INRIA, Evaluation of Theme Réseaux et Télécoms

Project-team PLANETE

March 2012

Project-team title: PLANETE - Protocols and Applications for the Internet

Scientific leader: Walid Dabbous

Research center: Sophia Antipolis-Méditerranée and Grenoble Rhône-Alpes

1 Personnel

Personnel (November 2007)

	Misc.	INRIA	CNRS	University	**Total**
DR / Professors		2			**2**
CR / Assistant Professor		4			**4**
Permanent Engineer		1			**1**
Temporary Engineer		5			**5**
PhD Students	3	6		1	**10**
Post-Doc.					
Total	**3**	**18**		**1**	**22**
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 2012)

	Misc.	INRIA	CNRS	University	**Total**
DR / Professors		3			**3**
CR / Assistant Professor		4			**4**
Permanent Engineer		1			**1**
Temporary Engineer		6			**6**
PhD Students	1	8		1	**10**
Post-Doc.		3			**3**
Total	**1**	**25**		**1**	**27**
External Collaborators					
Visitors (> 1 month)					

Changes in staff

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

Comments:

Mohamed Ali Kaafar was recruited as Chargé de recherche in 2008, Thierry Turletti has been promoted as Directeur de Recherche in 2010.

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

- Walid Dabbous, DR1, Scientific Leader, Sophia Antipolis
- Claude Castelluccia, DR2, permanent leader for Grenoble
- Thierry Turletti, DR2, permanent leader for Sophia Antipolis
- Chadi Barakat, CR1, Sophia Antipolis
- Arnaud Legout, CR1, Sophia Antipolis
- Vincent Roca, CR1, Grenoble
- Mohamed Ali Kaafar, CR1, Grenoble
- Thierry Parmentelat, Senior Engineer, Sophia Antipolis
- Daniel Camara, Experienced Engineer, Sophia Antipolis
- Jonathan Detchart, Expert Engineer, Grenoble
- Fabrice Schuler, Expert Engineer, Grenoble
- Fréderic Urbani, Expert Engineer, Sophia Antipolis
- Julien Tribino, Associate Engineer, Sophia Antipolis
- Alina Quereilhac, Expert Engineer and PhD student, Sophia Antipolis
- Gergely Acs, Post doctoral fellow, Grenoble
- Damien Saucez, Post doctoral fellow, Sophia Antipolis
- Young Hwan Kim, Post doctoral fellow, Sophia Antipolis
- Sana Ben Hamida, PhD Student, Grenoble
- Abdelberi Chaabane, PhD Student, Grenoble
- Ludovic Jacquin, PhD Student, Grenoble
- Ferdaouss Mattoussi, PhD Student, Grenoble
- Ashwin Rao, PhD Student, Sophia Antipolis
- Anshuman Kalla, PhD Student, Sophia Antipolis

- Dong Wang, PhD Student, Grenoble
- Lukasz Olejnik, PhD Student, Grenbole
- Min-Dung Tran, PhD Student, Grenoble
- Wunan Gong, PhD Student, Sophia Antipolis

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

- Shafqat ur-Rehman, PhD defended in January 2012, looking for a Post Doctoral researcher position.
- Imed Lassoued, PhD defended in December 2011, now R&D Engineer at OneAccess, Sophia Antipolis.
- Pere Manils, PhD abandoned, now Expert Engineer at Inria Grenoble, France.
- Mathieu Lacage, PhD defended in November 2011, now with Alcmeon start-up.
- Mohamad Jaber, PhD defended in October 2011, now Post Doctoral fellow at Inria Reso project-team, Lyon, France.
- Daniele Perito, PhD defended in October 2011, now Post Doctoral fellow at UC Berkeley, USA in collaboration with the Planéte Inria project-team.
- Stevens Le Blond, PhD defended in April 2011, now Post Doctoral fellow at MPI-SWS, Kaiserslautern, Gremany.
- Naveed Bin Rais, PhD defended in February 2011, now Assistant Professor at COMSATS Institue of Information Technology, Lahore, Pakistan.
- Mohamed Karim Sbai, PhD defended in September 2010, now Post Doctoral fellow at TELECOM Bretagne, France.
- Amine Ismail, PhD defended in June 2010, now R&D Engineer at OneAccess, Sophia Antipolis, France.
- Mathieu Cunche, PhD defended in May 2010, now Researcher at NICTA, Sydney, Australia.
- Aurélien Francillon, PhD defended in October 2009, now Assistant Professor at Eurecom, Sophia Antipolis.
- Mate Soos, PhD defended in September 2009, now at Researcher Security Research Labs, Berlin, Germany.
- Diego Dujovne, PhD defended in May 2009, now Assistant Professor at Universidad Diego Portales, Santiago, Chile.
- Rodrigue Imad, was Post Doctoral fellow until June 2011, now Assistant Professor at University of Balamand, Lebanon.
- Roberto Cascella, was Post Doctoral fellow until January 2011, now Research Engineer at INRIA Rennes - Bretagne Atlantique.

- Angelo Spognardi, was Post Doctoral fellow until December 2008, now Assistant Professor at Università di Roma ”La Sapienza”.
- Amir Krifa, was Expert Engineer until July 2011, now Engineer at Sopra Group, Sophia Antipolis, France.
- Talip Baris Metin, was Expert Engineer until January 2011, now Engineer at Verivue, USA.
- Giovanni Gherdovich, was Expert Engineer until December 2010, now Engineer at SII group in Sophia Antipolis, France.
- Faker Moatamri, was Expert Engineer until March 2010, now Engineer in France (no information on current employer).
- Anil Kumar Vengalil, was Expert Engineer until February 2010, now Engineer in India (no information on current employer).
- Bilel Bin Romdhane, was Associate Engineer until January 2010, now PhD student at Eurecom, Sophia Antipolis, France.
- Mathieu Lacage, Inria Permanent Engineer & PhD student until November 2009, now Engineer at Alcmeon (start-up), Sophia Antipolis, France.
- Mads Hansen, was Associate Engineer until August 2009, now in Engineer Denmark (no information on current employer).
- Lionel Giraud, was Associate Engineer until January 2009, now Engineer at Comarch R&D, Grenoble, France.
- Jahanzeb Farooq, was Associate Engineer until September 2008, now Software Engineer at Siemens A/S, Copenhagen, Denmark.
- Mohamed Amine Chaoui, was Expert Engineer until June 2008, now Engineer at Experian, Monaco.

Last INRIA enlistments

- Mohamed Ali Kaafar was recruited as CR2 in the project-team in November 2008.

Other comments:

The Planète project-team was created in January 2001. It is expected that the two branches split and that two new project-teams are proposed. The process will start after the INRIA evaluation and should (hopefully) be finalized before the end of the year. More information about the future directions of these two project-team proposal will be provided in section 5.

2 Work progress

2.1 Keywords

Network Protocols, Wireless Networks, Security, Privacy, Monitoring, Content Centric Networking, Peer-to-Peer.

2.2 Context and overall goal of the project

The Planète project-team, located both at INRIA Sophia Antipolis - Méditerranée and INRIA Grenoble - Rhône Alpes research centers, conducts research in the domain of networking, with an emphasis on designing, implementing, and evaluating Internet protocols and applications. The main objective of the group is to propose and study new architectures, services and protocols that will enable efficient and secured communication through the Internet.

In order to cope with the growth of the Internet, the simple, original Internet architecture has accreted several hundred additional protocols and extensions. Networks based upon this significantly more complex architecture are increasingly difficult to manage in a way to ensure efficient connectivity to the over 2 billion users. The increasing, and implicit, reliance on the Internet has stimulated a major debate amongst experts as to whether the current architecture and protocols can continue to be patched, or whether they will collapse under the demands of future applications. There are signs that the current suite of protocols and solutions are becoming inadequate to cope with some common Internet trends: mobility of users and devices, unusual but legitimate traffic load (e.g., flash crowds), large heterogeneity in terms of devices and applications, delivery of real-time high-bandwidth video services, requirements for episodic connectivity, scalability in terms of number of nodes and users, complexity related to network, service and security management. Additionally, the original Internet was designed and built in an era of mutual trust, probably due to the small size of the "ARPANet" research community. Many of the protocol additions/extensions have required (still require) protection mechanisms to increase their robustness in the Internet of today where mutual trust is not granted. The volume and types of attempts to subvert the Internet will continue to increase, further stressing the current architecture. Current solutions for security are added a posteriori as a patch to overcome the limitations encountered, instead of being embedded in the system functionality. Furthermore, mobile network hosts are rapidly becoming the norm for the devices with which users access the Internet. An increasing number of the protocol additions/extensions deal with adding mobility support into the (initially wireline-focused) Internet architecture.

