INRIA

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team rap

Réseaux, Algorithmes et Probabilités

Paris - Rocquencourt

THEME COM

Activity Report

2007

Table of contents

1. Team

Head of project-team
Philippe Robert [Research Director (DR) Inria, HdR]

Administrative Assistant
Virginie Collette [(TR) Inria]

Research scientists (Inria)
Christine Fricker [Research Associate (CR) Inria]

Research scientists (external)
Fabrice Guillemin [France Télécom R&D, Lannion]
Stéphanie Moteau [France Télécom R&D, Lannion]

Associate Professor
Danielle Tibi [Délégation de l'Université Paris VII]

Ph. D. Students
Yousra Chabchoub [INRIA]
Hanène Mohamed [INRIA, until August]
Florian Simatos [Détachement du Corps des Télécoms]

2. Overall Objectives

2.1. Overall Objectives

The research team RAP (Networks, Algorithms and Communication Networks) created in 2004 is issued from a long standing collaboration between engineers at France Telecom R&D in Lannion and researchers from INRIA-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 networks management concerning traffic measurements on ADSL networks for example.

At INRIA-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 the efforts of RAP project will focus on few dedicated domains of application over a period of three or four years. The general goal of the collaboration is to develop, analyze and optimize algorithms for communication networks. For the moment, the current projects are the following.

1. Mathematical Models of Traffic Measurements of ADSL traffic.
2. Design of Algorithms to Sample TCP flows.

The RAP project also aims at 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 a 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 good 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. See Section 3.2.
2. Analysis of scaling methods for Markov processes : fluid limits and functional limit theorems. See Section 3.3.

3. Scientific Foundations

3.1. Measurements and Mathematical Modeling

Keywords: *Passive measurements*, *TCP traces*.

3.1.1. Traffic Modeling

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 the dynamics of Internet traffic is essential for network operators in order to offer quality of service and to supervise their networks. Since the celebrated 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, different hypotheses and assumptions have been explored to explain the reasons why and how Internet traffic should be self-similar.

A common approach to describing traffic in a backbone network consists of observing the bit rate process evaluated over fixed length intervals, say a few hundreds of milliseconds. 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 to 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 aims at describing the global traffic on a link, it is now usual to distinguish short transfers (referred to as mice) and long transfers (referred to as elephants) [21]. This dichotomy was not totally clear up to a recent past (see for instance network measurements from the MCI backbone network). Yet, the distinction between mice and elephants become 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 by adopting a flow based approach and more precisely the mice/elephants dichotomy. The intuitive definition of a mouse is that such a flow comprises a small number of packets so that it does not leave or leaves slightly the slow start regime. Thus, a mouse is not very sensitive to the bandwidth sharing imposed by TCP. On the contrary, elephants are sufficiently large so that one can expect that they share the bandwidth of a bottleneck according to the flow control mechanism of TCP. As a consequence, mice and elephants have a totally different behavior from a modeling point of view.

In our approach, we think that describing statistical properties of the Internet traffic at the packet level is not appropriate, mainly because of the strong dependence properties noticed above. It seems to us that, at this time scale, only signal processing techniques (wavelets, fractal analysis, ...) can lead to a better understanding of Internet traffic. It is widely believed that at the level of users, independence properties (like for telephone networks) can be assumed, just because users behave quite independently. Unfortunately, there is not, for the moment, a stochastic model of a typical user activity. Some models have been proposed, but their number of parameters is too large and most of them cannot be easily inferred from real measurements. We have chose to look at the traffic of elephants and mice which is an intermediate time scale. Some independence properties seem to hold at that level and therefore the possibility of Markovian analysis. Note that despite they are sometimes criticized, 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, ...).

