



Activity Report 2016

Project-Team RAP

Networks, Algorithms and Probabilities

RESEARCH CENTER
Paris

THEME
Networks and Telecommunications

Table of contents

1. Members	1
2. Overall Objectives	1
3. Research Program	2
3.1. Scaling of Markov Processes	2
3.2. Design and Analysis of Algorithms	2
3.3. Structure of random networks	3
4. New Results	3
4.1. Random Graphs	3
4.2. Resource Allocation in Large Data Centres	3
4.3. Ressource allocation in vehicle sharing systems	5
4.4. Scaling Methods	5
4.4.1. Fluid Limits in Wireless Networks	5
4.4.2. Large Unreliable Stochastic Networks	6
4.5. Stochastic Models of Biological Networks	6
4.5.1. Stochastic Modelling of self-regulation in the protein production system of bacteria.	6
4.5.1.1. Feedback model	6
4.5.1.2. Models with Cell Cycle	7
4.5.1.3. Multi-protein Model	7
4.5.2. Stochastic Modelling of Protein Polymerization	8
5. Bilateral Contracts and Grants with Industry	8
6. Partnerships and Cooperations	9
6.1.1. Visits of International Scientists	9
6.1.2. Visits to International Teams	9
7. Dissemination	9
7.1. Promoting Scientific Activities	9
7.1.1. Scientific events selection	9
7.1.2. Conferences	9
7.1.3. Scientific expertise	10
7.2. Teaching - Supervision - Juries	10
7.2.1. Teaching	10
7.2.2. Juries	10
8. Bibliography	10

Project-Team RAP

Creation of the Project-Team: 2004 February 01, end of the Project-Team: 2016 December 31

Keywords:

Computer Science and Digital Science:

- 1.3. - Distributed Systems
- 6.1.2. - Stochastic Modeling (SPDE, SDE)
- 6.1.4. - Multiscale modeling
- 7.10. - Network science

Other Research Topics and Application Domains:

- 1.1. - Biology
- 1.1.2. - Molecular biology
- 1.1.10. - Mathematical biology
- 6. - IT and telecom
- 6.5. - Information systems

1. Members

Research Scientists

Philippe Robert [Team leader, Inria, Senior Researcher, HDR]
Nicolas Broutin [Inria, Researcher, HDR]
Christine Fricker [Inria, Researcher]

PhD Students

Renaud Dessalles [INRA]
Sarah Eugene [Inria, until Aug 2016, granted by FP7 BANGMD-ERCSKIPPERAD project]
Wen Sun [Inria]
Guilherme Thompson [Inria, granted by Brazilian Grant]
Veronica Quintuna Rodriguez [ORANGE Labs]

Post-Doctoral Fellow

Davit Martirosyan [Inria, from Sep 2016]

Visiting Scientist

Ravi Rasendra Mazumdar [Inria]

Administrative Assistant

Nelly Maloisel [Inria]

Others

Yousra Chabchoub [ISEP, Associate Professor]
Hanene Mohamed [Univ. Paris X, Associate Professor]
Thanh-Huy Nguyen [Inria, Intern, from Mar 2016 until Jul 2016]
Othmane Safsafi [ENS Paris, PhD Student]

2. Overall Objectives

2.1. Overall Objectives

The research team RAP (Networks, Algorithms and Communication Networks) was created in 2004 on the basis of a long standing collaboration between engineers at Orange Labs in Lannion and researchers from Inria Paris. The initial objective was to formalize and expand this fruitful collaboration.

At Orange Labs in Lannion, the members of the team are experts in the analytical modeling of communication networks as well as on some of the operational aspects of network management concerning traffic measurements on ADSL networks, for example.

At Inria Paris, the members of RAP have a recognized expertise in modeling methodologies applied to stochastic models of communication networks.

RAP also has the objective of developing new fundamental tools to investigate *probabilistic* models of complex communication networks. We believe that mathematical models of complex communication networks require a deep understanding of general results on stochastic processes. The two fundamental domains targeted are:

1. Design and analysis of algorithms for communication networks.
2. Analysis of scaling methods for Markov processes: fluid limits and functional limit theorems.