The network must therefore evolve to overcome the current limitations and allow adding new capabilities and services. Future networks should be available anytime and anywhere, be accessible from any communication device, support efficiently data dissemination, cope with mobility, require little or no management overhead, be resilient to failures, malicious attacks and natural disasters, and be trustworthy for all types of communication traffic. It is therefore important to address a balance of theoretical and experimental researches that expand the understanding of large, complex, heterogeneous networks, design of access and core networks based on emerging wireless and optical technologies, and keep evolving the Internet. These challenging researches include, but are not limited to, the following examples:

- Intermittent connectivity networks;
- New paradigms for data dissemination;
- Network security and privacy;
- Network measurements;
- Networking evaluation platforms.

In Planète, we have decided to tackle these challenges focusing on the following research directions that represent essential building blocks to the future Internet architecture.

Towards Data-Centric Networking

Today's Internet is characterized by high node and link heterogeneity. Nodes may vary substantially in terms of their processing, storage, communication, and energy capabilities. They may also exhibit very different mobility characteristics, from static nodes to nodes that are considerably mobile (e.g., vehicles). One of the challenges for future networks is to provide access to data anytime anywhere in the presence of high degree of heterogeneity and disconnections resulting in an *intermittent connectivity*. The disconnections may last longer than what "traditional" routing protocols (e.g., MANET routing) can handle. Several new routing paradigms (often referred to as Delay Tolerant Networking) have been proposed to handle possibly frequent, long-lived disconnections. However, a number of challenges remain, including: (1) The support of scalable and transparent integration with "traditional" routing mechanisms. (2) The study of heuristics for selecting forwarding strategies. (3) The design of unicast and multicast transmission algorithms with congestion and error control algorithms tailored for episodically connected networks and taking into account the intrinsic characteristics of flows. (4) The design of incentive-based mechanisms to ensure that nodes forward packets while preventing or limiting the impact of possible misbehaving nodes.

On the other hand, it is important to provide *content dissemination* systems even if the network does not provide efficient support for multicast. Indeed, this is an optimal data dissemination technology, that enables the creation of new commercial services, like IPTV over the Internet and multimedia content distribution on DVB-H/SH networks. This is also an efficient way to share information in WiFi, WiMax, sensor networks, or mobile ad hoc infrastructures. However, some challenges are still to be addressed namely: (1) The protocols and applications that enable the high level control of broadcasting sessions are currently missing. The goal is to enable the content provider to securely control the underlying broadcasting sessions, to be able to launch new sessions if need be, or prematurely stop an existing session and to have feedback and statistics on the past/current deliveries. (2) An Application level Forward Error Correction (AL-FEC) building block remains the cornerstone on which the whole broadcasting system relies. It is therefore important to design and evaluate new codes, capable of producing a large amount of redundancy (thereby approaching rateless codes), over very large objects, while requiring a small amount of memory/processing in order to be used on lightweight embedded systems and terminals. (3) The security building blocks and protocols that aim at providing content level security, protocol level security, and network level security must be natively and seamlessly integrated. Many components already exist. It is therefore important to identify them, know how to optimally use them, and to design/adapt the missing components, if any. (4) It is important also to seamlessly integrate these broadcasting systems to the Internet, so that users can benefit from the service, no matter where and how he is attached to the network. For instance there is a major discrepancy when considering flow control aspects, since broadcasting networks are using a constant bit rate approach while the Internet is congestion controlled.

The peer-to-peer (P2P) paradigm is another alternative for efficient data dissemination by exploiting the principle of sharing among interested users/hosts. Among the P2P protocols, BitTorrent has been largely used and his performance studied extensively, in particular its "ISP friendliness". Indeed, whereas P2P content distribution enables financial savings for the content providers, it dramatically increases the traffic on inter-ISP links. To solve this issue, the idea to keep a fraction of the P2P traffic local to each ISP was introduced a few years ago. Since then, P2P solutions exploiting locality have been introduced. However, several fundamental issues on locality still need to be explored. In particular, how far can we push locality, and what is, at the scale of the Internet, the reduction of traffic that can be achieved with locality? On the other hand, it is also important to support P2P data exchange between ad-hoc communities. However, the wireless medium and the mobility of nodes completely change the properties of peer-to-peer protocols. The dynamics becomes even more complex as it is a function of the environment and of the relative position of peers. The challenge here is to design adequate mechanisms to enhance the performance of P2P protocols in wireless and mobile infrastructures.

In addition to the above challenges, an important body of the research deals with the Internet architecture itself and with finding an alternative for it. The terms *Content Centric Networking* or *Information Centric Networking* summarize these efforts. Indeed, a major trend nowadays is that users are only interested in data "content" and do not want anymore to explicitly address where those contents are. Finding content should be a service offered by the network. The basic idea in Content Centric Networks is that the universal object is not the IP packet anymore. It becomes the "named chunk of data". Instead of establishing a conversation with a specific server, the user expresses its interest in receiving named data. Any router that stores a copy of the data can therefore send a copy. Identifying data chunks by names relaxes therefore the constraint to access these data from a specific source server. Chunk content is typically cryptographically signed and the associated degree of security and trust can be modulated depending on the sophistication of this signature.

This architecture is inherently efficient, exploiting the low cost of memory to reduce bandwidth requirements. Economies also result from using the natural multicast/dissemination facilities of the content centric approach in place of P2P applications. Mobility and even intermittent connectivity is easily handled in since data are no longer tied to a given location. Note that this architecture alleviates some of the tension related to the network neutrality debate as ISPs would provide the caching service currently managed by the content providers. In short, Information-Centric Networking attempts to achieve scalability, security and performance by moving from a location-based to a name-based retrieval and routing model and by integrating content-based security. The challenges to consider here are the corresponding naming problem, routing and resource allocation, reliable transport, data security and authentication, and content storage management.

Network security and Privacy

One of the reason of the Internet success is that it provides ubiquitous inter-connectivity. This is also one of the its main weakness since it allows to launch attacks and to exploit vulnerabilities in a large-scale basis. The Internet is vulnerable to many different attacks, for example, Distributed Denial-of Service (DDoS) attacks, epidemic attacks (Virus/Worm), spam/phishing and intrusion attacks. The Internet is not only insecure but it also infringes users' privacy. Those breaches are due to the Internet protocols but also to new applications that are being deployed (VoIP, RFID,...). A lot of research is required to improve the Internet security and privacy. For example, more research work is required to understand, model, quantify and hopefully eliminate (or at least mitigate) existing attacks. Furthermore, more and more small devices (RFIDs or sensors) are being connected to the Internet. Current security/cryptographic solutions are too expensive and current trust models are not appropriate. New protocols and solutions are required : security and privacy must be considered in the Internet architecture as an essential component. The whole Internet architecture must be reconsidered with security and privacy in mind.

We are focusing on three main domains: (1) embedded system security, where the challenge here is to make sensor network applications resilient and survivable under hostile attacks; (2) privacy leakages related to different Internet services, such as Google, BitTorrent, Skype and Tor; (3) protecting user data from unauthorized access, where the challenge is to maintain the user control and ownership of its data in an 'Informartion centric network'.

Network Measurements

In the absence of an advanced management and control plan in the Internet, and given the simplicity of the service provided by the core of the network and the increase in its heterogeneity, it is nowadays common that users experience a service degradation. This can be in the form of a pure disconnection, a decrease in the bandwidth or an increase in the delay or loss rate of packets. Service degradation can also be caused by protocol anomalies, an attack, an increase in the load, or simply a problem at the source or destination machines. Actually, it is not easy to diagnose the reasons for service degradation. Basic tools exist as ping and traceroute, but they are unable to provide detailed answers on the source of the problem nor on its location. From operator point of view, the situation is not better since an operator has only access to its own network and can hardly translate local information into end-to-end measurements. The increase in the complexity of network traffic is not easing the life of users and operators. The challenge in this direction is to come up with efficient and robust solutions for network monitoring both at the edge and in the core. On one side, it is necessary to reduce the cost of measurements without compromising accuracy (e.g., by resorting to time and space sampling and correlating measurements). On the other hand, there is a continuous need to better understand the characteristics of Internet traffic and the quality of service perceived by end users at the edge. In case of anomalies, the challenge is to detect them in a fast and accurate way and to understand their root causes. Measurements are also needed to deploy new solutions and services and tune them optimally (e.g. overlay optimization).

