy, *Dynamics of internet routing information*, in SIGCOMM, 1993, pp. 45–52.

[17] J. Hawkinson and T. Bates, *Guidelines for creation, selection, and registration of an Autonomous System (AS)*. RFC 1930 (Best Current Practice), March 1996.

[18] V. Paxson, *Measurements and Analysis of End-to-End Internet Dynamics*, PhD thesis, University of California, April 1997.

[19] E. W. Zegura, K. L. Calvert, and M. J. Donahoo, *A quantitative comparison of graph-based models for internet topology*, IEEE/ACM Trans. Netw., 5 (1997), pp. 770–783.

[20] D. J. Watts, *The structure and dynamics of small-world systems*, PhD thesis, 1997. Copyright - Copyright UMI - Dissertations Publishing 1997; Last updated - 2014-03-12; First page - n/a.

[21] D. J. Watts and S. H. Strogatz, *Collective dynamics of 'small-world'networks*, nature, 393 (1998), pp. 440–442.

[22] P. S. Eriksson, E. Perfilieva, T. Björk-Eriksson, A.-M. Alborn, C. Nordborg, D. A. Peterson, and F. H. Gage, *Neurogenesis in the adult human hippocampus*, Nature medicine, 4 (1998), pp. 1313–1317.

[23] A.-L. Barabási and R. Albert, *Emergence of scaling in random networks*, science, 286 (1999), pp. 509–512.

[24] M. Faloutsos, P. Faloutsos, and C. Faloutsos, *On power-law relationships of the internet topology*, in Proceedings of the conference on Applications,

technologies, architectures, and protocols for computer communication, SIGCOMM '99, New York, NY, USA, 1999, ACM, pp. 251–262.

[25] J. M. Kleinberg, R. Kumar, P. Raghavan, S. Rajagopalan, and A. S. Tomkins, *The web as a graph: Measurements, models, and methods*, in Computing and combinatorics, Springer, 1999, pp. 1–17.

[26] D. S. Roselli, J. R. Lorch, T. E. Anderson, et al., *A comparison of file system workloads.*, in USENIX Annual Technical Conference, General Track, 2000, pp. 41–54.

[27] A. Medina, I. Matta, and J. Byers, *On the origin of power laws in internet topologies*, SIGCOMM Comput. Commun. Rev., 30 (2000), pp. 18–28.

[28] W. Aiello, F. Chung, and L. Lu, *A random graph model for massive graphs*, in Proceedings of the thirty-second annual ACM symposium on Theory of computing, STOC '00, New York, NY, USA, 2000, ACM, pp. 171–180.

[29] A. Broido and k. claffy, *Internet topology: Connectivity of ip graphs*, in ITCom 2001: International Symposium on the Convergence of IT and Communications, International Society for Optics and Photonics, 2001, pp. 172–187.

[30] P. J. Cameron, *The random graph revisited*, in European Congress of Mathematics, Springer, 2001, pp. 267–274.

[31] M. E. Newman, *Assortative mixing in networks*, Physical review letters, 89 (2002), p. 208701.

[32] A. Vázquez, R. Pastor-Satorras, and A. Vespignani, *Large-scale topological and dynamical properties of the internet*, Phys. Rev. E, 65 (2002), p. 066130.

[33] A. Vázquez, R. Pastor-Satorras, and A. Vespignani, *Large-scale topological and dynamical properties of the internet*, Physical Review E, 65 (2002), p. 066130.

[34] M. Coates, R. Castro, R. Nowak, M. Gadhiok, R. King, and Y. Tsang, *Maximum Likelihood Network Topology Identification from Edge-based Unicast Measurements*, in Proc. ACM SIGMETRICS, 2002.

[35] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, *Network topology generators: degree-based vs. structural*, in Proceedings of ACM SIGCOMM, 2002.

[36] Q. Chen, H. Chang, R. Govindan, and S. Jamin, *The origin of power laws in internet topologies revisited*, in Proceedings of IEEE INFOCOM, vol. 2, 2002, pp. 608 – 617 vol.2.

[37] N. Alon, *Testing subgraphs in large graphs*, Random Structures & Algorithms, 21 (2002), pp. 359–370.

[38] F. Chung and L. Lu, *The average distances in random graphs with given expected degrees*, Proceedings of the National Academy of Sciences, 99 (2002), pp. 15879–15882.

[39] S. Saroiu, K. P. Gummadi, and S. D. Gribble, *Measuring and analyzing the characteristics of napster and gnutella hosts*, Multimedia systems, 9 (2003), pp. 170–184.

[40] A. Lakhina, J. Byers, M. Crovella, and P. Xie, *Sampling biases in IP topology measurements*, in INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies, vol. 1, march 2003, pp. 332 – 341 vol.1.

[41] M. E. Newman, *The structure and function of complex networks*, SIAM review, 45 (2003), pp. 167–256.