From the very beginning, it has been decided that RAP would focus on a number of particular issues over a period of three or four years. The general goal of the collaboration with Orange Labs is to develop, analyze and optimize algorithms for communication networks. The design of algorithms to allocate resources in large distributed systems is currently investigated in the framework of this collaboration:

3. Research Program

3.1. Scaling of Markov Processes

The growing complexity of communication networks makes it more difficult to apply classical mathematical methods. For a one/two-dimensional Markov process describing the evolution of some network, it is sometimes possible to write down the equilibrium equations and to solve them. The key idea to overcome these difficulties is to consider the system in limit regimes. This list of possible renormalization procedures is, of course, not exhaustive. The advantages of these methods lie in their flexibility to various situations and to the interesting theoretical problems they raised.

A fluid limit scaling is a particularly important means to scale a Markov process. It is related to the first order behavior of the process and, roughly speaking, amounts to a functional law of large numbers for the system considered.

A fluid limit keeps the main characteristics of the initial stochastic process while some second order stochastic fluctuations disappear. In “good” cases, a fluid limit is a deterministic function, obtained as the solution of some ordinary differential equation. As can be expected, the general situation is somewhat more complicated. These ideas of rescaling stochastic processes have emerged recently in the analysis of stochastic networks, to study their ergodicity properties in particular.

3.2. Design and Analysis of Algorithms

Data Structures, Stochastic Algorithms

The general goal of the research in this domain is of designing algorithms to analyze and control the traffic of communication networks. The team is currently involved in the design of algorithms to allocate bandwidth in optical networks and also to allocate resources in large distributed networks. See the corresponding sections below.

The team also pursues analysis of algorithms and data structures in the spirit of the former Algorithms team. The team is especially interested in the ubiquitous divide-and-conquer paradigm and its applications to the design of search trees, and stable collision resolution protocols.

3.3. Structure of random networks

This line of research aims at understanding the global structure of stochastic networks (connectivity, magnitude of distances, etc) via models of random graphs. It consists of two complementary foundational and applied aspects of connectivity.

RANDOM GRAPHS, STATISTICAL PHYSICS AND COMBINATORIAL OPTIMIZATION. The connectivity of usual models for networks based on random graphs models (Erdős–Rényi and random geometric graphs) may be tuned by adjusting the average degree. There is a *phase transition* as the average degree approaches one, a *giant* connected component containing a positive proportion of the nodes suddenly appears. The phase of practical interest is the *supercritical* one, when there is at least a giant component, while the theoretical interest lies at the *critical phase*, the break-point just before it appears.

At the critical point there is not yet a macroscopic component and the network consists of a large number of connected component at the mesoscopic scale. From a theoretical point of view, this phase is most interesting since the structure of the clusters there is expected (heuristically) to be *universal*. Understanding this phase and its universality is a great challenge that would impact the knowledge of phase transitions in all high-dimensional models of *statistical physics* and *combinatorial optimization*.

RANDOM GEOMETRIC GRAPHS AND WIRELESS NETWORKS. The level of connection of the network is of course crucial, but the *scalability* imposes that the underlying graph also be *sparse*: trade offs must be made, which required a fine evaluation of the costs/benefits. Various direct and indirect measures of connectivity are crucial to these choices: What is the size of the overwhelming connected component? When does complete connectivity occur? What is the order of magnitude of distances? Are paths to a target easy to find using only local information? Are there simple broadcasting algorithms? Can one put an end to viral infections? How much time for a random crawler to see most of the network?

NAVIGATION AND POINT LOCATION IN RANDOM MESHES. Other applications which are less directly related to networks include the design of improved navigation or point location algorithms in geometric meshes such as the Delaunay triangulation build from random point sets. There the graph model is essentially fixed, but the constraints it imposes raise a number of challenging problems. The aim is to prove performance guarantees for these algorithms which are used in most manipulations of the meshes.

4. New Results

4.1. Random Graphs

Participant: Nicolas Broutin.

Self-similar real trees defined as fixed-points [15]: Random trees that are fixed points of some random decompositions are ubiquitous: the essential building blocks of the scaling limits of graphs, but also various other trees associated to combinatorial models are such trees. We study a general class of fixed-points equations in spaces of measure metric spaces that yield such objects, and study the existence/uniqueness of the fixed-points in the natural spaces of interest. We also obtain geometric information such as fractal dimension or estimates about the degrees directly from the equations. This is joint work with Henning Sulzbach.

4.2. Resource Allocation in Large Data Centres

Participants: Christine Fricker, Philippe Robert, Guilherme Thompson, Veronica Quintana Rodriguez.

