



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team rap

Réseaux, Algorithmes et Probabilités

Paris - Rocquencourt

Theme : Networks and Telecommunications

Activity
R *eport*

2009

Table of contents

1. Team	1
2. Overall Objectives	1
3. Scientific Foundations	2
3.1. Measurements and Mathematical Modeling	2
3.2. Design and Analysis of Algorithms	2
3.3. Scaling of Markov Processes	3
4. New Results	4
4.1. Algorithms: On-line Algorithms for Traffic Measurements	4
4.2. Algorithms: Bandwidth Allocation in Optical Networks	5
4.3. Algorithms: Internet without Congestion Control Algorithms	5
4.4. Scaling Methods: Stochastic Models of peer-to-peer networks	5
4.5. Scaling Methods: Interaction of TCP Flows	5
4.6. Scaling Methods: Cognitive radio networks	6
4.7. Miscellaneous	6
4.7.1. Metastability and communication networks	6
4.7.2. Continuous-state branching process in random environment	7
4.7.3. Load balancing via random local search.	7
4.7.4. The Evolution of a Spatial Stochastic Network	7
4.7.5. A Scaling analysis of a Cat and Mouse Markov chain	8
5. Contracts and Grants with Industry	8
6. Other Grants and Activities	8
7. Dissemination	8
7.1. Leadership within scientific community	8
7.2. Teaching	9
7.3. Conference and workshop committees, invited conferences	9
8. Bibliography	9

1. Team

Research Scientist

Christine Fricker [Research Associate (CR) Inria]
Fabrice Guillemin [France Telecom R&D, Lannion]
Stéphanie Moteau [France Telecom R&D, Lannion]
Philippe Robert [Team Leader, Research Director (DR) Inria, HdR]
James Roberts [Research Director (DR) Inria]

Faculty Member

Danielle Tibi [Université Denis Diderot]

PhD Student

Yousra Chabchoub [INRIA]
Mathieu Feuillet [Détachement du Corps des Mines]
Florian Simatos [Détachement du Corps des Télécoms]

Visiting Scientist

Nelson Antunes [Visiting Scientist, University of Algarve]

Administrative Assistant

Virginie Collette [(AI) Inria]

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 France Telecom R&D in Lannion and researchers from INRIA Paris — Rocquencourt. The initial objective was to formalize and expand this fruitful collaboration.

At France-Telecom R&D 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 — Rocquencourt, the members of RAP have a recognized expertise in modeling methodologies applied to stochastic models of communication networks.

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 is to develop, analyze and optimize algorithms for communication networks. Projects currently in progress are the following.

1. Mathematical models of traffic measurements of ADSL traffic.
2. Design of algorithms to allocate bandwidth in optical 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. It could be argued that, since stochastic networks are “applied”, general results concerning Markov processes (for example) are not of real use in practice and, therefore, that ad-hoc results are more helpful. Recent developments in the study of communication networks have shown that this point of view is flawed. Technical tools such as scaling methods, large deviations and rare events, requiring a sound understanding of some fundamental results concerning stochastic processes, are now used in the analysis of these stochastic models. Two domains are currently investigated

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

3. Scientific Foundations

3.1. Measurements and Mathematical Modeling

3.1.1. Traffic Modeling

The characterization of Internet traffic has become over the past few years one of the major challenging issues in telecommunications networks. As a matter of fact, understanding the composition and dynamics of Internet traffic is essential for network operators in order to control quality of service and supervise their networks. Since the well-known paper by Leland *et al.* on the self-similar nature of Ethernet traffic in local area networks, a huge amount of work has been devoted to the characterization of Internet traffic. In particular, a number of different hypotheses on the origins and reasons for the self-similarity of Internet traffic have been explored.

A common approach to describing traffic in a backbone network consists in observing the bit rate process evaluated over fixed length intervals of a few hundreds of milliseconds, say. Long range dependence as well as self-similarity are two basic properties of the bit rate process, which have been observed through measurements in many different situations. Different characterizations of the fractal nature of traffic have been proposed in the literature (see for instance Norros on the monofractal characterization of traffic). An exhaustive account of the fractal characterization of Internet traffic can be found in the book by Park and Willinger. Even though long range dependence and self-similarity properties are very intriguing from a theoretical point of view, their significance in network design has recently been questioned.