Networking Evaluation Platforms

Networking protocols are evaluated either by simulation, emulation or experiments on a real platform. Simulations allow a fast evaluation process, fully controlled scenarios, and reproducibility. However, they lack realism and the accuracy of the models implemented in the simulators is hard to assess. Emulation allows controlled environment and reproducibility, but it also suffers from a lack of realism. Experiments allow more realistic environment and implementations, but they lack reproducibility and ease of use. Therefore, each evaluation technique has strengths and weaknesses.

There was a general feeling in the community few years ago that simulators are not adequate tools to validate new protocols any more due to lack of realism. It was therefore important to enhance simulation on this aspect so that it becomes more realistic, and easy to interface with real platforms as this would position the simulations again as a major tool in the overall validation process.

Testing on PlanetLab has become a nearly obligatory step for an empirical research paper on a new network application or protocol to be accepted into a major networking conference or by the most prestigious networking journals. As an overlay, PlanetLab allows large scale testing and deployment without network support. However, it was built based on a centralized model and didn't have specific support for wireless links. In addition, experimenters did not have the knowledge about the networking conditions (topology, traffic) during the experiment. There was also a lack of tools ease repeating the experiments and processing and storing the experiments results. All these elements are however important to allow a more scientific approach of network protocol validation on experimental platforms. Focusing on the integration of wireless experimentation testbeds, it is important to enhance the PlanetLab testbed adding diversity with the support of wireless links on PlanetLab and shifting to a federation testbeds model.

Experimentation is evolving as a viable and realistic performance analysis approach in wireless networking research. Realism is provisioned by deploying real software (network stack, drivers, OS), and hardware (wireless cards, network equipment, etc.) in the actual physical environment. However, the experimenter is more likely to be dogged by tricky issues because of calibration problems and bugs in the software/hardware tools. This, coupled with difficulty of dealing with multitude of hardware/software parameters and unpredictable characteristics of the wireless channel in the wild, poses significant challenges in the way of experiment repeatability and reproducibility. Furthermore, experimentation has been impeded by the lack of standard definitions, measurement methodologies and full disclosure reports that are particularly important to understand the suitability of protocols and services to emerging wireless application scenarios. Lack of tools to manage large number experiment runs, deal with huge amount of measurement data and facilitate peer-verifiable analysis further complicates the process. It is therefore important to define a scientifically rigorous experimentation methodology that enables protocols benchmarking in wireless networks.

Methodology

Based on a practical view, the Planète approach to address the above research topics is to design new communication protocols or mechanisms, to implement and to evaluate them either by simulation or by experimentation on real network platforms. We also desgin and develop the adequate evaluation (simulation or experimentation) envrionments. Our work includes therefore a substantial technological component since we implement an deploy our mechanisms in pre-operational systems.

In addition to our experimentation and deployment specificities, we closely work with researchers from various domains to broaden the range of techniques we can apply to networks. In particular, we apply techniques of the information and queuing theories to evaluate the performance of protocols and systems. The collaboration with physicists and mathematicians is, from our point of view, a promising approach to find solutions that will build the future of the Internet.

In order to carry out our approach as well as possible, it is important to attend and contribute to IETF (Internet Engineering Task Force) and other standardization bodies meetings on a regular basis, in order to propose and discuss our ideas in the working groups related to our topics of interests.

2.3 Planète objectives as presented in November 2007

Here follows (in subsections 2.3.1 to 2.3.4) the description of our project-team objectives that were presented in November 2007 on the occasion of the last evaluation. Then we detail in sections 2.4 to 2.7 the achievements for each of these objectives during the 2008-2011 period.

The project-team Planète will focus on new Internet architectures for the next four years. We will tackle this problem on four different axis:

- Data-centric Networking;
- Network security;
- Network monitoring;
- Network evaluation platforms.

2.3.1 Objective 1: Data centric Networking

From the Internet design, back to 1970, the resources to be addressed and localized are computers. Indeed, at that time there were few machines interconnected, and nobody believed this number the ever be larger that a few tens of thousand of machines. Moreover, those machines where static machines with well identified resources (e.g., a given hierarchy of files) that were explicitly requested by the users. Today, the legacy of this architecture is the notion of URLs that explicitly address specific resources on a specific machine. Even if modern architectures use caches to replicate contents with DNS redirection to make those caches transparent to the end-users, this solution is only an hack that do not solve the today real problem: Users are only interested in data and do not want anymore to explicitly address where those data are. Finding data should be a service offered by the network.

In this context of data-centric network, which means that the network architecture is explicitly built to transparently support the notion of content, a data can be much more than a simple content. In such a network you can, of course, request a specific file without specifying explicitly its location, the network will transparently return with closest instance of the content. You can also request a specific service to a person without knowing its explicit network location. This is in particular the case of a VoIP or an instant messaging conversation.

We propose to work on such data centric architectures as a follow-up and federating axe for three of our current activities (adaptive multimedia transmission protocols for heterogeneous networks, data dissemination paradigms and peer-to-peer systems).

A data-centric architecture is much more than a simple modification in the naming scheme currently used in the Internet. It requires a major rethinking a many fundamental building blocks of the current Internet. The project-team Planète will focus on three specific problems in a data-centric architecture that are related to the transport of the data:

- the adaptive multimedia transmission protocols for heterogeneous networks;

- the data dissemination paradigms in multicasting/broadcasting environments;
- the data dissemination on a peer-to-peer architecture.

Adaptive multimedia transmission protocols for heterogeneous networks Today's Internet is characterized by high node and link heterogeneity. Nodes may vary substantially in terms of their processing, storage, communication, and energy capabilities. They may also exhibit very different mobility characteristics, from static nodes to nodes that are considerably mobile (e.g., vehicles). Links may be wired or wireless and thus operate at widely varying rates and exhibit quite different reliability characteristics. One of the challenges of data-centric architecture is to provide access to data anytime anywhere in the presence of high degree of heterogeneity. This means that due to a number of factors such as node mobility, link instability, power-aware protocols that, for example, turn nodes off periodically, etc., the network will not be connected all the time. Additionally, disconnections may last longer than what "traditional" routing protocols (e.g., MANET routing) can handle. These types of network, a.k.a, intermittently connected networks, or even episodically connected networks have recently received considerable attention from the networking research community. Several new routing paradigms have been proposed to handle possibly frequent, long-lived disconnections. However, a number of challenges remain, including:

- The support of scalable and transparent integration with "traditional" routing mechanisms including wired infrastructure, infrastructure-based wireless and MANET routing.
- The study of heuristics for selecting forwarding nodes (e.g., based on node's characteristics such as node's speed, node's resources, sociability level, node's historic, etc.
- The design of unicast and multicast transmission algorithms with congestion and error control algorithms tailored for episodically connected networks and taking into account the intrinsic characteristics of flows.
- The design of incentive-based mechanisms to ensure that nodes forward packets while preventing or limiting impact of possible misbehaving nodes.

The solutions proposed, which are likely to extensively use cross-layer mechanisms, will be evaluated using the methodology and the tools elaborated in our new *Experimental Platform* research direction.

Data dissemination paradigms in multicasting/broadcasting environments The multicast/broadcast content delivery systems are playing an increasingly important role in data-centric networking. Indeed, this is an optimal dissemination technology, that enables the creation of new commercial services, like IPTV over the Internet, satellite-based digital radio and multimedia transmission to vehicles, electronic service guide (ESG) and multimedia content distribution on DVB-H/SH networks. This is also an efficient way to share information in WiFi, WiMax, sensor networks, or mobile ad hoc infrastructures.

Our goal is to take advantage of our strong background in the domain to design an *efficient, robust (in particular in case of tough environments) and secure (since we* *believe that security considerations will play an increasing importance) broadcasting system.*

During this period, we will focus on the following activities:

- the protocols and applications that enable the high level control of broadcasting sessions (like the FLUTE/ALC sessions) are currently missing. The goal is to enable the content provider to securely control the underlying broadcasting sessions, to be able to launch new sessions if need be, or prematurely stop an existing session and to have feedback and statistics on the past/current deliveries.

- the AL-FEC building block remains the cornerstone on which the whole broadcasting system relies. The goal is to design and evaluate new codes, capable of producing a large amount of redundancy (thereby approaching rateless codes), over very large objects, while requiring a small amount of memory/processing in order to be used on lightweight embedded systems and terminals.

  Note that we are the leader of the CAPRI-FEC ANR/RNRT project that focuses on high performance AL-FEC codes. This project, which has been launched in March 2007 for a total duration of three years, is a major asset that will help us to reach these goals.