Efficient resource allocation in large data centers has become crucial matter since the expansion in volume and in variety of the internet based services and applications. Everyday examples, such as Video-on-Demand and Cloud Computing are part of this change in the internet environment, bringing new perspectives and challenges with it. Resource pooling (gathering resources to avoid idleness) and resource decentralization (to bring the service “closer” to the user) are too an important topic in service design, specially because of the inherent dichotomy presented in this discussion. Understanding and assessing the performance of such systems ought enable to better resource management and, consequently, better quality of service.

Currently, most systems operate under decentralized policies due to the complexity of managing data exchange on large scale. In such systems, customer demands are served respecting their initial service requirements (a certain video quality, amount of memory or processing power etc.) until the system reaches saturation, which then leads to the blockage of subsequent customer demands. Strategies that rely on the scheduling of tasks are often not suitable to address this load balancing problem as the users expect instantaneous service usage in real time applications, such as video transmission and elastic computation. Our research goal is to understand and redesign its algorithms in order to develop decentralized schemes that can improve global performance using local instantaneous information. This research is made in collaboration with Fabrice Guillemin, from Orange Labs.

In a first approach to this problem, we examined offloading schemes in fog computing context, where one data centers are installed at the edge of the network. We analyze the case with one data center close to user which is backed up by a central (usually bigger) data center. In [10], when a request arrives at an overloaded data center, it is forwarded to the other data center with a given probability, in order to help dealing with saturation and reducing the rejection of requests. In [17], we studied another scheme, where requests are systematically forwarded by the small data to a larger one, but with some trunk reservation to ensure service performance in the second one. We have been able to demonstrate the behavior and performance of these systems, using the invariant distribution of a random walks in the quarter plane, and obtaining explicit expressions for both schemes. Those two papers shed some light in the effectiveness of this fog computing design, by investigating two basic and intuitive policies, whose advantages can now be compared.

In [11] and [16], we investigated allocation schemes which consist in reducing the bandwidth of arriving requests to a minimal value. In the first, this process is initiated when the system is saturated and in the second when the system is close to saturation. We analyzed the effectiveness of such a downgrading policies. In the case of downgrading at saturation, we were able to find an explicit expression of the key performance metrics when two types of customers share a resource and type two asks for the double of resources compared to type one. And, for the second case, we could show that if the system is correctly designed then we can stop losing clients. We developed a mathematical model which allows us to predict system behavior under such a policy and calculate the optimal threshold (in the same scale as the resource) after which downgrading should be initiated. We proved the existence of a unique equilibrium point, around which we have been able to determine the probability a customer receives service at requested quality. We have also shown that system blockage becomes indeed negligible. This policy finds a natural application in the framework of video streaming services and other real time applications. Notably, we are able to derive explicit and simple expressions for many aspects of this system, giving special predictability the outcome of such policy.

Recently, we started to investigate the framework of network function virtualization, another emergent stream stream of research in resource allocation. We start by considering the execution of Virtualized Network Functions (VNFs) in data centers whose capacities are limited and service execution time is constrained by telecommunication protocols. Virtualization practices play a crucial role in the evolution of telecommunications network architectures, since the service providers can reduce the investment on the edge and share resource more efficiently. Macrofunctions are virtualized into micro ones and treat individually. Through simulations and basic mathematical models, we aroused the discussion of three different prioritization policies and their *trade-offs*. They have shown that in for parallelizable macrofunctions (i.e. no order of execution), the greedy algorithm ensures the best performance in terms of execution delay. For chained ones, macrofunctions whose microfunctions need to be run in a certain order, this algorithm is not suitable, the Round Robin and the Dedicated Core policies perform with the same level.

With these results in mind, we have extend our research towards more complex systems, investigating the behaviour of multiple resource systems (such as a Cloud environment, where computational power is provided using unities of CPU and GB of RAM). We analyzed cooperation between data centers offering multiple resources and under imbalanced loads, a problem that naturally arises from the decentralization of resources. Again, we consider instantaneous service. By forwarding some clients across the system, we could design a policy that is allows cooperation between system and preserves service quality at both data centers. We consider two types of demands asking for two types of resources; particularly, type one clients demand more

of type one resource (and symmetrically for type two). We have shown that under our forwarding scheme, which offloads clients requiring most of the saturated resource locally at each data center, we can eliminate losses (in a well design system). Some other interesting properties that can help systems designers are as well derived, such as the minimum threshold for the sustainability of such scheme and the offloading rates. A document is being written to further publication.

4.3. Ressource allocation in vehicle sharing systems

Participants: Christine Fricker, Hanene Mohamed, Thanh-Huy Nguyen.