3.1.2. Sampling the Internet Traffic

Traffic measurement is an issue of prime interest for network operators and networking researchers in order to know the nature and the characteristics of traffic supported by IP networks. The exhaustive capture of traffic traces on high speed backbone links, with rates larger than 1 Gigabit/s, however, leads to the storage and the analysis of huge amounts of data, typically several TeraBytes per day. A method of overcoming this problem is to reduce the volume of data by sampling traffic. Several sampling techniques have been proposed in the literature (see for instance [17], [20] and references therein). In this paper, we consider the deterministic $1/N$ sampling, which consists of capturing one packet every other N packets. This sampling method has notably been implemented in CISCO routers under the name of NetFlow which is widely deployed nowadays in commercial IP networks.

The major issue with $1/N$ sampling is that the correlation structure of flows is severely degraded and then any digital signal processing technique turns out very delicate to apply in order to recover the characteristics of original flows [20]. An alternative approach consists of performing a statistical analysis of flow as in [17], [18]. The accuracy of such an analysis, however, greatly depends on the number of samples for each type of flows, and may lead to quite inaccurate results. In fact, this approach proves efficient only in the derivation of mean values of some characteristics of interest, for instance the mean number of packets or bytes in a flow.

3.1.3. Algorithms of Sampling

Deriving the general characteristics of the TCP traffic circulating at some edge router has potential applications at the level of an ISP. It can be to charge customers proportionally to their use of the network for example. It can be also to detect what is now called “heavy users”.

Another important application is to detect the propagation of worms, attacks by denial of service (DoS). And, once the attack is detected, to counter it with an appropriate algorithmic approach. Due to the natural variation of the Internet traffic, such a detection (through sampling !) is not obvious. Robust algorithms have to be designed to achieve such an ambitious goal. An ultimate (and ambitious !) goal would be of having an automatic procedure to counter this kind of attacks.

3.1.4. Goals

- Propose a fairly *simple and accurate* estimation of the traffic circulating in an ADSL network. A limited number of parameters should characterize the traffic at the first order. Note that ADSL traffic is significantly different from the usual academic traffic analyzed up to now (more than 80% of the ADSL traffic is from Peer to Peer networks).
- *Infer* through sampling the parameters of the model proposed to describe the ADSL traffic.
- *Design and analyze algorithms* to detect in sampled traffic attacks by worms or DoS and more generally unusual events.

3.2. Design and Analysis of Algorithms

Keywords: *Data Structures*, *Stochastic 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 at 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.

3.2.1. Goals

- *Static case*. Generalize and simplify the results currently proved by using analytic tools. Prove limit theorems for *distributions* instead of averages as it is currently the case.
- *Dynamic case*. Study renormalization techniques to analyze tree algorithms under heavy traffic. The understanding of the fundamental features of these algorithms with a traffic of requests is a major issue in this domain. Because of the quite complex technical framework, the partial results obtained up to now with analytical tools hide, in some way, the general phenomena.

3.3. Scaling of Markov Processes

Keywords: *Fluid Limits*, *Functional Limit Theorems*, *Statistical Physics*.

As the complexity of communication networks increases (and, consequently, the algorithms regulating them), the classical mathematical methods used to estimate the stationary behavior, the transient behavior show more and more their limitations. 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, this kind of approach is not, in general, possible. The key idea to overcome these difficulties is to consider limiting procedures for the system.

- By considering the asymptotic behavior of the probability of some events like it is done for large deviations at a logarithmic scale or for heavy tailed distributions, or looking at Poisson approximations to describe a sequence of events associated to 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 $\eta \to \eta_c$: some of the nodes grow according to some classical limit theorem and the rest of the nodes reach some equilibrium which can be described.
- by changing the time scale and the space scale with a common normalizing factor N and let N goes to infinity. This leads to functional limit theorems, see below.

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

3.3.1. An Example of Scaling Methods : TCP

In our past work, the Congestion Avoidance Algorithm of the TCP protocol has been analyzed by using such a technique. The equilibrium of the *one*-dimensional Markov chain associated to this algorithm is not known for the moment. A large number of papers have been written on this famous AIMD Algorithm. But either it was, in some way, idealized or approximations were used without justifications. In a series of papers, Dumas *et al.* [2], Guillemin *et al.* [3], a conveniently rescaled (time and space) Markov process has been analyzed in the limit when the loss rate of packets of some long connection was converging to 0. It provided a *rigorous* analysis to the scaling properties of this important algorithm of TCP.