- the security building blocks and protocols that aim at providing content level security, protocol level security, and network level security must be natively and seamlessly integrated. This is also true of the associated protocols that enable the initialization of the elementary building blocks (e.g. in order to exchange security parameters and keys). Many components already exist. The goal here is to identify them, know how to optimally use them, and to design/adapt the missing components, if any.

- it is important that these broadcasting systems be seamlessly integrated to the Internet, so that users be able to benefit from the service, no matter where and how he is attached to the network. More precisely we will study the potential impacts of a merge of the broadcasting networks and the Internet, and how to address them. For instance there is a major discrepancy when considering flow control aspects, since broadcasting network are using a constant bit rate approach while the Internet is congestion controlled.

Data dissemination on a peer-to-peer architecture When a native broadcasting service is not enabled by the network, data should still be able to be disseminated to a large population in a scalable way. A peer-to-peer architecture support such an efficient data dissemination.

We have gain a fundamental understanding of the key algorithms of BitTorrent on the Internet. We plan to continue this work in two directions. First, we want to study how a peer-to-peer architecture can be natively supported by the network. Indeed, the client-server architecture is not robust to increase in load. The consequence is that when a site becomes suddenly popular, it usually becomes unreachable. The peer-to-peer architecture is robust to increase in load. However, a native support in the network of this architecture is a hard problem as it has implications on many components of the network (naming, addressing, transport, localization, etc.)

Second, we want to evaluate the impact of wireless and mobile infrastructures on peer-to-peer protocols. This work has recently started with the COLOR PURPURA project (in collaboration with University of Avignon) and with the European project Expeshare. The wireless medium and the mobility of nodes completely change the properties of peer-to-peer protocols. The dynamics becomes even more complex as it is a function of the environment and of the relative position of peers.

2.3.2 Objective 2: Network security

The objectives of the team for the next four years is to continue our research on wireless security, and more specifically on WSN and RFID security. There are still many research challenges to be solved and the emergence of the "Internet of things" is just in its infancy. Our past and current work on WSN and RFID security is very promising. We intend to continue this effort and focus the next four year to the design of real and deployable systems. We will also develop a new research topic on the security of the Next-Generation Internet. The important goal of this new task is to rethink about the architecture of the Internet with security as a major design requirement, instead of an after-thought.

WSN security A lot of work has been done in the area of WSN security in the last year, but we believe that this is still the beginning and a lot of research challenges need to be solved.

On the one hand it is widely believed that the sensor networks carry a great promise: Ubiquitous sensor networks will allow us to interface the physical environment with communication networks and the information infrastructure, and the potential benefits of such interfaces to society are enormous, possibly comparable in scale to the benefits created by the Internet. On the other hand, as with the advent of the Internet, there is an important associated risk and concern: How to make sensor network applications resilient and survivable under hostile attacks? We believe that the unique technical constraints and application scenarios of sensor networks call for new security techniques and protocols that operate above the link level and provide security for the sensor network application as a whole. Although this represents a huge challenge, addressing it successfully will result in a very high pay-off, since targeted security mechanisms can make sensor network operation far more reliable and thus more useful. This is the crux of our work.

The goal of our team is to design new security protocols and algorithms for constrained devices and to theoretically prove their soundness and security. Furthermore, to complement the fundamental exploration of cryptographic and security mechanisms, we will simulate and evaluate these mechanisms experimentally. As already said, this work has already started and a lot of results were already produced. However some topics still require attention. Amongst them:

- Secure Group Communications: Due to the wireless broadcast environment, group communication is common WSN. Many important services in wireless sensor networks are performed by (dynamic) groups. For example, a group of nodes might be responsible for jointly tracking a vehicle through the network. Consequently, secure protocols for group management are required, securely admitting new group members and supporting secure group communication. While significant effort has been spent to study key management for devices that have high capability or reasonably limited capabilities, group keying in resource-constraint environment remains a fundamental problem to be resolved. A major remaining challenge is to develop techniques resilient against misbehavior by protocol participants, i.e., insider attacks. Such attacks are particularly relevant in hostile settings (e.g., battlefield and law enforcement) where potential compromise or capture of WSN nodes must be taken into account. Therefore, further research is needed in order to fortify existing techniques (or devise new ones) to detect, and recover from, insider misbehavior. Developing efficient solutions that are reliable, resource-efficient (if not optimal) and self-healing is of primary concern in the context of digital battlefield. We believe that the concept of self-healing wireless sensor networks, developed by our project-team, is very promising and deserve to be extended.

- Integration within the Internet: Some of the WSN will actually be connected to the Internet and will be accessible by any authorized remote users. While most of the work so far have considered stand-alone WSN, we believe that new security protocols are required to support Internet-connected WSN. For example, in order to protect against DoS attacks, the WSN needs to verify, in a resource efficient way, that the user is authorized to query it. Furthermore, if end-to-end security is required between the user and the sensors, new encryption and security protocols need to be developed since standard protocols, such as IPSecM will, certainly, be too costly for the WSN part.

RFID security As already mentioned, the ubiquitous use of RFID tags and the development of what has become termed "the Internet of things" will lead to a variety of security threats, many of which are quite unique to RFID deployment. Already industry, government, and citizens are aware of some of the successes and some of the limitations or threats of RFID tags, and there is a great need for researchers and technology developers to take up some of daunting challenges that threaten to undermine the commercial viability of RFID tags on the one hand, or to the rights and expectations of users on the other.

The proposed research project will focus on two important issues in the use of RFID tags.

- *Device Authentication*: allows us to answer several questions such as: Is the tag legitimate? Is the reader a tag interacts with legitimate?

- *Privacy*: is the feature through which information pertaining to a tag's identity and behavior is protected from disclosure by unauthorized parties or by unauthorized means by legitimate parties such as readers.

In a public library, for example, the information openly communicated by a tagged book could include its title or author. This may be unacceptable to some readers. Alternatively, RFID- protected pharmaceutical products might reveal a person's pathology. Turning to authenticity, if the RFID tag on a batch of medicines is not legitimate, then the drugs could be counterfeit and dangerous.

Authentication and privacy are concepts that are relevant to both suppliers and consumers. Indeed, it is arguable that an RFID deployment can only be successful if all parties are satisfied that the integrity between seller and buyer respects the twin demands of authentication and privacy.

The main goal of the project, therefore, is to propose and to prototype the design of cryptographic algorithms and secure protocols for RFID deployment. These algorithms and protocols may be used individually or in combination, and we anticipate that they will aid in providing authentication or privacy. One particular feature of the research in the RFID-AP project is that the work must be practical.

Many academic proposals can be deeply flawed in practice since too little attention has been paid to the realities of implementation and deployment. This project will therefore be notable for the way theoretical work will be closely intertwined with the task of development and deployment.

The challenges to be addressed in the project are considerable. In particular there are demanding physical limits that apply to the algorithms and protocols that can be implemented on the cheapest RFID tags. While there often exist contemporary security solutions to issues such as authentication and privacy, in an RFID-based deployment they are not technically viable. And while one could consider increasing the technical capability of an RFID-tag to achieve a better range of solutions, the solution is not economically viable.

Most of this work will be done within the context of the national funded ANR RFIDAP project (2007-2010), in collaboration with the CEA Leti, France Telecom and the Eurecom institute.

Future Internet Security The current Internet has reached its limits; a number of research groups around the world are already working on future Internet architectures. The new Internet should have built-in security measures and support for wireless communication devices, among other things. A new network design is needed to overcome unwanted traffic, malware, viruses, identity theft and other threats plaguing today's Internet infrastructure and end hosts. This new design should also enforce a good balance between privacy and accountability. Several proposals in the area have been made so far, and we expect many more to appear in the near future.

Some mechanisms to mitigate the effects of security attacks exist today. However, they are far from perfect and it is a very open question how they will behave on the future Internet. Cyber criminals are very creative and new attacks (e.g. VoIP spam, SPIT) appear regularly. Furthermore, the expectation is that cyber criminals will move into new technologies as they appear, since they offer new attack opportunities, where existing countermeasures may be rendered useless.

The ultimate goal of this research project is to contribute to the work on new Internet architecture that is more resistant to today's and future security attacks. This goal is very challenging, since some of future attacks are unpredictable. We will analyze some of the established and some of the new architectural proposals, attempting to identify architectural elements and patterns that repeat from one architectural approach to another, leading to understanding how they impact the unwanted traffic issue and other security issues. Some of the more prominent elements will be rather easy to identify and understand, such as routing, forwarding, end-to-end security, etc. Others may well be much harder to identify, such as those related to data-oriented networking, e.g., caching.