While self-similar models introduced so far in the literature aim to describe the overall traffic on a link, it is now usual to distinguish short transfers (referred to as mice) and long transfers (referred to as elephants) [31]. This dichotomy was not totally clear up to the recent past (see, for instance, network measurements from the MCI backbone network). However, the distinction between mice and elephants is becoming more and more evident with the emergence of peer-to-peer (p2p) applications, which give rise to a large amount of traffic on a small number of TCP connections. The above observation leads us to analyze ADSL traffic through a flow based approach intended, in particular, to clarify the mice/elephants dichotomy. The intuitive definition of a mouse is that it has such a small number of packets that TCP hardly leaves the slow start regime. Thus, a mouse is not sensitive to the bandwidth sharing imposed by TCP congestion avoidance. On the other hand, elephants are sufficiently large that one can assume they share bottleneck bandwidth through TCP flow control. As a consequence, mice and elephants have a totally different behavior from a modeling point of view.

We consider that describing statistical properties of Internet traffic at packet level is not appropriate, mainly because of the strong dependence properties noted above. At this time scale, signal processing techniques (wavelets, fractal analysis, ...) can lead to a better description of Internet traffic but do not bring the insights necessary for traffic control. It is widely believed that, at user level, independence properties can be assumed (as for telephone networks), just because users behave quite independently. Unfortunately, there is still no agreed stochastic model of typical user activity. Some models have been proposed, but their number of parameters is too large and most cannot be easily inferred from real measurements. We have chosen to look at the traffic of elephants and mice which defines an intermediate time scale. Some independence properties seem to hold at this level and allow the possibility of Markovian analysis. Note that despite some criticism, Markovian techniques are basically the *only* tools that can give a sufficiently precise description of the evolution of various stochastic models (average behavior, distribution of the time to overflow buffers, ...).

3.2. Design and Analysis of Algorithms

The stochastic models of a class of generic algorithms with an underlying tree structure, the splitting algorithms, have a wide range of applications. To classify the massive data sets generated by traffic measurements, these algorithms turn out to be fundamental. Hashing mechanisms such as Bloom filters are currently investigated in the light of these new applications. These algorithms have also been used for now more than 30 years in various areas, among which

- Data structures. Fundamental algorithms on data structures are used to sort and search. They are sometimes referred to as divide and conquer algorithms.
- Access Protocols. These algorithms are used to give a distributed access to a common communication channel.
- Distributed systems. Recently, algorithms to select a subset of a group of identical communicating components like ad hoc networks, sensor networks and more generally mobile networks are using a related approach.

This class of algorithms is fundamental, their range of applications is very large and, moreover, they have a nice underlying mathematical structure. Trees are the main mathematical objects to describe them. The associated stochastic processes can be seen as a discrete version of fragmentation processes which have been recently thoroughly investigated by Bertoin, Pitman and others. They are also related to random recursive decompositions of intervals introduced by Mauldin and Williams and their developments in fractal geometry by Falconer, Lapidus, etc...

A very large subset of the literature has been devoted to the static case analysis, mainly because of its applications in theoretical computer science. In the dynamic case, i.e., when the shape of the tree changes according to some random events, little work has been done for this class of algorithms. Their analysis has been, for the moment, mainly achieved by using analytical methods with functional transforms, complex analysis techniques and inversions theorems. Curiously, despite of the intuitive underlying stochastic structures, probabilistic studies of these algorithms are quite scarce.

3.3. 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. When the number of nodes is more than 3, however, this kind of approach is not possible in general. The key idea to overcome these difficulties is to consider the system in limit regimes.

- By considering the asymptotic behavior of the probability of some events as in the case of large deviations at a logarithmic scale or for heavy tailed distributions, or by looking at Poisson approximations to describe a sequence of events associated with them.
- By taking some parameter η of the model and look at the behavior of the system when it approaches some critical value η_c . In some cases, even if the model is complicated, its behavior simplifies as $\eta \rightarrow \eta_c$: some of the nodes grow according to a classical limit theorem while the remainder reach an equilibrium which can be described.
- By changing the time scale and space scale with a common normalizing factor N and letting N go to infinity. This leads to functional limit theorems as discussed below.