3.3.2. Fluid Limits

A fluid limit scaling is a particular important way of scaling a Markov process. It is related to the first order behavior of the process, roughly speaking, it 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 complicates a lot any investigation : Analytical tools such as Wiener-Hopf techniques for dimension 1 cannot be easily generalized to higher dimensions. It is possible nevertheless to get some insight on the behavior of these processes through some limit theorems. The limiting procedure investigated 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 stochastic fluctuations of second order vanish with this procedure. In “good cases”, a fluid limit is a deterministic function, solution of some ordinary differential equation. As it 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 [22] for example. In statistical physics, these methods are quite classical, see Comets [16].
Multi-Class Networks. The state space of the Markov processes encountered up to now were embedded into some finite dimensional vector space. For $J \in \mathbb{N}$, $J \geq 2$ and $j = 1,...J$, λ_j and μ_j are positive real numbers. It is assumed that J Poissonnian arrivals flows arrive at a single server queue with rate λ_j for $j = 1,..., J$ and customers from the jth flow require an exponentially distributed service with parameter μ_j. All the arrival flows are assumed to be independent. The service discipline is FIFO.

A natural way to describe this process is to take the state space of the finite strings with values in the set $\{1, ..., J\}$, i.e. $\mathcal{S} = \cup_{n \geq 0}\{1, ..., J\}^n$, with the convention that $\{1, ..., J\}^0$ is the set of the null string. If $n \geq 1$ and $x = (x_1, ..., x_n) \in \mathcal{S}$ is the state of the queue at some moment, the customer at the kth position of the queue comes from the flow with index x_k, for $k = 1$, ..., n. The length of a string $x \in \mathcal{S}$ is defined by $\|x\|$. Note that $\| \cdot \|$ is not, strictly speaking, a norm. For $n \geq 1$, there are J^n vectors of length n ; the state space has therefore an exponential growth with respect to that function. Hence, if the string valued Markov process $(X(t))$ describing the queue is transient then certainly the length $\|X(t)\|$ converges to infinity as t gets large. Because of the large number of strings with a fixed length, the process $(X(t))$ itself has, a priori, infinitely many ways to go to infinity. Bramson [15] has shown that complicated phenomena could indeed occur. It turns out that the “classical” fluid limits methods of the finite dimensional case cannot be used in such a setting. This is probably one of the most challenging question in the domain to be able to propose new methods to tackle the problems due to the infinite dimension of the state space. Dantzer and Robert [1] derives results in this direction. See also the corresponding chapter of Robert [5].

3.3.3. Goals

The general goals are, in some way, contained in the previous sections. They will consist in developing scaling techniques in the various cases encountered in sampling problems or tree algorithms where the traffic will be supposed to be close to saturation. The following fundamental questions will be analyzed :

— Study the impact of randomness in fluid limit processes. This has been already partially investigated in Dantzer and Robert [1].

— Develop techniques to investigate metastability phenomena observed in some models of networks in the scaling limit due to mean field approach. See Kelly [19].

4. New Results

4.1. Algorithms for Traffic Measurements

Participants: Yousra Chabchoub, Christine Fricker, Philippe Robert.

4.1.1. Counting algorithms

The aim of the study is to design an algorithm to count long flows (elephants) in Internet traffic and to get their characteristics such as IP addresses and sizes. The algorithm must be adapted to very different traffic characteristics and also to traffic variations.

The algorithm is based on Bloom filters : for each incoming packet, the counters of some of the entries of the filter are incremented. The selection of the corresponding entries is done with some hashing values of the flow characteristics. Ideally for a given flow, the values of its entries in the filter should give the number of packets of the flow. But because of collisions of hashing values of different flows (mainly due to the numerous small TCP flows in IP traffic), the filters has to be emptied for time to time. Estan and Varghese proposed a scheme in this domain in 2003. It turns out that the efficiency of this algorithm is highly dependent on the traffic. We proposed a new algorithm which is, in our opinion (and also with the experiments we did) more robust. Refreshment of a filter occurs every time the filling rate reaches some threshold r.