[42] S.-T. Park, D. M. Pennock, and C. L. Giles, *Comparing static and dynamic measurements and models of the Internet's AS topology*, in Proceedings of IEEE Infocom, 2004.

[43] C. Faloutsos, K. S. McCurley, and A. Tomkins, *Fast discovery of connection subgraphs*, in Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining, ACM, 2004, pp. 118–127.

[44] A. Narula-Tam, E. Modiano, and A. Brzezinski, *Physical topology design for survivable routing of logical rings in wdm-based networks*, Selected Areas in Communications, IEEE Journal on, 22 (2004), pp. 1525–1538.

[45] B. Zhang, R. Liu, D. Massey, and L. Zhang, *Collecting the internet as-level topology*, ACM SIGCOMM Computer Communication Review, 35 (2005), pp. 53–61.

[46] J.-L. Guillaume and M. Latapy, *Complex network metrology*, Complex Systems, (2005), pp. 83–94.

[47] J. Xia, L. Gao, and T. Fei, *Flooding attacks by exploiting persistent forwarding loops*, in Proceedings of the ACM SIGCOMM Internet Measurement Conference, 2005.

[48] B. Quoitin and S. Uhlig, *Modeling the Routing of an Autonomous System with CBGP*, in IEEE Network Magazine, 2005.

[49] A. Bonato, *A survey of models of the web graph*, in Combinatorial and algorithmic aspects of networking, Springer, 2005, pp. 159–172.

[50] I. F. Akyildiz, X. Wang, and W. Wang, *Wireless mesh networks: a survey*, Computer Networks, 47 (2005), pp. 445 – 487.

[51] Y. W. Teh, D. Newman, and M. Welling, *A collapsed variational bayesian inference algorithm for latent dirichlet allocation*, in NIPS'06, 2006, pp. 1353–1360.

[52] J. Quittek and K. White, *Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations*. RFC 4560 (Proposed Standard), June 2006.

[53] Y. Koren, S. C. North, and C. Volinsky, *Measuring and extracting proximity in networks*, in Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining, ACM, 2006, pp. 245–255.

[54] M. Stojmenovic, A. Nayak, and J. Zunic, *Measuring linearity of a finite set of points*, in IEEE Conference on Cybernetics and Intelligent Systems, June 2006, pp. 1–6.

[55] M. Lad, D. Massey, and L. Zhang, *Visualizing Internet Routing Changes*, IEEE Transactions on Visualization and Computer Graphics, special issue on Visual Analytics, (2006).

[56] V. D. Blondel, J.-L. Guillaume, J. M. Hendrickx, and R. M. Jungers, *Distance distribution in random graphs and application to network exploration*, Phys. Rev. E, 76 (2007), pp. 066–101.

[57] B. Donnet and T. Friedman, *Internet topology discovery: a survey*, Communications Surveys & Tutorials, IEEE, 9 (2007), pp. 56–69.

[58] J.-J. Pansiot, *Local and Dynamic Analysis of Internet Multicast Router Topology*, Annales des télécommunications, 62 (2007), pp. 408–425.

[59] R. Oliveira, B. Zhang, and L. Zhang, *Observing the Evolution of Internet AS Topology*, in Proceedings of ACM SIGCOMM, 2007.

[60] M. Latapy and C. Magnien, *Complex network measurements: Estimating the relevance of observed properties*, in Proc. of INFOCOM 2008, 2008, pp. 1660–1668.

[61] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira, *Traceroute Anomalies: Detection and Prevention in Internet Graphs*, Computer Networks, 52 (2008), pp. 998–1018.

[62] H. van den Esker, R. van der Hofstad, and G. Hooghiemstra, *Universality for the distance in finite variance random graphs*, J. Stat. Phys., 133 (2008), pp. 169–202.

[63] E. Bullmore and O. Sporns, *Complex brain networks: graph theoretical analysis of structural and functional systems*, Nature Reviews Neuroscience, 10 (2009), pp. 186–198.

[64] C. Magnien, F. Ouédraogo, G. Valadon, and M. Latapy, *Fast dynamics in internet topology: Observations and first explanations*, in Proceedings of the Fourth IEEE International Conference on Internet Monitoring and Protection, 2009, pp. 137–142.

[65] D. Achlioptas, A. Clauset, D. Kempe, and C. Moore, *On the bias of traceroute sampling*, Journal of the ACM, 56 (2009).

[66] M. Rubinov and O. Sporns, *Complex network measures of brain connectivity: Uses and interpretations*, NeuroImage, 52 (2010), pp. 1059 – 1069. Computational Models of the Brain.

[67] D. Cvetković, P. Rowlinson, and S. Simić, *An introduction to the theory of graph spectra*, Cambridge-New York, (2010).