The motivation for this work is that the clean slate architectures provide a unique opportunity to provide built in security capabilities that would enable the prevention of phenomenon like unwanted traffic. New architectures will most likely introduce additional name-spaces for the different fundamental objects in the network and in particular for routing objects. These names will be the fundamental elements that will be used by the new routing architectures and security must be a key consideration when evaluating the features offered by these new name-spaces.

2.3.3 Objective 3: Network Monitoring and Topology Inference

The Planète project-team will keep exploring the area of network monitoring. In addition to the work on extensions for what we have already proposed, our focus will be on the monitoring of the Internet for the purpose of problem detection and troubleshooting. Indeed, in the absence of an advanced management and control plan in the Internet, and given the simplicity of the service provided by the core of the network and the increase in its heterogeneity, it is nowadays common that users experience a service degradation. This can be in the form of a pure disconnectivity or a decrease in the bandwidth or an increase in the delay or loss rate of packets. Service degradation can be caused by protocol anomalies, an attack, an increase in the load, or simply a problem at the source or destination machines. Actually, it is not easy to diagnose the reasons for service degradation. Basic tools exist as ping and trace-route, but they are unable to provide detailed answers on the source of the problem nor on its location. From operator point of view, the situation is not better since an operator has only access to its own network and can hardly translate local information into end-to-end measurements. The increase in the complexity of networks as is the case of wireless mesh networks will not ease the life of users and operators.

The purpose of our work in this direction will be to study to which extent one can troubleshoot the current Internet either with end-to-end solutions or core network solutions. Our aim is to propose an architecture that allows end-users by collaborating together to infer the reasons for service degradation. This architecture can be purely end-to-end or can rely on some information from the core of the network as BGP routing information. We will build on this study to understand the limitations in the current Internet architecture and propose modifications that will ease the troubleshooting and make it more efficient in future network architectures. One possible direction could be the proposition of a two-layer signaling protocol a la ICMP in which edge routers are probed on end-to-end basis to collect local information on what is going on inside each network along the path. Our future contributions in this direction will be the subject of validation over large scale experimental platforms as PlanetLab and OneLab. On the funding side, our target is French and European initiatives on experimental facilities and future Internet research.

2.3.4 Objective 4: Network Evaluation Plaforms

Over the next couple of years, we are considering to build on top of our expertise of the PlanetLab software, within a co-development effort with Princeton University, and to bring the following enhancements to both the software and the PlanetLab Europe platform:

- Scalability: The PlanetLab public platform has grown up to 800+ nodes and 400+ sites over a four-year period, which is quite a success story. However the current centralized management structure is not likely to scale up to a much more important scale; this is the reason why we have started to run the PlanetLab Europe platform in a **federation** with PlanetLab Central at Princeton. The model, that is totally transparent to users/experimenters, basically allows to aggregate several experimental platforms into a unified one, offering the user the ability to use the whole platform.

  We have been active in designing and implementing this federation paradigm, and plan to go on improving it so as to support a much wider cooperation framework, by breaking the barriers between various experimental testbeds.

- Heterogeneity: On a parallel track, we are planning on studying the impact that a reservation model could have on the current PlanetLab resource allocation scheme. The objectives here are to provide a platform that would be appropriate for both continuous services – and there are quite a few of those that run on PlanetLab on a daily basis – and with experiments that need only a finite amount of time but that, on the other hand, need for various reasons to gain full access to the hardware. This would open the door to much more controlled experimental environments, that are for instance required in the wireless arena.

  This leads at last to the aspects related to virtualization. Although we do not have specific skills in the virtualization techniques per se, we are in a position to take good advantage of the many paradigms and implementation that have become increasingly fashionable recently. We believe that supporting a wider spectrum of these techniques to build the underlying experiment abstraction is likely to widely improve the capabilities of the experimental platform as a whole.

Recall that the evaluation of new network protocols and architectures is at the core of networking research. This evaluation is usually performed using simulations (e.g., NS2), emulations (e.g., Emulab), or in the wild experimental platforms (e.g., PlanetLab). Simulations allow a fast evaluation process, fully controlled scenarios, and reproducibility. However, they lack realism and the accuracy of the models implemented in the simulators is hard to assess. Emulation allows controlled environment and reproducibility, but it also suffers from a lack of realism. Experiments allow more realistic environment and implementations, but they lack reproducibility and ease of use. Therefore, each evaluation technique has strengths and weaknesses. However, there is currently no way to combine them in a scientific experimental workflow. Typical evaluation workflows are split into four steps: topology description and construction, traffic pattern description and injection, trace instrumentation description and configuration, and, analysis based on the result of the trace events and the status of the environment during the experimentation. To achieve the integration of experimental workflows among the various evaluation platforms, the two following requirements must be verified:

- Reproducibility: A common interface for each platform must be defined so that a same script can be run transparently on different platforms. This also implies a standard way to describe scenarios, which includes the research objective of the scenario, topology description and construction, the description of the traffic pattern and how it is injected into the scenario, the description and configuration of the instrumentation, and the evolution of the environment during the experimentation

- Comparability: As each platform has different limitations, a way to compare the conclusions extracted from experiments run on different platforms, or on the same platform but with different conditions (this is in particular the case for in the wild experimental platforms) must be provided.

Benchmarking is the function that provides a method of comparing the performance of various subsystems across different environments. Both reproducibility and comparability are essential to benchmarking. In order to facilitate the design of a general benchmarking methodology, we plan to integrate and automate a networking experiments workflow within the OneLab platform. This requires that we:

- Automate the definition of proper scenario definition taking in consideration available infra-structure to the experiment.
- automate the task of mapping the experimentation topology on top of the available OneLab topology. We propose to first focus on a simple one-to-one node and link mapping the beginning.
- define and provide extensive instrumentation sources within the OneLab system to allow users to gather all interesting trace events for offline analysis
- measure and provide access to ”environment variables” which measure the state of the OneLab system during an experimentation
- define an offline analysis library which can infer experimentation results and comparisons based on traces and ”environment variables”.

To make the use of these components transparent, we plan to implement them within a simulation-like system which should allow experiments to be conducted within a simulator and within the OneLab testbed through the same programming interface. The initial version will be based on the NS3 programming interface.

2.4 Work Progress on Objective 1: Towards Data Centric Networking

A worldwide effort is underway in the network community to re-design the network architecture for the future Internet, and several solutions propose changing the architecture from a host-centric to a data centric or information centric design. Indeed, the Internet was originally designed for connecting point-to-point fixed terminals using their IP addresses, which is known to lead to an inefficient behavior when either terminals or content move. Another weakness is that the Internet of today fails in finding contents by itself and in leveraging the widespread of contents for a more efficient dissemination. With the proliferation of contents of all types and the facility with which users move and share content, these problems become an impediment for the evolution and survivability of the Internet.

Furthermore, today’s Internet is characterized by high node and link heterogeneity. Nodes may vary substantially in terms of their processing, storage, communication, and energy capabilities and links may be wireless or wired and thus operate at widely varying rates and exhibit quite different reliability characteristics. This means that the network may not be connected all the time, due to a number of factors such as node mobility, link instability, etc. Additionally, disconnections may last longer than what “traditional” routing protocols (e.g., MANET routing) can handle. Such challenging environments, a.k.a, intermittently connected networks, or Delay Tolerant Networks (DTN), have recently received considerable attention from the networking research community. One of the challenges of data-centric architecture we have tackled is to provide efficient access to data anytime anywhere in the presence of high degree of heterogeneity and challenging environments.

On the other hand, multicast/broadcast content delivery systems are playing an increasingly important role in data-centric networking. Indeed, this is an optimal dissemination technology, that enables the creation of new commercial services, like

IPTV over the Internet, satellite-based digital radio and multimedia transmission to vehicles and multimedia content distribution on DVB-H/SH networks. This is also an efficient way to share information in WiFi, WiMax, sensor networks, or mobile ad hoc infrastructures. Our goal here is to take advantage of our strong background in the domain to design efficient, robust (in particular in case of challenging environments) and secure broadcasting system.

Our main contributions on this objective are summarized in Section 2.4.3.

2.4.1 Personnel

Chadi Barakat (CR1), Rao Naveed Bin Rais (PhD student), Mathieu Cunche (PhD student), Walid Dabbous (DR1), Amir Krifa (PhD student), Ferdaouss Mattoussi (PhD student), Vincent Roca (CR1), Karim Sbai (PhD student), Thierry Turletti (DR2).