This algorithm has been both theoretically analyzed and with simulations.

First for an analytical point of view, the model has been simplified in order to evaluate the impact of refreshment on the proportion of short flows detected as long flows (false positives). The problem is related to the statistics of the number of balls in m urns at the time where a fixed proportion r of them are non zero when m is large. In a close context, it is also the number of coupons of different kinds for the coupon collector when its collection reaches a proportion r of the total collection. The stationary proportion of counters at say j converges when m tends to ∞ to the invariant probability measure at j of the length of a $M/G/1/C$ queue with service times one at departure times, for an arrival rate related to r. Analytical expressions are obtained even if the convergence proofs are difficult. The main conclusion of the analysis is that the threshold r must not reach some critical value, otherwise the proportion of false positives is very high. Nevertheless, the threshold r must remain enough high in order not to miss long flows. For that purpose, asymptotics when m is large for the mean and the variance of the stationary time between two refreshments are derived. The numbers of false negatives is related to this quantity.

In practice, the algorithm has been tested on both commercial France Telecom traces and academic Abilene traces.

4.1.2. Application to detection of attacks

These algorithms are also used to detect Denial of Service attacks in Internet traffic. Two main kinds of attacks are targeted : Syn-flooding and attacks by volume. For example, a Syn-flooding consists in neutralizing a node by sending to it a huge number of SYN packets from different sources. The node then cannot face to this demand, waiting for acknowledgments to the SYN packets it receives and therefore collapses quickly. During the attack, the traffic is composed of a very huge number of SYN packets to a given destination. These packets can be considered as *flows*, i.e. a sequence of packets with common characteristics, which are the SYN tag and the same destination address.

The idea is to detect these few very long flows using a Counting Filter Algorithm (or Bloom Filters Algorithm). The parameters of the algorithm must be adapted to have a very short response time, less than a minute. It takes into account intensive attacks (a huge amount of SYN packets in a short time) and progressive attacks (a moderated number of SYN packets on a long duration time). Moreover the destination of the attack can be returned by the algorithm.

The adaptive mechanism to refresh filters is more drastic than for counting long flows : the filters are emptied at each refreshing time. The main idea is to compare a new (in the current period) maximum flow size to a maximum flow size which represents the normal traffic and which is recomputed at each period with some part dedicated to the past values to take into account the variations of the normal traffic.

Simulations have been done on two traces given by FranceTelecom for the research project Oscar : a OTIP trace of 65 hours with DoS attacks and a ADSL trace of 3 hours. It shows that the threshold must be maintained to 20%. Then the algorithm has been tested together with different algorithms proposed in the Oscar contract. Experiments has been done on line with ADSL traffic generated by the mean of Planetlab from Planete team project, with the collaboration of the partners. It allows us to improve our algorithm in order to adapt it better to the high variability of traffic.

4.2. Analysis of Algorithms

4.2.1. Analysis of Splitting Algorithms

Participants: Hanène Mohamed, Philippe Robert.

Algorithms with an underlying tree structure are quite common in computer science and communication networks. Splitting algorithms are examples of such algorithms.

A splitting algorithm is a procedure that divides recursively into subgroups an initial group of n items until each of the subgroups obtained has a cardinality strictly less than some fixed number D. A common problem is, given an initial number n of requests, to estimate the time it takes to complete the algorithm. In the language of trees, it amounts to give an asymptotic expression of the number R_n of nodes of the corresponding tree.

Hanène defended her PhD on July 13.

4.2.2. Analysis of File Sharing Algorithms

Participants: Florian Simatos, Philippe Robert.

We study one of the core feature of peer-to-peer networks, namely the fact that a client offers a file that it has downloaded. This mechanism makes it possible for such networks to scale with the load, because a high demand eventually leads to a high capacity. In contrast with earlier works, our goal is to derive rigorous mathematical results on such systems, and not to rely on simulation. This choice implies to study a simplified version of peer-to-peer networks, where the focus is set on few specific features.