[68] R. Mozzachiodi and J. H. Byrne, *More than synaptic plasticity: role of nonsynaptic plasticity in learning and memory*, Trends in Neurosciences, 33 (2010), pp. 17 – 26.

[69] L. Lovász and B. Szegedy, *Testing properties of graphs and functions*, Israel Journal of Mathematics, 178 (2010), pp. 113–156.

[70] C. Crespelle and F. Tarissan, *Evaluation of a new method for measuring the internet degree distribution: Simulation results*, Comput. Commun., 34 (2011), pp. 635–648.

[71] I. Cunha, R. Teixeira, and C. Diot, *Measuring and Characterizing End-to-End Route Dynamics in the Presence of Load Balancing*, in Proceedings of Passive and Active Measurement Conference, 2011.

[72] M. Latapy, C. Magnien, and F. Ouédraogo, *A radar for the internet*, Complex Systems, 20 (2011), pp. 23–30.

[73] L. Benamara and C. Magnien, *Removing bias due to finite measurement of dynamic systems: case study on p2p systems*, arXiv preprint arXiv:1104.3694, (2011).

[74] F. de Montgolfier, M. Soto, and L. Viennot, *Treewidth and hyperbolicity of the internet*, in Network Computing and Applications (NCA), 2011 10th IEEE International Symposium on, Aug 2011, pp. 25–32.

[75] X. Wang, M. Latapy, and M. Soria, *Deciding on the type of the degree distribution of a graph from traceroute-like measurements*, International Journal of Computer Networks & Communications, 4 (2012).

[76] B. Huffaker, M. Fomenkov, and k. Claffy, *Internet Topology Data Comparison*, tech. rep., Cooperative Association for Internet Data Analysis (CAIDA), May 2012.

[77] D. Krioukov, M. Kitsak, R. S. Sinkovits, D. Rideout, D. Meyer, and M. Boguñá, *Network cosmology*, Scientific reports, 2 (2012).

[78] K. Wehmuth, E. Fleury, and A. Ziviani, *A New Model for Time-Varying Graphs*, in Temporal and Dynamic Networks: From Data to Models, Copenhagen, Denmark, Jun 2013.

[79] S. Kirgizov, C. Magnien, and F. Tarissan, *Studying the impact of measurement frequency on the ip-level routing topology dynamics*, in 24-ème édition du colloque Gretsi, Septemer 2013. http://kirgizov.complexnetworks.fr/publications/gretsi.pdf.

[80] C. Magnien, A. Medem, S. Kirgizov, and F. Tarissan, *Towards realistic modeling of IP-level routing topology dynamics*, Networking Science, (2013), pp. 1–10.

[81] J. Viard and M. Latapy, *Identifying roles in an IP network with temporal and structural density*, in Sixth IEEE International Workshop on Network Science for Communication Networks (NetSciCom 2014), Canada, May 2014, p. 1.

[82] S. Kirgizov and C. Magnien, *Peaks and valleys in the size distribution of shortest path subgraphs*. http://www-complexnetworks.lip6.fr/~kirgizov/publications/peaks-and-valleys.pdf, 2014.

[83] S. Kirgizov and F. Queyroi, *Stochastic process estimation from partial observations: Poisson case*, January 2014. http://kirgizov.complexnetworks.fr/notes/1-point-poisson.pdf.

[84] A. Agresti, *An Introduction to Categorical Data Analysis (Wiley Series in Probability and Statistics)*, Wiley-Interscience, 2 ed., Mar 2007.

[85] C. Berge, *Théorie des graphes et ses applications*, Collection universitaire de mathématiques, Dunod, Paris, 1958.

[86] B. Bollobás, *Random Graphs*, vol. 73 of Cambridge Studies in Advanced Mathematics, Cambridge University Press, 2001.

[87] S. N. Dorogovtsev and J. F. Mendes, *Evolution of networks: From biological nets to the Internet and WWW*, Oxford University Press, 2013.

[88] R. Durrett, *Random graph dynamics*, Cambridge series in statistical and probabilistic mathematics, Cambridge University Press, 2006.

[89] F. Harary, *Graph Theory*, Addison-Wesley, Reading, MA, 1969.

[90] D. Knuth, *The Art of Computer Programming: Seminumerical algorithms*, Addison-Wesley series in computer science and information processing, Addison-Wesley, 1981.

[91] D. Kőnig, *Theorie Der Endlichen Und Unendlichen Graphen*, 1936.

[92] L. Lovász, *Large Networks and Graph Limits*, American Mathematical Society colloquium publications, American Mathematical Society, 2012.

[93] J. Postel, *RFC 791 Internet Protocol - DARPA Inernet Programm, Protocol Specification*, September 1981.

[94] E. Prisner, *Graph dynamics*, vol. 338, CRC Press, 1995.

[95] S. Wasserman, *Social network analysis: Methods and applications*, vol. 8, Cambridge university press, 1994.