INRIA, Evaluation of Theme Networks and Telecommunications

Project-team RESO

March 21, 2012

Project-team title: RESO: Optimized Protocols and Softwares for High-Performance Networks

Scientific leader: Paulo Gonçalves

Research center: Grenoble Rhône-Alpes

Common project-team with: Laboratoire de l'Informatique du Parallélisme, CNRS, ENS Lyon and the UCB Lyon 1 (UMR CNRS - ENS Lyon - UCB Lyon 1 - INRIA 5668)

1 Personnel

Personnel (March 27, 2008)

	Misc.	INRIA	CNRS	University	**Total**
DR (1) / Professors		1		1	**2**
CR (2) / Assistant Professors		2		2	**4**
Permanent Engineers (3)		0,5	0,5	0,5	**1,5**
Temporary Engineers (4)			4		**4**
PhD Students	2	2		4	**8**
Post-Doc.	1				**1**
Total	**3**	**5,5**	**4,5**	**7,5**	**20,5**
External Collaborators					
Visitors ($>$ 1 month)	1				**1**

(1) "Senior Research Scientist (Directeur de Recherche)"
(2) "Junior Research Scientist (Chargé de Recherche)"
(3) "Civil servant (CNRS, INRIA, ...)"
(4) "Associated with a contract (Ingénieur Expert or Ingénieur Associé)"

Personnel (March 21, 2012)

	Misc.	INRIA	CNRS	University	**Total**
DR / Professors				1	**1**
CR / Assistant Professor		2		3	**5**
Permanent Engineer		0,5	0,5		**1**
Temporary Engineer		3			**3**
PhD Students	3	3		1	**7**
Post-Doc.	2	1			**3**
Total	**5**	**9,5**	**0,5**	**5**	**20**
External Collaborators					
Visitors ($>$ 1 month)					

Changes in staff

DR / Professors CR / Assistant Professors	Misc.	INRIA	CNRS	University	total
Arrival				1	1
Leaving		1			1

Comments: Since June 2010, Pascale Vicat-Blanc (DR1 INRIA), former head of the RESO team, is on leave of absence to create and to lead her start-up Lyatiss.

Current composition of the project-team (March 21, 2012):

- Permanent personnel
  - Paulo Gonçalves, Research Associate (CR1) INRIA [Habilité], team leader
  - Laurent Lefèvre, Research Associate (CR1) INRIA
  - Isabelle Guérin-Lassous, Professor University of Lyon (UCBL) [Habilitée]
  - Thomas Begin, Associate Professor University of Lyon (UCBL)
  - Jean-Patrick Gelas, Associate Professor University of Lyon (UCBL)
  - Olivier Glück, Associate Professor University of Lyon (UCBL)
  - Mathieu Imbert, Permanent Research Engineer INRIA – 50%
  - Jean-Christophe Mignot, Permanent Research Engineer CNRS – 50%
  - Laetitia Lecot, Assistant, INRIA – 50%
- PhD students
  - Thiago Abreu, ENS Lyon, RESCUE grant, [2011-2014]
  - Doreid Ammar, INRIA Alcatel Lucent common Lab. [2009-2012]
  - Mohammed Diouri, UCBL scholarship, MENRT [2010-2013]
  - Ines Doghri, PhD student, with Orange CIFRE [2009-2012]
  - Van Dan Nguyen, (co-advisement with Univ. of Hanoi), AUF fellowship [2009-2012]
  - Shubhabrata Roy, INRIA SAIL grant [2010-2013]
  - Marina Sokol, INRIA Alcatel Bell Labs joint Lab. [2009-2012] (with RESO until Oct. 2011, now with MAESTRO)
  - Ghislain Landry Tsafack, INRIA HEMERA grant [2011-2014]
- Post-Doctoral fellows
  - Ashley Chonka, ENS Lyon, PETAFLOW grant [2011-2012]
  - Mohamad Jaber, INRIA Alcatel Bell Labs joint Lab. [2011-2012]

  - Anne-Cécile Orgerie, ENS Lyon, PETAFLOW grant [2011-2012]

- Contract Expert Engineers

  - Julien Carpentier, INRIA COMPATIBLE ONE grant [2012]
  - Maxime Morel, INRIA COMPATIBLE ONE grant [2011-2012]
  - Olivier Mornard, HIPCAL [2008-2009], AUTONOMIC INTERNET [2009-2010], CNRS [2010-2011], COMPATIBLE ONE [2011-2012]

Current position of former project-team members (including PhD students during the 2008-2011 period):

- Permanent personnel

  - Pascale Vicat-Blanc (DR1 INRIA, until June 2010): CEO of Lyatiss

- PhD students

  - Fabienne Anhalt (INRIA GEYSERS, 2008-2011): Engineer Lyatiss
  - Narjess Ayari (France Telecom / INRIA): Engineer Orange Labs
  - Dinil Mon Divakaran (INRIA Alcatel Lucent joint Lab, 2007-2010): Assistant Professor Indian Institute of Technology Mandi
  - Pierre-Solen Guichard (INRIA Alcatel Lucent joint Lab, 2008-2010, PhD aborted in 2nd year): Engineer EDF
  - Romaric Guillier (INRIA CARRIOCAS 2006-2009, Engineer INRIA HIPCAL & SAIL, 2009-2010): Engineer Lyatiss
  - Ludovic Hablot, (ENS Lyon, UCBL scholarship, 2006-2009), Engineer Service Society
  - Guilherme Koslovski, (INRIA SAIL, 2008-2011): Engineer (self-entrepreneur) Brazil
  - Patrick Loiseau, (ENS Lyon scholarship, 2006-2009), Assistant Professor Eurecom
  - Dino Lopez Pacheco (Mexican Government grant, 2004-2008): Assistant Professor University of Nice
  - Anne-Cecile Orgerie (ENS Lyon, UCBL scholarship, 2008-2011): Post-Doc position with RESO at University of Melbourne
  - Sébastien Soudan (ENS Lyon, UCBL scholarship, 2006-2009): Engineer Lyatiss
  - Rémi Vannier, (INRIA AEOLUS 2006-2009, ATER UCBL 2009-2010): Professor High School

- Post-Doctoral fellows and temporary teaching assistants

  - Manoj Dahal (INRIA CARRIOCAS 2008-2009): Professor at SMIT, India.
  - Marcos Dias de Assunção (INRIA GREEN-NET, 2009-2010): Engineer IBM Brazil
  - Emmanouil Dramitinos (INRIA AEOLUS 2008-2009): unknown position
  - Olivier Grémillet (INRIA Alcatel Lucent joint Lab., 2008-2009): Engineer at Viveris Technologies, Toulouse.
  - Ibrahim Mouhamad (INRIA Alcatel Lucent joint Lab, 2009): unknown position
  - Lucas Nussbaum (Teaching Assistant (ATER UCBL), 2008-2009): Assistant Professor University Nancy 2

- Contract Expert Engineers

  - Damien Ancelin (INRIA EC-GIN, 2008-2009): unknown position
  - Aurelien Cedeyn (ENS Lyon, 2008-2009): unknown position

- Abderhaman Cheniour (INRIA AUTONOMIC INTERNET, 2008-2010): unknown position
- Oana Goga (INRIA ADT ALADDIN, 2008-2009), PhD student LIP6
- Philippe Martinez (INRIA DTI, 2008-2010): Engineer Sysfera
- Marcelo Pasin (INRIA EC-GIN, 2007-2008), Assistant Professor University of Lisbon
- Augustin Ragon (INRIA OGF EUROPE, 2009-2010): Engineer Sysfera
- Armel Soro (INRIA ADT ALADDIN, 2009-2010), Engineer Lyatiss
- Sébastien Soudan (INRIA EC-GIN, 2009-2010): Engineer Lyatiss

Last INRIA enlistments

There was no INRIA researcher recruitment during the evaluation period.

Other comments:

After P. Vicat-Blanc's departure in june 2010, P. Gonçalves replaced her as scientific leader of the project-team RESO.

2 Work progress

2.1 Keywords

Virtualization, clouds, autonomic and programmable networks, energy efficiency, dynamic network management, probabilistic resource provisioning, Quality of Service, performance evaluation, traffic and queueing models, large scale experiments, high performance computing, communication Middleware, wireless networks.

2.2 Context and overall goal of the project

The INRIA team-project RESO has been created in december 1st, 2003 and was leaded by Pascale Vicat-Blanc until june 1, 2010 and by Paulo Gonçalves since then. In addition to its INRIA affiliation, RESO is also a research team at the "Laboratoire de l'Informatique du Parallélisme" (LIP) - Unité Mixte de Recherche (UMR) between CNRS, INRIA, ENS Lyon and Université Claude Bernard of Lyon 1. Presently, RESO comprises 6 permanent researchers and its overall size and composition (viz. permanent / non permanent ratio) is roughly the same as it was at the end of 2007 (although it reached its highest peak in 2008-2009 with 37 people).

RESO is focusing on resource management and communication protocols in the context of high speed networks and virtualized infrastructures (clouds). We apply our results to the domain of high demanding applications and Future Internet.

Objectives. As the use of networks for on-demand computing is gaining in the large Internet, new bandwidth and network resource sharing paradigms are necessary to enable a customized traffic delivery that is dependent of the originating application. The coordination of networking, computing and storage capacities requires the design, development and deployment of new resource management mechanisms able to autonomously identify, reserve, co-allocate and reconfigure the available resources. Eventually, this adaptive administration should turn possible to control the allocation of shared resources to meet the volatility of the applications' demand while rationalizing their usage in terms of energy efficiency. All this, combined with the constantly increasing dimensions of interconnected

networks, prompts the emergence of probabilistic approaches whenever deterministic or static methods fail at facing scalability.

The core of RESO's research activity lies in the development of such networking service, with special emphasis on:

- application-wise Quality of Service,
- traffic-aware transport protocols,
- energy consumption efficiency,
- network metrology.

All these interacting aspects have in common the same target, which is that of an efficient and dynamic network resource provisioning and scheduling.

One of the key challenge in this direction, is the provisioning of a secure, flexible, transparent and adapted transport infrastructure for high performance data access and processing. The resulting concept of *Network As a Service* aims at proposing context-aware solutions that autonomously adapt the allocated IT infrastructures to sustain the current needs or constraints of the Future Internet. This is one of the main motivations for RESO's outcomes.

Methodology. The RESO approach relies on the theoretical and practical analysis of the limitations of current network practices, to motivate the exploration of new paradigms that show more adapted to the emerging infrastructures. This research framework at the interface of a specific network context and a challenging application domain, induces a close interaction between these two layers (viz. network and application). Our methodology is based on a deep evaluation of the functionalities and performance of high speed infrastructures. We study the high end and the original requirements before we design and we evaluate the proposed solutions.
Schematically, our approach is organized into four major research axes :

Axis 1: Advanced protocol implementation and networking equipment. Since several years, virtualization of the operating system is used in end system to improve security, isolation, reliability and flexibility of the environments. These mechanisms becoming incontrovertible in large scale distributed systems, we explore how the same can also be adapted and used in data transport networks and specifically in switching and routing equipments.

Axis 2: End-to-end Quality of Service and Transport layer. The goal of this axis is to guarantee quality of service in machine/user to machine/user communication while using efficiently the resources of the future networks.

Axis 3: High-Speed Network's traffic metrology, analysis and modeling. This activity consists in measuring along time, the nature and the amount of exchanged information between the constituents of a system. It is then a matter of using the collected data to forecast the network load evolution, so as to anticipate congestion, and more widely, to guarantee a certain (probabilistic) Quality of Service, rationalizing the resources usage.

Axis 4: Application to Grid Middleware and Applications. This axis focuses on the application of specific solutions to the grid context and on their implementation in a real environment such as the national research instrument GRID'5000.

From the beginning, the main activity of RESO is anchored in communication softwares, services and protocols in the context of high performance networks, with privileged application to the Grids' environment. During the last time span though, the center of

mass of our research has sensibly moved. It was so, primarily to match the evolving needs of emergent infrastructures and applications: hence our recent focusing on network virtualization (with the noticeable development of a software suite for managing virtual infrastructures – groundwork of Lyatiss spinoff) and the reinforcement of our activity on energy efficiency in large scale distributed systems. But our research axes repositioning also occurred because the internal competence of RESO changed. This mutation started with the incorporation of P. Gonçalves in 2006 and then with the recruitment of T. Begin in 2009. Although they work in two distinct fields (signal processing and performance evaluation, respectively), they leveraged this complementarity to consolidate the RESO's component that deals with metrology and probabilistic approaches. More recently, the departure of P. Vicat-Blanc entailed the transfer to the Lyatiss start-up of RESO's knowledge on network virtualization techniques. Naturally, this shift had a significant impact on the contours of axes 1 and 4 in the last two years.

2.3 Objectives for the evaluation period

In our previous evaluation report, we proposed to pursue our research activity towards *the design of a new protocol architecture for Future Internet* and to contribute *to solve some of the critical issues faced by the current architecture.* Indeed, we believed that *the Internet re-design would raise the opportunity to better understand and to better assess higher-level system requirements that could then be used as guidelines for the lower layer architecture.* One main reasoning of RESO was thus to claim that certain *mechanisms, usually implemented as part of applications, might conceivably migrate into the network itself.*

Because their *deployment undergo a constant spreading (thanks e.g. to the fiber-to-the-home) and also because their technical specificities are considerably evolving*, we decided *to keep our focus on high speed networks* and to maintain *a methodological approach combining theory and large scale experiments on real testbeds.* As an evolution with respect to the previous period, we also foresaw *to emphasize the theoretical and the upstream aspects of our activity adopting a clean slate thinking as advocated by the network community.* Regarding applications, we deemed wise *to leverage our expertise and our long-term involvement in the* GRID'5000 *platform to continue investigating grid middleware and applications.*

To be more precise on our objectives, we decomposed our line of thought into concrete targets aimed at addressing a series of challenges relating to each of our four research axes.

2.3.1 Axis 1: Advanced protocol implementation and networking equipment

In this direction, we proposed to address the following main questions:

- *Where and how do we need to integrate the required autonomy to manage and to control high speed networks at large scale?*
- *How can we obtain both high performance and efficient communications along with controlled energy consumption and reliability in virtualized infrastructures context?*

To tackle the first issue, we envisioned to *explore the concept of "dynamic programming enablers" for dynamic service driven configuration of communication resources.* While we planed to start by *studying the mechanisms and infrastructures required to support these elements*, our longer term ambition was really to *provide services that use Internet facilities, with new functionalities able to handle self-management operations in differentiated and integrated services.* In this perspective of integrating context-awareness in Internet services and infrastructures, our interest focused on two particular concerns:

– *Session awareness: service-aware dependable systems are required to warrant client and server transparency, low cost during failure-free periods and performance sustainability during failures.*

– *Energy awareness: to cope with energy limitations in the future Internet generation, we proposed to develop a dynamic monitoring system of energy consumption. The resulting information will play as an additional resource for distributed large-scale systems and should allow current jobs (OAR - a batch scheduler) and network (BDTS - bulk data transfers scheduling) schedulers to adapt their decisions to energy efficiency.*

Recalling the second main issue we intended to address in this axis, it is a fact that virtualizing techniques of the Operating System improve security, isolation, reliability and flexibility of the environment, and that such mechanisms should become incontrovertible in large-scale distributed systems, noticeably in data transport networks and routers. However, the inherent overhead needs to be integrated into the models used to describe the performance of virtualized system, in order to forecast their behavior. Following up our earlier study based on XEN virtual routers, we planed to continue the *experimental evaluation of virtualization impact on high-speed network performance, and to start the viability analysis of virtual infrastructures approaches regarding end systems performance.* More particularly, we evoked *network processors based on NIC-embedded virtualization designs* to help us devising *efficient solutions for the HIPerNet model relying on HIP and virtualization paradigms.*

2.3.2 Axis 2: End-to-end Quality of Service and Transport layer

To cope with the drastic evolution of networks' technologies and concepts as well as to address the changes in Internet practices, we had initiated in this axis a series of works revolving around the following main concepts: *i) dynamic bandwidth sharing and congestion control in the Future Internet and ii) control and flow management in semantic networks.* We then proposed to follow up these directions and to specifically address the questions:

- *Which congestion control and which transport protocol for high speed networks in large scale contexts (FTTH for example)?*
- *How to efficiently share and to dynamically provision a network bandwidth dedicated to computing tasks?*
- *Is the "flow-aware" approach a reliable solution to solve end to end quality of service issues raised by the very high speed Future Internet?*

Regarding the first item, we committed to *complete our contribution to the collective methodological effort towards a benchmark design for high speed transport protocol comparison.* In addition, we intended to *set up an attractive and open testbed for these studies using* GRID'5000. For the numerous protocols and congestion control mechanisms recently proposed for high-speed (optical) networks, our idea was to assess their ability at handling highly dynamic traffic loads, their versatility w.r.t different environments and applications and also, their interoperability strength.

Our second interrogation is tightly related to the multiscale nature of network resource control. Basically, we can consider three different timescales corresponding to three aggregation levels: "data scale" corresponds to the packet scale, "control scale" to that of flows and "management scale" which deals with the aggregate of flows. These three timescales

should be considered simultaneously as all have an impact on the system. While TCP behaves at the data timescale and some routing protocols (e.g. BGP and OSPF) at the management timescale, we had undertaken the development of Bulk Data Transfer Service which in turn operates at the control timescale. We thus proposed to *pursue this work on BDTS, whose goal is to schedule and to forward packet aggregates in such a way that it improves the predictability of periods of massive data set transfer.* Originality of our approach stemmed from the exploitation of time dimension, the definition of a fine-grain user control plane and from the implementation of the virtual network resource reservation paradigm. We then believed that *exploring this type of service-oriented network resource management at large scale and in heterogeneous environments would help to test the limits of these protocol architectures and to validate the potential of alternative abstractions. It should also build better insight on fundamental issues concerning the interactions between control- and data-planes in the one hand, and between abstraction layers, session, transport and network in the other hand.*

In a different direction, but still considering dynamic bandwidth sharing, we announced our intention *to resort to Lagrangian optimization theory to propose resource allocation strategies particularly well suited to dynamic features either due to node mobility (in wireless local networks) or to the intrinsic traffic volatility.*

Finally, considering our last investigation item of this axis, end-to-end QoS, one important motivation at the origin of RESO, took a new lease with the *Semantic Networking* theme of INRIA-Alcatel Lucent Common Lab. This latter shed a new light by considering the flow abstraction on its own. In this direction, our main objective was *to question the fundamental paradigm of packet switching networks in high-speed networks and to shift towards coarser grain approaches like flow-oriented strategies. We had in mind to revisit flow admission control, flow routing, flow scheduling and interaction with transport protocols.*

2.3.3 Axis 3: High-Speed Network's traffic metrology, analysis and modeling

This research axis is the most recently introduced theme in the RESO project (2006) and the course of action we proposed for the evaluation period was to address the two following generic questions:

- *How do the traffic statistical properties really impact the Quality of Service?*
- *How to identify and to classify, in real time, transiting flows, according to a sensible typology?*

Beyond these two mere questionings though, the aim of this metrology activity was manifold: Instrumentally speaking, we had started *to conceive a metrology platform able to capture and to timestamp traffic at packet level on high-speed links.* From a statistical analysis viewpoint, we intended *to determine the effect and the limits of sampling resolution on the estimation of important traffic properties*; we intended *to work out a more complete understanding of the scale invariant properties of network traffic and of their origins*; and we also envisioned *to formalize and to quantify the impact of these diverse statistical traffic features on network performance (i.e. QoS).* Finally from a management perspective, we planed *to build upon this flow traffic characterization to come up with data-centric solutions for dynamic resource allocation, flow admission control, routing or scheduling...*

The recruitment in 2009 of T. Begin, assistant professor, brought in a new constituent in this research axis: performance evaluation and queueing theory. This expertise area was missing in RESO to get soundly based models for the arrival processes of requests

and to obtain accurate estimates of the different waiting and service times generated by our admission control processes.