This list of possible renormalization procedures is, of course, not exhaustive. This methodology has become more and more developed over the last ten years. Its advantages lie in its flexibility to various situations and to the interesting theoretical problems it has raised since then.

3.3.1. Fluid Limits

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.

It is in general quite difficult to have a satisfactory description of an ergodic Markov process describing a stochastic network. When the dimension of the state space d is greater than 1, the geometry significantly complicates any investigation: analytical tools such as Wiener-Hopf techniques for dimension 1 cannot be easily generalized to higher dimensions. It is nevertheless possible to gain insight into the behavior of these processes through some limit theorems. The considered limiting procedure consists in speeding up time and scaling appropriately the process itself with some parameter. The behavior of such rescaled stochastic processes is analyzed when the scaling parameter goes to infinity. In the limit, one gets a sort of caricature of the initial stochastic process which is defined as a *fluid limit*.

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. See Rybko and Stolyar [32], for example. In statistical physics, these methods are quite classical, see Comets [27].

4. New Results

4.1. Algorithms: On-line Algorithms for Traffic Measurements

Participants: Yousra Chabchoub, Christine Fricker.

4.1.1. Analysis of a counting algorithm

Joint work with H. Mohamed, University of Nanterre. This work is related to the design and performance evaluation of a probabilistic algorithm to identify long flows (elephants) in Internet traffic.

The algorithm, originally proposed by Azzana [26] is based on Bloom filters: the ID of each packet is sent via independent hashing functions to filters (hash tables). The idea is that, since each packet of a flow is hashed to the same counters, long flows (more than $C = 20$ packets) will then be identified. Because of the accumulation of short flows, filters must be refreshed from time to time to cope with repeated collisions in hash tables. Specifically, the non-null counters are decremented by one every time the filling rate of the filter reaches some threshold r .

This algorithm has been widely experimented, it has been tested on both commercial Orange traces and academic Abilene traces. The algorithm engenders both false positives (mice detected as elephants) and false negatives (missed elephants). To estimate the number of false negatives, a simple model has been proposed, first for mice with size one, in terms of an urn and ball model. It has been used to evaluate the impact of refreshment on the proportion of false positives, i.e., short flows detected as long flows with more than C packets. Limit theorems of the empirical distribution of the filter counters (mean field limit) when the filter size is large have been obtained. The limiting stationary distribution is deterministic and has a nice interpretation in terms of queues. The proof is based on the convergence to a dynamical system and uses a queueing interpretation of the fixed point of the system. The convergence of the invariant measure is completely proved for $C = 2$, where a Lyapunov function is exhibited.

In the considered model, all counters associated with the ID of a given packet are incremented. In Chabchoub *et al.* [16] we investigate an improvement of this policy which consists in incrementing only the lowest valued counters. It is clear that this algorithm will allow a more accurate estimate of the number of long flows. Using the shortest queue drastically reduces the tail of queue size distribution. This problem has been extensively analyzed recently under the general heading of the power of (two) choice(s). In our setting, a new version of the algorithm has been studied: only one counter is incremented among d counters provided by d hashing functions. This algorithm has been tested on traffic traces and shown to perform well, especially when d is small (2 or 4). The previous results partly extend to this framework: the limiting time between two refreshes is analytically tractable. The convergence to a dynamical system is obtained, though it is more complicated to describe analytically. The existence of a fixed point is proved but not uniqueness. The global stability of the dynamical system, which could lead to the same limit result remains an open question. It is under investigation.

4.2. Algorithms: Bandwidth Allocation in Optical Networks

Participants: Nelson Antunes, Yousra Chabchoub, Christine Fricker, Philippe Robert, James Roberts.

The development of dynamic optical switching is widely recognized as an essential requirement to meet anticipated growth in Internet traffic. Since September 2009, RAP has begun an investigation into the traffic management and performance evaluation issues that are particular to this technology. A first analysis of passive optical networks used for high speed Internet access has led to the proposal of an original dynamic bandwidth allocation algorithm and to an evaluation of its traffic capacity [18]. Our activity on optical networking is carried out in collaboration with Orange Labs with whom we have a research contract. We have also established contacts with Alcatel Bell Labs and look forward to collaborating with Iraj Saniee and his team on their proposed time-domain wavelength interleaved networking architecture (TWIN).

4.3. Algorithms: Internet without Congestion Control Algorithms

Participant: Mathieu Feuillet.