Vehicle sharing systems are becoming an urban mode of transportation, and launched in many cities, as Velib' and Autolib' in Paris. One of the major issues is the availability of the resources: vehicles or free slots to return them. These systems became an important topic in Operation Research and now the importance of stochasticity on the system behavior is commonly admitted. The problem is to understand the system behavior and how to manage these systems in order to provide both resources to users.

Our stochastic model is the first one taking into account the finite number of spots at the stations.

Equivalence of ensembles We used limit local theorems to obtain the asymptotic stationary joint distributions of several station states when the system is large (both numbers of stations and bikes), in the case of finite capacities of the stations. This gives the asymptotic independence property for node states. This widely extends the existing results on heterogeneous bike-sharing systems.

Load balancing policies. Recently we investigated some load balancing algorithms for stochastic networks to improve the bike sharing system behavior. We focus on the choice of the least loaded station among two to return the bike. In real systems, this choice is local. Thus the main challenge is to deal with the choice between two neighboring stations.

For that, a set of N queues, with a local choice policy, is studied. When a customer arrives at queue i , he joins the least loaded queue between queues i and $i + 1$. When the load tends to zero, we obtain an asymptotic for the stationary distribution of the number of customers at a queue. It allows to compare local choice, no choice and choice between two chosen at random.

For a bike-sharing homogeneous model, we study a deterministic cooperation between the stations, two by two. Analytic results are achieved in an homogeneous bike-sharing model. They concern the limit as the system is large, the so-called mean-field limit, and its equilibrium point. Results on performance mainly involve an original closed form expression of the stationary blocking probability in the classical join-the-shortest-queue model. These results are compared by simulations with the policy where the users choose the least loaded station between two stations to return close to their destination. It turns out that, because of randomness, the choice between two neighbours gives better performance than grouping stations two by two.

Bike-sharing model with waiting In real systems, if the customer does not find the resource (a bike or an place to return), he can either leave, or search in a neighbouring station, or wait. We extend a basic model to take into account waiting.

4.4. Scaling Methods

Participants: Philippe Robert, Wen Sun.

4.4.1. Fluid Limits in Wireless Networks

This is a collaboration with Amandine Veber (CMAP, École Polytechnique). The goal is to investigate the stability properties of wireless networks when the bandwidth allocated to a node is proportional to a function of its backlog: if a node of this network has x requests to transmit, then it receives a fraction of the capacity proportional to $\log(1 + x)$, the logarithm of its current load. This year we completed the analysis of a star network topology with multiple nodes. Several scalings were used to describe the fluid limit behaviour.

4.4.2. Large Unreliable Stochastic Networks

The reliability of a large distributed system is studied. The framework is a system where files have several copies on different servers. When one of these servers breaks down, all copies stored on it are lost. These copies can be retrieved afterwards if there is another copy of the same files stored on other servers. In the case where no other copy of a given file is present in the system, it is definitely lost. We study two math models on this problem.

In the first model, it is assumed that the duplication process is local, any server has a capacity to make copies to another server, but the capacity can only be used for the copies present on this server. We have studied the asymptotic behavior of this system, i.e. the number of servers is large, via mean field methods. We have shown that asymptotically, the load of each server can be described by a non-linear Markov process. This limiting process can also give an exponential decay of the number of files. This is a joint work with Reza Aghajani, Brown University.

In the second model, two policies for the reassignment of files are studied. It is assumed that each server has a neighborhood, that consists of a set of servers in the system. When a server breaks down, it restarts immediately but empty. Copies on it are reassigned to other servers in the neighborhood, following “Random Choice” (RC) policy or “Power of choices” (PoC) policy.

- (RC) Each copy joins a server in the neighborhood at random.
- (PoC) Each copy chooses several servers in the neighborhood at random, and joins the least loaded one.

The asymptotic behaviors of these two policies are investigated through mean field models. We have shown that when the number of servers gets large, the load of each server can be approached by a linear (resp. non-linear) Markov process for RC (resp. PoC) policy. The equilibrium distributions of these asymptotic processes are also given. This is a joint work with Inria/UPMC Team Regal.

4.5. Stochastic Models of Biological Networks

Participants: Renaud Dessalles, Sarah Eugene, Philippe Robert, Wen Sun.

4.5.1. Stochastic Modelling of self-regulation in the protein production system of bacteria.

This is a collaboration with Vincent Fromion from INRA Jouy-en-Josas, which started in December 2013.