2.3.4 Axis 4: Application to Grid Middleware and Applications

This transverse axis seeks at applying and testing the solutions we propose in a grid context. We then had identified three main directions to investigate during the evaluation period:

- *Efficient MPI implementation for the grid;*
- *Optical resource provisioning and optimal bandwidth sharing for the grid;*
- *Network-aware dynamic grid composition.*

Message Passing Interface standard is often used in parallel applications for communication needs. While prioritarly designed for homogeneous clusters, MPI are not so well adapted to the high-speed interconnects heterogeneity in grids nor to the WAN structures realizing inter-site communications. In a nutshell, our declared goal was *to significantly improve efficiency of MPI and performance execution of MPI applications on the grid.* The following is a list of specific challenges we have strived to address: *(i) How to maintain good performance execution when long-way communications are involved due to the distributed nature of applications? (ii) Study the impact of TCP protocol for WAN communications and its interactions with MPI applications. Subsequently, propose adjustments of TCP usage to cope with MPI applications running in grid platforms. (iii) Propose new optimized transactions at communication layers and elaborate on a framework that allows execution of a multiple tasks MPI application to take into account the grid topology and network state.*

Within the emerging context of dynamically configurable optical networks, RESO undertook *the studying and the prototyping of interactions between the different components needed to perform user-specified tasks like bandwidth reservation, path computation and network signaling.* In parallel and to integrate this technological breakthrough, we wanted to keep learning *how to configure, to control and to tune optical components with computer ends.*

In the framework of advanced self-organizing, scalable and fault-tolerant computational grids, we had proposed a global design for network-aware resource discovery adapted to large scale distributed environments. To evaluate the feasibility of the proposed solution and to asses its pros and cons, we had planed *to develop a prototype relying on an independent P2P overlay enabling reliable communication between network peers.* This prototype was also expected *to be embedded in the HIPerNET and the SRV prototypes under development at that time.*

2.4 Objective 1: Optimized protocol implementation and networking equipment

2.4.1 Personnel

J.-P. Gelas, P. Gonçalves, O. Gluck, I. Guérin-Lassous, L. Lefèvre, P. Vicat-Blanc – F. Anhalt, N. Ayari, M. Diouri, D. Divakaran, R. Guillier, G. Koslovski, A.-C. Orgerie, D. Pacheco, S. Soudan, G. Tsafack – M. Dias de Assunçao – J. Carpentier, A. Cheniour, M. Morel, O. Mornard

2.4.2 Project-team positioning

Several projects, prototypes, and commercial solutions have been proposed to provide Virtual Resources, Virtual Networks, and Virtual Infrastructures (VIs) within different contexts, in Grids and Clouds. Grids have explored decoupling services from physical infrastructure with virtualization techniques. Rezmerita et al (2006) proposed PVC, combining various Grid, P2P and VPN approaches to create instant Grids.
Existing commercial products, such as Amazon's Elastic Compute Cloud (EC2), Enomaly's Elastic Computing Platform (ECP), and GOGRID, allow users to reserve a set of resources, choose an operating system, and customize them. Other products are less configurable, like 3Tera's AppLogic where the operating system cannot be chosen. Some frameworks for Cloud orchestration, such as OpenNebula and Eucalyptus, provide tools for Infrastructure-as-a-Service provisioning. All these Cloud solutions are suitable for performing computation and storage but are not aimed to host virtual routers, nor they are intended to control the network or its communication performance.
Some authors have explored virtualization for enabling end-to-end control while using physical equipments from different physical infrastructure providers. Among other solutions, we will retain VINI (Bavier et al., 2006), GENI, Trellis (Bhatia et al., 2008), DaVinci (He and co-workers, 2008)... Recent european projects (e.g. SAIL and GEYSERS) are specifically dedicated to VI provisioning. GEYSERS works out solutions for implementing virtual networks on top of distributed optical providers, while SAIL addresses the provisioning of virtual networks to interconnect distributed Cloud data centers, proposing the dynamic creation of Flash Network Slice. RESO is full partner of both projects.

Concerning energy efficiency in large scale systems, several french teams (from INRIA: ASCOLA, AVALON - from CNRS: SEPIA (IRIT)), and international teams (BSC Barcelona, University of Sevilla, Virginia Tech, University of Melbourne) investigate this topic. Through international initiatives like the European COST action IC804 (L. Lefèvre, WP leader) and the GreenTouch initiative (L. Lefèvre, executive board member and WG co-chair), RESO exchanges and cooperates with some of these teams.

Fighting the ossification of large scale networks is a hot topic adressed by numerous network research teams in the world. Autonomic and network aware equipments research ranges from experimental and lab approach to some operational experiments by network providers. University of Passau (H. de Meer) has explored the aspect ol large networks virtualization while programmability and autonomicity in networks are explored by the team of A. Galis in Univ. College of London. Within the FP7 AUTONOMIC INTERNET, we collaborated with these two teams on the design and the validation of some large scale frameworks.

2.4.3 Scientific achievements

Virtual Infrastructures (VIs) combine on-demand provisioning of Information Technology (IT) resources and dynamic virtual networks. This combination, sought to allow dynamic

service provisioning, is transforming the Internet into a world-wide reservoir of interconnected resources, where computational, storage, and communication services are available on-demand for different users and applications. It is in this context that we derived a set of notable achievements contributing to the manageability and operability of these concepts.

−**Autonomic and network aware equipments.** We explore the programmability of future networks infrastructures by adding enhanced or autonomic funtionnalities within network equipments (routers, gateways). During the PhD of D. Lopez Pacheco [27], we proposed interoperable solutions for the XCP protocol (eXplicit Control Protocol) which improves the robustness and the interoperabiltiy of router assisted based protocols like XCP with classical equipments and transport procotols (TCP) [139]. By fighting some ossification of Internet infrastructures, autonomic approaches can allow large scale management of networks without human intervention. We proposed some programmatic enablers for autonomic network equipments like internet routers [103, 119, 160] [68] or like lightweight equipments [52] [118]. We validated these approaches through their deployment on platforms (GRID'5000, DSLLAB, Temic) and proposed some software framework (ANPI, OVNI).
−**Session awareness.** Based on our previous work [1], we continued to explore and proposed session aware distributed network solutions which support the reliability mandatory to operators services [34]. During the PhD of Narjess Ayari (joint PhD with France Telecom Labs, Lannion) [17] we proposed and developped mechanisms that improve the QoS of Internet servers [33] [90–92] and stateful devices like firewalls [58,59] [141–144] subjected to next generation services, under overload or across failures.
−**Energy efficiency.** During the PhD of Anne-Cécile Orgerie [26], we explored the design of energy aware and energy effcicient reservation infrastructures frameworks [53, 62, 63] for distributed systems like Grids [151, 152], Clouds [148] [54, 60] and future Networks [149,150,155] [61,65]. Based on theoretical models, simulations [156], and real experimental validations [107, 107, 154] through software proposal, our approach allows important energy savings while supporting a respected quality of experiment. Our approach in the context of virtualized systems is currently under validation with the development of an energy aware open cloud stack (FUI COMPATIBLE ONE project) involving some industrial partners (BULL, eNovance) and applied to some experimental platforms (local clusters and GRID'5000).
Within the PhD of Mehdi Diouri, we have started to explore the next challenging target of HPC systems, that is to say the design of green exascale services able to support ultra large deployed application [108] [209]. Some part of this activity concerns some fault tolerance exascale services which result with some collaboration with Franck Cappello (INRIA / University of Illinois Urbana Champaign Joint Laboratory). Within the Phd of Landry Tsafack (joint with IRIT Lab, Toulouse), we support the profiling of HPC applications through some DNA like approaches which allow us to propose prediction models able to widely coordinate Green leverages (DVFS, LPI, On/Off approaches).
− **Performance and sharing in virtualized networks.** As a groundwork towards routers virtualization, we analyzed and evaluated the impact of virtualization mechanisms on communication performance.

Our conclusion is that software optimization is a promising approach for experimentation but dedicated hardware is required for production networks [31], yet remaining a vivid debate in the community.. We then conceived VxSwitch[2], a virtualized switching fabric

[1] "Procédés de gestion de sessions multi-flux", N. Ayari, D. Barbaron, L. Lefevre, France Telecom R&D Patent, June 2007

[2] VXSwitch Patent - INPI: No 10/00368, 2010. LYaTiss, INRIA, ENS Lyon

that offers a wide range of assets: it allows flexible sharing of hardware switch resources among several virtual switches; it enables users to set up and configure virtual switches in terms of number of ports, capacity per port and buffer sizes; each virtual switch can have its own packet-scheduling and queuing mechanism [83]. We also proposed a virtual network service for controlling the amount of resources conferred to each virtual network, while ensuring performance isolation. This service includes virtual routers and links, both entirely configurable, and it has been implemented in two different ways: using software virtual resources (HIPerNet in [30]) or using OpenFlow switches. As a proof of concept, we finally implemented this virtual network service in the GRID'5000 testbed. A detailed assessment of this promising approach is reported in [73] [170].
This entire work was performed during F. Anhalt's PhD.
– **Virtual Infrastructure Description Language (VXDL).** We designed this descriptive and declarative language to allow users and systems to exhaustively describe the components of a virtual infrastructure, including virtual resources, network topology, and organization of the internal timeline [134]. This fully abstracted view of infrastructure and requests increases considerably the application portability and open up a promising field for dynamic VI provisioning. A open forum dedicated to its improvement and usage has been created (`http://www.vxdlforum.org/`). In [133], we used this language to develop a mechanism that translates workflows of a distributed application into VI specifications. We then tackled the NP-hard problem of resource allocation to VIs, and we came up with an effective mapping solution onto a reduced search space. This solution[3] directly leverages VXDL attributes and yields a good trade-off between users and Internet providers interests [132]. We finally proposed and developed a mechanism where users can delegate the reliability expectations of their applications to the InP, contracting a transparent service. Thanks to VXDL, when a failure happens the service automatically adapts and keeps the entire reliability management transparent to the running application [135].
VXDL language and applications described above correspond to G. Koslovski PhD work.

2.4.4 Collaborations

Relating to the activities of Axis 1, we have past or current collaborations with :

- Pablo Neira Ayuso from University of Sevilla on the topic of stateful firewalls, this collaborations gives rise to numerous high quality publications;
- RESO has hosted Alejandro Fernandez-Montes (from University of Sevilla, Spain, Nov 2008 and Sept-Nov 2009, RESO , Lyon) on the work on "Prediction Models for Energy Efficiency in Large Scale Distributed Systems";
- Permanent members of the IRIT Lab, Sepia Team (Jean-Marc Pierson and Georges Da Costa) on the energy efficiency issues in large scale systems;
- Denis Barbaron (France Telecom R&D) on session awareness;
- University college of London (A. Galis, S. Clayman) and UPC (J. Loyola and J. Serrat) teams from the AUTONOMIC INTERNET European project on programmatic enablers in autonomic networks. This collaboration leads to joint software development and publications.
- The Joint laboratory for Peta Scale Computing between INRIA and the University of Illinois at Urbana Champaign (Franck Cappello)

[3]VXAlloc Patent - INPI: 10/01626, 2010, LYaTiss, INRIA, ENS Lyon — VXCap Patent - INPI: 10/01624, 2010, LYaTiss, INRIA, ENS Lyon.

- The INRIA team-project PLANETE (Vincent Roca) and the I3S (Johan Montagnat) within the context of HIPCAL project and specifically on dynamic virtual infrastructures [171, 172] [176] and the VXDL language [133] [73].
- Alcatel-Lucent Bell Labs (Dominique Verchère, Olivier Audouin) with whom we were full partners in the CARRIOCAS project. Our joint work relied on the virtualization of optical network infrastructures [168] [32, 75]
- Alcatel-Lucent Bell Labs (Peter Vetter) on virtualization of virtual home gateways services in the context of the wireline working group in GreenTouch

2.4.5 External support

The achievements obtained in the area of *Autonomic networks*, *Session-* and *Energy awareness* were supported by:

- AUTONOMIC INTERNET (EU FP7) [A. Cheniour, O. Mornard]
- PRIMEENERGYIT (EU Intelligent Energy in Europe) [M. Dias de Asuncao, Engineer]
- COMPATIBLE ONE (FUI OSEO) [O. Mornard, M. Morel, J. Carpentier, Engineers]
- GREEN-NET (ARC) [M. Dias de Asuncao, Postdoc]
- DSLLAB (ANR)
- HEMERA (large-scale INRIA initiative) [G. Tsafack, PhD grant]
- the Mexican government, which granted a PhD scholarship to D. Pacheco for his work on *Next generation networks*
- the French ministry of education through the financing of A.-C. Orgerie PhD on *energy efficiency in large scale distributed systems*
- European Action COST IC804 on "Energy efficiency for large scale distributed systems"
- GreenTouch

The works conducted in the *Virtualization* domain were essentially financed by:

- HIPCAL (ANR)
- GEYSERS (EU FP7) [F. Anhalt, PhD grant]
- SAIL (EU FP7) [G. Koslovski, PhD grant]

2.4.6 Self assessment

Most of the results relating to Virtualization (VxSwitch and VXDL notably) gave rise to patents licensing. It is this outstanding break-through that partly motivated the creation in 2010 of the Lyatiss start-up by P. Vicat-Blanc (CEO) and S. Soudan (CTO). Now, an important part of RESO's outcomes obtained during the evaluation period stands at the core of Lyatiss activity and development. This technological transfer was awarded several national and international prizes to Lyatiss and we believe this is the striking sign of a high quality upstream research. As a counterpart though, the leaving of P. Vicat-Blanc and of all her former PhD students and collaborators has now drained RESO of its main know-how on networking virtualization techniques.

Current activities on Green networking are embedded and partly supported by the GreenTouch initiative. Activities on energy efficiency have helped to establish some research communities on this topic through the ARC GREEN-NET and the Action Interfaces, European COST. These activities will continue in the future.

2.5 Objective 2 : End-to-end Quality of Service and Transport Layer for Future Networks

2.5.1 Personnel

T. Begin, P. Gonçalves, I. Guérin-Lassous, L. Lefèvre, P. Vicat-Blanc – D. Divakaran, I. Doghri, P.-S. Guichard, R. Guillier, L. Hablot, V. N. Nguyen, S. Soudan, R. Vannier – E. Dramitinos, O. Grémillet, M. Pasin – P. Bozonnet, Ph. Martinez

2.5.2 Project-team positioning

– *High-speed transport protocols.* New TCP variants or other alternatives have been explored to cope with high-speed networks. Proposals such as HighSpeed TCP (Floyd) and Scalable TCP (Kelly) increase the aggressiveness in high-throughput situations while staying fair to standard TCP flows in legacy contexts. FAST (Caltech) or compound TCP (Microsoft) leverage the queueing information provided by round-trip time variations, in order to efficiently control buffering in routers and manage IP congestion optimally. These proposals are actively analyzed and experimented by the international community. Few groups in INRIA have been working on TCP and its high speed variants (MAESTRO, PLANETE, TREC). RESO is recognized at national level for its expertise in this domain.
– *Dynamic bandwidth sharing and congestion control.* Bandwidth (and more generaly resource) management is a major challenge to the networking community and it generates an intense worldwide activity. G-Lambda[4] (a joint research group between AIST, KDDI R&D Labs, NTT and NICT) proposed the *Grid Network Service - Web Services Interface* (GNS-WSI) to perform bandwidth reservation in networks. In the same spirit, the CARRIOCAS project developed a *Scheduling and Reconfiguration Service. StarPlane Management Plane* (SPMP) is a DRAC-based solution to dynamically configure the optical network supporting DAS-3, the Dutch grid. As for the PHOSPHORUS european project, it invented a *Network Service Plane* (NSP) for grids environments, which manages both computing resources and network. To end our non exhaustive list, the *Dynamic Resource Allocation via GMPLS Optical Networks* (DRAGON) project develops technologies allowing for dynamic and in-advance reservation of networks resources in an heterogeneous and multi-domain network.
– *Bandwidth sharing in wireless networks.* The problem of bandwidth sharing in wireless networks is becoming more and more prominent due to the wireless traffic increase along with the complicated sharing on a wireless medium. Many works have been proposed for a fair share in multihop wireless networks (teams of K. Narhstedt or J. Walrand for instance), but most of these studies are based on complicated contention models that are hard to use in pratice. Concerning the available bandwidth estimation, many solutions have been proposed recently like for instance AAC, IAB. Our solutions were the first to integrate collisions, retransmissions and induced waiting mechanisms in this estimation.
– *Flow-aware networks.* Flow-aware networking has been proposed at the beginning of the 2000's to have a better control of the network and to limit signaling (team of J. Roberts, France Télécoms R&D). Solutions, like Openflow or Anagran, allow the traffic control on a per-flow basis. The Semantic Networking paradigm, proposed in the common laboratory INRIA Alcatel-Lucent BellLabs, aims at combining control at flow level, traffic-awareness and self-management. One originality of our approach lies in the absence of *a priori* knowledge about the network but rather in its learning through regular updates on a knowledge plane.

[4] `http://www.g-lambda.net`

2.5.3 Scientific achievements

– **High-speed transport protocol.** We carried out an exhaustive and fair comparative evaluation of transport protocols in large-scale high-speed networks. To this end, an operational and strict procedure NXE[5] [126] has been set up – along with corresponding tools – to highlight the characteristic behavior of transport solutions in a real and utility context. This work comprises a variety of payoffs among which: active participation to the standardization effort of an evaluation test-suite [82]; conception of a end-user-oriented transport protocol benchmark dedicated to bulk data transfers [127]; a user-friendly tool to detect the bottlenecks existing between two end-hosts [6]. This work was at the core of R. Guillier's PhD work.

– **Dynamic bandwidth sharing and congestion control.** Concerning this other challenge raised by the Future Internet, we consolidated the algorithmic development and strengthened the effectiveness of our BDTS[7] and FLOC[8] solutions. In these latest versions, the network resource allocation service allocates network bandwidth based on the users' utilities while taking care of the network status, meaning that we condition the bandwidth delivery to the data plane so as to enable regular applications to efficiently use the allocated resources [78] [71]. Then, as reconfigurable networks naturally prompt the elaboration of dynamic provisioning, we formalized the corresponding provisioning optimization problem in a tiered architecture implying the client, the service provider and the network provider [75] [164]. In a related domain, we also elaborated on the extension of routing games with convex cost functions to consider time as a tool to quantify. The service provider solution we exposed shows to improve the social cost of allocation as compared to the one of the Wardrop equilibrium reached by selfish decisions [200]. BDTS was at the core of S. Soudan's PhD work.

– **Flow-based networks.** The last point above is overlapping with the PhD dissertation of D. Mon Divakaran, which dealt with the design aspects of a flow-based network for the Future Internet. Focusing on the two basic entities of the Internet architecture (routers and datagrams), we aimed at improving the completion time of small flows (with negligible performance degradation for larger flows), while analyzing how the other two entities (users and operator) will interact in the presence of such QoS mechanisms. The scalable Flow-Scheduler architecture we ended up to, not only instanciates such QoS mechanisms but also reduces the time-out faced by small flows [112]. We implemented this scheduler using a Combined-Input-Crosspoint-Queued (CICQ) architecture [111]. We moreover studied the robustness of size-based scheduling systems when user-behavior fails because large flows have an incentive to behave like the small ones. We tackled this issue by analyzing the existence of equilibria as well as the constraints for stability in such a dynamical system, where the number of misbehaving users can become large [199]. In parallel, but still in order to improve the efficiency of a flow-based network, we explored the resulting gain when transmitting (the protocols remaining unchanged) standard packets embedded in a meta structure that we termed XLFrames [109, 110]. The works on flow-based networks have been extended to the admission control problem. We first studied and compared several admission control solutions within the semantic networks context where small flows are

[5]NXE: Definition, configuration, deployment, run and analysis of a large scale experiment for protocol evaluation. APP: NXE V1.0 No IDDN.FR.001.030005.000.S.P.2009.000.10800.

[6]PATHNIF: Systematic analysis and evaluation of the capacity of potential bottlenecks and of end to end network paths. APP: version 1.0, march 2009 : IDDN.FR.001.260002.000.S.P.2009.000.10800

[7]BDTS: Dynamic network bandwidth allocation, bulk data transfers scheduling. APP : jBDTS Engine V2.0 No IDDN.FR.001.220025.000.S.P.2008.000.10700.