2.4.2 Project-team positioning

A large number of groups around the world are contributing on data centric networking, such as Parc, Cambridge University, University of California, Helsinki IIT, KTH, SICS and ALU Bell Labs. For activities related to DTN, several groups are carrying out relevant research to ours such as UMASS, University of Waterloo, USC, UCSC, University of Cambridge, RPI, ETH, EPFL and LIP6. Concerning the AL-FEC domain, it is a highly competitive domain and people at Digital Fountain (now Qualcomm) have patented several key techniques since 1997. If this situation does not prevent research, it prevents a free usage of the research outcomes for practical systems. Therefore we have decided to stay away from techniques known to be protected, no matter how efficient they may be and to push free solutions as far as possible. This unusual standpoint turned out to be rather fruitful (our solutions are not so far from Qualcomm's AL-FEC codes' performance) and is a clear distinctive feature. For instance the work being carried out at DLR (German Aerospace Center, Germany) and at University of Bologna (Italy) on AL-FEC consider techniques that we deliberately ignored. Within Inria, there are some AL-FEC activities within the ASAP EPI (Rennes), more specifically on coordinated regenerating codes.

The originality of our approach is to combine our expertise on system, modeling and protocol design to propose new algorithms with the objective to improve current applications and enable more exciting ones. We also insist on the importance of validating our mechanisms over extensive simulations fed by realistic traces, and when possible on using experimentations made with real applications on real conditions.

2.4.3 Scientific achievements

File sharing in wireless ad hoc networks File sharing protocols, typically BitTorrent, are known to perform very well over the wired Internet where end-to-end performances are almost guaranteed. However, in wireless ad-hoc networks the situation is different due to topology constraints and the fact that nodes are at the same time peers and routers and have limited resources. We have designed a solution that minimizes the average download finish time per peer while encouraging peers to collaborate by enforcing a fair sharing of data. We validated it by simulations and extensive experiments over the well known ORBIT platform, see [SSB10, SB09, SSB09b] for details. We have implemented the solution in an application called BitHoc, which stands for BitTorrent for wireless ad hoc networks[1]. It includes two principal components: an innovative distributed membership management service and a content sharing service. BitHoc has been demonstrated at two conferences [KSBT09a, KSBT09b] and has profited from the support of the ITEA European Project Expeshare.

Routing strategies and resource management in Delay Tolerant Networks A large number of opportunistic routing strategies for DTN have been proposed by the network community. We have identified generic network characteristics that are relevant to the routing process (e.g., network density, node heterogeneity, mobility patterns) and dissected different challenged wireless networks or applications based on these characteristics. We proposed a taxonomy for intermittently connected networks with the objective to help in selecting the most appropriate routing protocol for the application or network in hand [BRTO11, SBRT$^+$11]. In many envisioned applications, participating nodes include handhelds, vehicles, sensors, etc. These various classes have diverse characteristics and mobility patterns, and can contribute quite differently to the routing process. We have shown that proposed routing solutions, which perform well in homogeneous scenarios, are not as competent in an heterogeneous setting. To this end, we proposed a class of routing schemes that can identify the nodes of highest utility for routing, improving the delay and delivery ratio by 4-5 times [STO09].

In order to enable communication in an internet connecting heterogeneous networks prone to network disruptions, we have designed a message delivery framework, called MeDeHa [RTO08, BRTO11]. MeDeHa benefits from network heterogeneity (e.g., nodes supporting more than one network and nodes having diverse resources) to improve message delivery and employs opportunistic routing to support nodes with episodic connectivity. It can bridge the connectivity gap between infrastructure-based and multi-hop infrastructure-less networks. We showcased MeDeHa interconnection of MANETs with infrastructure-based networks, allowing network coverage to be extended to regions where infrastructure deployment is sparse or nonexistent, see [BRMTO11]. We have deployed and evaluated the MeDeHa framework on a real network testbed and conducted experiments in "hybrid" scenarios running partly on ns-3 simulation and partly on real nodes. The performance obtained show significant improvement in average delivery ratio and a significant decrease in average delivery delay in the face of episodic connectivity. The MeDeHa framework was demonstrated at the ACM Mobicom conference 2010 [BRMTO10] and at the ACM Sigcomm conference in 2011 [KMR$^+$11].

We have also designed an optimal scheduling and drop policy that can optimize different performance metrics, such as the average delivery rate and the average delivery delay. We were able to associate to each message inside the network a utility value that can be calculated locally, and that allows to compare it to other messages upon scheduling and buffer congestion. Our solution called HBSD (History Based Scheduling and Drop) [KBS08a, KBS08b] integrates methods to reduce the overhead of the history-collection plane and to adapt to network conditions. [KBS12] provides an extension to a heterogeneous mobility scenario in addition to refinements to the history collection algorithm. An implementation is proposed for the DTN2 architecture as an external router and experiments have been carried out by both real trace driven simulations and experiments[2] over the SCORPION testbed at the

[1] See http://planete.inria.fr/bithoc

[2] See http://planete.inria.fr/HBSD_DTN2/

University of California Santa Cruz.

Then, we have worked on an extension of HBSD for one-to-many communications. The new framework, called MobiTrade, provides a utility driven trading system for efficient content dissemination on top of a disruption tolerant network. The trading mechanism allows a node (*merchant*) to buy, store, and carry content for other nodes (its *clients*) so that it can later trade it for content it is personally interested in. MobiTrade achieves up to 2 times higher query success rates compared to other content dissemination schemes, and can successfully isolates selfish devices[3] [KBS11].

Application-Level Forward Error Correction Codes (AL-FEC) and their applications to Broadcast/Multicast Systems The AL-FEC building block is a cornerstone on which many data broadcasting systems rely, both for file datacasting and streaming services. Our activities in the domain followed two goals.

A first goal was to design and evaluate high performance AL-FEC codes and codecs (both aspects are closely linked), capable of producing a large amount of redundancy (A.K.A. small-rate codes), over very large objects, potentially offering Unequal Erasure Protection (UEP) capabilities, while requiring a small amount of memory and processing power so that they can be used on lightweight terminals. We have demonstrated that simple LDPC-Staircase codes can achieve erasure recovery performance close to ideal codes while offering encoding and decoding times over 1Gbps, which makes them interesting solutions for practical systems [CSR$^+$08, CR09b]. We have also introduced several new AL-FEC codes to go further, like GLDPC-Staircase codes [CSR$^+$08] that easily achieve our small-rate goal with performance even closer to that of ideal codes, and a new form of QC-LDPC codes [CSR10] that enables a precise control of the recovery performance versus decoding complexity tradeoff. Finally we have introduced an efficient solution for UEP [RRSI11, RRS11a, RRS11b] that greatly simplifies their practical realization when compared to some state of the art approaches (e.g. PET and its derivatives) while offering similar or improved performance.

A second goal was to standardize these codes and the way they can be used in the file datacasting and streaming services, so that they can be used in practical systems, and to produce and distribute high quality codecs in order to promote our activities and provide proof of concepts. Three IETF RFCs have been published in the period on the subject [WBR11, RNF08, LRPP09] (additional IETF documents in progress). LDPC-Staircase codes have been recently adopted as the AL-FEC solution for ISDB-Tmm (Integrated Services Digital Broadcasting, Terrestrial Mobile Multimedia), a Japanese standard for push VOD services. The fact these codes have been preferred to a major AL-FEC competitor is the recognition of their intrinsic qualities and is a great achievement. We have also launched the http://openfec.org project, a unique initiative for the promotion of free AL-FEC codes and codecs, and we derived a high performance version that is now commercialized.

2.4.4 Collaborations

The research on file sharing in wireless networks started within a collaboration with the laboratory LIA in Avignon (the PURPURA Color Action) then profited from the support of the ITEA European project Expeshare. We had several interactions with the other partners of Expeshare, in particular with the University of Every that

[3] See http://planete.inria.fr/MobiTrade/

choose BitHoc as a reference model for the development of its video stream component. As for the DTN opportunistic routing and resource management problems, it is profiting from two main collaborations. One with Thrasyvoulos Spyropoulos from Eurecom who was previously at ETH and another one from Katia Obraczka at UCSC. This latter collaboration is part of the COMMUNITY associated team with UCSC. The activities on broadcast/multicast systems has benefited from collaborations with the Expway French company (that commercializes both our FLUTE/ALC protocol stack and our LDPC-Staircase advanced codec) and with ALU-Bell Labs (in the context of the Inria-ALU Bell Labs joint laboratory).

2.4.5 External support

The INRIA PURPURA Color Action has funded the activities on file sharing in wireless networks, and has been followed up by the ITEA European project Expeshare[4]. The activities on routing and resource management in Delay Tolerant networks is profiting from the collaboration with UCSC within the COMMUNITY associated team[5]. The contributions related to broadcast/multicast systems have been made possible thanks to ANR/CAPRI-FEC project (leaded by V. Roca) on AL-FEC techniques, the ANR/ARSSO project follow-up on robust streaming techniques, as well as an INRIA grant (2-year engineer support).