In prokaryotic cells (e.g. E. Coli. or B. Subtilis) the protein production system has to produce in a cell cycle (i.e. less than one hour) more than 10^6 molecules of more than 2500 kinds, each having different level of expression. The bacteria uses more than 67% of its resources to the protein production. Gene expression is a highly stochastic process: bacteria sharing the same genome, in a same environment will not produce exactly the same amount of a given protein. Some of this stochasticity can be due to the system of production itself: molecules, that take part in the production process, move freely into the cytoplasm and therefore reach any target in the cell after some random time; some of them are present in so much limited amount that none of them can be available for a certain time; the gene can be deactivated by repressors for a certain time, etc. We study the integration of several mechanisms of regulation and their performances in terms of variance and distribution. As all molecules tend to move freely into the cytoplasm, it is assumed that the encounter time between a given entity and its target is exponentially distributed.

4.5.1.1. Feedback model

We have also investigated the production of a single protein, with the transcription and the translation steps, but we also introduced a direct feedback on it: the protein tends to bind on the promoter of its own gene, blocking therefore the transcription. The protein remains on it during an exponential time until its detachment caused by thermal agitation.

The mathematical analysis aims at understanding the nature of the internal noise of the system and to quantify it. We tend to test the hypothesis usually made that such feedback permits a noise reduction of protein distribution compared to the “open loop” model. We have made the mathematical analysis of the model (using a scaling to be able to have explicit results), it appeared that reduction of variance compared to an “open loop” model is limited: the variance cannot be reduced for more than 50%.

We proposed another possible effect of the feedback loop: the return to equilibrium is faster in the case of a feedback model compared to the open loop model. Such behaviour can be beneficial for the bacteria to change of command for a new level of production of a particular protein (due, for example, to a radical change in the environment) by reducing the respond time to reach this new average. This study has been mainly performed by simulation and it has been shown that the feedback model can go 50% faster than the open loop results.

4.5.1.2. *Models with Cell Cycle*

Usually, classical models of protein production do not explicitly represent several aspects of the cell cycle: the volume variations, the division and the gene replication. Yet these aspects have been proposed in literature to impact the protein production. We have therefore proposed a series of “gene-centered” models (that concentrates on the production of only one type of protein) that integrates successively all the aspects of the cell cycle. The goal is to obtain a realistic representation of the expression of one particular gene during the cell cycle. When it was possible, we analytically determined the mean and the variance of the protein concentration using Marked Poisson Point Process framework.

We based our analysis on a simple model where the volume changes across the cell cycle, and where only the mechanisms of protein production (transcription and translation) are represented. The variability predicted by this model is usually assimilated to the “intrinsic noise” (i.e. directly due to the protein production mechanism itself). We then add the random segregation of compounds at division to see its effect on protein variability: at division, every mRNA and every protein has an equal chance to go to either of the two daughter cells. It appears that this division sampling of compounds can add a significant variability to protein concentration. This effect directly depends on the relative variance (Fano factor) of the protein concentration: this effect is stronger as the relative variance is low. The dependence on the relative variance can be explained by considering a simplified model. With parameters deduced from real experimental measures, we estimate that the random segregation of compounds can double the variability of the genes with the lowest relative variance.

Finally, we integrate the gene replication to the model: at some point in the cell cycle, the gene is replicated, hence doubling the transcription rate. We are able to give analytical expressions for the mean and the variance of protein concentration at any moment of the cell cycle; it allows to directly compare the variance with the previous model with division. We show that gene replication has little impact on the protein variability: an environmental state decomposition shows that the part of the variance due to gene replication represents only at most 2% of the total variability predicted by the model.

In the end, these results are compared to the real experimental measure of protein variability. It appears that the models with cell cycle presented above tend to underestimate the protein variability especially for highly expressed proteins.

4.5.1.3. *Multi-protein Model*

In continuation of the previous models, we propose a model that still considers the division and the gene replication but which also integrates the sharing of common resources: the different genes are in competition for the limited quantity of RNA-polymerases and ribosomes in order to produce the mRNAs and proteins. The goal is to examine if fluctuations in the availability of these macromolecules have an important impact on the protein variability, as it has been suggested in literature. As the model considers the interaction between the different protein productions, one needs to represent all the genes of the bacteria altogether: it is therefore a multi-protein model.

As this model is too complex to be studied analytically, we have developed a procedure to estimate the parameters so that they correspond to real experimental measures. We then perform simulations in order to determine the variance of each protein and compare them with the one predicted by the models with cell cycle

previously presented. It appears that the common sharing of RNA-polymerases and ribosomes has a limited impact on the protein production: for most of proteins the variance increases of at most 10%.

