INRIA

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team rap

Réseaux, Algorithmes et Probabilités

Paris - Rocquencourt

THEME COM

Activity Report

2008

Table of contents

1. Team

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

Faculty Member
Danielle Tibi [Délégation de l'Université Paris VII until August]

PhD Student
Yousra Chabchoub [INRIA]
Florian Simatos [Détachement du Corps des Télécoms]

Administrative Assistant
Virginie Collette [(AI) Inria]

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 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 networks 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 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 aim 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) [30]. 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 chosen 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 for 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 [25], [29] and references therein). 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 [29]. An alternative approach consists of performing a statistical analysis of flow as in [25], [26]. 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 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.

3.2.1. 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*.

Because of the growing complexity of communication networks, the classical mathematical methods of investigation 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 lie 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 Guillemin *et al.* [28], 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 [32] for example. In statistical physics, these methods are quite classical, see Comets [23].
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 [22] 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 [24] derives results in this direction. See also the corresponding chapter of Robert [31].

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 [24].
- Develop techniques to investigate metastability phenomena observed in some models of networks in the scaling limit due to mean field approach. See Kelly [27].

4. New Results

4.1. Algorithms : On-line Algorithms for Traffic Measurements

Participants: Yousra Chabchoub, Christine Fricker, Fabrice Guillemin, Philippe Robert, Danielle Tibi.

4.1.1. A counting algorithm based on Bloom filters

The aim of the study is to design a probabilistic 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 packet, the IP address is sent via independent hashing functions to filters. The idea is that, if the cases of the filters have counters, while each packet of one flow is hashed to the same cases, the long flows (more than 20 packets) could be approximately counted. For that, the filters must be cleared for time to time, in order to avoid accumulation of short flows, very numerous, which pollute the filters. Contrary to the original scheme devised by Estan and Varghese where filters are cleared periodically independently of the traffic, the refreshment proposed in our algorithm occurs every time the filling rate of the filter reaches some threshold r.

In practice, the algorithm has been tested on a wide set of real traces : commercial Orange traces and academic Abilene traces. The algorithm used for elephants identification was adapted to a quite different domain, attacks detection. We are particularly interested in SYN flooding and volume flooding which are the most common DoS (Denial of Service) attacks. Using a convenient definition of a flow, these attacks can be considered as big flows. However, in the context of attacks detection, the refreshing mechanism of the multistage filter has a new purpose. It should eliminate quickly all normal flows, for this reason a more aggressive refreshing method has to be used.

Moreover attacks are always defined as a notable deviation from a standard behavior. To describe the standard traffic, the main idea of the algorithm is to evaluate a varying average denoted by mn of the largest flow in several sliding time windows. The quantity mn is periodically actualized in order to adapt to varying traffic conditions. It is a weighted average that takes into account all its past values to follow carefully traffic variations. A flow is considered as an attack if it deviates significantly from the actual varying average.

The algorithm was tested against several traces from France Telecom IP backbone network. The alarms detected by the algorithm coincide very often with anomalous behaviors. A distributed version of the algorithm at the level of a network has been proposed.

4.1.2. Analytical Study of the Algorithm

This work [11] is a collaboration with Frédéric Meunier (ENPC). The model has been simplified in order to evaluate the impact of refreshment on the proportion of short flows detected as long flows, with more than say C packets (false positives). When short flows have just one packet, a simplified model has been analyzed. Limit theorems of the empirical distribution of the filter have been obtained when the filter size m is large. The limit is deterministic and has a nice interpretation in terms of queues. They are completely proved for C=2, where a Lyapounov function is exhibited. Its existence is conjectured for $C > 2$. The proof is based on the convergence to a dynamical system and widely uses a queueing interpretation of the fixed point of the system.

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 high enough 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.

Work in progress : The previous model concerns one hashing function or several hashing functions but incrementing all the hit counters. A further study is focused on the fact that incrementing the lowest counters can maintain the highest counters much lower. Such an algorithm should have better performance. This question is related to the so-called power of choices in Bloom filters or in queues which has attracted much attention recently. Results have been obtained and several open problems are being investigated. An analysis of this version of the algorithm has been proposed. For the moment, it suggests a new version of the algorithm, where just one counter is incremented at each packet. Such a version seems to have good performances, moreover its analysis seems tractable.