It is commonly accepted that the observed robustness of the Internet in the last two decades, despite an exponential growth in traffic, is largely due to the use of congestion control algorithms. Since the introduction of congestion control mechanisms in TCP in the late 1980s, no congestion collapse has been observed. Unfortunately, it is not clear if the IETF will be able to enforce the use of TCP in the future. More and more applications do not comply with the “TCP-friendliness” principle. With the generalization of broadband access through optical fiber (FTTH), the impact of such applications cannot be neglected.

In this project we analyze Internet behavior at flow level in the absence of any control congestion. We assume all sources send at their maximum access rate and recover from packet loss by the use of some source coding or retransmissions independent of the congestion control. The bandwidth allocation is then fully determined by the buffer management policy implemented in routers. We study two possible policies: *Fair Dropping* and *Tail Dropping*. In both cases, we study the efficiency of resource utilization in terms of the maximum load the network can sustain. For *Fair Dropping*, utilization is optimal. This is not the case for *Tail Dropping* although efficiency is shown to be very high in most topologies of practical use. For the latter, the order of magnitude difference between the maximum access rates and link capacity plays a crucial role. The complete model and first results are presented in [15].

4.4. Scaling Methods: Stochastic Models of peer-to-peer networks

Participants: Philippe Robert, Florian Simatos.

File-sharing networks are distributed systems used to disseminate files among a subset of Internet end nodes. A file is split into several pieces called chunks. The general simple principle is that once a node has retrieved a chunk, it in turn becomes a server for this chunk. In Leskela *et al.* [22], we have considered a stochastic model for chunk arrival times and download durations. We determined the maximal arrival rate that such a network can accommodate, i.e., the conditions under which the Markov process describing this network is ergodic. Technical estimates related to the survival of interacting branching processes are key ingredients to establish the stability of these systems. Several cases are considered: networks with one and two chunks where a complete classification is obtained and several cases of networks with n chunks.

4.5. Scaling Methods: Interaction of TCP Flows

Participant: Philippe Robert.

This is a collaboration with Carl Graham (CMAP, École Polytechnique). Mathematical modeling of data transmission in communication networks has been the subject of intense activity for some time now. For data transmission, the Internet can be described as a very large distributed system with self-adaptive capabilities to the different congestion events that regularly occur at its numerous nodes. Various approaches have been used in this respect: control theory, ordinary differential equations, Markov processes, optimization techniques, ...

The coexistence of numerous connections in a network with a general number of nodes is analyzed in this work. The mean-field limit of a Markovian model describing the interaction of several classes of permanent connections in a network is analyzed. As with TCP, each connection has a self-adaptive behavior in that its transmission rate along its route depends on the level of congestion of the nodes of the route. Since several classes of connections going through the nodes of the network are considered, an original mean-field result in a multi-class context is established. It is shown that, as the number of connections goes to infinity, the behavior of the different classes of connections can be represented by the solution of an unusual set of non-linear stochastic differential equations depending not only on the sample paths of the process, but also on its *distribution*. Existence and uniqueness results for the solutions of these equations are derived. Properties of their invariant distributions are investigated and it is shown that, under some natural assumptions, they are determined by the solutions of a fixed point equation in a finite dimensional space. See Graham and Robert [10].

The uniqueness of the solutions of the associated fixed point equation have been investigated in Graham *et al.* [17]. Uniqueness results of such solutions are proved for different topologies: rings, trees and a linear network and with various configurations for routes through nodes.

4.6. Scaling Methods: Cognitive radio networks

Participants: Philippe Robert, Florian Simatos.

Joint work with Ed Coffman, Shuzo Tarumi and Gil Zussman from Columbia University, USA. In light of current and projected demands of wireless communications, techniques are needed to make more efficient use of the wireless spectrum. In one class of applications, certain bands of the wireless spectrum are not licensed to specific users, but instead are reserved for users that compete for channel allocation on a dynamic basis; this dynamic spectrum allocation is also envisioned for cognitive radio networks in which unlicensed users compete for bandwidth within the temporarily unused channels of licensed users. Reconfigurability is a key property of the opportunistic users: Their channels may consist of a number of disjoint sub-bands allocated to them dynamically; a channel is not simply a fixed, single continuous band of frequencies. Exploiting this property leads to new and intriguing fragmentation issues.