[8]FLOC: Limitation and triggering of flow rate. APP: FLOC V0.12 No IDDN.FR.001.290009.000.S.P.2009.000.10200.

always accepted. This comparison has shown that a good parameterization of the tested solutions is very dependent on the traffic and the scenarios. To bypass this issue, we have proposed a new admission control solution based on a knowledge plane computed on line and the modeling of the communication link by a single server queue. We obtained promising results on the simulator NS3. This latter study is part of the ongoing PhD work of D. Ammar.

– **Bandwidth sharing in wireless networks.** Wireless communications are now extensively used. IEEE 802.11-based networks are also widely used in different contexts including classic wireless LAN, mesh networks and multihop wireless networks. Its simplicity, its current development and its availability in commercial products increase the success of 802.11-based networks. The use of IEEE 802.11, and specifically in multihop networks, raises issues, like efficiency and/or fairness issues. We have proposed a distributed and dynamic rate allocation solution that is based on a simple sharing model. Due to its simplicity, we can derive a network protocol, called *Profiterole*, that can practically be used in multihop wireless networks. This protocol provides a fair bandwidth sharing between end-to-end flows while maintaining an efficient overall throughput in the network. This solution has been implemented in NS2 and evaluated with extensive simulations [167]. This work was at the core of R. Vannier's PhD thesis. Some applications may require specific performance that can not be achieved by Profiterole, or by any fair solution, since their goal is to provide a fair share of the radio medium. For these applications, QoS mechanisms may be more adapted. To this end, we have proposed a new available bandwidth estimation for multihop wireless networks called RABE (Retransmission-based Available Bandwidth). This method integrates the average number of retransmission attempts in the available bandwidth estimation. RABE is evaluated by simulation and the obtained results show that RABE can achieve a mean error ratio less than 20% in comparison with the real measurement. Furthermore RABE is between two times and ten times more accurate than existing solutions [145]. This work is at the core of V.N. Nguyen's PhD thesis.

Finally, we have introduced an assessment framework, based on *history-dependent utility functions* that can be used as a holistic performance evaluation tool of these fairness schemes. These functions quantify the satisfaction that the users obtain from the way their long-lived service sessions are allocated bandwidth, due to the behavior of the fair schemes. This way we can unambiguously compare the performance of various fair solutions whose maximization goals are inherently different (max-min fairness, proportional fairness, etc.) [116].

2.5.4 Collaborations

With regards to these themes, we had the opportuniy to work with:

- Our colleagues from the joint lab between INRIA and Alcatel Lucent Bell Labs on the Semantic Networking axis. More precisely, Eitan Altman (INRIA team-project MAESTRO) co-advised the PhD thesis of D. Mon Divakaran on flow-based networking [109–112, 163]. Philippe Nain (head of INRIA team-project MAESTRO) is PhD Director of M. Sokol and P. Gonçalves is co-advising with Konstantin Avrachenkov (MAESTRO) her work on semi-supervised learning and traffic awareness [86, 87, 124]. The ensemble of these works is accompanied by Ludovic Noirie (Alcatel Lucent) [80, 81, 109, 110].

- Victor Moraru from IFI (University of Hanoi, Vietnam) and Tahiry Razafindralambo (INRIA team-project POPS) on the QoS mechanisms for multihop wireless networks;

this collaboration gives rise to common publications [125, 145, 157–159] [66, 67] and a PhD student supervision.

- RESO has hosted Andrei Agapi, a PhD student from the Vrije University (June-July 2008), to work on flow scheduling within the Internet [78].

2.5.5 External support

The results relating to *High-speed transport protocols* and to *Bandwidth control* were obtained in the framework of the projects:

- CARRIOCAS (System@tic Competitive pole) [R. Guillier, PhD grant]
- EC-GIN (EU FP6)
- the French ministry of education with S. Soudan PhD grant on *dynamic Bandwidth sharing*.

The work on *Flow-based networks* was essentially supported by the:

- Common laboratory between INRIA and Alcatel-Lucent Bell Labs, with in particular the PhD grants of D. Mon Divakaran and of D. Ammar.

The *Wireless* activities are conducted under the auspices of the programs:

- AEOLUS (EU FP6) [R. Vannier, PhD grant]
- RESCUE (ANR)
- MISSION (INRIA ARC) [T. Abreu, PhD grant].

2.5.6 Self assessment

As for some results of Objective 1, the works on BDTS and FLOC gave rise to patents licensing and the technological transfer with Lyatiss was awarded several national and international prizes. The leaving of P. Vicat-Blanc and of all her former PhD students and collaborators has amputated RESO from part of its expertise on dynamic bandwidth sharing and on congestion control techniques for high-speed networks.

Regarding the works on bandwidth sharing in multihop wireless networks, we think that we have significantly improved the existing solutions in terms of performance, feasibility and evaluation. Our works have been published in highly selective conferences and journals (e.g., IEEE TMC, ACM MSWiM).

Regarding the works on flow-based networks, this activity is more recent but yet has led to first promising results which raised interests from a company like Alcatel-Lucent Bell-Labs. Moreover, in a more general setting, these works, as the ones carried out for wireless networks, have raised very interesting problems on the probabilistic resource management of networks that we plan to investigate in the future.

2.6 Objective 3 : High Speed Network's traffic metrology and statistical analysis

2.6.1 Personnel

T. Begin, P. Gonçalves, I. Guérin-Lassous, P. Vicat-Blanc – T. Abreu, D. Ammar, P. Loiseau, S. Roy, M. Sokol – M. Imbert, O. Grémillet

2.6.2 Project-team positioning

Network measurement is a mature domain. However, as the current trend in network design is to integrate more knowledge in the system management, it fosters the development of innovative information treatments. More specifically, regarding the analysis and the modeling of network traffic, our interest lies in a topical domain of activity federating (sometimes dividing) a large number of research groups at the national and international scale. We only list here some of the prominent actors whose work is tightly connected to our results.
– *Statistical properties of aggregated traffic.* First results related to self similarity in packet switching networks goes back in time with the pioneer works of Taqqu et al. (1986) on renewal models, Leland et al. (1993), Paxson and Floyd (1994), Beran (1995), Garrett and Willinger (1994), Crovella (1998), Abry et al. (1998) on long range dependence and its origins, and more recently of Veres (2000), Roughan et al. (2002), Hohn and Veitch (2002), Figueiredo at al. (2002), Doukhan (2003) or MacIntyre (2008) on the role of TCP regarding these properties. As for Dovrolis et al (2005), Barakat (2002), Riedi and Lévy-Véhel (1997), Feldmann (1998) or Sarvotham et al (2001), they studied scaling properties at the packet scale.
– *QoS.* Fewer works studied the impact of traffic statistical properties on the network performance. Among them, we can notice the interesting results obtained by Erramilli (since 1996), Park and Willinger (2000), Norros (1994), Massoulié et al. (1999) or the series of articles by Mandjes (2000)...

At INRIA level, several team-projects are concerned with stochastic modeling in the networking area. That is clearly the case of MAESTRO (Ph. Nain) and MESCAL (B. Gaujal) on the modeling, the performance evaluation, the optimization and the control of networks and their applications; of RAP (Ph. Robert) that studies probabilistic models for large communication networks; of TREC (F. Baccelli) focused on random geometry to improve networks performance, of HIPERCOM (P. Minet), which devotes one of its research axis to statistical traffic modeling; or of PLANETE (W. Dabbous) for its metrology and traffic classification components. Moreover, the common laboratory between INRIA and Alcatel Lucent Bell Labs is a rich and pertinent framework that fosters inner and outer INRIA collaborations on traffic awareness topics.

Regarding cooperative programs, several projects around the world are providing measurement infrastructures, tools and data repositories: CAIDA (`http://www.caida.org/`) is certainly the most popular but M-Lab (`http://www.measurementlab.net/`), WIDE MAWI WorkingGroup (`http://mawi.wide.ad.jp/`), RIPE Atlas (`http://atlas.ripe.net/`), Predict (`https://www.predict.org/`), Grenouille (`http://cmon.grenouille.com`) are other very active organizations. There are also a number of projects targeting the large-scale deployment of measurement instruments inside home networks (Microsoft HomeMaestro, SamKnows, `http://www.SamKnows.com`). In the same vein, RESO is partner of to two proposal submissions:

– **Metroscope** is a french initiative involving several research operators (INRIA, CNRS, Institut Telecom), a private company (Technicolor) and national authorities (RENATER, ARCEP, AFNIC). Metroscope is aimed at deploying an Internet observatory.
– **Probman** has been submitted as a STREP proposition to the 7th Framework Program of the European Commission on the probabilistic management of networks. It involves a consortium of wide-world reckoned teams (SICS, Technion Israel Inst. of Tech. , PTIN, IT Portugal, INRIA) and Alcatel-Lucent Bell Labs France.

2.6.3 Scientific achievements

– **Metrology.** As regards the metrology instrumentation of GRID'5000, we upgraded our packet capture probe so that we can now monitor a 10Gbps link capacity. As for the 1Gbps version, we checked that capture is loss free and bidirectional, time stamping precision is close to the μs and storage capacity permits several hours of capture at full speed [138] [224,226]. Since then, we have initiated the design of a synchronized distributed metrology service in order to monitor with the same timestamp several ingress-egress nodes of a large scale network. To allow affordable large scale deployment, we turned to a different technology combining simple Network Interface Cards with RAD-Clock, an experimental time protocol developed by the CUBINlab.

From a more theoretical side, we derived a series of interesting results that contribute to better characterizing the statistical features of traffic processes and to getting a deeper insight on their originating mechanisms.
– **Heavy Tailed Distribution and Long Range Dependence.** Thanks to our large scale testbed, we experimentally validated the functional bond that theoretically exists between the heavy tailed distribution of flow size and the long range dependence property of aggregated traffic [55]. For the first time, we were able to demonstrate that this theoretical relation holds true on real traffic traces and over a wide range of experimental conditions. Then, our interest in this relation led us to extend it to a model that takes into account the correlations at flow scale [56]. Such correlations exist and will become more prominent in the future Internet with the emergence of flow-aware control mechanisms correlating transmission rates to flow characteristics (size, duration, etc). In another work, we then studied the impact of this flow size heavy tail index (so of LRD) on the network performance. Considering the case of TCP protocol, we showed that the impact is far more intricate than what was claimed in the literature, and more surprisingly, that sometimes it can even increase the network performance [123]. Finally, we proposed a maximum likelihood estimator of the heavy-tail exponent of a random variable from a doubly censured dataset realization [137]. This doubly-censured situation is representative of a sampled sequence of multiplexed packets leading to undersized flows and, when the sampling rate gets too small, to unobserved (small) flows. Our solution was shown to outperform, under very general conditions, all state-of-the-art estimators.
– **Traffic and Large Deviations.** Focusing on long-lived TCP flows (representing a large portion of Internet traffic), we extended the scope of the famous Padhye's relation and demonstrated that deviations of a long TCP flow's throughput from its almost-sure mean value are rigorously quantifiable. This result relies on an ergodic large deviations principle proved to hold on almost every single realization of a large class of stochastic processes. Applied to a Markov chain modeling the time evolution of a TCP congestion window, it means also that it is possible in practice to probabilistically bound the throughput variations at the different scales of interest [136].
All aforementioned results where obtained during the PhD work of P. Loiseau.
– **Probabilistic Resource Management.** We reused the same Large Deviation prin-

ciple mentioned above, to characterize the workload volatility of a content delivery server when the users' behavior can reliably be described by a Markovian based epidemiological model [49]. This is an ongoing PhD work (S. Roy) whose ultimate objective is to propose probabilistic resource management strategies to cope with the inherent elasticity in cloud networking.
– **Semi-Supervised Machine Learning.** In a different domain, we proposed a canonical formalism for graph-based semi-supervised learning that encompasses standard or normalized Laplacian and PageRank methods [87]. We apply the resulting flexibility of these clustering techniques to user and content classification in P2P networks [86]. This work is being carried out in M. Sokol's PhD.
– **Use of clustering techniques.** Interestingly, RESO has recently integrated classical clustering techniques, namely k-means, in the aim of building a knowledge plane for management task. This clustering techniques allows us to partition numerous measurements into a limited number clusters and so to miniminze the amount of data to deal with. This work is being carried out in D. Ammar's PhD.
– **Numerical solutions to queueing systems.** Since 2009, RESO has a collaboration with Pr. Brandwajn (University of California Santa Cruz). This joint work mainly addresses queueing theory and modeling issues. More precisely, we developed semi-numerical methods to solve the steady-state distribution for the number of requests in the queue. The main originality of this work stems from the use of conditionnal probabilities.

2.6.4 Collaborations

This metrology axis gave rise to several fruitful national and international collaborations.

- With Julien Barral (Prof. Université Paris 13, also member of INRIA team-project Sisyphe headed by M. Sorine), we collaborated on the Large Deviations Principles applied to scale invariant properties of traffic [136] and also within the framework of the ANR DMASC project [35].
- Our work published in [55] is the outcome of a study carried out with the SiSyPhe team at ENS-Lyon (Patrice Abry and Pierre Borgnat). With P. Abry and Patrick Flandrin (SiSyPhe, ENS-Lyon) we have a long term collaboration on Multi-scale analysis [45, 46, 50] [76, 77] [3].
- The proposed maximum likelihood estimator we derived for flow size distributions was performed in cooperation with the INRIA project-team MISTIS headed by Florence Forbes [137].
- RESO has hosted Dr Tomohiro Kudoh for 2 weeks and Dr Yuetsu Kodama for 1 week as invited researchers to work on the design of a packet capture feature within GtrcNet10 equipment integrated in the GRID'5000 cluster [138] [224]. We also worked together on the OGF Network Service Interface use case document [6, 13].
- In Spring 2010, Pr. Brandwajn (University of California Santa Cruz) spent 3 months as an invited researcher at RESO. This visit provided the opportunity for RESO to address issues related to queueing theory and performance evalution interests. Associated publications include a journal paper [39] and articles in conferences [100–102].

2.6.5 External support

The development of our metrology system *Metroflux* along with the theoretical achievements on traffic scaling properties, were co-financed by the projects:

- GridNet-FJ (INRIA associated team)
- EC-GIN (EU FP6)
- GRID'5000-Aladdin (ADT INRIA)

P. Loiseau who authored most of these results, was granted a PhD ENS-Lyon scholarship. The *Probabilistic management* work is being financed by two main sources:

- SAIL (EU FP7) [S. Roy, PhD grant]
- PETAFLOW (ANR) [M. Jaber and A. Chonka, post-doctoral fellowships]

The PhD works of M. Sokol on *Semi-supervised learning* and of D. Ammar on *Admission control* are financed by the:

- Common Laboratory between INRIA and Alcatel-Lucent Bell Labs.

Finally, the collaboration with A. Brandwajn, which yielded the results on queueing theory, benefited from the Visiting programs of both INRIA and ENS-Lyon.

2.6.6 Self assessment

Regarding the characterization of statistical properties of network traffic and the identification of certain of their generating mechanisms, we believe that we have significantly contributed to a better comprehension of these complex systems. Our work has been reckoned in some of the best conferences (ACM SIGMETRICS, IFIP PERFORMANCE) and transactions (IEEE TON, STOCHASTIC MODELS) of the domain and it is something we are quite proud of. We believe that this work holds ground in a deep knowledge in statistical signal processing and in a strong expertise with experimental data analysis and modeling, two assets of RESO that we wish to reinforce in the future.
The design and the deployment of our packet capture system based on the Gtrc-Net technology, provided us with a flexible and high performance metrology tool. This achievement represents a tremendous effort during the last four years, and notwithstanding its success (all our experimental studies relied on MetroFlux), we found very difficult to valorize its outcome in the network community. By inexperience also, we considerably underestimated the time spent on technological pitfalls, which drastically hampered the progression of our scientific advances.

Not surprizingly, RESO worked towards a rapprochement between the objectives of the present axis and that of axis 2 (end-to-end QoS). Indeed, we initiated our metrology activity (in 2006) having in mind to integrate traffic awareness (from measurements and analysis) to network management in order to better rationalize resources usage. This coupling is becoming effective with the PhD work of S. Roy on probabilistic resource management in elastic environments and with that of D. Ammar on the notion of knowledge plane and admission control. It is this approach combining measurement, analysis, modeling and decision that we want to emphasize in one direction of RESO's follow-up.

2.7 Objective 4 : Application to Grid Middleware and Applications

2.7.1 Personnel

J.-P. Gelas, P. Gonçalves, O. Glück, L. Lefèvre – R. Guillier, L. Hablot, P. Loiseau, S. Soudan – M. Imbert, J.-C. Mignot, O. Mornard

2.7.2 Project-team positioning

Many MPI implementations are available for the grid like MPICH2, OpenMPI, GridMPI or MPICH-G2. Some of them need a fine tuning to be efficiently executed with TCP. Others like MetaMPICH, PACX-MPI or GridMPI also introduce proxies at the LAN/WAN interfaces, but at a proxy level or usage different from ours. Others provide optimized algorithms for collective operations, for instance MPICH2, MagPIe, GridMPI or MPICH-G2. GridMPI also attempts to avoid burst by proposing a mechanism of packets spacing, which therefore reduces the number of TCP retransmissions. The INRIA RUNTIME team has proposed MPICH-Madeleine that efficiently manages communications between heterogeneous high speed cluster networks. MPICH-V designed by the INRIA GRAND LARGE team focusses on fault tolerant mechanism for MPI applications. Finally, even if we have focalized on MPI applications, our proposition is not specific to MPI applications because it is at the socket level and thus, it would apply to any grid application executed above TCP.

2.7.3 Scientific achievements

– Design and development of MPI5000. Firstly, we studied the interactions between TCP and MPI applications executed on grids, and we identified two main problems: (1) a high latency on the grid (the application waits a long time for duplicate acknowledgments and for the TCP congestion window to increase); (2) the communication profile of MPI applications can substantially increase the number of RTO timeouts and idle times during WAN communications. To circumvent these, we then proposed to insert a new transparent layer between MPI and TCP, which takes into account the grid topology to split TCP connections between MPI processes. This layer, called MPI5000, is based on LAN/WAN proxies that better manage long-distance traffics and splits each LAN-WAN-LAN TCP connection in three distinct ones: LAN-LAN, WAN- WAN and LAN-LAN. The proposed architecture turns more reactive to congestion and aggregates bursty MPI traffic into larger flows over long-distance links. The advantage is manyfold: *(i)* it maintains the congestion window value closer to the actually available WAN bandwidth, *(ii)* it eliminates some RTO timeouts, and *(iii)* it retransmits duplicate acknowledgments faster. The TCP waiting time of MPI applications is thus reduced and ergo, global performances are improved. Our experimental validations of MPI5000 on GRID'5000, demonstrated that the expectable gain in terms of execution time can go up to 35% for certain applications, and that compared to MPICH-G2, MPI5000 can outperform a grid enabled MPI implementation.

Finally, our architecture allowed us to use different TCP implementations depending on wether we deal with local or distant communications. We finally implemented on each relay, a priority variable scheduling policy for MPI messages and we experimentally demonstrated that favoring the control messages (w.r.t. collective operations and small messages) could by itself decrease the execution time of some applications by almost 10%.

– Optical resource provisioning and optimal bandwidth sharing for the grid. After the fact, it deemed artificial to isolate in this 4th axis, the results related to these themes as they are intimately connected to those of axes 1 and 2, already described in the

corresponding sections. Let us simply recall that to a large extent, they were obtained in the framework of the PhD works by S. Soudan, F. Anhalt and G. Koslovski.

2.7.4 Collaborations

- We had fruitful exchanges with Masatoshi Ishikawa (AIST at University of Tokyo) who developed GridMPI, a grid-oriented MPI implementation similar to the system of relays that we used in MPI5000. In partnership with RESO, AIST tested and validated their approach on the realistic environment provided by GRID'5000.
- As already listed in the collaborations of Objective 3 (Metrology), it is thanks to the Gridnet-FJ associated team with AIST (T. Kudoh and Y. Kodama) that we succeeded in the development of our packet capture system MetroFlux [138] [224]. The Gtrc-Net card, at the core of this metrology probe, was also used to emulate latency in the TCP experiments conducted both by S. Soudan and by R. Guillier in their respective PhD works.