2.4.6 Self assessment

We have proposed several mechanisms that highly enhance performance of applications in future data centric networks. Important efforts have been made to validate those solutions using both simulations and experimentations. Here we would like to emphasize the difficulty to experiment in real our algorithms (time necessary to build new applications and testbeds). This research objective will be strengthened in the coming years, in particular in the area of routing and congestion control for content-oriented networks. We believe that by combining routing and congestion control, we can optimize resource consumption. We will also study the implications of using information centric networking from an economical perspective. Finally, the validation of data centric approaches, either by simulations or experimentations, need more advanced traffic and mobility models, in particular ones that take into account the social relations between devices and together with their locations. We will continue our effort in this direction for a more faithful validation of data centric networking approaches.

2.5 Work Progress on Objective 2: Network Security and Privacy

During these last 4 years, the Planète project-team had significant activities in security and privacy. In security, the team worked on embedded system (wireless sensor networks, medical devices,...) and Internet security (password security, botnet detection,...). It had major contributions in network, system security and applied cryptography with articles in top conferences.

The project also contributed to the area of online privacy and had major results. It studied privacy leakages in different Internet services, including Google, Skype, BitTorrent, Facebook, and Tor, and were able to detect many different privacy breaches. It also worked in privacy-preserving systems and architectures and

[4] See http://virtual.vtt.fi/virtual/expeshare/

[5] See http://inrg.cse.ucsc.edu/community/

proposed new solutions users to better control their data on the Internet and on in smart environments, such as smart grids or RFID systems.

2.5.1 Personnel

Claude Castelluccia (DR2), Walid Dabbous (DR1), Arnaud Legout (CR1), Mohamed Ali Kaafar (CR1), Stevens Le Blond (Phd), Pere Manils (Phd), Abdelberi Chaabane (Phd), Daniele Perito (Phd).

2.5.2 Project-team positioning

They are many research groups working in the area of security and privacy. Planète is primarily focussing on network applications. It mostly designs network attacks and secure/private networking systems. Planète works on applied cryptography, protocol design and metrology, and is interested in real applications. However, Planète does not consider all research areas. For example, Planète does not have any research activity in malware, intrusion detection and formal methods.

Planète is one of the rare groups at Inria that works on security and privacy. Madynes (Loria) works on VoIP and network management security. However, they have a system approach. Other groups work on cryptography. However, no group work on network security and privacy.

The closest group in France is the Eurecom security group (headed by Refik Molva). We have many common research interests, and actually collaborate quite intensively.

In Europe, EPFL (JP Hubaux), ETHZ (S.Capkun), BME (L.Buttyan), Leuven (C.Diaz), Cased (A.Sadegui), TUB (JP Seifert), MPI (P.Francis) are groups that work on similar research problems and with which we collaborate one way or the other (either through EU projects or by visits/exchanges).

Outside of Europe, we collaborate with UCI, Berkeley,NICT and ICT (China).

University of Washington (Tom Anderson group) and Telefonica research (Pablo Rodriguez group) are the two main leading research groups that work also on the problems address by our Bluebear project. In particular, both groups have been active in the field of large scale BitTorrent measurements and its privacy implications. However, we significantly differ in the focus. To the best or our knowledge, the bluebear project is the only one to address the issue of large scale privacy infringements from individuals without any dedicated infrastructure and privileged information.

2.5.3 Scientific achievements

Embedded System Security. The security of wireless sensor network devices relies essentially on the security of their operating systems. Most WSN devices use an Harvard architecture, where data and program memories are physically separated. It was commonly assumed that, because of this physical separation, these devices were immune to code injection attacks. However, we showed, by designing and implementing the first worm on Harvard architecture devices, that it is in fact possible to inject code on Harvard architecture devices [CFSP09]. This work utilizes a technique derived from ROP (Return Oriented Programming) that alters control flow in order to execute pre-existing instructions in the device's program memory. Those instructions are then chained to inject the malicious code in the program memory. This contribution is fundamental and had a huge impact since it invalidates a commonly used assumption. Our second important contribution in this area is related to code attestation techniques. Given that it is often impossible to guarantee the immunity of a WSN device, it is of a prominent importance to be able to verify the integrity of the device state. A solution, code attestation, is to check their integrity in order to verify that a device do not run malicious code. We showed that most of existing solutions attest only program memory and are therefore vulnerable to several attacks [RCHBC09]. Among other things we developed a stealth malicious code (rootkit) that de-installs and re-installs itself before and after attestation. This makes the attestation procedure useless. Those results shows that it is mandatory to verify all the memories of a device (RAM, ROM, EEPROM) of a WSN device in order to reinforce the security of embedded systems. We then considered the problem of Implantable Medical Devices security. IMD manufacturers have started adding wireless capabilities to many implantable medical devices, including pacemakers and cardioverter defibrillators. This allows doctors to access vital information and send commands to these devices quickly, but raises security concerns. This wireless link can in fact be used to glean personal information from such a device, to drain its batteries remotely, and to make it malfunction in dangerous ways. Securing IMDs is challenging since the solution should protect medical devices by preventing unauthorized access but still allow ease access by medical staff. We designed a new and innovative solution that relies on the physical proximity of the communicating device [RCHBC09]. A device will always be accessible from up to 10 meters away, and will normally enforce a series of authentication steps before allowing access. In an emergency, however, when the device detects that the patient using it is in trouble, it will grant access to anyone who is physically close to the patient (within about three centimeters). We finally considered the problem of secure data aggregation in wireless sensor networks. Aggregation is challenging if end-to-end privacy between sensors and the sink (or aggregate integrity) is required. We proposed a simple and provably secure encryption scheme that allows efficient additive aggregation of encrypted data [CC11]. Only one modular addition is necessary for ciphertext aggregation. The security of the scheme is based on the indistinguishability property of a pseudorandom function (PRF), a standard cryptographic primitive. We also performed a formal treatment to the security of concealed data aggregation (CDA) and the more general private data aggregation (PDA) is given. While there exist a handful of constructions, rigorous security models and analyses for CDA or PDA are still lacking. Standard security notions for public key encryption, including semantic security and indistinguishability against chosen ciphertext attacks, are refined to cover the multi-sender nature and aggregation functionality of CDA and PDA in the security model. The proposed security model is sufficiently general to cover most application scenarios and constructions of privacy-preserving data aggregation. An impossibility result on achieving security against adaptive chosen ciphertext attacks in CDA/PDA is shown. A generic CDA construction based on public key homomorphic encryption is given, along with a proof of its security in the proposed model. The security of a number of existing schemes is analyzed in the proposed model.

Impact of this work: Those results were published in prestigious conferences (e.g. ACM CCS), prestigious journals (ToSN) and received a lot of media attention. ACM CCS is undoubtedly the most prestigious conference in the field of computer security (more specifically the security of networks and systems). Our work received important media attention (articles in several newspaper, MIT Techreview, ACM news,...). Finally, the IMD solution is being patented, and discussions are in progress to transfer it to a pacemaker manufacturer.

Information Leakage on the Internet. As the amount of personal information stored at remote service providers increases, so does the danger of private information leakage. In this work, we studied privacy leakages related to different Internet services, including Google, Skype and Tor. Google records all searches made by a Google signed-in user in a Web History. The Web History is used to provide personalized results and keyword suggestions for searches that a user has already made. We designed the Historiographer, a novel inference attack that reconstructs the web search history of Google users, even though this service is supposedly protected from session hijacking by a stricter access control policy [CDCP10]. The Historiographer uses a reconstruction technique to infer search history from the personalized suggestions fed by the Google search engine. Its validity is confirmed through experiments conducted over real network traffic. We then showed that all BitTorrent downloads can be monitored from a single machine in real time [LBLL$^+$10]. Over a three months experiment, we have shown that we can collect 148 million IP addresses, downloading 1.2 million contents. In addition, we also developed an attack to identify by their IP address 70% of all BitTorrent content providers, that is the first peer who inserted a content in BitTorrent. Using an anonymizing network like Tor does not help [LBMC$^+$11] . Indeed, not only we can find the public IP address of a BitTorrent users on top of Tor, but we can also de-anonymize all the Tor traffic of this user, as long as he uses BitTorrent at least once. This work exhibits some fundamental issues in the Tor architecture, and in particular in the management of the exit nodes and in the notion of circuits. We then considered the Skype service. By leveraging on this service, we designed an attack that enables to map a social identity to an IP address [LBZL$^+$11]. This attack works for all Skype users (more than 500 millions) and cannot be detected. In addition, we also show that we can track the mobility and the BitTorrent downloads of Skype users. Finally, we studied the problem of linking online profiles using only usernames (best paper-[PCKM11]). Two family of techniques were introduced. The first one estimates the uniqueness of a username to link profiles that have the same username. We gather from language model theory and Markov-Chain techniques to estimate uniqueness. Usernames gathered from multiple services have been shown to have a high entropy and therefore might be easily traceable. We extend this technique to cope with profiles that are linked but have different usernames and tie our problem to the well known problem of record linkage. All the methods we tried have high precision in linking username couples that belong to the same users. Ultimately we show a new class of profiling techniques that can be exploited to link together and abuse the public information stored on online social networks and web services in general.