In Coffman *et al.* [20], we study a baseline mathematical model of these issues and arrive at a number of important insights that need to be borne in mind in system design. We adopt the most basic model in which a spectrum is shared by unlicensed users only, each characterized by a desired total bandwidth and the duration of a time interval over which it is needed. As users come and go, gaps of available bandwidth develop randomly in both size and position. When allocating bandwidth to a user's channel, the spectrum is searched for gaps in a linear scan; gaps are allocated to the channel until the full requested bandwidth has been provided. Fundamental questions include: Is the number of fragments (sub-bands) into which a user's channel is divided a stable process (e.g., can fragmentation increase indefinitely within a continuous model)? Is there a relation between the numbers of users and gaps similar to the 50 storage allocation? Are there normal limit laws similar to those of other fragmentation problems? Affirmative answers to these and similarly posed questions are observed experimentally; the results are rather surprising at first glance, but plausible derivations are given for each.

4.7. Miscellaneous

Participants: Florian Simatos, Philippe Robert, Danielle Tibi.

4.7.1. Metastability and communication networks

Danielle Tibi. In Antunes *et al.* [25], in the context of mobility, a model of a network that converges, as the number N of nodes gets large, to a mean-field limit with several stable equilibrium points has been exhibited. A similar phenomenon had been observed for a network under a rerouting procedure in Gibbens *et al.* [30]. The analysis in [25] relies on the derivation of an explicit Lyapunov function for the limiting dynamical system. Such a function is not available for the rerouting procedure.

The question of metastability - meaning exponential growth of exit times from neighborhoods of stable points - is not addressed in these papers, though this phenomenon is highly suspected to occur for both models (particularly in view of simulation results). But the non-reversibility is here an obstacle to using standard approaches to metastability.

These models are now re-examined in the light of Freidlin and Wentzell's large deviations approach of randomly perturbed dynamical systems [28]. It can be shown that both models essentially fit the scheme developed in this reference for a certain class of Markov jump processes. For these processes, it provides large deviations results for exit times, points and paths from neighborhoods of stable points. The central quantity involved in the action functionals is the *quasi-potential*, that can be interpreted as the energy function associated to the system in the limit $N \rightarrow \infty$.

These results can be adapted to our context. For the model in [25], it can also be proved that the above mentioned Lyapunov function is equal to the quasi-potential of a simplified process, which, contrary to the model in [29], has the particularity of being asymptotically reversible.

4.7.2. *Continuous-state branching process in random environment*

Florian Simatos, joint work with Vincent Bansaye (University Pierre et Marie Curie). Continuous state branching processes (CSBP) are continuous time, continuous state space Markov processes which satisfy the branching property. Until now, mainly time-homogeneous CSBP have been studied since they naturally appear as limits of Galton-Watson branching processes.

On the other hand, a growing interest has been recently devoted to Galton-Watson branching processes in a random environment: these processes are relevant from a modeling perspective, and it turns out that they have many interesting features from a mathematical standpoint as well. The goal of this collaboration is twofold: define CSBP in a varying environment as the solution of a certain martingale problem. Once these processes are defined, we want to let the environment be random and study the basic properties of CSBP in a random environment. Performance properties of interest are the extinction probability and the existence of a genealogy, for instance.

4.7.3. *Load balancing via random local search.*

Florian Simatos, joint work with A. Ganesh, S. Lilenthal, D. Manjunath and A. Proutière, Microsoft Cambridge. In traditional approaches to load balancing in parallel server systems, e.g., Join-the-Shortest-Queue and variants thereof, clients make smart server selections upon arrival. In this work we analyze a different approach, where load balancing is performed by random load resampling and migration strategies. Clients initially attach to an arbitrary server, but may switch server independently at random instants of time in an attempt to improve their service rate. Load resampling is particularly relevant in scenarios where clients cannot predict the load of a server before being actually attached to it. An important example is in wireless spectrum sharing where clients try to share a set of frequency bands in a distributed manner.

In [21] we analyze closed systems, where we derive tight estimates of the time it takes to balance the load across servers. We also study open systems where clients arrive according to a random process and leave the system upon service completion. In this scenario, we analyze how client migrations within the system interact with the system dynamics induced by client arrivals and departures.