Initially, in our model, N peers are interested in downloading a file, and each one of them initiates a random counter. When a counter expires, the corresponding peer enters the system and is queued at the less loaded server, where it is served according to the Processor Sharing discipline. Upon completion of its download, the peer becomes a server in turn.

Since we do not consider external arrivals of new peers, two different regimes exist : in the first one, the capacity offered by the system is very small compared to the demand, leading to a highly overloaded system. But eventually, many servers are created, and the situation is reversed. We have looked at the duration of the first regime : It turns out that a good indicator to know whether the system is in the first regime or not is to investigate the number of arrivals between the creation of two successive servers. The system is in the first regime as long as this number is high, and so it is natural to look at the first time when only few peers arrive in such a time interval.

This problem can be cast in terms of urns and balls in the following way. Let T_i be the time at which the ith server is created, and A_j the time at which the jth peer arrives. Then the random point process $(T_i, i \geq 1)$ divides the half real line $\mathbb{R}^+$ into random intervals, such that each point A_j falls in some interval. If one sees the points A_j as balls, and the intervals (T_{i-1}, T_i) as random urns, the problem becomes to know the index of the first empty urn in an urns and balls problem with infinitely many random urns. Considering this problem, many interesting phenomena appear. First, it turns out that the number of empty urns in some proper range converges in distribution to a Poisson random variable with a random intensity. Moreover, when this intensity is not integrable, the behavior of the average of this random variable is dictated by rare events, that we have been able to explicitly write down.

Possible directions for further work on this subject include the generalization of the urns and balls problem to more general distributions, and a study of the second regime of the peer-to-peer system.

4.3. Large Stochastic Networks

4.3.1. Stability Properties of Loss Networks

Participants: Nelson Antunes, Christine Fricker, Philippe Robert, Danielle Tibi.

A new class of stochastic networks has been introduced and analyzed. Their dynamics combine the key characteristics of the two main classes of queueing networks : loss networks and Jackson type networks.

1. Each node of the network has finite capacity so that a request entering a saturated node is rejected as in a loss network.
2. Requests visit a subset of nodes along some (possibly) random route as in Jackson or Kelly's networks.

This class of networks is motivated by the mathematical representation of cellular wireless networks. Such a network is a group of base stations covering some geographical area. The area where *mobile users* communicate with *a base station* is referred to as *a cell*. A base station is responsible for the bandwidth management concerning mobiles in its cell. New calls are initiated in cells and calls are handed over (transferred) to the corresponding neighboring cell when mobiles move through the network. A new or a handoff call is accepted if there is available bandwidth in the cell, otherwise, it is rejected.

Paper [7] has been accepted in the Annals of Probability.

4.3.2. Large networks with mean-field dynamics

Participant: Danielle Tibi.

Several examples of networks consisting of N exchangeable nodes exhibit a limiting dynamics, as N tends to infinity, given by a “master equation" : $\dot{y} = L_y^* y$ where y is the empirical distribution of the N nodes and (L_y) is a y-indexed set of reversible infinitesimal generators. This means that the evolution of a typical node at time t in the (instantaneous) global environment $y(t)$ is (approximately for N large) that of a Markov process with generator $L_{y(t)}$. This conveys the mean-field property. (For these networks, the state of one individual node evolves together with the global state of the network, as represented by $y(t)$, on the same normal time scale.) Our objective is to associate to such a dynamics an energy function (or Lyapunov function), permitting to classify the equilibrium states as stable/unstable (then obtaining a feedback on the original random network stationary distribution for large N). Such a function was obtained in Antunes *et al.* [7] ; it involved the relative entropy $H(y/\pi_y)$ of y with respect to the L_y-invariant probability π_y. This result can be straightforwardly generalized only to a very particular class of (L_y)'s. The general problem could be related to minimizing $y' \mapsto H(y/y') + G(y')$ for fixed y on some variety V and for some well chosen G. In this direction, we proved (by a very simple argument) a strict convexity property for $y' \mapsto H(y/y')$ when y' is represented in some multi-parameter Gibbs form adapted to the “model" (L_y). It then gives the existence and uniqueness of a minimum for $y' \mapsto H(y/y') + G(y')$ on any affine variety V for proper G's. (Incidentally, we also obtain a nice “geometrical" property for the relative entropy.) The function $y \mapsto H(y/y'_{min}) + G(y'_{min})$ would then be a Lyapunov function if we could ensure (by a proper choice of G and V) that the minimum y'_{min} is on the same side as π_y of some y-depending hyperplane . (In the example of Antunes *et al.* [7]) y'_{min} is just equal to π_y and V is the “natural" set of measures π_y.)