Finally, we have investigated other possible sources of variability by presenting other simulations that integrate some specific aspects: variability in the production of RNA-polymerases and ribosomes, uncertainty in the division and DNA replication decisions, etc. None of the considered aspects seems to have a significant impact on the protein variability.

4.5.2. Stochastic Modelling of Protein Polymerization

This is a collaboration with Marie Doumic, Inria MAMBA team.

The first part of our work focuses on the study of the polymerization of protein. This phenomenon is involved in many neurodegenerative diseases such as Alzheimer's and Prion diseases, e.g. mad cow. In this context, it consists in the abnormal aggregation of proteins. Curves obtained by measuring the quantity of polymers formed in in vitro experiments are sigmoids: a long lag phase with almost no polymers followed by a fast consumption of all monomers. Furthermore, repeating the experiment under the same initial conditions leads to somewhat identical curves up to translation. After having proposed a simple model to explain this fluctuations, we studied a more sophisticated model, closer to the reality. We added a conformation step: before being able to polymerize, proteins have to misfold. This step is very quick and remains at equilibrium during the whole process. Nevertheless, this equilibrium depends on the polymerization which is happening on a slower time scale. The analysis of these models involves stochastic averaging principles.

We have also investigated a more detailed model of polymerisation by considering the the evolution of the number of polymers with different sizes ($X_i(t)$) where $X_i(t)$ is the number of polymers of size i at time t . By assuming that the transitions rates are scaled by a large parameter N , it has been shown that, in the limit, the process ($X_i^N(t)$) is converging to the solution of Becker-Döring equations as N goes to infinity. For another model including nucleation, we have given an asymptotic description of the lag time at the first and second order. These results are obtained in particular by proving stochastic averaging theorems.

The second part concerns the study of telomeres. This work is made in collaboration with Zhou Xu, Teresa Teixeira, from IBCP in Paris.

In eukaryotic cells, at each mitosis, chromosomes are shortened, because the DNA polymerase is not able to duplicate one ending of the chromosome. To prevent loss of genetic information- which could be catastrophic for the cell- chromosomes are equipped with telomeres at their endings. These telomeres do not contain any genetic information; they are a repetition of the sequence T-T-A-G-G-G thousands times. At each mitosis, there is therefore a loss of telomere. As it has a finite length, when the telomeres are too short, the cell cannot divide anymore: they enter in replicative senescence. Our model tries to captures the two phases of the shortening of telomeres: first, the initial state of the cells, when the telomerase is still active to repair the telomeres. Second, when the telomerase is inhibited, we try to estimate the senescence threshold, when the replication of the cells stops. See [8].

5. Bilateral Contracts and Grants with Industry

5.1. Bilateral Contracts with Industry

- Contrat de recherche externalisé avec ORANGE SA "Scheduling Global OS". Duration three years 2014-2016.
- *Christine Fricker* is the leader of PGMO project "Systèmes de véhicules en libre-service: Modélisation, Analyse et Optimisation" with G-Scop (CNRS lab, Grenoble) and Ifsttar. From 1 to 3 years. From 1/10/2013 to 30/9/2016.
- PhD grant CJS (Contrat Jeune Scientifique) Frontières du vivant of INRA for Renaud Dessalles.
- PhD grant from Fondation Sciences Mathématiques de Paris for Wen Sun.

- PhD grant from Brazilian Government for Guilherme Thompson.
- CELTIC+ Contract “SENDATE”.

6. Partnerships and Cooperations

6.1. International Research Visitors

6.1.1. Visits of International Scientists

- Louigi Addario-Berri (McGill)
- Gabor Lugosi (ICREA and Pompeu Fabra)
- Christina Goldschmidt (Oxford)
- Hui He (Beijing Normal University)
- Elie de Panafieu (Bell Labs France)
- Minmin Wang (Buenos Aires)
- Neil Walton (UK)

6.1.1.1. Internships

- Thanh-Huy, Nguyen, M1 student at Polytechnique (Avril-July 2016). Internship on “Resource allocation in vehicle-sharing systems”.

6.1.2. Visits to International Teams

- *Nicolas Broutin* has visited the mathematics department of the University of Bath, the School of Computer Science at McGill University, the computer science laboratory in Bordeaux (LaBRI) and the NYU-ECNU institute for mathematical sciences at NYU Shanghai.
- *Philippe Robert* (17-27 April 2016) and *Wen Sun* (25-29 April 2016) have visited the Department of Applied Mathematics at Brown University.