4.7.4. *The Evolution of a Spatial Stochastic Network*

Philippe Robert. Motivated by bandwidth allocation algorithms in wireless networks, the asymptotic behavior of a stochastic network represented by a birth and death processes of particles has been analyzed in [24]. Particles are created at rate λ_+ and their location is independent of the current configuration. Deaths are due to negative particles arriving at rate λ_- . The death of a particle occurs when a negative particle arrives in its neighborhood and kills it. Several killing schemes are considered. The arriving locations of positive and negative particles are assumed to have the same distribution. By using a combination of monotonicity properties and invariance relations it is shown that the configurations of particles converge in distribution for several models. The problems of uniqueness of invariant measures and of the existence of accumulation points

for the limiting configurations have been investigated. It has been shown for several natural models that if $\lambda_+ < \lambda_-$ then the asymptotic configuration has a finite number of points with probability 1. It has also been shown that systems with $\lambda_+ < \lambda_-$ and an infinite number of particles in the limit exist.

4.7.5. A Scaling analysis of a Cat and Mouse Markov chain

Philippe Robert, joint work with Nelly Litvak (University of Twente, Holland). An original on-line page-ranking algorithm has been investigated in Litvak and Robert [23]: starting from an arbitrary Markov chain (C_n) on a discrete state space \mathcal{S} , a Markov chain (C_n, M_n) on the product space \mathcal{S}^2 , the cat and mouse Markov chain, is constructed. The first coordinate of this Markov chain behaves like the original Markov chain and the second component changes only when both coordinates are equal. The asymptotic properties of this Markov chain are investigated. A representation of its invariant measure is in particular obtained. When the state space is infinite it is shown that this Markov chain is in fact null recurrent if the initial Markov chain (C_n) is positive recurrent and reversible. In this context, the scaling properties of the location of the second component, the mouse, are investigated in various situations: simple random walks in \mathbb{Z} and \mathbb{Z}^2 , reflected simple random walk in \mathbb{N} and also in a continuous time setting. For several of these processes, a time scaling with rapid growth gives an interesting asymptotic behavior related to limit results for occupation times and rare events of Markov processes.

5. Contracts and Grants with Industry

5.1. Contracts

Participation to the ANR Projet Blanc “SADA” on Discrete Random Structures, three year contract starting from 2005. Participants: University of Bordeaux, University of Caen, Computer science department of Ecole Polytechnique, INRIA Algo and Rap projects, University of Versailles.

CRE contract “Algorithmes d’allocation de ressources dans les réseaux optiques” with Orange Labs on bandwidth allocation algorithm in optical networks. Two years starting from 2009.

6. Other Grants and Activities

6.1. Visiting scientists

RAP team has received the following people:

Vincent Bansaye (LPMA, Paris) , Guillaume Vu-Brugier (Orange Labs, Lannion) , Florent Malrieu (Université Rennes 1, France) , Vincent Benezech (ENS, France) , Mérouane Debbah (Alcatel-Lucent chair on Flexible Radio, Supelec, France) , Iraj Saniee (Alcatel Bell Labs), Alexandre Proutière (Microsoft, Cambridge), Patrick Loiseau (ENS Lyon), Pierre Del Moral (Université de Bordeaux), Bert Zwart (CWI, Amsterdam).

7. Dissemination

7.1. Leadership within scientific community

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”. He is member of the scientific council of EURANDOM. He is also associate Professor at the École Polytechnique in the department of applied mathematics where he is in charge of lectures on mathematical modeling of networks.

*Jim Roberts*co-organized the workshop [EuroNFTraf09](#) on Traffic management and traffic engineering for the future Internet (7-8 December). He was named of Fellow (membre émérite) of the [SEE](#) in November 2009.

7.2. Teaching

Philippe Robert gives Master2 lectures “Stochastic Networks” in the laboratory of the Probability of the University of Paris VI. He is also giving lectures in the “Programme d’approfondissement de Mathématiques Appliquées et d’Informatique” on Networks and Algorithms at the École Polytechnique.

7.3. Conference and workshop committees, invited conferences

Yusra Chabchoub gave a talk at ITC21 conference in Paris.

Mathieu Feuillet gave a talk at the Young European Queueing Theorists conference, November 19–21, in Eindhoven.