2.7.5 External support

During his Phd work on *MPI5000*, L. Hablot was granted a French ministry scholarship and his activity received research support from:

- GRID'5000-Aladdin (ADT INRIA).

As for the other outcomes of this transverse axis, they did not get specific financing but overlapped with several already mentioned programs:

- GEYSERS (EU FP7)
- SAIL (EU FP7)
- CARRIOCAS (System@tic competitive pole)
- EC-GIN (EU FP6)
- HIPCAL (ANR)
- PETAFLOW (ANR)
- GridNet-FJ (INRIA associate team)

2.7.6 Self assessment

Thanks to MPI5000, we have shown that splitting TCP connections for WAN communications of MPI-like applications improves significantly application performances. However, MPI5000 is nothing more but a prototype. The software should be improved to be used at very large scale. The MPI5000 proxy overhead should be reduced by implementing a kernel version of the proxies and hence avoid two extra copies in user space. Finally, we should implement and evaluate further optimizations now available thanks to MPI5000: bandwidth reservation between proxies, parallel long-distance connections, use of an optimized/modified TCP protocol on the WAN.

More generally though, this research axis is certainly the one that was the most impacted by the orientations changes in RESO. In our initial thematic apportionment, the existence of this axis was fully justified when we used to study grid applications on GRID'5000 (e.g. MPI5000, CARRIOCAS, IGTMD), but the current evolution of our activities led us to rather consider GRID'5000 as an experimental testbed to deploy metrology equipments for traffic capture or for energy consumption measurements. Presently, this axis contains only few specific themes, and to a large extent, it is already implicit in the activities carried out in the other branches of RESO. If RESO had to continue per se, we would certainly recommend to split this 4th axis and to merge it with the objectives of axes 1, 2 and 3.

3 Knowledge dissemination

3.1 Publications

	2008	2009	2010	2011	Total
PhD Thesis	2	5	1	3	11
H.D.R (*)			1		1
Journal	9	8	11	11	39
Conference proceedings (**)	24	34	17	21	96
Book chapter	1	2	1	3	7
Book (written)				1	1
Book (edited)		4			4
Proceeding & special issues (edited)	4	1	2		7
Patent		1	4		5
General audience papers					
Technical report	3	7	5	1	16
Deliverable	3		3	2	8

(*) HDR Habilitation à diriger des Recherches
(**) Conference with a program committee

Major journals in the field with RESO publications in 2008-2012

1. *Future Generation Computer Systems*: 4 publications [52, 71, 72] [11].
2. *Journal of Supercomputing*: 3 publications [54, 65] [7]
3. *IEEE Transactions on Mobile Computing*: 1publication [69]
4. *IEEE/ACM Transactions on Networking*: 1 publication [55]
5. *Journal of Applied Probability*: 1 publication [39]
6. *Stochastic Models*: 1 publication [56]

Major conferences in the field with RESO publications in 2008-2012

1. *IFIP Networking*: 4 publications [109, 112, 139, 158]
2. *ACM/IEEE MSWiM*: 3 publications [116, 145, 167]
3. *IFIP Performance*: 2 publications [123, 136]
4. *ACM Sigmetrics*: 1 publication [137] and a best demonstration award [226]

3.2 Software

Lyatiss Weaver Suite

1. Software characterization: A-5, SO-4, SM-4, EM-4, SDL-5
2. Own Contribution characterization: DA-4, CD-4, MS-4, TPM-4
3. Short Description. The following list of softwares, whose development was initiated at RESO , are in constant evolution as they represent a direct outcome of the research carried out at RESO until Aug. 2011. They also form the core of the technological transfer to Lyatiss; embedded in the *Lyatiss Weaver Suite*, they implement the solutions for virtual resources orchestration and infrastructure services.

BDTS Dynamic network bandwidth allocation, bulk data transfers scheduling.
APP : JBDTS version 1, Dec. 15, 2007: IDDN.FR.001.220025.000.S.P.2008.000.10700

FLOC Limitation and triggering of flow rate.
APP: version 0.12, Feb. 17, 2009: IDDN.FR.001.290009.000.S.P.2009.000.10200

NXE Definition, configuration, deployment, run and analysis of a large scale experiment for protocol evaluation.
APP: version 1.0, Nov. 2008: IDDN.FR.001.030005.000.S.P.2009.000.10800

VXcore Resource temporal database manager.
APP: version 1.0, March 15, 2009 : IDDN.FR.001.290012.000.S.P.2009.000.10800

VXtopology Resource spacial database manager.
APP: version 1.0, March 15, 2009 : IDDN.FR.001.290012.000.S.P.2009.000.10800

VXScheduler Adaptation of virtual infrastructure request and scheduling.
Patent version 1.0, March 15, 2009 : IDDN.FR.001.290010.000.S.P.2009.000.10800

VXDL Parser interpretation and XML traduction of virtual infrastructures specifications.
Patent version 2.0, March 20, 2009: IDDN.FR.001.260009.000.S.P.2009.000.10800)

SRVdemonstrator Scheduling, Reconfiguration and Virtualization of Network resources for intensive computing environment.

PATHNIF Systematic analysis and evaluation of the capacity of potential bottlenecks and of end to end network paths.
APP: version 1.0, March 2009 : IDDN.FR.001.260002.000.S.P.2009.000.10800

HIPerNet v0.5 Cloud solution HIPerNet engine is a software implementing discovery, selection, allocation, scheduling and management of virtual private execution infrastructures over the Internet. HIPerNET v0.5 is focusing on virtual end-resource deployment and configuration.
Distributed under GPL license.

4. Associated URLs:
   `http://www.ens-lyon.fr/LIP/RESO/web/`
   `http://www.lyatiss.com`

***G/M/c*-like queue solver**

1. Software characterization: A-3, SO-2, SM-3, EM-3, SDL-4
2. Own Contribution characterization: DA-4, CD-4, MS-4, TPM-4
3. Short description.
   This tool provides a simple web application to promote the use of our algorithms for solving classical queueing systems. It currently features the numerical solution to the steady-state distribution for the number of customers in the system and other customary performance parameters for a queue with multiple servers, general arrivals, exponential service times and a possibly finite buffer, (i.e., $G/M/c$-like and $G/M/c/N$-like queue). The steady-state solution to this queue is based on a simple and stable recurrence whose leverages the use of conditonal probabilities [39]. This tool is a joint work with Pr. Brandwajn (University of California Santa Cruz). We

will include new models to this tool in the near future. As of beginning of 2012, this site is averaging over 15 unique visits per day.

4. Associated URL: `http://queueing-systems.ens-lyon.fr`
5. Associated publications: [39]

MPI5000

1. Software characterization: A-3, SO-2, SM-2, EM-1, SDL-2
2. Own Contribution characterization: DA-4, CD-4, MS-4, TPM-4
3. MPI5000 is a communication layer between the application (MPI for example) and the transport protocol (TCP) which improves communications of distributed applications over wide area network in grids. For instance, MPI5000 reduces the impact of retransmissions and the impact of congestion window in such a context. MPI5000 is automatically and transparently executed without modifying the application. The general principle is to introduce proxies at the interface between the local network and the long- distance network to differentiate communications. These proxies allows to put forward the split of TCP connections in order to avoid losses and retransmissions on the long-distance links. This mechanism also allows to keep the congestion window closer to available throughput on the long-distance network.
4. Associated publication: [122]

ShowWATTS

1. Software characterization: A-3, SO-2, SM-2, EM-3, SDL-4
2. Own Contribution characterization: DA-4, CD-4, MS-4, TPM-4
3. ShowWATTS is an open source software framework. The aim of ShowWATTS is to provide Grid/Cloud users and developers with tools to deal with power consumption of such infrastructures. ShowWATTS provides (i) a unificated interface to interact with heterogeneous power measure hardware, (ii) an efficient solution to store and fetch logged measures, and (iii) web interfaces or GUI client tools to expose live measures to the end users (and infrastructure administrator) and to extract data power consumption of users' applications for post-mortem experiment analysis.
4. Associated publication: [154]

3.3 Valorization and technology transfer

The creation of the Lyatiss[9] start-up by P. Vicat-Blanc (former scientific leader of RESO) and by S. Soudan (former PhD student at RESO) constitutes the main technological transfer and a veritable success-story of our scientific activity. The by-product of RESO research already yield five patents (applying to the EU zone and soon extended to the US and Canada).

From the institutional side, RESO participated to several national and european projects that also implied many different industrial partners (Alcatel Lucent, Ericsson, HP, BULL, ...). Directly or indirectly, these latter benefited from and valorized the technical developments carried out by RESO within these frameworks. The joint lab between INRIA and Alcatel Lucent is also a transfer vector that not only enhances our impact on the industrial world, but reciprocally and to some extent, it orients our upstream research to better address the forthcoming challenges relevant for the society.

[9] `http://www.lyatiss.com/`

3.4 Teaching

As 2/3 of the permanent staff of RESO are professor or assistant professors at University of Lyon, RESO is highly involved in teaching activities, with a total amount of about 4000 class hours for the evaluation period.

- **Isabelle Guérin Lassous** ($\sim$ 200 hrs/year)

  2008-2012: University Lyon 1 - Master; Quality of Service and Multimedia Systems - Networking - Ad Hoc Networks - Autonomic Networks

  2010-1012: ENS Lyon - Master; Network algorithms

  2011-2012: University Lyon 1 - Master; Wireless Networks

  2011-2012: ESIROI (La Réunion) - 5th year; Quality of Service - Wireless Networks

- **Thomas Begin** ($\sim$ 200 hrs/year)

  2009-2012: University Lyon 1 - Master; Computer Networks - Advanced Topics in Communication Networks I & II

  2011-1012: ENS Lyon - Master; Network Architecture

  2009-2012: University Lyon 1 - Under graduate; Computer Networks

  2010-2011: IFI (Hanoi, Vietnam) - Master; Advanced Topics in Communication Networks I

- **Jean-Patrick Gelas** ($\sim$ 250 hrs/year)

  2008-2012: University Lyon 1 - Master : Computer Networks (protocols)

  2009-2012: University Lyon 1 - Master : Network Architecture (hardware)

  2008-2012: University Lyon 1 - Master : Operating System

  2008-2012: University Lyon 1 - Master : Embedded systems and software

  2008-2009: IFI (Hanoi, Vietnam) - Master : Computer Networks (protocols)

- **Olivier Glück** ($\sim$ 250 hrs/year)

  2008-2012: University Lyon 1 - Licence; Computer Networks - Database and network initiation

  2008-2012: University Lyon 1 - Master; Internet protocols and applications

  2011-1012: ENS Lyon - Master; Network Architecture

  2011-2012: University of Sciences HCMUS (Ho Chi Minh ville, Vietnam) - Licence; Computer Networks

- **Paulo Gonçalves** ($\sim$ 50 hrs/year)

  Since 2008 ENS Lyon - Responsible for the cursus axis "Models and optimization for emergent infrastructures", Master *Informatique Fondamentale*.

  2008-2011 ENS Lyon - Master; "Models and Processes for Network Traffic" (24h).

  2009-2011 ENS Lyon - Master; organization of four research schools (4×24h).

  2011-2012 CPE Lyon - Master; "Introduction to Compressive Sensing" (8h).

- **Laurent Lefevre** ($\sim$ 50 hrs/year)

  2011-1012: ENS Lyon - Master; Network Architecture (8h)

  2011-2012 University of Addis Abeba - Ethiopia - Master; Advanced networks : performant, intelligent and green (48h)

3.5 General Audience Actions

General audience papers or talks:

- Evolution du standard pour réseau sans fil: IEEE 802.11, I. Guérin Lassous, Techniques de l'Ingénieur, TE7376, October 2009.

Videos

- Semantic Networking project: `http://www.dailymotion.com/video/xgudmh_laboratoire-commun-alcatel-lucent-inria-demonstration_tech`

Interviews

- Laurent Lefevre : Papier : Dossier pour la Science 2010, La Recherche 2010, 01 Informatique Business et Technologies 2010, Okapi 2011 - Radio : Radio Méditerranée Internationale 2010, RCF Isère 2010

3.6 Visibility

Editorial activities

- Isabelle Guérin Lassous is in the editorial board of Computer Communications (Elsevier), Ad Hoc Networks (Elsevier) and Discrete Mathematics and Theoretical Computer Science journal.
- Laurent Lefevre has edited a special theme on "Towards Green ICT", ERCIM News, October 2009 and proceedings of CCGrid2008, PDCAT 2008, Special issue of Journal of Supercomputing 2010.
- Paulo Gonçalves is with the Editorial board of EURASIP (officer of the Local Liaison Board).

Organization of workshops/conferences

- Isabelle Guérin Lassous
  - Program co-Chair of ACM International Symposium on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks (PE-WASUN) in 2008, 2009, 2010 and 2011 and General Chair in 2012.
  - Poster Chair of ACM MobiHoc in 2009.
- Laurent Lefevre
  - General (co) chair of IEEE CCGrid2008, IEEE HPCC2009.
  - Program (co) chair of GADA2009, IEEE/ACM Grid2010, (vice) GreenCom 2010, IEEE HPCC2012, Cloud&Grid2012.
  - Workshop chair of PDCAT08, HPPAC09, E2GC2 2009, SUNSET 2011.
  - (Co) Organizer of events : Parco2009, Entretiens Jacques Cartier 2011, GreenDays 2010, 2011 and 2012.

- (Co) Organizer of INRIA events : INTECH seminary 2010, Rencontres INRIA Industrie 2012, INRIA booth in Supercomputing 2009, 2011 and 2012.

Involvement and responsibilities in the scientific community

- Isabelle Guérin Lassous
  - co-Vice-Chair of the Expertise Committe of the program SINI2 (Computer Science) at the ANR.
  - member of the CNU 27 (2008-2011).
  - co-leader of the project Semantic Networking of the commun laboratory INRIA-BellLabs.
- Laurent Lefevre
  - Member of the CNU 27 (2010-2011).
  - INRIA Representative, Executive Board member and Working Group Vice chair in the GreenTouch organization (since 2009).
  - Steering committee member of conferences : IEEE CCGrid (2008-2012), ICPS (2008-2009), IEEE/ACM Grid2011.
  - Management Committee member and French representative in the European COST Action IC804.
  - TPC member of various international conferences (IEEE CCGrid, Supercomputing 2009,2011, ICCS, NSS, UCC, E-Energy, PDCAT, HPDC08...)
- Paulo Gonçalves
  - Member of the AERES evaluation committee of LTSI Lab., Univ. of Rennes 1 (2010).
  - ANR expert for the "Software Technologies" program SIMI3 (2008-2010).
  - PC member of *International Workshop on High Speed Network and Computing Environments for Scientific Applications* (2010, 2012); *International Workshop on World Wide Computing Environments for Knowledge Discovery* (2011).
- Jean-Patrick Gelas
  - PC member of the Energy Efficient Grids, Clouds and Clusters Workshop (E2GC2 2010 and 2009).
  - TPC member of the Fifth Workshop on High-Performance, Power-Aware Computing (HPPAC 2009).
  - Poster chair of the 8th IEEE International Symposium on Cluster Computing and the Grid (CCGrid2008).
  - PC member of the Mardi Gras 2008 conference, 15th conference in cooperation with ACM SIGAPP.

4 External Funding

(k euros)	2008	2009	2010	2011
INRIA Research Initiatives				
ARC† GREEN-NET	20	20	–	–
ARC† MISSION	–	–	7	7
ADT ALADDIN	7	–	–	–
National initiatives				
ANR IGTMD	32,3	6,5	–	–
ANR DSLLAB	8,4	2	–	–
ANR HIPCAL	90,8	8,8	–	–
ANR DMASC	–	2,5	–	–
ANR RESCUE	–	–		28,6
ANR PETAFLOW	–	–	5	56
COMPATIBLE ONE	–	–	–	43,6
ACTION INTERFACES	–	–	4	–
CARRIOCAS	74,4	120,5	–	–
European projects				
FP6 (STREP) EC-GIN	96	102,9	159,4	–
FP6 (IP) AEOLUS	57,9	60,9	–	–
FP7 (IP) AUTONOMIC INTERNET	52	70,7	143,6	–
FP7 (IP) GEYSERS	–	–	84,5	35,1
FP7 (IP) STREP SAIL	–	–	43	80
FP7 OGF EUROPE	5	42,7	–	–
FP7 (IC) COST	–	–	–	–
FP7 PRIME ENERGY IT	–	–	25,8	38,8
FP7 (NoE) EURO NF	5	5	–	10
Associated teams				
GRID-NET France-Japon	6,6	5		
Industrial contracts				
CIFRE FRANCE TELECOM R&D	8	–	–	–
CIFRE ORANGE	4,5	4,5	4,5	4,5
INRIA-ALBL COMMON LAB	100,3	106	122,5	89
ANAGRAN	–	–	–	–
Scholarships				
PhD *	130,9	28,9	57,7	33,6
Post Doc*	–	–	–	–
AI+	–	43	43,3	–
ODL#	26,2	26,2	26,2	26,2
ENSL/CNRS Engineer suport	26,2	26,2	26,2	26,2
Other funding				
INRIA internship	–	–	6,7	3,6
INRIA invited researcher	–	–	6	–
ENS invited researcher	–	–	13,1	–
INRIA allocation	30	25	26	16,5
ENS Lyon allocation	6,6	2,9	4,2	1,5
CNRS allocation	2,9	2	1,6	1,5
Region support to start-up creation (Lyatiss)	–	11,4	15	–
Total	**791**	**723,6**	**823,8**	**503,2**

† INRIA Cooperative Research Initiatives
‡ Large-scale Initiative Actions
∗ other than those supported by one of the above projects
+ junior engineer supported by INRIA
# engineer supported by INRIA

ARCs

Green-Net[10] (INRIA, 2008-2009) This project explores the design of energy-aware software frameworks dedicated to large scale distributed systems. Laurent Lefèvre is leading the INRIA ARC Green-Net on "Power aware software frameworks for high performance data transport and computing in large scale distributed systems" which involved 4 partners : INRIA RESO and MESCAL (Grenoble), IRIT (Toulouse), Virginia Tech (USA).

Mission (INRIA, 2010-2011) stands for MobIle SubStitutIOn Networks and is focused on the performance study and the feasibility to deploy a fleet of mobile wireless routers to help a wired network that can not offered its services anymore (e.g., due to a failure). Other partners are: lip6 (Paris) and INRIA Lille.

National initiatives

igtmd (ANR, ENSL, 2006-2008) The aim of this project is to design, develop and validate mechanisms that concretely make the interoperability of heterogeneous grids a reality. The project concentrates on the following topics: a) Bulk data transfers, b) replication and referring mechanisms, c) information system and job management interoperability, d) grid control and monitoring, e) usage of statistics and accounting data.

DslLab (ANR, ENSL, 2006-2008) aims at building and using an experimental platform about distributed systems running on DSL Internet. The objective is twofold: *(i)* to provide with accurate and customized measures of availability, activity and performances in order to characterize and tune the models of the ADSL resources; *(ii)* to provide with a validation and experimental tool for new protocols, services and simulators and emulators for these systems. In this project, RESO has mainly collaborated with LRI (Orsay) and INRIA MESCAL team.

Hipcal[11] (ANR, INRIA, 2007-2010). The goal of this project was to explore an approach in a break with current services-oriented principles developed in grids, to enhance the application portability, the communications performance control and their security, simultaneoulsy. Hipcal studied a new paradigm (grid substrate) based on confined virtual private execution infrastructure for resource control in grids. In particular, we proposed and implemented new approaches for bandwidth sharing and end to end network quality of service guarantees. Use-cases in biomedical applications deployed on grid'5000, served as proof-of-concept. Joint project between INRIA (RESO, grand large, planete) and CNRS (IBPC, I3S).

dmasc (ANR, INRIA, 2008-2012) The main objective is to develop advanced multifractal analysis tools, from mathematically ground results to efficient estimators. We apply these methods to the analysis, to the modeling and to the classification (for non invasive diagnoses) of cardiovascular systems. This project, leaded by J. Barral

[10] http://www.ens-lyon.fr/LIP/RESO/Projects/GREEN-NET/
[11] http://www.ens-lyon.fr/LIP/RESO/Projects/HIPCAL/ProjetsHIPCAL.html

(Univ. Paris 13), is a partnership between INRIA (SISYPHE and RESO), university of Paris 12 and Paris 13 and Paris Sud (équipe d'accueil EA 4046 Service de Réanimation Médicale CHU de Bicêtre).