7. Dissemination

7.1. Promoting Scientific Activities

7.1.1. Scientific events selection

7.1.1.1. Member of the editorial boards

- *Nicolas Broutin* is member of the steering committee of the international meeting on analysis of algorithms (AofA).
- *Philippe Robert* is Associate Editor of the Book Series “Mathématiques et Applications” edited by Springer Verlag and Associate Editor of the journal “Queueing Systems, Theory and Applications”.

7.1.2. Conferences

- *Nicolas Broutin* has given lectures at annual meeting of the ALEA working group of the GDR-IM, the workshop GRAAL that took place at IHES, at the workshop organized in honor of Svante Janson’s 60’s birthday in Sweden, and the workshop on geometric networks at the Symposium on Computational Geometry that was held in Eindhoven. He gave a plenary lecture at the conference on analysis of algorithms that was held in Strobl in June. He has also exposed his results at the probability seminar in Grenoble, the seminar of algorithmic and combinatorial geometry at IHP, and at the probability seminar of the NYU-ECNU Math institute in Shanghai.

Nicolas Broutin has given the following lectures:

- "Almost optimal sparsification of random geometric graphs" on Jan 30 at the Journées Présage at the Computer Science department of the University of Marne-la-Vallée,
- "The scaling limit of inhomogeneous random graphs" on June 29 at the Paris-Bath workshop that took place at the IHP.
- "Combinatorial structures and fixed-points arguments" on Nov 1 at NYU-ECNU Mathematics institute in Shanghai
- "Random graphs and their scaling limits" on Nov 24 at the Hong Kong University of Science and Technology

He has visited the mathematics department of Beijing Normal University (China), the computer science department of McGill University (Canada), the Mathematics Institute at the Hong-Kong University of Science and Technology, the NYU-ECNU institute for mathematical sciences at NYU Shanghai (China).

- *Renaud Dessalles* has given talks at the conference "Les probabilités de demain" (May 2016) and "Journées Modélisation Aléatoire et Statistique" (August 2016).
- *Christine Fricker* and GT have given talks at the Sigmetrics and MAMA workshop, Juan-les-Pins, 14-18 /6/2016 and at the 2nd ECQT (European Conference on Queueing Theory) Toulouse, 18-20/7/2016.
- *Veronica Quintana Rodriguez* gave a talk at the conference "17th International Network Strategy and Planning Symposium" in september (Montreal).
- *Philippe Robert* gave a talk at the department of mathematics of Brown University (April), at the stochastic networks conference in San Diego, and at the conference "Modern Problems in Theoretical and Applied Probability" in Novosibirsk in August.

7.1.3. Scientific expertise

- *Christine Fricker* is member of the jury of agrégation.
- *Philippe Robert* is member of the scientific council of EURANDOM.

7.2. Teaching - Supervision - Juries

7.2.1. Teaching

Master : *Nicolas Broutin* Master Parisien de Recherche en Informatique (MPRI), in the course 2.15 on Analysis of Algorithms.

Master: *Nicolas Broutin* , Analysis of Algorithms, NYU Shanghai.

Master: *Philippe Robert* , Master, Probabilités et Applications, UPMC.

7.2.2. Juries

- *Philippe Robert* has been member of the jury of PhD defense of Sarah Eugène, Claudio Imbrenda and Tanguy Cabana.

8. Bibliography

Publications of the year

Doctoral Dissertations and Habilitation Theses

- [1] S. EUGENE. *Stochastic modelling in molecular biology: a probabilistic analysis of protein polymerisation and telomere shortening*, UPMC LJLL, September 2016, <https://hal.inria.fr/tel-01377561>