Christine Fricker gave talks at the SMAI meeting, May 25–29th in La-Colle-sur-Loup, France and at the 15th INFORMS Applied Probability Conference on July 12–15th at Cornell University.

Philippe Robert gave talks at Lycée Jeanne d’Albret (Saint Germain en Laye) in May, Bedlewo (Poland) in May, Bell-Laboratories in June, ENS-Cachan in November, NETCOOP conference in Eindhoven in November. *Philippe Robert* was TPC member for ITC 21.

Jim Roberts was TPC member for Infocom 2010, Infocom WiP workshop, ONDM 2010, ITC Specialist Seminar on Multimedia Applications. He gave invited talks at the conferences ITC 21 (September) and DNAC09 (November) and at a seminar at Orange Labs (September).

Florian Simatos gave talks at Université de Rennes 1 in January 2009, Imperial College, London in January, Université de Versailles Saint-Quentin-en-Yvelines in May, University of Toulouse in October.

Florian Simatos visited Moez Draief, Imperial College (London), March, 9–16, Saint-Flour Probability Summer School, July 5–18, Bert Zwart, CWI (Amsterdam), September 28 – October 1, Université de Toulouse (Toulouse), October 26–28, Balakrishna Prabhu, LAAS (Toulouse), October 29.

8. Bibliography

Major publications by the team in recent years

- [1] N. ANTUNES, C. FRICKER, P. ROBERT, D. TIBI. *Stochastic Networks with Multiple Stable Points*, in "Annals of Probability", vol. 36, n^o 1, 2008, p. 255–278.
- [2] J.-F. DANTZER, P. ROBERT. *Fluid limits of string valued Markov processes*, in "Annals of Applied Probability", vol. 12, n^o 3, 2002, p. 860–889.
- [3] F. GUILLEMIN, P. ROBERT, B. ZWART. *AIMD algorithms and exponential functionals*, in "Annals of Applied Probability", vol. 14, n^o 1, 2004, p. 90–117.
- [4] H. MOHAMED, P. ROBERT. *A probabilistic analysis of some tree algorithms*, in "Annals of Applied Probability", vol. 15, n^o 4, November 2005, p. 2445–2471.
- [5] P. ROBERT. *Stochastic Networks and Queues*, Stochastic Modelling and Applied Probability Series, vol. 52, Springer, New-York, June 2003.

Year Publications

Doctoral Dissertations and Habilitation Theses

- [6] Y. CHABCHOUB. *Analyse et modélisation du trafic Internet*, Université Pierre et Marie Curie, October 2009, Ph. D. Thesis.

- [7] F. SIMATOS. *Étude de modèles probabilistes de réseaux pair-à-pair et de réseaux avec mobilité*, École Polytechnique, December 2009, Ph. D. Thesis.

Articles in International Peer-Reviewed Journal

- [8] C. FRICKER, F. GUILLEMIN, P. ROBERT. *An identification problem in an urn and ball model with heavy tailed distributions*, in "Probability in the Engineering and Informational Sciences", 2009, <http://fr.arxiv.org/abs/0812.2546>, To Appear.
- [9] C. FRICKER, F. GUILLEMIN, P. ROBERT. *Perturbation analysis of an $M/M/1$ queue in a diffusion random environment*, in "Queueing Systems, Theory and Applications", vol. 61, n^o 1, 2009, p. 1-36, <http://fr.arxiv.org/abs/0812.2543>.
- [10] C. GRAHAM, P. ROBERT. *Interacting Multi-Class Transmissions in Large Stochastic Networks*, in "Annals of Applied Probability", vol. 19, n^o 6, 2009, p. 2334–2361, <http://arxiv.org/abs/0810.0347>.
- [11] F. GUILLEMIN, P. ROBERT. *Analysis of Steiner subtrees of Random Trees for Traceroute Algorithms*, in "Random Structures and Algorithms", vol. 35, n^o 2, September 2009, p. 194–215, <http://www.arxiv.org/abs/cs.NI/0702156>.
- [12] H. MOHAMED, P. ROBERT. *Dynamic Tree Algorithms*, in "Annals of Applied Probability", 2009, <http://arxiv.org/abs/0809.3577>, To Appear.
- [13] P. ROBERT, F. SIMATOS. *Occupancy Schemes Associated to Yule Processes*, in "Advances in Applied Probability", vol. 41, n^o 2, 2009, p. 600–622, <http://arxiv.org/abs/0810.3079>.
- [14] F. SIMATOS, D. TIBI. *Spatial homogenization in a stochastic network with mobility*, in "The Annals of Applied Probability", 2009, <http://arxiv.org/abs/0807.1205>, to appear.