PETAFLOW[12] (ANR, ENSL, 2009-2012) Generation or processing of peta-scale data benefits from the emergence of adequate *Information and communication technologies* with respect to high performance *computing-networking-visualization* and their mutual *awareness*. In this project, RESO aims at proposing network solutions to guarantee the Quality of Service (in terms of reliability level and of transfer delay properties) of a high speed, transnational long-distance connection used in an interactive, high performance computing application. Another specificity of this application is the peta-scale volume of the treated data corresponding to the upper airway flow modeling. Collaborative project between the GIPSA Lab (Grenoble), MOAIS (INRIA Grenoble), RESO , the University of Osaka (the Cybermedia Center and the Department of Information Networking) and the University of Kyoto (Visualization Laboratory).

RESCUE[13] (ANR, ENSL, 2010-2013). RESCUE follows the ARC MISSION project. In RESCUE, we investigate both the underlying mechanisms and the deployment of a substitution network, aimed at overcoming local failures on a base network, maintaining a satisfactory level of service to the users. Unlike many projects and other scientific works that consider mobility as a drawback, in RESCUE we use the controlled mobility of the substitution network (composed of a fleet of dirigible wireless mobile routers) to help the base network reduce contention or to create an alternative network in case of failure. The advantages of an on-the-fly substitution network are manifold: Reusability and cost reduction; Deployability; Adaptability. Other partners of this project are the INRIA team POPS, LIP6, LAAS, and France Telecom.

COMPATIBLE ONE (FUI, INRIA, 2010-2012). This project, funded by the *Fonds Unique Interministériel* (FUI) is dealing with the building of a Cloud architecture open software stack. In this project, RESO is focused on the design and provisioning of energy aware and energy efficient components in order to include energy aspects in QoS, SLAs and billing in clouds architectures. In this project, INRIA works in close cooperation with the project leader BULL and some partners SMEs : Prologue, Activeon, Xwiki, Mandriva and ENovance.

ALADDIN GRID'5000[14] (ADT, INRIA). GRID'5000 testbed is an experimental Grid platform gathering ten sites geographically distributed in France. The GRID'5000 site of Lyon comprises now around 300 processors interconnected with the 10 Gbit per second network. Lyon site is nationally recognized to gather the *networking expertise* with skilled research and engineering teams focused on networking equipments (Metroflux, Gtrc-Net10...). Lyon site also hosts an important part of the Green GRID'5000 infrastructure by hosting a set of 1500 wattmeters and exposing energy measurements to the GRID'5000 community.

CARRIOCAS[15] (Competitive pole System@tic, INRIA, 2006-2009). In this collaborative work we studied and implemented an ultra high bit rate (up to 40 Gbps per wavelength) network interconnecting super computers, storage servers and high resolution visualization devices to support data and computing intensive applications in industrial and scientific domains. This testbed was intended to be the experimental first

[12] http://petaflow.gforge.inria.fr/
[13] http://rescue.lille.inria.fr/
[14] https://www.grid5000.fr/mediawiki/index.php/Grid5000:Home
[15] http://www.systematic-paris-region.org/fr/projets/carriocas

step towards a transition from local to external storage and computing systems. More specifically, RESO was in charge of the design and the prototyping of the "Resource Scheduling Reconfiguration and Virtualization - SRV" component. As a regional initiative, this project gathered under the coordination of Alcatel-Lucent, more than 20 academic and industrial IT actors from *Ile de France* (participation of RESO stemmed from our close collaboration with the INRIA team GRAND LARGE).

INTERFACES (GRID'5000, Inst. of Grids, INRIA, 2010-2011) RESO is leading this *Action Interfaces Recherches en Grilles*: "Energy efficiency in Grids from experimental to operational platforms". The goal of this action is to propose and evaluate energy aware software components able to be deployed in production data centers of Grids. Some energy monitoring infrastructure will be deployed and validated in operational centers (like the IN2P3 Data center). Main partners of this project are RESO, ASCOLA and MESCAL INRIA teams and IRIT.

European projects

SAIL[16] (FP7, INRIA, 2009-2012) Sought objective is the research and the development of novel networking technologies using proof-of-concept prototypes to lead the way from current networks to the Network of the Future. SAIL leverages state of the art architectures and technologies, extends them as needed, and integrates them using experimentally-driven research, producing interoperable prototypes to demonstrate utility for a set of concrete use-cases. RESO contributes to workpackage D on *Cloud Networking*. The consortium comprises 25 academic and industrial partners from Europe (Sweden, Germany, Finland, UK, France, Spain, Portugal, Ireland) and from Israel and Australia.

GEYSERS[17] (FP7, INRIA, 2009-2012) We envision to qualify optical infrastructure providers and network operators with a new architecture, to enhance their traditional business operations. Following this objective, GEYSERS will specify and implement a novel optical-network architecture able to support "Optical Network+Any-IT" resource provisioning seamlessly and efficiently. Energy-consumption metrics for the end-to-end service routing are part of this efficiency. The consortium involves more than 20 academic and industrial partners from Europe (Italy, Switzerland, Germany, Poland, The Netherlands, Greece, Belgium, UK, Spain) and India.

AUTONOMIC INTERNET[18] (FP7, INRIA, 2008-2010) proposed a transition from a service agnostic Internet to service-aware network, managing resources by applying autonomic principles. In order to achieve the objective of service-aware resources and to overcome the ossification of the current Internet AUTONOMIC INTERNET developed a self-managing virtual resource overlay that can span across heterogeneous networks that can support service mobility, security, quality of service and reliability. We emphasized co-operation of the following activities: autonomic control principles, resource virtualization, enhanced control algorithms, information modeling, policy based management and programmability. RESO was mainly involved in the programmability of the AUTONOMIC INTERNET overlay by proposing an Autonomic Network Programming Interface that supports large scale service deployment.

[16] `http://www.sail-project.eu/`
[17] `https://www.geysers.eu/`
[18] `http://ist-autoi.eu`

RESO has mainly collaborated with University College of London (UK), University of Passau (Germany), UPC (Spain) and Hitachi.

AEOLUS (FP6, INRIA, 2005-2008). Acronym for "Algorithmic Principles for Building Efficient Overlay Computers", AEOLUS investigated the principles and developed the algorithmic methods for building an overlay computer that enables an efficient and transparent access to the resources of an Internet-based global computer. The university of Patras (Greece) was the prime contractor.

EC-GIN[19] (FP6, INRIA, 2006-2009) Based on a number of properties that make Grids unique from the network perspective, the project EC-GIN has developed a tailored network technology in dedicated support of Grid applications. These technical solutions have been supplemented with a secure and incentive-based Grid Services network traffic management system, which balanced the conflicting performance demand and the economic use of resources in the network and within the Grid. EC-GIN outcomes stemmed from a close academic collaboration between Europe (INRIA, UIBK, UniZH, ULANC, UniS, UIO) and China (BUPT, CTTL, CMDI).

OGF EUROPE (FP7, INRIA, 2008-2010). RESO participates to the OGF-Europe to reinforce the french participation to Open Grid Forum in standardization activities. Our contribution mainly concentrates on Telco interaction and Energy-efficiency in Grid context.

COST (IC, INRIA, 2009-2013). The Action "Energy efficiency for large scale distributed systems" has for main objective to foster original research initiatives addressing energy awareness/saving and to increase the overall impact of European research in the field of energy efficiency in distributed systems. The Action objectives can be summarized on scientific and societal points of view: sharing and merging existing practices will lead the Action to propose and disseminate innovative approaches, techniques and algorithms for saving energy while enforcing given Quality of Service (QoS) requirements. Laurent Lefèvre is Management Committee member and French representative in this COST action. This project is leaded by University of Toulouse (France) and gathers a total of 20 partner countries.

PRIMEENERGYIT[20] (IEE Intelligent Energy in Europe, INRIA, 2010-2012). The proposed project is designed to further enforce energy efficient market development for central IT equipment based on the previous initiatives and with a focus on so far largely uncovered IT hardware including storage and network equipment as well as new power management technologies. RESO is mainly involved in energy efficiency criteria in the context of storage for small and medium datacenters. Other partners are: Electricity of Austria (coordinator), Berlin Energy Agency, Berlin Institute of Technology, BIO Intelligence Service, Politecnico di Milano, GAIA, ICLEI, University of Coimbra, Seven.

EURO-NF (Specific Joint Project, INRIA, 2008-2010). This consortium agreed to initiate a specific joint research project focused on the impact of security and privacy on emerging and evolving ICT energy efficiency mechanisms. The addressed problems are of high importance and span multiple domains of technology, including energy efficiency, virtualization and IT security/privacy, and consequently have a deep impact on the evolution of next generation networks and future Internet. This specific joint research project is leaded by university of Passau (Germany) in collaboration with university of Wien (Austria), the CERTH-ITI (Greece) and INRIA (France).

[19] http://www.ec-gin.eu/corpsite/display/main.asp
[20] http://www.efficient-datacenter.eu/

Associated teams and other international projects

GRID-NET (Ass. team INRIA-AIST Japan, 2007-2009). Thanks to this collaboration, we investigated four main directions: 1) High speed transport protocol over very high speed links, 2) Bandwidth allocation and control in Grids, 3) Optimisation of MPI communications in Grids, and 4) Co-design of GtrcNET-packet capture functionality. The integration of the Gtrc-Net 1 and 10 (developed at AIST) into the GRID'5000 infrastructure led us to a series of noticeable results that are described in the scientific achievements of axes 2, 3 and 4 (e.g BDTS, Metroflux, MPI5000).

NEGST (JSPT-CNRS). The objective of this project was to promote the collaborations between Japan and France on grid computing technology. We considered three main lines of investigation: 1) Grid interoperability and applications; 2) GridMetrics and 3) Instant Grid and virtualization of grid computing resources. RESO mainly participates to the Grid Metrics topic.

Industrial contracts

INRIA ALCATEL LUCENT BELL LABS[21] (Common laboratory, 2008-2012)
RESO participates in the Research Action *Semantic Networking* (SEM- NET) which advocates a new paradigm for the networks of the future bringing together *flow-based networking*, *traffic-awareness* and *self-management* concepts to get *plug-and-play* networks. The natural traffic granularity is the flow. RESO's task is to elaborate on the admission control of flows in routers having in mind the current status of the network and the underlying applications. In a different work, we also consider the problem of graph-based semi-supervised approaches applied to content- and user-based classifications in networks.
Besides Alcatel Lucent, we are mostly working with the INRIA team MAESTRO.

FRANCE TELECOM R&D (Cifre, INRIA, 2005-2008). The subject of this industrial contract was "Network load balancing on layer 7 switching for high performance and high available Linux based platforms".

ANAGRAN (INRIA, 2008). We have designed and ran experimentations of the ANAGRAN FR router within GRID'5000.

5 Objectives for the next four years

5.1 Context

After a 8 years lifespan, the scientific contours of the INRIA team-project RESO have naturally evolved, due notably to three major transformations:

- the cessation of the virtualization activity with the departure of P. Vicat-Blanc to create the Lyatiss, a Inria RESO spinoff;
- the rising power of the metrology axis impulsed by pooling the internal competence in signal processing, performance evaluation and distributed algorithmic;

[21] `http://inria.bell-labs.commonlab.homeip.net/`

- the recognition of energy efficiency as an economical and societal challenge prompting the idea of energy efficient large scale distributed software infrastructures.

We believe that the last two items are now sufficiently mature to split and to participate to the creation of two new INRIA initiatives that will emphasize the visibility of our themes. One, centered around the *analysis of dynamic networks*, will involve from the RESO side, P. Gonçalves (INRIA), T. Begin (U. Lyon 1) and I. Guérin Lassous (U. Lyon 1). The other initiative relates to the activities of L. Lefèvre (INRIA), J.-P. Gelas (U. Lyon 1) and O. Glûck (U. Lyon 1) on the *energy efficiency in networks*. Both projects will be hosted at the LIP, ENS Lyon, and their respective goals are summarized in the following sections (with full texts in annex).

5.2 Dynamic Networks: Temporal and Structural Capture Approach

The scientific lines of this proposition stem from a joint reflexion between a part of RESO (T. Begin, P. Gonçalves and I. Guérin Lassous) and D-Net (E. Fleury and C. Crespelle) a companion INRIA project at ENS Lyon. Although the two had distinct motivations at their origin, we meet on the systems we analyze and their nature, and we both endeavor at understanding the complex mechanisms that cause their transformations. It is this convergence that led us to pool our efforts and to combine our competences in *Graph Theory*, *Signal Processing*, *Performance Evaluation* and *Distributed Algorithms* to tackle the study of "Dynamic networks". The now widely spread notion of "Network Science" covers a variety of organized systems where the agents can communicate together, as it is the case for instance with the technological networks such as Internet (thoroughly studied by RESO) or with social, collaborative and communication networks or even epidemiological contact nets (as monitored and modeled by D-Net)... One originality of our proposition dwells in our ability to collect large sets of interactive data, or footprints, allowing us for inferring the relations that bond the different network's agents. The considerable size of those systems has fostered the development of statistical / stochastic methods able to reveal and to describe their (often) multi-scale organization and from there to deduce significant information. This approach stands at the core of Network Science and has proven particularly efficient at identifying universal organizing principles and at clarifying the relationship between the topology and the function of the systems. So far, yet, Network Science has neglected the dynamic aspect, which is characteristic of several social, biological of information networks. Our main objective is thus to work out, based on a multi-disciplinary foundation, a theoretical and methodological framework devoted to time-space networks, able to connect structural features (viz. vertices / entities or edges / links) with dynamic properties. To this end, we necessarily have to consider the four modules of a processing chain and to address the scientific issues therein:
Monitoring – Which information is pertinent for dynamic networks? How to measure it? Which space and time resolutions and which refreshing rate?
Analysis – Is there a causality between structural organization, exchanges density and dynamic properties? Can we expect universal features common to dynamic systems. How to quantify them? How to estimate them?
Modeling – What is the correct level to be modeled (agents' behavior, connections randomness)? Can we discard some complexity aspects? If so, what are the consequences? Are the models always identifiable? easy to simulate?
Algorithms – How to conceive appropriate algorithms? The control plane and the performance evaluation of the underlying dynamic processes, have to react and to adapt to the changes in topology and to the interactions variability. Overhead of the distributed

algorithm deployed to monitor the system state must strive not to degrade its good functioning.

With their own tools, D-NET and RESO have acquired a specific expertise with respect to empirical data analysis. We wish to pursue in this direction, and to leverage the complementarity of our skills to develop in a broaden theoretical framework, generic methods for detecting structural changes in large scale networks. Our proposal stands a the cross-road of the mathematics and computer science domains and will profit by our implication in the Complex Institute (IXXI) of ENS Lyon.

5.3 Large Algorithms and Software Architectures for Service Oriented Platforms

The other half of RESO permanent members (J.-P. Gelas, O. Gluck and L. Lefevre) and associated colleagues (PhDs students and Engineers) will join the new INRIA initiative AVALON and continue to explore their activities on energy efficiency.

The fast evolution of hardware capabilities in terms of wide area communication as well as of machine virtualization leads to the requirement of another step in the abstraction of resources with respect to applications. Those large scale platforms based on the aggregation of large clusters (Grids), huge datacenters (Clouds), collections of volunteer PC (Desktop computing platforms), or high performance machines (Supercomputers) are now available to researchers of different fields of science as well as to private companies. This variety of platforms and the way they are accessed have also an important impact on how applications are designed (i.e. the programming model used) as well as how applications are executed (i.e. the runtime/middleware system used). The access to these platforms is driven through the use of different services providing mandatory features such as security, resource discovery, virtualization, load-balancing, monitoring, etc. Platform as a Service (PaaS) and Software as a Service (SaaS) have thus to play an important role in the future development of large scale applications.

The overall idea of the AVALON research team is to consider the whole system, ranging from resources to applications, to be able to design adequate programming and resources abstractions to provide simple to use abstractions to programmers while enabling efficient exploitation of resources. More precisely, the team is going to focus on large-scale, heterogeneous, and elastic resources – ranging from supercomputers to Clouds and Grids, and to services, component and workflow models – including some domain specific models such as GridRPC or MapReduce. In addition to well-known metrics such as makespan, resource utilisation, etc, the team will consider other metrics, such as energy efficiency or consumption for example.

RESO members will provide their expertise in energy efficiency in the context of application and resources models because they will in particular require to deal with energy concerns, as energy price is becoming a major limiting factor for large scale infrastructures.

Physically monitoring the energy consumption of few resources is now becoming a reality; injecting such local measurements as a new parameter in multi-objective optimization models is also more and more common. However, dealing with energy consumption and energy efficiency at large scale is still a real challenge. This activity, initiated in the RESO team since 2008, will be continued by the AVALON team by investigating energy consumption and efficiency on large scale (external, internal) monitoring of resources. Also, while physical resources start to be well mastered, another challenge is to deal with virtualized resources and environments.

Moreover, in recent years, we have seen the apparition of hardware based green leverages (on/off, idle modes, dynamic frequency and speed scaling, etc) applied to various

kinds of physical resources (cpu, memory, storage and communications). To exploit them, these facilities must be incorporated into middle- ware software layers (schedulers, resource managers, etc). AVALON will explore the benefits of such lever- ages, for example with respect to elasticity, to improve the energy efficiency of distributed applications and services and to limit the energy consumption of platforms. The goal will be to provide the needed amount of physical and virtual resources to fulfill the needs of applications. Such provision is greatly influenced by a large set of contextual choices (hardware infrastructures, software, location, etc).

6 Bibliography of the project-team

Books and Monographs

[1] Pascale Vicat-Blanc Primet. *Cluster and Computing Networks*. Wiley-ISTE, June 2011.

[2] Pascale Vicat-Blanc Primet, Sebastien Soudan, Romaric Guillier, and B. Goglin. *Réseaux de calcul - Des grappes aux nuages de calcul*. Hermes Sciences Publications, 2011.

Books or proceedings edition

[3] P. Abry, P. Gonçalves, and J. Lévy Véhel. *Scaling, Fractals and Wavelets*. Digital signal and image processing series. ISTE – John Wiley & Sons, Inc., London (UK), 2009.

[4] L. Bononi and I. Guérin Lassous, editors. *Fifth ACM International Symposium on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks (PE-WASUN)*, Vancouver, Canada, October 2008. ACM.

[5] A. Boukerche and I. Guérin Lassous, editors. *Sixth ACM International Symposium on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks (PE-WASUN)*, Canary Islands, Spain, October 2009. ACM.

[6] I. Guérin Lassous, editor. *Seventh ACM International Symposium on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks (PE-WASUN)*, Bodrum, Turkey, October 2010. ACM.

[7] Zhiyi Huang, John Hine, Laurent Lefèvre, Tony McGregor, Yi Pan, Hong Shen, and Zhiwei Xu. *Special Issue on "Emerging Research in Parallel and Distributed Computing" - Journal of Supercomputing*, volume 51. Journal of Supercomputing - Springer Netherlands eds., March 2010.

[8] Zhiyi Huang, Zhiwei Xu, Nathan Rountree, Laurent Lefèvre, Hong Shen, John Hine, and Yi Pan. *Proceedings of PDCAT 2008 : The Ninth International Conference on Parallel and Distributed Computing, Applications and Technologies*. IEEE Computer Society, Dunedin, New Zealand, December 2008.

[9] Laurent Lefèvre and Jean-Marc Pierson. *Special theme ERCIM News : Towards Green ICT*, volume 79. ERCIM EEIG, October 2009.

[10] Thierry Priol, Laurent Lefèvre, and Rajkumar Buyya. *Proceedings of CCGrid2008 : Eight IEEE International Symposium on Cluster Computing and Grid*. IEEE, Lyon, France, May 2008.