4.2. Algorithms : Analysis of Splitting Algorithms

Participant: Philippe Robert.

This is a collaboration with Hanène Mohamed (Université de Versailles-Saint-Quentin). 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.

A general tree algorithm processing a random flow of arrivals has been analyzed. Capetanakis-Tsybakov-Mikhailov's protocol in the context of communication networks with random access is an example of such an algorithm. In computer science, this corresponds to a trie structure with a dynamic input. Mathematically, it is related to a stopped branching process with exogeneous arrivals (immigration). Under quite general assumptions on the distribution of the number of arrivals and on the branching procedure, it is shown that there exists a *positive* constant λ_c so that if the arrival rate is smaller than λ_c, then the algorithm is stable under the flow of requests, i.e., that the total size of an associated tree is integrable. At the same time a gap in the earlier proofs of stability of the literature is fixed. When the arrivals are Poisson, an explicit characterization of λ_c is given. Under the stability condition, the asymptotic behavior of the average size of a tree starting with a large number of individuals is analyzed. The results are obtained with the help of a probabilistic rewriting of the functional equations describing the dynamic of the system. The proofs use extensively this stochastic background throughout the paper. In this analysis, two basic limit theorems play a key role : the renewal theorem and the convergence to equilibrium of an auto-regressive process with moving average. See Mohamed and Robert [19].

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

Participants: Fabrice Guillemin, Philippe Robert, Florian Simatos.

4.3.1. Transient behavior

Peer-to-peer networks are complex systems with a variety of purposes and behaviors. Although their impact on real-world networks is important (it amounts to 80% of the traffic in some cases), their mathematical modeling is still unsatisfactory. In [14], we looked at a special case of such networks.

We were interested in the initialization phase, when a popular file has just appeared so that many customers rush to try and download it. We more precisely investigated the time needed for the network to cope with the high demand. The network indeed starts highly overloaded since many customers want to download the file and it has just appeared ; but after a sufficiently long period of time, the situation will be reversed since only few customers will still want the file, and it will be offered by the numerous customers of have previously downloaded it. The problem was to estimate the time at which the network shifts between these two extreme situations.

Technically, the problem of estimating this time turned out to be equivalent to a bins and balls problem in random environment : given a *random* probability distribution (P_i) such that, conditionally on (P_i), N balls are thrown independently, each into bin i with probability P_i, what is the index of the first empty bin ? We analyzed the annealed version of this problem in [14], and devoted [20] to the more complex case of the quenched version. By comparing the two different results, we argued why the answer of by the annealed problem was the correct one, thus giving a good insight into the behavior of the network.

4.3.2. Bins and balls in random environment

Although the annealed case of the above bins and balls problem could be analyzed with standard techniques, the full treatment of the quenched case in [20] was more involved. Bins and balls problems with a deterministic probability distribution (p_i) have already been extensively studied in the literature, but the instance of this problem in a random environment is fairly new. It has been considered in the particular "stick-breaking" case, i.e., when $P_i = W_1 \cdots W_{i-1}(1 - W_i)$ and (W_i) are i.i.d. and each lies in $(0, 1)$. Because of the peer-to-peer motivation, we studied a different case : if (t_i) is the sequence of split times of a Yule process, then $P_i = e^{-\rho t_{i-1}} - e^{-\rho t_i}$ where $\rho > 0$ is a parameter of the problem.

For our results for bins and balls problems, especially in random environment, *global* quantities, such as the number of bins receiving at least one ball were investigated. Here we gave a *geometric* description of the locations of empty bins. More precisely, if the point process $\mathbb{N}_n$ has Dirac masses at suitably rescaled indices of empty bins when n balls are thrown :

$$\mathbb{N}_n = \left\{ \frac{i}{n^{1/(\rho+2)}} : \text{ bin } i \text{ empty when } n \text{ balls are thrown} \right\},$$

then we proved weak convergence of this sequence of point processes to an original point process that we explicitly characterized.

This result shows at the same time the originality of the methods employed : point processes make it possible to describe the landscape of empty bins both concisely and precisely.