International Peer-Reviewed Conference/Proceedings

- [15] T. BONALD, M. FEUILLET, A. PROUTIERE. *Is the Law of the Jungle sustainable for the Internet?*, in "Proceedings of INFOCOM", 2009.
- [16] Y. CHABCHOUB, C. FRICKER, H. MOHAMED. *Analysis of a Bloom filter algorithm via the supermarket model*, in "Proceedings of ITC21", September 2009.
- [17] C. GRAHAM, P. ROBERT, M. VERLOOP. *Stability Properties of Networks with Interacting TCP Flows*, in "Proceedings of NET-COOP, Eindhoven", Springer Verlag, November 2009.

Other Publications

- [18] N. ANTUNES, C. FRICKER, P. ROBERT, J. ROBERTS. *Upstream traffic capacity of a WDM EPON under online GATE-driven scheduling*, November 2009, Preprint.
- [19] Y. CHABCHOUB, C. FRICKER, F. GUILLEMIN, P. ROBERT. *Adaptive algorithms for the identification of large flows in IP traffic*, June 2009, <http://arxiv.org/abs/0901.4846>, Preprint.

- [20] E. COFFMAN, P. ROBERT, F. SIMATOS, S. TARUMI, G. ZUSMAN. *Channel Fragmentation in Dynamic Spectrum Access Systems*, November 2009, Preprint.
- [21] A. GANESH, S. LILENTHAL, D. MANJUNATH, A. PROUTIERE, F. SIMATOS. *Load Balancing via Random Local Search in Closed and Open systems*, 2009.
- [22] L. LESKELÄ, P. ROBERT, F. SIMATOS. *On the Stability Properties of File-Sharing Networks*, July 2009, <http://arxiv.org/abs/0907.0375>, Preprint.
- [23] N. LITVAK, P. ROBERT. *A Scaling analysis of a Cat and Mouse Markov chain*, May 2009, <http://arxiv.org/abs/0905.2259>, Preprint.
- [24] P. ROBERT. *The Evolution of a Spatial Stochastic Network*, August 2009, <http://arxiv.org/abs/0908.3256>, Preprint.

References in notes

- [25] N. ANTUNES, C. FRICKER, P. ROBERT, D. TIBI. *Stochastic Networks with Multiple Stable Points*, in "Annals of Probability", vol. 36, n^o 1, 2008, p. 255-278.
- [26] Y. AZZANA. *Mesures de la topologie et du trafic internet*, Universite Paris 6, July 2006, <http://www-c.inria.fr/wiki/pub/RAP/FormerMembers/Azzana-PhD.pdf>, Ph. D. Thesis.
- [27] F. COMETS. *Limites hydrodynamiques*, in "Astérisque", n^o 201-203, 1991, p. Exp. No. 735, 167–192 (1992), Séminaire Bourbaki, Vol. 1990/91.
- [28] M. I. FREIDLIN, A. D. WENTZELL. *Random perturbations of dynamical systems*, Second, Springer-Verlag, New York, 1998, Translated from the 1979 Russian original by Joseph Szücs.
- [29] R. J. GIBBENS, P. J. HUNT, F. P. KELLY. *Bistability in communication networks*, in "Disorder in physical systems, New York", Oxford Sci. Publ., Oxford Univ. Press, 1990, p. 113–127.
- [30] R. J. GIBBENS, P. J. HUNT, F. P. KELLY. *Bistability in communication networks*, in "Disorder in physical systems, New York", Oxford Sci. Publ., Oxford Univ. Press, 1990, p. 113–127.
- [31] V. PAXSON, S. FLOYD. *Wide area traffic: The failure of the Poisson assumption*, in "IEEE/ACM Trans. on Networking", 1995, p. 226-244.
- [32] A. N. RYBKO, A. L. STOLYAR. *On the ergodicity of random processes that describe the functioning of open queueing networks*, in "Problems on Information Transmission", vol. 28, n^o 3, 1992, p. 3–26.