[11] Piero Spinnato, Pascale Vicat-Blanc Primet, Chris Edwards, and Michael Welzl, editors. *Special Section on Networks for Grid Applications*. Elsevier, april 2009. International Journal on Future Generation Computer Systems.

[12] Pascale Vicat-Blanc Primet, Tomohiro Kudoh, and Joe Mambretti, editors. *Networks for Grid Applications*. Springer, Beijing, China, june 2009.

[13] Pascale Vicat-Blanc Primet, Joe Mambretti, and Tomohiro Kudoh. *Proceedings of the 2nd International ACM/ICST GRIDNETS conference on Networks for Grid Applications*. oct 2008.

Doctoral dissertations and "Habilitation" theses

[14] Fabienne Anhalt. Analyse et évaluation d'un routeur virtuel basé sur XEN. Master's thesis, INSA de Lyon, Jun 2008.

[15] Fabienne Anhalt. Intégration de la ressource routeur virtuel au logiciel HIPerNET de composition et d'exploitation d'infrastructures virtuelles à la demande. Master's thesis, INSA de Lyon, Sep 2008.

[16] Fabienne Anhalt. *Network virtualization: performance, sharing and applications*. PhD thesis, ENS-Lyon, Université de Lyon, July 2011.

[17] Narjess Ayari. *Contributions to Session Aware Frameworks for Next Generation Internet Services*. PhD thesis, École Normale Supérieure de Lyon, October 2008.

[18] Dinil Mon Divakaran. *Dealing with Elephants in the Internet: Towards a Flow-Based Internet Architecture*. PhD thesis, ENS-Lyon, Université de Lyon, July 2010.

[19] P. Gonçalves. *Analyse temps-fréquence et modèles d'invariance d'échelle : Contribution à l'étude de systèmes complexes*. Habilitation à diriger des recherches, École Normale Supérieure de Lyon, Mathématiques Informatique Fondamentale, 2010.

[20] Romaric Guillier. *Methodologies and Tools for the Evaluation of Transport Protocols in the Context of Highspeed Networks*. PhD thesis, ENS-Lyon, Université de Lyon, September 2009.

[21] Ludovic Hablot. *Réseau longue distance et application distribuée dans les grilles de calcul : étude et propositions pour une interaction efficace*. PhD thesis, ENS Lyon, Université de Lyon, 2009.

[22] Guilherme Koslovski. *Dynamically provisioned virtual infrastructures: specification, allocation and execution*. PhD thesis, ENS-Lyon, Université de Lyon, July 2011.

[23] Guilherme Koslovski, Pascale Vicat-Blanc Primet, and Andrea Schwertner Charão. VXDL: Uma Linguagem para Descrićão de Interconexões e Recursos em Grades Virtuais. Master's thesis, Programa de Pós-Graduaćão em Informática, Brasil, August 2008.

[24] Patrick Loiseau. *Contributions to the analysis of scaling behavior and quality of service in networks: experimental and theoretical aspects.* PhD thesis, École Normale Supérieure de Lyon, December 2009.

[25] Anne-Cécile Orgerie. Energy-Aware Frameworks for High-Performance Data Transport and Computations in Large-Scale Distributed Systems. Master's thesis, Ecole Normale Supérieure de Lyon, Jun 2008.

[26] Anne-Cécile Orgerie. *An Energy-Efficient Reservation Framework for Large-Scale Distributed Systems.* PhD thesis, École Normale Supérieure de Lyon - France, Sept. 2011.

[27] Dino M. Lopez P. *Propositions for a robust and inter-operable eXplicit Control Protocol on heterogeneous high speed networks.* PhD thesis, École Normale Supérieure de Lyon, June 2008.

[28] Vannier Rémi. *Partage de la bande passante équitable dans les réseaux ad hoc.* PhD thesis, ENS Lyon, INRIA, 2009.

[29] Sebastien Soudan. *Bandwidth Sharing and Control in High-Speed Networks: Combining Packet- and Circuit-Switching Paradigms.* PhD thesis, ENS-Lyon, Université de Lyon, 2009.

Articles in referred journals and book chapters

[30] Fabienne Anhalt, Guilherme Koslovski, and Pascale Vicat-Blanc Primet. Specifying and provisioning Virtual Infrastructures with HIPerNET. *ACM International Journal of Network Management (IJNM) - special issue on Network Virtualization and its Management*, 2010.

[31] Fabienne Anhalt and Pascale Vicat-Blanc Primet. Analysis and experimental evaluation of network data-plane virtualization mechanisms. *International Journal On Advances in Intelligent Systems*, 2009.

[32] O. Audouin, D. Barth, M. Gagnaire, C. Mouton, P. Vicat-Blanc Primet, D. Rodrigues, L. Thual, and D. Verchère. Carriocas project: Towards converged internet infrastructures supporting high performance distributed applications. *IEEE/OSA Journal of Lightwave Technology*, 27(12):1928–1940, 2009.

[33] Narjess Ayari, Denis Barbaron, Laurent Lefèvre, and Pascale Vicat-Blanc Primet. Fault tolerance for highly available internet services: Concepts, approaches, and issues. *IEEE Communications Surveys and Tutorials*, 10(2):34–46, July 2008.

[34] Narjess Ayari, Laurent Lefèvre, and Denis Barbaron. Design and evaluation of a session-aware admission-control framework for improving service providers profitability. *Journal of Internet Engineering - Special issue on Service Oriented Infrastructures*, 4(1), December 2010.

[35] J. Barral and P. Gonçalves. On the estimation of the large deviations spectrum. *Journal of Statistical Physics*, 144(6):1256–1283, 2011.

[36] Thomas Begin, Bruno Baynat, Alexandre Brandwajn, and Francis Sourd. A DFO technique to calibrate queueing models. *Computers & Operations Research*, 37(2):273 – 281, February 2009.

[37] Thomas Begin and Alexandre Brandwajn. Higher-order distributional properties in closed queueing networks. *Performance Evaluation*, 66(11):607 – 620, November 2009.

[38] Thomas Begin, Alexandre Brandwajn, Bruno Baynat, Bernd Wolfinger, and Serge Fdida. High-level approach to modeling of observed system behavior. *Performance Evaluation*, 67(5):386 – 405, May 2010.

[39] Alexandre Brandwajn and Thomas Begin. A recurrent solution of $Ph/M/c/N$-like and $Ph/M/c$-like queues. *Journal of Applied Probability*, 49(1), March 2012.

[40] H. Carrão, A. Araújo, P. Gonçalves, and M. Caetano. Multitemporal meris images for land cover mapping at national scale: the case study of portugal. *International Journal of Remote Sensing*, 31(8, 2063), 2008.

[41] H. Carrão, P. Gonçalves, and M. Caetano. Contribution of multispectral and multitemporal information from modis images to land cover classification. *Elsevier, Remote Sensing of Environment*, 112(3):986–997, 2008.

[42] H. Carrão, P. Gonçalves, and M. Caetano. A nonlinear harmonic model for fitting satellite image time series: Analysis and prediction of land cover dynamics. *IEEE Trans. in Geosciences and Remote Sensing*, 48(4 -part 2):1919–1930, 2010.

[43] C. Chaudet and I. Guérin Lassous. *Réseaux de capteurs*, chapter Principes et protocoles d'accès au médium. ISTE/Wiley, 2009.

[44] Marcos Dias de Assunção, Alexandre di Costanzo, and Rajkumar Buyya. A cost-benefit analysis of using Cloud computing to extend the capacity of clusters. *Cluster Computing*, 13(3):335–347, September 2010.

[45] M. Doret, H Helgason, P. Abry, P. Gonçalves, C. Gharib, and P. Gaucherand. Multifractal analysis of fetal heart rate variability in fetuses with and without severe acidosis during labor. *American Journal of Perinatology*, 28(4):259–266, Nov. 2010.

[46] P. Flandrin, P. Gonçalves, and P. Abry. Scale Invariance and Wavelets. In P. Abry, P. Gonçalves, and J. Lévy Véhel, editors, *Scaling, Fractals and Wavelets*, pages 71–102. ISTE – John Wiley & Sons, Inc., 2009.

[47] Brice Goglin, Olivier Glück, and Pascale Vicat-Blanc Primet. Interaction efficace entre les réseaux rapides et le stockage distribué dans les grappes de calcul. *Technique et Science Informatiques*, 27(7/2008):911–940, September 2008.

[48] P. Gonçalves, J.-Ph. Ovarlez, and R. Baraniuk. Quadratic Time-Frequency Analysis III: The Affine Class and Other Covariant Classes. In F. Hlawatsch and F. Auger, editors, *Time-Frequency Analysis: Concepts and Methods*, pages 193–226. ISTE – John Wiley & Sons, Inc., 2008.

[49] P. Gonçalves, S. Roy, and T. Begin. Dynamic resource management in clouds: A probabilistic approach. *IEICE Transactions on Communications, special session on Networking Technologies for Cloud Services*, 2012. To appear.

[50] H. Helgason, P. Abry, P. Gonçalves, C. Gharib, P. Gaucherand, and M. Doret. Adaptive multiscale complexity analysis of fetal heart rate. *IEEE, Transactions on Biomedical Engineering*, 58(8):2186–2193, August 2011.

[51] Laurent Lefèvre. Le cout écologique d'internet - ecological cost of internet (in french). *Pour la Science*, 66:40–41, January 2010.

[52] Laurent Lefèvre and Jean-Patrick Gelas. IAN2 : Industrial Autonomic Network Node architecture for supporting personnalized network services in the industrial context. *The International Journal of Future Generation Computer Systems (FGCS) - Grid Computing: Theory, Methods and Applications*, 24(1):58–65, January 2008.

[53] Laurent Lefèvre and Anne-Cécile Orgerie. Towards energy aware reservation infrastructure for large-scale experimental distributed systems. *Parallel Processing Letters - Special Issue on Clusters and Computational Grids for Scientific Computing*, 19(3):419–433, September 2009.

[54] Laurent Lefèvre and Anne-Cécile Orgerie. Designing and evaluating an energy efficient cloud. *The Journal of SuperComputing*, 51(3):352–373, March 2010.

[55] P. Loiseau, P. Gonçalves, G. Dewaele, P. Borgnat, P. Abry, and P. Vicat-Blanc. Investigating self-similarity and heavy-tailed distributions on a large scale experimental facility. *IEEE/ACM Transactions on Networking*, 18(4):1261–1274, August 2010.

[56] P. Loiseau, P. Vicat-Blanc, and P. Gonçalves. A long-range dependent model for network traffic with flow-scale correlations. *Stochastic Models*, 27:333–361, 2011.

[57] N. Mitton, B. Sericola, S. Tixeuil, E. Fleury, and I. Guérin Lassous. Self-stabilization in Self-organized Wireless Multihop Networks,. *Ad Hoc and Sensor Wireless Networks Journal*, 11(1-2), 2011.

[58] Pablo Neira Ayuso, Rafael M. Gasca, and Laurent Lefèvre. Demystifying cluster-based fault-tolerant firewalls. *IEEE Internet Computing : Special Issue on Unwanted Traffic*, 13(6):30–37, November 2009.

[59] Pablo Neira Ayuso, Rafael M. Gasca, and Laurent Lefèvre. Communicating between the kernel and user-space in linux using netlink sockets. *Journal Software: Practice and Experience*, 40(9):797–810, August 2010.

[60] Anne-Cécile Orgerie, Marcos Dias de Assuncao, and Laurent Lefèvre. *Energy Aware Clouds*, chapter Chapter on book "Grids, Clouds and Virtualization" - Massimo Cafaro and Giovanni Aloisio editors, pages 145–170. Springer Book, October 2010.

[61] Anne-Cécile Orgerie and Laurent Lefèvre. *Energy-efficient data transfers in large-scale distributed systems*, chapter Chapter in Handbook of Energy-Aware and Green Computing, Chapman & Hall. 2011.

[62] Anne-Cécile Orgerie and Laurent Lefèvre. Eridis: Energy-efficient reservation infrastructure for large-scale distributed systems. *Parallel Processing Letters*, 21:133–154, June 2011.

[63] Anne-Cécile Orgerie and Laurent Lefèvre. *Energy-Efficient Reservation Infrastructure for Grids, Clouds and Networks*, chapter Chapter in Energy Aware Distributed Computing Systems, Wiley Series on Parallel and Distributed Computing, John Wiley & Sons. John Wiley & Sons, 2012. To appear.

[64] Anne-Cécile Orgerie, Laurent Lefèvre, and Jean-Patrick Gelas. Etudier l'usage pour économiser l'énergie dans les systèmes distribués à grande echelle : l'approche eari. *TSI : Technique et Science Informatiques*, 30(5):515–538, May 2011.

[65] Anne-Cécile Orgerie, Laurent Lefèvre, and Isabelle Guérin-Lassous. Energy-efficient bandwidth reservation for bulk data transfers in dedicated wired networks. *The Journal of SuperComputing - Special issue on Green Networks*, March 2011.

[66] T. Razafindralambo and I. Guérin Lassous. Increasing Fairness and Efficiency using the MadMac Protocol in Ad Hoc Networks. *Ad Hoc Networks*, 6(3), May 2008.

[67] T. Razafindralambo, I. Guérin Lassous, L. Iannone, and S. Fdida. Dynamic Packet Aggregation to Solve Performance Anomaly in 802.11 Wireless Networks. *Computer Networks*, 52(1), January 2008.

[68] J. Rubio-Loyola, A. Galis, A. Astorga, J. Serrat, L. Lefevre, A. Fischer, A. Paler, and H. de Meer. Scalable service deployment on software defined networks. *IEEE Communications Magazine*, 49(12):84–93, December 2011.

[69] S. Sarr, C. Chaudet, G. Chelius, and I. Guérin Lassous. Bandwidth Estimation for IEEE 802.11-based Ad Hoc Networks. *IEEE Transactions on Mobile Computing*, 7(10), 2008.

[70] S. Sarr, S. Khalfallah, and I. Guérin Lassous. Dynamic Bandwidth Management in IEEE 802.11-Based Multihop Wireless Networks. *Annals of Telecommunications*, 2011.

[71] Sebastien Soudan, Binbin Chen, and Pascale Vicat-Blanc Primet. Flow scheduling and endpoint rate control in gridnetworks. *Future Generation Computer Systems*, 2008.

[72] Piero Spinnato, Pascale Vicat-Blanc Primet, Chris Edwards, and Michael Welzl. Editorial special section on networks for grid applications. *International Journal on Future Generation Computer Systems*, april 2009.

[73] Tram Truong Huu, Guilherme Koslovski, Fabienne Anhalt, Pascale Vicat-Blanc Primet, and Johan Montagnat. Joint elastic cloud and virtual network framework for application performance optimization and cost reduction. *Journal of Grid Computing (JoGC)*, 2010.

[74] Pascale Vicat-Blanc, Sergi Figuerola, Xiaomin Chen, Giada Landi, Eduard Escalona, Chris Develder, Anna Tzanakaki, Yuri Demchenko, Joan A. García Espín, Jordi Ferrer, Ester López, Sébastien Soudan, Jens Buysse, Admela Jukan, Nicola Ciulli, Marc Brogle, Luuk van Laarhoven, Bartosz Belter, Fabienne Anhalt, Reza Nejabati, Dimitra Simeonidou, Canh Ngo, Cees de Laat, Matteo Biancani, Michael Roth, Pasquale Donadio, Javier Jiménez, Monika Antoniak-Lewandowska, and Ashwin Gumaste. *Bringing Optical Networks to the Cloud: an Architecture for a Sustainable future Internet*, chapter Chapter on FIA book. Springer Lecture Notes, April 2011.

[75] Pascale Vicat-Blanc Primet, Sebastien Soudan, and Dominique Verchere. Virtualizing and scheduling optical network infrastructure for emerging it services. *Optical Networks for the Future Internet (special issue of Journal of Optical Communications and Networking (JOCN))*, 1(2):A121–A132, 2009.

Publications in Conferences and Workshops

[76] P. Abry, H. Helgason, P. Gonçalves, E. Pereira de Souza Neto, M. Doret, and C. Gharib. Multifractal analysis for in partum fetal-ECG diagnosis. In *IEEE Int. Conf. on Acoust. Speech and Sig. Proc.*, Dallas, TX (US), March 2010.

[77] P. Abry, H. Wendt, S. Jaffard, H. Helgason, P. Gonçalves, E. Pereira, Cl. Gharib, P. Gaucherand, and M. Doret. Methodology for multifractal analysis of heart rate variability: From lf/hf ratio to wavelet leaders. In *32nd Int. Conf. of the IEEE Engineering in Medicine and Biology Society*, Buenos Aires (Argentina), Sept. 2010.

[78] Andrei Agapi, Sebastien Soudan, Marcelo Pasin, Pascale Pascale Vicat-Blanc Primet, and Thilo Kielmann. Optimizing deadline-driven bulk data transfers in overlay networks. In *ICCCN 2009 Track on Pervasive Computing and Grid Networking (PCGN)*, San Francisco, USA, 2009.

[79] Doreid Ammar, Thomas Begin, and Isabelle Guérin-Lassous. A new tool for generating realistic internet traffic in $NS-3$. In *Proceedings of the 4th International ICST Conference on Simulation Tools and Techniques, IEEE SIMUTools 2011*, page 3, Barcelona, Spain, March 2011.

[80] Doreid Ammar, Thomas Begin, Isabelle Guérin-Lassous, and Ludovic Noirie. Contrôles d'admission basés sur des mesures : Evaluation et comparaison de solutions. In *Proceedings of the 15th Colloque Francophone sur l'Ingénierie des Protocoles, CFIP11*, page 12, Sainte Maxime, France, May 2011. Hermès.

[81] Doreid Ammar, Thomas Begin, Isabelle Guérin-Lassous, and Ludovic Noirie. Evaluation and comparison of mbac solutions. In *Proceedings of the 36th Conference on Local Computer Networks, IEEE LCN 2011*, page 4, Bonn, Germany, October 2011.

[82] Lachlan Andrew, Cesar Marcondes, Sally Floyd, Lawrence Dunn, Romaric Guillier, Wang Gang, Lars Eggert, Sangtae Ha, and Injong Rhee. Towards a common tcp evaluation suite. In *PFLDnet 2008*, March 2008.

[83] Fabienne Anhalt, Dinil Mon Divakaran, and Pascale Vicat-Blanc Primet. A virtual switch architecture for hosting virtual networks on the internet. In *11th International Conference on High Performance Switching and Routing*, Dallas, Texas, USA, 6 2010.

[84] Fabienne Anhalt, Guilherme Koslovski, Marcelo Pasin, Jean-Patrick Gelas, and Pascale Vicat-Blanc Primet. Les Infrastructures Virtuelles à la demande pour un usage flexible de l'Internet. In *JDIR 09 : Journees Doctorales en Informatique et Reseaux*, Belfort, France, February 2009.

[85] Fabienne Anhalt and Pascale Vicat-Blanc Primet. Analysis and experimental evaluation of data plane virtualization with Xen. In *ICNS 09 : International Conference on Networking and Services*, Valencia, Spain, April 2009.

[86] K. Avrachenkov, P. Gonçalves, A. Legout, and M. Sokol. Graph based classification of content and users in graph based classification of content and users in bittorrent. In *NIPS Conference, Big Learning Workshop*, 2011.

[87] K. Avrachenkov, P. Gonçalves, A. Mishenin, and M. Sokol. Generalized optimization framework for graph-based semi-supervised learning. In *SIAM Data Mining*, 2012.

[88] Narjess Ayari, Denis Barbaron, and Laurent Lefèvre. Evaluating session aware admission control strategies for improving the profitability of service providers. In *The 3rd IEEE Workshop on Enabling the Future Service-Oriented Internet: Towards Socially-Aware Networks - Held in conjunction with IEEE GLOBECOM 2009*, Honolulu, USA, December 2009.

[89] Narjess Ayari, Denis Barbaron, Laurent Lefèvre, and Pascale Vicat-Blanc Primet. On improving the reliability of cluster based voice over ip services. In *FastAbstract : DSN 2008 : The 38th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, Anchorage, Alaska, USA, June 2008.