4.3.3. Fluid limits for a processor sharing network with mobile users

Participants: Florian Simatos, Danielle Tibi.

We determine the stability region for a network consisting of a finite set of nodes operating as processor sharing servers, between which users follow Markovian routes during their service. A fluid limit approach requires to get estimates on certain hitting times, especially, in this case, the time when the relative occupancies of the different nodes get close (or deviate from) the equilibrium distribution of the Markovian routing. This is obtained through the construction of a martingale associated to the multidimensional Markov process describing the network. Beyond determining the stability region we get a precise description of how transience occurs.

5. Contracts and Grants with Industry

5.1. Contracts and Grants with Industry

Participation to the CRE with France Telecom “Mathematics of Internet Measurements”. Two years contract starting from 2005.

Participation to the RNRT project “OSCAR” on the attack detection in the Internet. Two years contract starting from April 2006.

Participation to the ACI Masse de données “FLUX” on the probabilistic counting methods of large data sets occurring in traffic measurements, biological sequences, dictionaries. Participants : INRIA (Algo project), INRIA (Rap project) and University of Montpellier. Three years contract starting from 2004.

Participation to the ANR Projet Blanc “SADA” on the 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.

6. Other Grants and Activities

6.1. Visiting scientists

RAP team has received the following people :
Sindo Nunez-Queija (CWI, Amsterdam), Matthieu Jonckheere (CWI, Amsterdam), Nelly Litvak (University of Twente), Sem Borst (University of Eindhoven and Lucent), Djalil Chafai (Université de Toulouse), Andreas Löpker (University of Eindhoven).

7. Dissemination

7.1. Leadership within scientific community

Philippe Robert is the Chairman of the Project Committee of INRIA-Rocquencourt.

Philippe Robert is Associate Editor of the Book Series “Mathématiques et Applications” edited by Springer Verlag.

Philippe Robert is associate Professor at the École Polytechnique in the department of applied mathematics. He is in charge of lectures on mathematical modeling of networks.

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 “Majeure de Mathématiques Appliquées et d’Informatique” on Networks and Algorithms at the École Polytechnique. He gave lectures on algorithms at the University of Casablanca from April 10 to 13.

7.3. Conference and workshop committees, invited conferences

Yousra Chabchoub gave a talk at the ITC’20 Conference (Ottawa) June 17/21.

Yousra Chabchoub and *Philippe Robert* were at the RNRT project Oscar meeting at Lyon in February 1 and 2

Christine Fricker , *Hanène Mohamed* , *Philippe Robert* and *FLorian Simatos* gave talks at Informs conference (Eindhoven) July 9/11.

Christine Fricker was from April 18th to 21th at the FranceTelecom R&D center at Lannion.

FLorian Simatos gave talks at France Telecom February 19/20, ALEA (Luminy) March 19/23, University of Nice April 23/24, CWI November 19/21 and September 3/4 .

Philippe Robert gave talks at the University of Casablanca April 10/13, the University of Nice May 23/24, IST Conference in Nancy June 19, the University of Lisbonne December 4/7.

8. Bibliography

Major publications by the team in recent years

[1] J.-F. DANTZER, P. ROBERT. *Fluid limits of string valued Markov processes*, in "Annals of Applied Probability", vol. 12, n° 3, 2002, p. 860–889.