Articles in International Peer-Reviewed Journals

- [2] L. ADDARIO-BERRY, N. BROUTIN, C. GOLDSCHMIDT, G. MIERMONT. *The scaling limit of the minimum spanning tree of the complete graph*, in "Annals of Probability", 2016, 78 p. , to appear, <https://hal.inria.fr/hal-00773360>
- [3] N. BROUTIN, O. DEVILLERS, R. HEMSLEY. *Efficiently navigating a random Delaunay triangulation*, in "Random Structures and Algorithms", 2016, vol. 49, n^o 1, pp. 95–136 [DOI : 10.1002/RSA.20630], <https://hal.inria.fr/hal-00940743>
- [4] N. BROUTIN, L. DEVROYE, G. LUGOSI. *Almost optimal sparsification of random geometric graphs*, in "The Annals of Applied Probability : an official journal of the institute of mathematical statistics", 2016, vol. 26, pp. 3078-3109 [DOI : 10.1214/15-AAP1170], <https://hal.archives-ouvertes.fr/hal-01056127>
- [5] N. BROUTIN, J.-F. MARCKERT. *A new encoding of coalescent processes. Applications to the additive and multiplicative cases*, in "Probability Theory and Related Fields", 2016, vol. 166, n^o 1, pp. 515–552 [DOI : 10.1007/s00440-015-0665-1], <https://hal.inria.fr/hal-01092562>
- [6] N. BROUTIN, M. WANG. *Cutting down p-trees and inhomogeneous continuum random trees*, in "Bernoulli", 2016, 44 pages, 6 figures, <https://hal.archives-ouvertes.fr/hal-01056125>
- [7] M. DOUMIC, S. EUGENE, P. ROBERT. *Asymptotics of Stochastic Protein Assembly Models*, in "SIAM Journal on Applied Mathematics", November 2016, vol. 76, n^o 6, 20 p. [DOI : 10.1137/16M1066920], <https://hal.inria.fr/hal-01301266>
- [8] S. EUGENE, T. BOURGERON, Z. XU. *Effects of initial telomere length distribution on senescence onset and heterogeneity*, in "Journal of Theoretical Biology", January 2017, vol. 413, 8 p. , <https://hal.inria.fr/hal-01378596>
- [9] S. EUGENE, W.-F. XUE, P. ROBERT, M. DOUMIC-JAUFFRET. *Insights into the variability of nucleated amyloid polymerization by a minimalistic model of stochastic protein assembly*, in "Journal of Chemical Physics", May 2016, vol. 144, n^o 17, 12 p. [DOI : 10.1063/1.4947472], <https://hal.inria.fr/hal-01205549>
- [10] C. FRICKER, F. GUILLEMIN, P. ROBERT, G. R. THOMPSON. *Analysis of an offloading scheme for data centers in the framework of fog computing*, in "ACM Transactions on Modeling and Performance Evaluation of Computing Systems", September 2016, vol. 1, n^o 4, 18 p. [DOI : 10.1145/2950047], <https://hal.inria.fr/hal-01179206>
- [11] C. FRICKER, F. GUILLEMIN, P. ROBERT, G. R. THOMPSON. *Analysis of Downgrading for Resource Allocation*, in "Performance Evaluation Review", September 2016, vol. 44, n^o 2, 3 p. [DOI : 10.1145/3003977.3003986], <https://hal.inria.fr/hal-01406664>
- [12] P. ROBERT, J. TOUBOUL. *On the dynamics of random neuronal networks*, in "Journal of Statistical Physics", September 2016 [DOI : 10.1007/s10955-016-1622-9], <https://hal.inria.fr/hal-01075242>
- [13] W. SUN, M. FEUILLET, P. ROBERT. *Analysis of Large Unreliable Stochastic Networks*, in "Annals of Applied Probability", October 2016, vol. 26, n^o 5, 41 p. [DOI : 10.1214/15-AAP1167], <https://hal.inria.fr/hal-01203096>

Other Publications

- [14] R. AGHAJANI, P. ROBERT, W. SUN. *A Large Scale Analysis of Unreliable Stochastic Networks*, August 2016, working paper or preprint, <https://hal.archives-ouvertes.fr/hal-01359208>
- [15] N. BROUTIN, H. SULZBACH. *Self-similar real trees defined as fixed-points and their geometric properties*, October 2016, working paper or preprint, <https://hal.inria.fr/hal-01384309>
- [16] C. FRICKER, F. GUILLEMIN, P. ROBERT, G. R. THOMPSON. *Allocation Schemes of Ressources with Downgrading*, April 2016, working paper or preprint, <https://hal.inria.fr/hal-01301291>
- [17] F. GUILLEMIN, G. R. THOMPSON. *Analysis of a trunk reservation policy in the framework of fog computing*, October 2016, working paper or preprint, <https://hal.inria.fr/hal-01378533>
- [18] P. ROBERT. *Jacques Neveu et les modèles probabilistes de réseaux*, August 2016, working paper or preprint, <https://hal.inria.fr/hal-01382215>
- [19] P. ROBERT, A. VEBER. *A Scaling Analysis of a Star Network with Logarithmic Weights*, September 2016, working paper or preprint, <https://hal.inria.fr/hal-01377703>