Complementary to our first work on peer-to-peer networks where we have looked at the initialization phase, we are currently investigating a model for peer-to-peer networks with stationary arrivals. To avoid the combinatorial difficulties inherent to file sharing peer-to-peer networks, we decided to focus on a special case where peers download the chunks constituting the file sequentially. This gives raise to a nice system of tandem queues, and we aim at describing the stability condition, which has no closed-form expression, as precisely as possible. The case where two queues are in tandem has already attracted the attention of the community, but no formal results have been proved so far.

4.4. Scaling Methods : A Stability Analysis of Mobile networks

Participants: Florian Simatos, Danielle Tibi.

Mobile networks are a new class of stochastic networks which have emerged recently impelled by new technologies, such as web surfing on PDAs. The main originality of these models is that, in contrast with classical queueing networks like Jackson networks for instance, users in mobile networks move from one cell to another independently of the service they receive ; the fixed capacity of each cell is thus shared among the users present at any given time.

For a particular stochastic model, we established in [21] the stability condition by using scaling techniques of Markov processes. The main contribution of this work consisted in applying these techniques in an unusual framework. Indeed, these techniques are well suited when the sequence of rescaled processes is tight. But because of the co-existence of two different time-scales, one for the movements of customers within the network and one for the total number of customers, stochastic models for mobile networks are not amenable to this approach. Rather, we had to control the distance between, on one hand, the empirical measure which accounts for the distribution of customers across the network, and on the other hand a measure that describes the homogenization of the system. In the transient case, a striking phenomenon occurs : almost surely, no deviation from the homogenization measure ever occurs when the initial number of customers gets large.

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 an intense activity for some time now. For data transmission, the Internet network 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. In the same way as for the TCP algorithm, each of the connections 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 non-linear stochastic differential equation 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 [18].

4.6. Miscellaneous

Participants: Christine Fricker, Florian Simatos.

4.6.1. The infinite server queue with varying arrival rate

(C. Fricker). A stochastic model of sales on a period can be viewed as the area under the number of customers during this time period in an infinite server queue with varying arrival rate. The transient behavior of the target quantity for a queue with general service time is obtained by the representation of the quantity as a functional of a Poisson process in dimension 2. Our proof of the heavy traffic limit theorems for the $M/G/\infty$ queue can be extended to the varying arrival rate case.

4.6.2. Generation of multi-polymer systems

(F. Simatos). Consta *et al.* proposed in 1999 an efficient algorithm to approximately sample from a multi-polymer system according to an arbitrary probability distribution. The state of a d-dimensional multi-polymer system is a collection of self-avoiding walks on the torus in dimension d, with the additional constraint that no two polymers can intersect each other. Building on their work, we defined and analyzed a variant of this algorithm in [15]. We moreover proved some properties concerning the irreducibility of the Metropolis algorithm, which apply to the original algorithm and were not known.

5. Contracts and Grants with Industry

5.1. Contracts

Participation to the RNRT project “OSCAR” on the attack detection in the Internet. Two years contract starting from April 2006. Participants : ENS Lyon, Get, INRIA, France Telecom, Laas, LIP6.

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 :

Balakrishna Prabhu and Sindo Nunez-Queija (CWI, Amsterdam), Nelly Litvak (University of Twente), Ilka Norros and Hannu Reittu (VTT, Finland), Ruth Williams (UCSD), Lasse Leskela (University of Helsinki), Raouf Jaïbi (University of Tunis)

Maaike Verloop, a PhD student of CWI (Amsterdam) visited our group from October 15 to December 12.

7. Dissemination

7.1. Leadership within scientific community

Philippe Robert was the Chairman of the Project Committee of INRIA Paris — Rocquencourt until September 1. He is Associate Editor of the Book Series “Mathématiques et Applications” edited by Springer Verlag. 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.

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

Yousra Chabchoub gave a talk at ENS on attack detection (October).

Christine Fricker gave a talk at the “Discrete Mathematics and Theoretical Computer Science Conference” on September 22-26th, 2008 in Blaubeuren, Germany. She visited Orange R&D labs at Lannion from August 19th to 21th.

Philippe Robert gave talks at Lycée Camille See (Paris) in February, LIAFA (University Paris VII) in February, Dynamo Conference (Reykjavic) in July, the University of Santiago, Chile, in July, Journees MAS in Rennes, August 27-29, IECN (Nancy) in October.