[2] V. DUMAS, F. GUILLEMIN, P. ROBERT. *A Markovian analysis of Additive-Increase Multiplicative-Decrease (AIMD) algorithms*, in "Advances in Applied Probability", vol. 34, n° 1, 2002, p. 85–111.

[3] F. GUILLEMIN, P. ROBERT, B. ZWART. *AIMD algorithms and exponential functionals*, in "Annals of Applied Probability", vol. 14, n° 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° 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] H. MOHAMED. *Probabilistic Analysis of Tree Algorithms*, thèse de doctorat, Paris VI University, Probability Laboratory, July 2007.

Articles in refereed journals and book chapters

[7] N. ANTUNES, C. FRICKER, P. ROBERT, D. TIBI. *Stochastic Networks with Multiple Stable Points*, in "Annals of Probability", 2007, http://www.arxiv.org/abs/math.PR/0601296.

[8] C. FRICKER, R. JAIBI. *On the fluid limit of the M/G/infinite server queue*, in "Queueing Systems, Theory and Applications", vol. 56, 2007, p. 255–265.

[9] C. GROMOLL, P. ROBERT, B. ZWART. *Fluid Limits for Processor Sharing Queues with Impatience*, in "Mathematics of Operation Research", To Appear, 2007.

Publications in Conferences and Workshops

[10] Y. CHABCHOUB, C. FRICKER, F. GUILLEMIN, P. ROBERT. *Deterministic Versus Probabilistic Packet Sampling in the Internet*, in "Managing Traffic Performance in Converged Networks, 20th International Teletraffic Congress, ITC20 2007, Ottawa, Canada, June 17-21, 2007, Proceedings", L. MASON, T. DRWIEGA, J. YAN (editors), Lecture Notes in Computer Science, vol. 4516, Springer, 2007, p. 678–689, http://dx.doi.org/10.1007/978-3-540-72990-7_60.

Miscellaneous

[11] Y. CHABCHOUB, C. FRICKER, F. GUILLEMIN, P. ROBERT. *A Study of Flow Statistics of IP Traffic with Application to Sampling*, Preprint, July 2007, http://www-rocq.inria.fr/~robert/src/papers/2007-3.pdf.

[12] C. FRICKER, F. GUILLEMIN, P. ROBERT. *Perturbation analysis of an $M/M/1$ queue in a diffusion random environment*, Preprint, April 2007, http://www-rocq.inria.fr/~robert/src/papers/2007-2.pdf.

[13] F. GUILLEMIN, P. ROBERT. *On the Discovery of Random Trees*, February 2007, http://www-rocq.inria.fr/~robert/src/papers/2007-1.pdf.

[14] F. SIMATOS. *Mathematical analysis of the recoil growth algorithm to generate multi-polymer systems*, Preprint, August 2007, http://arxiv.org/abs/0708.1116v1.

References in notes

[15] M. BRAMSON. *Instability of FIFO queueing networks*, in "Annals of Applied Probability", vol. 4, n^o 2, 1994, p. 414–431.

[16] F. COMETS. *Limites hydrodynamiques*, in "Astérisque", Séminaire Bourbaki, Vol. 1990/91, n^o 201-203, 1991, p. Exp. No. 735, 167–192 (1992).

[17] N. DUFFIELD, C. LUND, M. THORUP. *Properties and Prediction of Flow Statistics Properties and Prediction of Flow Statistics*, in "ACM SIGCOMM Internet Measurement Workshop", November 2002, p. 6–8.

[18] N. DUFFIELD, C. LUND, M. THORUP. *Estimating Flow Distributions from sampled Flow statistics*, in "SIGCOMM", 2003, p. 25–29.

[19] 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.

[20] N. HOHN, D. VEITCH. *Inverting sampled traffic*, in "IMC", October 2003, p. 27–29.

[21] V. PAXSON, S. FLOYD. *Wide area traffic : The failure of the Poisson assumption*, in "IEEE/ACM Trans. on Networking", 1995, p. 226-244.

[22] 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.