[90] Narjess Ayari, Denis Barbaron, Laurent Lefèvre, and Pascale Vicat-Blanc Primet. A session aware admission control scheme for next generation ip services. In *Fifth annual IEEE Consumer Communications and Networking Conference (IEEE CCNC 2008)*, Las Vegas, USA, January 2008.

[91] Narjess Ayari, Laurent Lefèvre, and Denis Barbaron. On improving the reliability of internet services through active replication. In *PDCAT 2008 : The Ninth International Conference on Parallel and Distributed Computing, Applications and Technologies*, pages 259–262, Dunedin, New Zealand, December 2008.

[92] Narjess Ayari, Pablo Neira Ayuso, Laurent Lefèvre, and Denis Barbaron. Towards a dependable architecture for highly available internet services. In *ARES'08 : The Third International Conference on Availability, Reliability and Security*, pages 422–427, Barcelona, Spain, March 2008.

[93] S. Baykut, P. Gonçalves, P.-H. Luppi, P. Abry, E. Pereira de Souza Neto, and D. Gervasoni. Emd-based analysis of rat eeg data for sleep state classification. In *Biosignals*. Springer, January 2009.

[94] Thomas Begin and Alexandre Brandwajn. Note sur la simulation d'une file $M/G/1$ selon la distribution du temps de service. In *Proceedings of the 12th Rencontres Francophones sur les Aspects Algorithmiques de Télécommunications, ALGOTEL10*, page 4, Belle Dune, France, June 2010.

[95] Thomas Begin and Alexandre Brandwajn. Une solution approchée pour les files $Ph/Ph/1$ et $Ph/Ph/1/N$. In *Proceedings of the 13th Rencontres Francophones sur les Aspects Algorithmiques de Télécommunications, ALGOTEL11*, page 4, Cap Estérel, France, Mai 2011.

[96] Francoise Berthou, Robert Ferret, and Laurent Lefèvre. Consommation énergétique des datacentres : quand la politique européenne s'en mêle. In *JRES 2011 : 9 ème Journées Réseaux de l'enseignement supérieur et de la recherche*, Toulouse, France, November 2011.

[97] Mehdi Bezahaf, Thomas Begin, Bruno Baynat, and Serge Fdida. Note sur les performances de TCP dans un environnement sans-fil multisaut. In *Proceedings of the 14th Colloque Francophone sur l'Ingénierie des Protocoles, CFIP09*, page 2, Strasbourg, France, October 2009.

[98] Alexandre Brandwajn and Thomas Begin. A note on aspects of workload characterization in parallel access volumes. In *Proceedings of 19th the Computer Measurement Group, CMG09*, page 6, Dallais, US, December 2009.

[99] Alexandre Brandwajn and Thomas Begin. A note on the effects of service time distribution in the M/G/1 queue. In *Proceedings of the Standard Performance Evaluation Corporation Benchmark Workshop, SPEC09*, pages 138–144, Austin, Texas, January 2009.

[100] Alexandre Brandwajn and Thomas Begin. Preliminary results on a simple approach to $G/G/c$-like queues. In *Proceedings of the 16th International Conference on Analytical and Stochastic Modelling Techniques and Applications, ASMTA09*, page 15, Madrid, Spain, June 2009.

[101] Alexandre Brandwajn and Thomas Begin. Performance evaluation of a single node with general arrivals and service. In *Proceedings of the 18th International Conference on Analytical and Stochastic Modelling Techniques and Applications, ASMTA11*, page 15, Venice, Italy, June 2011.

[102] Alexandre Brandwajn and Thomas Begin. An approximate solution for $Ph/Ph/1$ and $Ph/Ph/1/N$ queues. In *Proceedings of the 3rd International Conference on Performance Engineering, ICPE12*, page 6, Boston, Massachusetts, April 2012.

[103] Lawrence Chang, Alex Galis, Bertrand Mathieu, Kerry Jean, Roel Ocampo, Lefteris Mamatas, Javier R. Loyola, Joan Serrat, Andreas Berl, Hermann de Meer, Steven Davy, Zeinab Movahedi, and Laurent Lefèvre. Self-organising management overlays for future internet services. In *MACE 2008 : 3rd IEEE International Workshop on Modelling Autonomic Communications Environments during the ManWeek 2008 : 4th International Week on Management of Networks and Services*, pages 74–89, Samos Island, Greece, September 2008.

[104] Georges Da-Costa, Marcos Dias de Assuncao, Jean-Patrick Gelas, Yiannis Georgiou, Laurent Lefèvre, Anne-Cécile Orgerie, Jean-Marc Pierson, Olivier Richard, and Amal Sayah. Multi-facet approach to reduce energy consumption in clouds and grids: The green-net framework. In *e-Energy 2010 : First International Conference on Energy-Efficient Computing and Networking*, pages 95–104, Passau, Germany, April 2010.

[105] Georges Da-Costa, Jean-Patrick Gelas, Yiannis Georgiou, Laurent Lefèvre, Anne-Cécile Orgerie, Jean-Marc Pierson, Olivier Richard, and Kamal Sharma. The green-net framework: Energy efficiency in large scale distributed systems. In *HPPAC 2009 : High Performance Power Aware Computing Workshop in conjunction with IPDPS 2009*, Rome, Italy, May 2009.

[106] Marcos Dias de Assuncao, Jean-Patrick Gelas, Laurent Lefèvre, and Anne-Cécile Orgerie. The green grid5000: Instrumenting a grid with energy sensors. In *5th International Workshop on Distributed Cooperative Laboratories: Instrumenting the Grid (INGRID 2010)*, Poznan, Poland, May 2010.

[107] Marcos Dias de Assuncao, Anne-Cécile Orgerie, and Laurent Lefèvre. An analysis of power consumption logs from a monitored grid site. In *IEEE/ACM International Conference on Green Computing and Communications (GreenCom-2010)*, pages 61–68, Hangzhou, China, December 2010.

[108] Mohammed Diouri, Olivier Gluck, and Laurent Lefevre. Vers des machines exaflopiques vertes. In *Renpar 20 : Rencontres francophones du Parallélisme*, May 2011.

[109] Dinil Mon Divakaran, Eitan Altman, Georg Post, Ludovic Noirie, and Pascale Vicat-Blanc Primet. Analysis of the effects of XLFrames in a network. In *IFIP/TC6 NETWORKING 2009*, pages 364–377, May 2009.

[110] Dinil Mon Divakaran, Eitan Altman, Georg Post, Ludovic Noirie, and Pascale Vicat-Blanc Primet. From packets to XLFrames: sand and rocks for transfer of mice and

elephants. In *IEEE INFOCOM 2009 Workshop on High-Speed Networks*, pages 1–6, Rio di Janerio, Brazil, Apr 2009.

[111] Dinil Mon Divakaran, Fabienne Anhalt, Eitan Altman, and Pascale Vicat-Blanc Primet. Size-Based flow scheduling in a CICQ switch. In *11th International Conference on High Performance Switching and Routing (HPSR 2010)*, pages 140–145, Dallas, Texas, USA, June 2010.

[112] Dinil Mon Divakaran, Giovanna Carofiglio, Eitan Altman, and Pascale Vicat-Blanc Primet. A Flow Scheduler Architecture. In *IFIP/TC6 NETWORKING 2010*, pages 122–134, Chennai, India, May 2010.

[113] I. Doghri, L. Reynaud, and I. Guérin Lassous. On The Recovery Performance of Single- and Multipath OLSR in Wireless Multi-Hop Networks. In *Third International ICST Conference on Ad Hoc Networks (ADHOCNETS),*, Paris, France, September 2011.

[114] Sébastien Doirieux, Bruno Baynat, and Thomas Begin. On finding the right balance between fairness and efficiency in WiMAX scheduling through analytical modeling. In *Proceedings of the 17th IEEE/ACM International Symposium on Modelling, Analysis and Simulation of Computer and Telecommunication Systems, MASCOTS09*, page 10, London, England, September 2009.

[115] E. Dramitinos and I. Guérin Lassous. A Bandwidth Allocation Mechanism for 4G. In *European Wireless Technology Conference*, Roma, Italy, September 2009.

[116] E. Dramitinos, I. Guérin Lassous, and R. Vannier. A Performance Evaluation Framework for Fair Solutions in Ad Hoc Networks. In *12-th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWIM)*, Tenerife, Canary Islands, Spain, October 2009.

[117] M. Dramitinos and I. Guérin Lassous. Un mécanisme d'allocation de bande passante dans les réseaux sans fil hétérogènes basé sur la notion d'enchère. In *Proceedings of the 12th Rencontres Francophones sur les Aspects Algorithmiques de Télécommunications, ALGOTEL10*, page 4, Belle Dune, France, June 2010.

[118] Gilles Fedak, Jean-Patrick Gelas, Thomas Hérault, Victor Iniesta, Derrick Kondo, Laurent Lefèvre, Paul Malecot, Lucas Nussbaum, Ala Rezmerita, and Olivier Richard. DSL-Lab: a Low-Power Lightweight Platform to Experiment on Domestic Broadband Internet. In *International Symposium on Parallel and Distributed Computing (ISPDC 2010)*, Istanbul, Turkey, July 2010.

[119] Alex Galis, Stuart Clayman, Laurent Lefevre, Andreas Fischer, Hermann de Meer, Javier Rubio-Loyola, Joan Serrat, and Steven Davy. Towards in-network clouds in future internet. In *The Future Internet - Future Internet Assembly 2011: Achievements and Technological Promises*, volume 6656, pages 19–33. Lecture Notes in Computer Science, Springer, May 2011. ISBN 978-3-642-20897-3.

[120] Alex Galis, Spyros Denazis, Alessandro Bassi, Pierpaolo Giacomin, Andreas Berl, Andreas Fischer, Hermann de Meer, John Srassner, Steven Davy, Daniel Macedo, Guy Pujolle, Javier R. Loyola, Joan Serrat, Laurent Lefèvre, and Abderhaman Cheniour. Management architecture and systems for future internet networks. In *FIA Book : "Towards the Future Internet - A European Research Perspective"*, pages 112–122, Prague, May 2009. IOS Press. ISBN 978-1-60750-007-0.

[121] Jean-Patrick Gelas, Olivier Mornard, and Pascale Vicat-Blanc Primet. évaluation des performances réseau dans le contexte de la virtualisation xen. In *CFIP 2008 : Colloque Francophone sur l'Ingénierie des Protocoles*, Les Arcs, France, March 2008.

[122] Olivier Glück and Jean-Christophe Mignot. Splitting TCP for MPI Applications Executed on Grids. In *2011 IEEE Ninth International Symposium on Parallel and Distributed Processing with Applications*, pages 207–212, Busan, Corée, République De, May 2011.

[123] O. Goga, P. Loiseau, and P. Gonçalves. On the impact of the flow-size distribution's tail index on network performance with tcp connections. In ACM, editor, *ACM, IFIP Performance Evaluation Review*, October 2011.

[124] O. Grémillet, P. Gonçalves, P. Vicat-Blanc, and A. Dupas. Traffic classification techniques supporting semantic networks. In *1st Int. Wireless Comm. and Mobile Comp. Conf., TRaffic Analysis and Classification W.*, Caen (France), Jun-jul 2010.

[125] I. Guérin Lassous, T. Razafindralambo, and R. Vannier. Evaluation d'un protocole de régulation de débit dans les réseaux sans fil multisauts. In *Conférence Francophone sur l'Ingénierie des Protocoles*, Strasbourg, France, October 2009.

[126] Romaric Guillier and Pascale Vicat-Blanc Primet. Methodologies and tools for exploring transport protocols in the context of high-speed networks. In *IEEE TCSC Doctoral Symposium*, May 2008.

[127] Romaric Guillier and Pascale Vicat-Blanc Primet. A User-Oriented Test Suite for Transport Protocols Comparison in DataGrid Context. In *ICOIN 2009*, January 2009.

[128] Ludovic Hablot, Olivier Glück, Jean-Christophe Mignot, and Pascale Vicat-Blanc Primet. Etude d'implémentations mpi dans une grille de calcul. In *Actes de Renpar'08*, Février 2008.

[129] Syed Hasan, Laurent Lefèvre, Zhiyi Huang, and Paul Werstein. Cross layer protocol support for live streaming media. In *AINA-08 : The IEEE 22nd International Conference on Advanced Information Networking and Applications*, pages 319–326, Okinawa, Japan, March 2008.

[130] Syed Hasan, Laurent Lefèvre, Zhiyi Huang, and Paul Werstein. Supporting large scale eresearch infrastructures with adapted live streaming capabilities. In *6th Australasian Symposium on Grid Computing and e-Research*, volume 82, pages 55–63, Wollongong, Australia, January 2008. Conferences in Research and Practice in Information Technology - Australian Computer Society.

[131] Helmut Hlavacs, Thomas Treutner, Jean-Patrick Gelas, Laurent Lefevre, and Anne-Cecile Orgerie. Energy consumption side-channel attack at virtual machines in a cloud. In *International Conference on Cloud and Green Computing (CGC 2011)*, Sydney, Australia, December 2011.

[132] G. Koslovski, S. Soudan, P. Gonçalves, and P. Vicat-Blanc. Locating virtual infrastructures: Users and inp perspectives. In *12th IEEE/IFIP International Symposium on Integrated Network Management - Special Track on Management of Cloud Services and Infrastructure*, Dublin (Ireland), May 2011.

[133] Guilherme Koslovski, Tram Truong Huu, Johan Montagnat, and Pascale Vicat-Blanc Primet. Executing distributed applications on virtualized infrastructures specified with the VXDL language and managed by the HIPerNET framework. In *First International Conference on Cloud Computing (CLOUDCOMP 2009)*, Munich, Germany, October 2009.

[134] Guilherme Koslovski, Pascale Vicat-Blanc Primet, and Andrea Schwertner Charão. VXDL: Virtual Resources and Interconnection Networks Description Language. In *GridNets 2008*, Oct. 2008.

[135] Guilherme Koslovski, Wai-Leong Yeow, Cedric Westphal, Tram Truong Huu, Johan Montagnat, and Pascale Vicat-Blanc Primet. Reliability support in virtual infrastructures. In *2nd IEEE International Conference on Cloud Computing Technology and Science (CloudCom 2010)*, Indianapolis, Indiana USA, Nov/Dec 2010. IEEE.

[136] P. Loiseau, P. Gonçalves, J. Barral, and P. Vicat-Blanc. Modeling tcp throughput: an elaborated large-deviations-based model and its empirical validation. In *IFIP, Performance*, Namur (Belgium), Nov. 2010.

[137] P. Loiseau, P. Gonçalves, S. Girard, F. Forbes, and P. Vicat-Blanc. Maximum likelihood estimation of the flow size distribution tail index from sampled packet data. In *ACM Sigmetrics*, June 2009.

[138] P. Loiseau, P. Gonçalves, R. Guillier, M. Imbert, Y. Kodama, and P. Vicat-Blanc. Metroflux: A high performance system for analyzing flow at very fine-grain. In *TridentCom*, April 2009.

[139] Dino M. Lopez Pacheco, Laurent Lefèvre, and Cong-Duc Pham. Lightweight fairness solutions for xcp and tcp cohabitation. In *IFIP/TC6 Networking 2008*, pages 715–726, Singapore, May 2008.

[140] Kashif Munir, Pascale Vicat-Blanc Primet, and Michael Welzl. Grid network dimensioning by modeling the deadline constrained bulk data transfers. In *11th IEEE International Conference on High Performance Computing and Communications (HPCC-09)*, Seoul, Korea, June 2009.

[141] Pablo Neira Ayuso, Laurent Lefèvre, and Rafael M. Gasca. hft-fw : Hybrid fault-tolerance for cluste-based stateful firewalls. In *ICPADS 2008 : The 14th IEEE International Conference on Parallel and Distributed Systems*, Melbourne, Australia, December 2008.

[142] Pablo Neira Ayuso, Rafael M. Gasca, and Laurent Lefèvre. Ft-fw: Efficient connection failover in cluster-based stateful firewall. In *PDP2008 : 16th Euromicro International Conference on Parallel, Distributed and network-based Processing*, pages 573–580, Toulouse, France, February 2008.

[143] Pablo Neira Ayuso, Rafael M. Gasca, and Laurent Lefèvre. Multiprimary support for the availability of cluster-based stateful firewalls using ft-fw. In *ESORICS 2008 : 13th European Symposium on Research in Computer Security*, pages 1–17, Malaga, Spain, October 2008.

[144] Pablo Neira Ayuso, Leonardo Maccari, Laurent Lefèvre, and Rafael M. Gasca. Stateful firewalling for wireless mesh networks. In *NTMS 2008 : The second IFIP International Conference on New Technologies, Mobility and Security*, Tangier, Morocco, November 2008.

[145] V.N. Nguyen, I. Guérin Lassous, C. Sarr, and V. Moraru. Retransmission-based Available Bandwidth Estimation in IEEE 802.11-based Multihop Wireless Networks. In *14th ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, Miami, Fl., USA, October 2011.

[146] Lucas Nussbaum. Rebuilding debian using distributed computing. In *6th International Workshop on Challenges of Large Applications in Distributed Environments (CLADE 2009), held in conjunction with the International ACM Symposium on High Performance Distributed Computing (HPDC 2009)*, June 2009.

[147] Lucas Nussbaum, Fabienne Anhalt, Olivier Mornard, and Jean-Patrick Gelas. Linux-based virtualization for hpc clusters. In *Linux Symposium 2009*, July 2009.

[148] Anne-Cécile Orgerie and Laurent Lefèvre. When clouds become green: the green open cloud architecture. In *Parco2009 : International Conference on Parallel Computing*, pages 228–237, Lyon, France, September 2009.

[149] Anne-Cécile Orgerie and Laurent Lefèvre. Energy-efficient framework for networks of large-scale distributed systems. In *ISPA 2011 : The 9th IEEE International Symposium on Parallel and Distributed Processing with Applications*, pages 250–255, Busan, Korea, May 2011.

[150] Anne-Cécile Orgerie and Laurent Lefèvre. Energy-efficient overlay for data transfers in private networks. In *IEEE International Conference on Networks (ICON 2011)*, Singapore, December 2011.

[151] Anne-Cécile Orgerie, Laurent Lefèvre, and Jean-Patrick Gelas. Chasing gaps between bursts : Towards energy efficient large scale experimental grids. In *PDCAT 2008 : The Ninth International Conference on Parallel and Distributed Computing, Applications and Technologies*, pages 381–389, Dunedin, New Zealand, December 2008.

[152] Anne-Cécile Orgerie, Laurent Lefèvre, and Jean-Patrick Gelas. Save watts in your grid: Green strategies for energy-aware framework in large scale distributed systems. In *ICPADS 2008 : The 14th IEEE International Conference on Parallel and Distributed Systems*, pages 171–178, Melbourne, Australia, December 2008.

[153] Anne-Cécile Orgerie, Laurent Lefèvre, and Jean-Patrick Gelas. Economies d'energie dans les systèmes distribués à grande echelle : l'approche eari. In *JDIR 09 : Journèes Doctorales en Informatique et Réseaux*, Belfort, France, February 2009.

[154] Anne-Cécile Orgerie, Laurent Lefèvre, and Jean-Patrick Gelas. Demystifying energy consumption in grids and clouds. In *The Work in Progress in Green Computing (WIPGC) Workshop, in conjunction with the first IEEE sponsored International Green Computing Conference*, pages 335–342, Chicago, USA, August 2010.

[155] Anne-Cécile Orgerie, Laurent Lefèvre, and Isabelle Guérin-Lassous. On the energy efficiency of centralized and decentralized management for reservation-based networks. In *IEEE Global Communications Conference (GLOBECOM 2011)*, Houston, USA, December 2011. IEEE Computer Society Press.

[156] Anne-Cécile Orgerie, Laurent Lefèvre, Isabelle Guérin-Lassous, and Dino Lopez Pacheco. Ecofen: an end-to-end energy cost model and simulator for evaluating power consumption in large-scale networks. In *SustaInet 2011 : First International*

Workshop on Sustainable Internet and Internet for Sustainability, Lucca, Italy, June 2011. IEEE Computer Society Press.