Florian Simatos gave talks at ENS, Paris (February), ACM Sigmetrics Conference, Annapolis, USA (June), The seventh world congress in statistics and probability, Singapore, (July), Fifth Colloquium on Mathematics and Computer Science, Blaubeuren, Germany, (September), Young European Queueing Theorists II Workshop, Eindhoven, Netherlands (December).

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

Articles in International Peer-Reviewed Journal

[6] 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, http://www.arxiv.org/abs/math.PR/0601296.

[7] Y. CHABCHOUB, C. FRICKER, F. GUILLEMIN, P. ROBERT. *Inference of flow statistics via packet sampling in the internet*, in "IEEE Communications Letters", To Appear, 2008, http://fr.arxiv.org/abs/0812.2544.

[8] 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", To Appear, 2008, http://fr.arxiv.org/abs/0812.2543.

[9] C. GROMOLL, P. ROBERT, B. ZWART. *Fluid Limits for Processor Sharing Queues with Impatience*, in "Mathematics of Operations Research", vol. 33, n^{o} 2, May 2008, p. 375–402, http://arxiv.org/abs/math/0604215.

[10] F. GUILLEMIN, P. ROBERT. *Analysis of Steiner subtrees of Random Trees for Traceroute Algorithms*, in "Random Structures and Algorithms", To appear, 2008, http://www.arxiv.org/abs/cs.NI/0702156.

International Peer-Reviewed Conference/Proceedings

[11] Y. CHABCHOUB, C. FRICKER, F. MEUNIER, D. TIBI. *Analysis of an algorithm catching elephants on the Internet*, in "Fifth Colloquium on Mathematics and Computer Science", DMTCS Proceedings Series, september 2008, p. 299-314.

[12] C. GRAHAM, P. ROBERT. *A multi-class mean-field model with graph structure for TCP flows*, in "Proceedings of European Consortium For Mathematics In Industry 2008, London", Springer, June 2008.

[13] N. LITVAK, P. ROBERT. *Analysis of an on-line algorithm for solving large Markov chains*, in "The 3rd International Workshop on Tools for solving Structured Markov Chains, Athens", ACM, October 2008.

[14] F. SIMATOS, P. ROBERT, F. GUILLEMIN. *A Queueing System for Modeling a File Sharing Principle*, in "ACM-Sigmetrics, Annapolis", ACM/IFIP WG 7.3, June 2008, http://fr.arxiv.org/abs/0803.1610.

[15] F. SIMATOS. *A variant of the Recoil-Growth algorithm to generate multipolymer systems*, in "Proceedings of the Fifth Colloquium on Mathematics and Computer Science, Blaubeuren", 2008.

Other Publications

[16] Y. CHABCHOUB, C. FRICKER, F. GUILLEMIN, P. ROBERT. *A Robust Statistical Estimation of Internet Traffic*, Preprint, June 2008, http://hal.inria.fr/inria-00347011/fr/.

[17] C. FRICKER, F. GUILLEMIN, P. ROBERT. *A statistics inference problem in an urn and ball model with heavy tailed distribution*, Preprint, 2008, http://fr.arxiv.org/abs/0812.2546.

[18] C. GRAHAM, P. ROBERT. *Interacting Multi-Class Transmissions in Large Stochastic Networks*, Preprint, October 2008, http://arxiv.org/abs/0810.0347.

[19] H. MOHAMED, P. ROBERT. *Dynamic Tree Algorithms*, Preprint, September 2008, http://arxiv.org/abs/0809.3577.

[20] P. ROBERT, F. SIMATOS. *Occupancy Schemes Associated to Yule Processes*, Preprint, October 2008, http://arxiv.org/abs/0810.3079.

[21] F. SIMATOS, D. TIBI. *Spatial homogenization in a stochastic network with mobility*, Preprint, 2008.

References in notes

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

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

[24] 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.

[25] 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.

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

[27] 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.

[28] 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.

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

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

[31] P. ROBERT. *Stochastic Networks and Queues*, Stochastic Modelling and Applied Probability Series, vol. 52, Springer, New-York, June 2003.

[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.