[157] R. Razafindralambo and I. Guérin Lassous. SBA: un algorithme simple de backoff pour les réseaux ad hoc. In *Algotel*, St Malo, France, May 2008.

[158] T. Razafindralambo and I. Guérin Lassous. SBA: a Simple Backoff Algorithm for Wireless Ad Hoc Networks. In *IFIP Networking*, Aachen, Germany, May 2009.

[159] Tahiry Razafindralambo, Thomas Begin, Marcelo Dias De Amorim, Isabelle Guérin-Lassous, Nathalie Mitton, and David Simplot-Ryl. Promoting quality of service in substitution networks with controlled mobility. In *Proceedings of the 10th International Conference on Ad Hoc Networks and Wireless, AdHocNow 2011*, page 14, Paderborn, Germany, July 2011.

[160] Javier Rubio-Loyola, Antonio Astorga, Joan Serrat, W. K. Chai, Lefteris Mamatas, Alex Galis, Stuart Clayman, Abderhaman Cheniour, Laurent Lefevre, Olivier Mornard, Andreas Fischer, Alexandru Paler, and Hermann de Meer. Platforms and software systems for an autonomic internet. In *IEEE Globecom 2010 - Next Generation Networking Symposium*, Miami, USA, December 2010.

[161] Javier Rubio-Loyola, Antonio Astorga, Joan Serrat, Laurent Lefevre, Abderhaman Cheniour, D. Muldowney, Steven Davy, Alex Galis, Lefteris Mamatas, Stuart Clayman, Daniel Macedo, Zeinab Movahedi, Guy Pujolle, Andreas Fischer, and Hermann de Meer. Manageability of future internet virtual networks from a practical viewpoint. In *FIA Book : "Towards the Future Internet - Emerging Trends from European Research"*, Valencia, April 2010.

[162] C. Sarr, S. Khalfallah, and I. Guérin Lassous. Gestion dynamique de la bande passante dans les réseaux ad hoc multi-sauts. In *9es Journées Doctorales Informatique et Réseau (JDIR)*, Belfort, France, January 2009.

[163] Sebastien Soudan, Dinil Mon Divakaran, Eitan Altman, and Pascale Vicat-Blanc Primet. Equilibrium in size-based scheduling systems. In *16th International Conference on Analytical and Stochastic Modelling Techniques and Applications*, pages 234–248, Madrid, Spain, Jun 2009.

[164] Sebastien Soudan and Pascale Vicat-Blanc Primet. Mixing malleable and rigid bandwidth requests for optimizing network provisioning. In *21st International Teletraffic Congress*, Paris, France, sept 2009.

[165] Ghislain Landry Tsafack, Laurent Lefevre, and Jean Patrick Gelas. On applying dtns to a delay constrained scenario in wired networks. In *International Workshop on Opportunistic and Delay/disruption-Tolerant Networking in conjunction with the 14th International Symposium on Wireless Personal Multimedia Communications*, Brest, France, October 2011.

[166] R. Vannier and I. Guérin Lassous. Partage équitable de la bande passante dans les réseaux ad hoc. In *CFIP*, Les Arcs, France, March 2008.

[167] R. Vannier and I. Guérin Lassous. Towards a Practical and Fair Rate Allocation for Multihop Wireless Networks based on a Simple Node Model. In *11th ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, Vancouver, Canada, October 2008.

[168] Dominique Verchere, Olivier Audouin, Bela Berde, Agostino Chiosi, Richard Douville, Helia Pouyllau, Pascale Vicat-Blanc Primet, Marcelo Pasin, Sebastien Soudan, Thierry Marcot, Veronique Piperaud, Remi Theillaud, Dohy Hong, Dominique Barth, Christian Cadéré, V. Reinhart, and Joanna Tomasik. Automatic network services aligned with grid application requirements in carriocas project. In *GridNets 2008*, Oct. 2008.

[169] P. Vetter, L. Lefevre, L. Gasca, K. Kanonakis, L. Kazovsky, A. Lee, C. Monney, X. Qiu, F. Saliou, and A. Wonfor. Research roadmap for green wireline access. In *Workshop on Green Communications and Networking during IEEE ICC'12*, Ottawa, Canada, June 2012. IEEE.

[170] Pascale Vicat-Blanc Primet, Fabienne Anhalt, and Guilherme Koslovski. Exploring the virtual infrastructure service concept in Grid'5000. In *20th ITC Specialist Seminar on Network Virtualization*, Hoi An, Vietnam, May 2009.

[171] Pascale Vicat-Blanc Primet, Jean-Patrick Gelas, Olivier Mornard, Guilherme Koslovski, Vincent Roca, Lionel Giraud, Johan Montagnat, and Tram Truong Huu. A scalable security model for enabling dynamic virtual private execution infrastructures on the internet. In *IEEE/ACM International Conference on Cluster Computing and the Grid (CCGrid2009)*, Shanghai, May 2009.

[172] Pascale Vicat-Blanc Primet, Jean-Patrick Gelas, Olivier Mornard, Guilherme Koslovski, Vincent Roca, Lionel Giraud, Johan Montagnat, and Tram Truong Huu. A scalable security model for enabling dynamic virtual private execution infrastructures on the internet. In *IEEE/ACM International Conference on Cluster Computing and the Grid*, Shangai (China), December 2008.

[173] Pascale Vicat-Blanc Primet, Sebastien Soudan, and Guilherme Koslovski. Lyatiss tuner. In *2nd IEEE Int. Conf. on Cloud Computing Technology and Science (CloudCom 2010)*. IEEE, 2010.

Internal Reports

[174] M. Abid, Andreas Berl, Z. Boudjemil, Abderhaman Cheniour, Steven Davy, Spyros Denazis, Alex Galis, Jean-Patrick Gelas, Claire Fahy, Andreas Fischer, Javier Rubio Loyola, Laurent Lefèvre, Daniel Macedo, Lefteris Mamatas, Hermann de Meer, Z. Movahedi, John Strassner, and H. Zimmermann. Autonomic internet initial framework - deliverable d6.1. Technical report, INRIA, August 2008.

[175] Andrei Agapi, Thilo Kielmann, Marcelo Pasin, Sebastien Soudan, and Pascale Vicat-Blanc Primet. Globally optimal bulk data transfers in overlay routing networks. Technical Report TR-0157, Institute on Grid Systems, Tools and Environments, CoreGRID - Network of Excellence, August 2008.

[176] Fabienne Anhalt, Christophe Blanchet, Romaric Guillier, Tram Truong Huu, Guilherme Koslovski, Johan Montagnat, Pascale Vicat-Blanc Primet, and Vincent Roca. HIPCAL: final report. Research Report, INRIA, 2010.

[177] Fabienne Anhalt and Pascale Vicat-Blanc Primet. Analysis and evaluation of a XEN based virtual router. Technical Report 6658, INRIA Rhône Alpes, Sep 2008.

[178] K. Avrachenkov, P. Gonçalves, A. Mishenin, and M. Sokol. Generalized optimization framework for graph-based semi-supervised learning. Technical Report 7774, INRIA, October 2011.

[179] Thomas Begin and Alexandre Brandwajn. A recurrent solution of $Ph/M/c/N$-like and $Ph/M/c$-like queues. Research Report RR-7321, INRIA, 2010.

[180] Abderhaman Cheniour, Jean-Patrick Gelas, Claire Fahy, Andreas Fischer, Javier Rubio Loyola, Laurent Lefèvre, Daniel Macedo, and Lefteris Mamatas. Autonomic internet initial service infrastructure - deliverable d5.1. Technical report, INRIA, December 2008.

[181] Abderhaman Cheniour, Jean-Patrick Gelas, Laurent Lefèvre, Lefteris Mamatas, Olivier Mornard, , and Javier Rubio Loyola. Autonomic internet autoi service plane and interfaces - deliverable d5.2. Technical report, INRIA, December 2010.

[182] Marcos Dias de Assuncao and Laurent Lefevre. Energy efficient it and infrastructure for data centres and server rooms. Technical report, Deliverable of the PrimeEnergyIT European project, October 2011.

[183] Dinil Mon Divakaran, Fabienne Anhalt, Eitan Altman, and Pascale Primet. Size-Based Flow Scheduling in a CICQ Switch. Research Report RR-7183, INRIA, 2010.

[184] Dinil Mon Divakaran, Giovanna Carofiglio, Eitan Altman, and Pascale Primet. A Flow Scheduler Architecture. Research Report RR-7133, INRIA, 2009.

[185] Dinil Mon Divakaran, Sebastien Soudan, Pascale Vicat-Blanc Primet, and Eitan Altman. A survey on core switch designs and algorithms. Research Report RR-6942, INRIA, 2009.

[186] Emmanouil Dramitinos, Rémi Vannier, and Isabelle Guérin Lassous. A utility-based framework for assessing fairness schemes in ad-hoc networks. Technical report, INRIA, 2009.

[187] M. Dramitinos and I. Guérin Lassous. Auction-based bandwidth allocation mechanisms for wireless future internet. Research Report RR-7188, INRIA, 2010.

[188] Gilles Fedak, Jean-Patrick Gelas, Thomas Herault, Victor Iniesta, Derrick Kondo, Laurent Lefevre, Paul Malecot, Lucas Nussbaum, Ala Rezmerita, and Olivier Richard. Dsl-lab: a platform to experiment on domestic broadband internet. Technical report, ANR, December 2008. Delivrable Projet ANR DSLLAB.

[189] Gilles Fedak, Jean-Patrick Gelas, Thomas Hérault, Victor Iniesta, Derrick Kondo, Laurent Lefèvre, Paul Malecot, Lucas Nussbaum, Ala Rezmerita, and Olivier Richard. DSL-Lab: a Platform to Experiment on Domestic Broadband Internet. Research Report RR-7024, INRIA, 2009.

[190] Jérôme Gallard, Adrien Lebre, Oana Goga, and Christine Morin. VMdeploy: Improving Best-Effort Job Management in Grid'5000. Research Report RR-6764, INRIA, 2008.

[191] Jean-Patrick Gelas, Olivier Mornard, and Pascale Vicat-Blanc Primet. Evaluation des performances reseau dans le contexte de la virtualisation xen. Technical Report 6409, INRIA, January 2008. "Also avalaible : LIP report RR2008-01".

[192] Romaric Guillier, Sebastien Soudan, and Pascale Vicat-Blanc Primet. Udt and tcp without congestion control for profile pursuit. Research Report 6874, INRIA, 03 2009. Also available as LIP Research Report RR2009-10.

[193] Ludovic Hablot, Olivier Glück, Jean-Christophe Mignot, Romaric Guillier, Sebastien Soudan, and Pascale Vicat-Blanc Primet. Interaction between mpi and tcp in grids. Research Report 6945, INRIA, 06 2009.

[194] P. Loiseau, P. Gonçalves, G. Dewaele, P. Borgnat, P. Abry, and P. Vicat-Blanc. Investigating self-similarity and heavy-tailed distributions on a large scale experimental facility. Technical Report 6472, INRIA, March 2008.

[195] Anne-Cécile Orgerie and Laurent Lefèvre. A year in the life of a large-scale experimental distributed system: usage of the Grid'5000 platform in 2007. Research Report 6965, INRIA, April 2009.

[196] Anne-Cécile Orgerie and Laurent Lefèvre. A year in the life of a large-scale experimental distributed system: usage of the Grid'5000 platform in 2008. Research Report 7481, INRIA, December 2010.

[197] Anne-Cécile Orgerie and Laurent Lefèvre. Greening the Clouds! Research report, Proceedings of the COST Action IC0804 on Energy Efficiency in Large Scale Distributed Systems - 1st Year, 2010.

[198] Anne-Cécile Orgerie and Laurent Lefèvre. Energy-Efficient Advance Bandwidth Reservation for Bulk Data Transfers in Dedicated Networks. Research report, Proceedings of the COST Action IC0804 on Energy Efficiency in Large Scale Distributed Systems - 2nd Year, 2011.

[199] Sebastien Soudan, Dinil Mon Divakaran, Eitan Altman, and Pascale Vicat-Blanc Primet. Equilibrium in size-based scheduling systems. Research Report 6888, INRIA, 03 2009. Also available as LIP Research Report RR2009-11.

[200] Sebastien Soudan, Dinil Mon Divakaran, Eitan Altman, and Pascale Vicat-Blanc Primet. Extending routing games to flows over time. Research Report 6931, INRIA, 05 2009.

[201] Rémi Vannier and Isabelle Guérin Lassous. Towards a practical and fair rate allocation for multihop wireless networks based on a simple node model. Research Report 6538, INRIA, 05 2008.

[202] Pascale Vicat-Blanc Primet, Jean-Patrick Gelas, Olivier Mornard, Guilherme Koslovski, Vincent Roca, Lionel Giraud, Roberto Podesta, and Victor Iniesta. Hipernet design document. Technical report, INRIA, July 2008. "Delivrable #3 : HIPCAL ANR-06-CIS-005".

[203] Pascale Vicat-Blanc Primet, Guilherme Koslovski, Fabienne Anhalt, Tram Truong Huu, and Johan Montagnat. Exploring the Virtual Infrastructure as a Service concept with HIPerNet. Research Report RR-7185, INRIA, 2010.

Miscellaneous

[204] W. Chai, Lefteris Mamatas, Alex Galis, Javier R. Loyola, Joan Serrat, Andreas Fischer, A. Paler, Y. Al-Hazmi, Andreas Berl, Hermann de Meer, Abderhaman

Cheniour, Laurent Lefèvre, Steven Davy, D. Muldowney, G. Koumoutsos, A. Bassi, and Zeinab Movahedi. Open source for future internet systems. FIA2009 : Future Internet Assembly in Stockholm - Poster, November 2009.

[205] Georges Da-Costa, Jean-Patrick Gelas, Yiannis Georgiou, Laurent Lefèvre, Anne-Cécile Orgerie, Jean-Marc Pierson, and Olivier Richard. The green-net approach for supporting energy efficient solutions in grids. Short paper and Poster in Renpar 2009 : French Meeting in Parallelism, September 2009.

[206] Marcos Dias de Assuncao, Laurent Lefèvre, and Anne-Cécile Orgerie. Investigating the energy consumption of virtual machines. Poster in E-Energy 2010 conference : The First International Conference on Energy-Efficient Computing and Networking, Passau, Germany, April 2010.

[207] Marcos Dias de Assuncao, Laurent Lefèvre, and Anne-Cécile Orgerie. Monitoring, collecting and exposing energy usage of large scale distributed infrastructures. Poster in E-Energy 2010 conference : The First International Conference on Energy-Efficient Computing and Networking, Passau, Germany, April 2010.

[208] Marcos Dias de Assuncao, Laurent Lefèvre, and Anne-Cécile Orgerie. Monitoring, collecting and exposing energy usage of large scale distributed infrastructures. Poster during Grid5000 Spring School, Lille, France, April 2010.

[209] Mohammed Diouri, Olivier Gluck, and Laurent Lefèvre. Green networking at exascale. Poster during Rescom 2011 Summer School, La Palmyre, France, June 2011.

[210] Alex Galis, Stuart Clayman, Andreas Fischer, Alexandru Paler, Yahya Al-Hazmi, Hermann de Meer, Jean-Patrick Gelas, Laurent Lefèvre, Olivier Mornard, Abderhaman Cheniour, Antonio Astorga, Joan Serrat, Javier R. Loyola, and Steven Davy. Future internet management platforms for network virtualisation and service clouds. Service Wave 2010 Demonstration, December 2010.

[211] I. Guérin Lassous. Standard pour les réseaux sans fil : IEEE 802.11. Journal Techniques de l'Ingénieur, 2009.

[212] Romaric Guillier. Méthodologies et outils pour l'évaluation des protocoles de transport dans les réseaux très haut débit. présentation, journées d'Automne RESCOM 2008, October 2008.

[213] Romaric Guillier and Pascale Vicat-Blanc Primet. Congestion collapse in grid5000. demo, Stanford Congestion Collapse workshop, April 2008.

[214] Romaric Guillier and Pascale Vicat-Blanc Primet. Pathnif. inria patent. 2009.

[215] Guilherme Koslovski and Pascale Vicat-Blanc Primet. VXDL Parser, 2009.

[216] Laurent Lefèvre. Towards green computing platforms - vers des plateformes de calcul vertes, May 2009.

[217] Laurent Lefèvre and Abderhaman Cheniour. A service enabler infrastructure for the future internet. Joint EMANICS, AutoI, Self-Net Workshop on Autonomic Management, London UK, April 2009.

[218] Laurent Lefèvre, Jean-Patrick Gelas, and Anne-Cécile Orgerie. How an experimental grid is used : the grid5000 case and its impact on energy usage. Poster CCGrid2008 : 8th IEEE International Symposium on Cluster Computing and the Grid, Lyon, France, May 2008.

[219] Laurent Lefevre, Olivier Mornard, Jean-Patrick Gelas, and Maxime Morel. Monitoring energy consumption in clouds: the compatibleone experience. Poster in Cloud and Green Computing conference, Sydney, Australia, December 2011.

[220] Laurent Lefèvre and Anne-Cécile Orgerie. EARI : Energy Aware Resource Infrastructure for Large Scale Distributed Systems. Poster INRIA Booth, Supercomputing 2008, Austin, USA, November 2008.

[221] Laurent Lefèvre and Anne-Cécile Orgerie. Green-net : Power aware software frameworks for high performance data transport and computing in large scale distributed systems. Poster National ARC Days - INRIA Sophia Antipolis, France, October 2008.

[222] Laurent Lefèvre and Anne-Cécile Orgerie. Green-net: Power aware software frameworks for high performance data transport and computing in large scale distributed systems. Poster National ARC Days, INRIA Sophia Antipolis, France, October 2008.

[223] P. Loiseau, P. Gonçalves, G. Dewaele, P. Borgnat, P. Abry, and P. Vicat-Blanc. Vérification du lien entre auto-similarité et distributions à queues lourdes sur un dispositif grande échelle, June 2008. 9 ième Atelier en Evaluation de Performances, Aussois, France.

[224] P. Loiseau, P. Gonçalves, Y. Kodama, and P. Vicat-Blanc. Metroflux: A fully operational high speed metrology platform, September 2008. Euro-NF workshop: New trends in modeling, quantitative methods and measurements, in cooperation with Net-coop, THOMSON Paris Research Labs, France.

[225] P. Loiseau, P. Gonçalves, and P. Vicat-Blanc. How TCP can kill self-similarity, September 2008. Euro-NF workshop: Traffic Engineering and Dependability in the Network of the Future, VTT, Finland.

[226] P. Loiseau, R. Guillier, O. Goga, M. Imbert, P. Gonçalves, and P. Vicat-Blanc. Automated traffic measurements and analysis in grid'5000. Best Demonstration award at ACM SIGMETRICS/PERFORMANCE (Seattle, US), June 2009.

[227] Anne-Cécile Orgerie and Laurent Lefèvre. Greening the clouds ! Poster during Rescom 2009 Summer School, La Palmyre, France, June 2009.

[228] Anne-Cécile Orgerie and Laurent Lefèvre. Towards a green grid5000. Best presentation award of the Grid5000 school, April 2009.

[229] Anne-Cécile Orgerie, Laurent Lefèvre, and Isabelle Guerin Lassous. Greening the wired networks. Poster during Rescom 2010 Summer School, Giens, France, June 2010.

[230] Anne-Cécile Orgerie, Laurent Lefèvre, and Isabelle Guerin Lassous. Greening the networks of large-scale distributed systems. Poster during Entretiens Jacques Cartier event, Colloquium on "Information and Communications Technologies: Are they Green?", Montreal, Canada, October 2011.

[231] Anne-Cécile Orgerie, Laurent Lefèvre, and Isabelle Guerin Lassous. Greening the networks of large-scale distributed systems. Poster during Rescom 2011 Summer School, La Palmyre, France, June 2011.

[232] Ghislain Landry Tsafack, Laurent Lefèvre, and Jean-Patrick Gelas. Dtn based management framework for green on/off networks. Poster during Rescom 2011 Summer School, La Palmyre, France, June 2011.