Impact of this work: All these contributions were published in top conferences (LEET, IMC, PETS) and received a large international press coverage (e.g., Le Monde, The New York Times Bits, The Register, NewScientist, Wired News, Slashdot, etc.)

Architecture of Privacy: The increasing amount of personal information disseminated over the Internet raises serious privacy concerns. Data may linger forever, and users often lose its control and ownership. This motivates the desire of binding availability of contents to expiration times set by the data owner. To this end, we formalized the notion of Ephemeral Data Systems (EDSs): EDSs protect privacy of past data and prevent malicious parties from accessing expired contents. We designed EphCom, a practical EDS using only a primary Internet service — the Domain Name Service (DNS) and its caching mechanism. EphCom does not rely on Trusted Platform Modules (TPM), centralized servers, peer-to-peer networks, or proactive actions of the users [CDCFAK11]. It is transparent to existing applications and services, and allows users to tightly control data lifetime. We developed Firefox extension that provides ephemeral email capabilities and a command line tool for ephemeral files.

Impact of this work: This work has been presented at ICNP2011. A prototype of EphCOM has been implemented and distributed.

2.5.4 Collaborations

We strongly collaborated with the research groups of:

- Prof. Gene Tsudik at UC Irvine, USA. Planète had an associated team with UCI until 2011. Several student visits/exchanges between Inria and UCI happened in the period 2005-2011.

- Prof. Keith Ross at NYU-Poly. Stevens Le Blond spent a few months at NYU-Poly for this collaboration and the result was an IMC'11 paper.

- Prof. Dawn Song at UC Berkeley since beginning of 2012. Planète has now a associated team with Gene Tsudik at UCI and Dawn Song at UC Berbeley. Daniele Perito spend several months at UCI and UCB.

2.5.5 External support

The following projects provided support to these activities:

- European project UbiSec&Sens
- European project WSAN4CIP
- ANR RFIDAP
- ANR ARESA2
- European project OneLab2

2.5.6 Self assessment

The project has major contributions and visibility in the field of security and privacy and published in the major conferences of these areas (ACM CCS, IEEE S&P, Esorics, LEET, IMC, PETS, ...).

We make several fundamental contributions in these areas and considered several important applications, such as medical device security, google/skype privacy. As a result, we received considerable media attention.

This project addresses an overlooked issue in the field of privacy, the privacy infringements made by individuals. We led this direction by showing that single individuals can severely infringe privacy at a large scale. We got worldwide press coverage, and prestigious publications for our work.

It is also a very hot research topic that we envision to get an even higher visibility in the future. We plan to continue working into that direction and to address architectural issues in order to solve some fundamental privacy issues that exist in the current Internet.

As individuals, we are all well-recognized researchers and collaborate with the best groups in the world (Berkeley, UCI, Nicta, MPI,..). We are also in the PC of the most prestigious conferences (such as CCS, Wisec,...).

Finally, we trained and graduated several doctoral and master students.

2.6 Work Progress on Objective 3: Network Measurements

The focus of our research is a better monitoring of the Internet and a better understanding of its main features. We gave a particular attention to four aspects. The first aspect is related to traffic measurements in the core. We worked on efficient solutions that allow an operator to perform measurements of traffic inside in its network while limiting the overhead and without compromising accuracy. The second aspect is related to the access performance as perceived by end users. The focus was on light-weight methods that allow scanning the access, detecting any shift in performances, and evaluating of the impact of this shift on the global access quality. The third aspect is related to applications where we tried on one hand to infer applications quickly and efficiently, and on the other hand to understand the behavior of some well known applications, in particular the video streaming applications. The fourth aspect is related to the topology of the Internet, its inference and the robustness of the proposed solutions.

2.6.1 Personnel

Chadi Barakat (CR1), Walid Dabbous (DR1), Arnaud Legout (CR1), Mohamed Ali Kaafar (CR2), Mohamad Jaber (PhD), Amir Krifa (PhD), Imed Lassoued (PhD), Ashwin Rao (PhD).

2.6.2 Project-team positioning

Network monitoring is gaining lot of interest in the research community and many teams across the world are active in this area. Among the groups that focus on the same aspects as us one can cite the followings. In France: LIP6, LAAS, Eurecom, Technicolor, Alcatel, Telecom Bretagne and ENS Lyon. In Europe, groups as those at ETH, MPI, Telefonica, UCL Louvain-La-Neuve, Univ of Liege, FTW, Politecnico di Torino, Fraunhofer, TU Berlin and Brescia University perform complementary research to ours. Over the world the list is large, it includes among others Boston University, Georgia Tech, CAIDA, University of Melbourne, Wide project, Microsoft research, Nicta, AT&T, University of Wisconsin and CMU. The particularity of our approach is that we merge experimentations with mathematical tools for a better understanding of the main network properties and a more efficient design of tools. We give a particular focus to the load caused by measurements and we try to reduce it as much as possible while not compromising accuracy. We also seek originality in the measurements themselves and in the observations made of them. In general, we try to be either original in the way we define our measurements and we design our tools, otherwise build over the contributions of others to make them more performing.

2.6.3 Scientific achievements

Herein we provide a summary of our main contributions in this area according to the four aspects we previously mentioned.

Traffic measurements in the core We have worked on the design of an adaptive centralized monitoring system based on NetFlow that provides visibility over the entire network of an ISP. Given a measurement task, the proposed system drives its own configuration, typically the packet and flow sampling rates in routers, in order to address the tradeoff between monitoring constraints (processing and memory cost, collected data) and measurement task requirements (accuracy, flexibility, scalability). We explained our architecture with an accounting application: estimating the number of packets per flow (e.g., Domain-to-Domain traffic, all traffic originated from a domain, destined to a domain). The architecture and the algorithms behind it are explained in [LB10] for the case of a proactive control and in [LBA11] for the case of a reactive control. In [LB11] the architecture and its algorithms are specified to a flow counting application. In all these works [Las11], the performances of our architecture are validated in typical scenarios over an experimental platform we developed for the purpose of the study [KLB10]. Our platform is called MonLab (Monitoring Lab) [6], it presents a new approach for the emulation of Internet traffic and for its monitoring across the different routers.

In parallel, we have also worked on the spectral properties of the sampled traffic. Herein, and within a collaboration with Politecnico di Bari, we have proposed a novel approach to predict the energy of the sampling error on the real time traffic volume estimation, based on a spectral analysis in the frequency domain [GB09, GB10, GBM11]. In [VGBB11], we applied our expressions of the error to design a real-time algorithm, that sets the IPFIX counter export timers in order to grant, to each flow, a target estimation accuracy.

Edge measurements of the Internet access performance Our research focused on evaluating the quality of the Internet access and this is within the framework of our ANR CMON (Collaborative Monitoring) project jointly with Grenouille.com (weather of Internet in France). We gave a particular attention to the detection of delay anomalies at the access and to the evaluation of the impact of these anomalies for the purpose of localizing their root causes. We refer to this as the *impact factor* of the anomaly, which models the percentage of affected destinations. We worked on finding estimates for the impact factor of network anomalies through a limited set of measurements to random nodes we call landmarks. Our first results showed that accurate estimates of the impact factor can be obtained with a finite set of randomly distributed landmarks [CB10].

We also worked on the exploitation of virtual coordinates for network monitoring. Our objective was to study the feasibility of using shifts in coordinates to detect any change in the network delay and to identify the origin of this shift. For this, we have studied the dynamics of Vivaldi coordinates in normal network conditions [JNB10]. The Vivaldi system is known to be one of the most interesting approaches for the calculation of Internet coordinates. We observed that, despite the instability of Vivaldi coordinates in their absolute values, there is still a stable internal structure that can better reflect the stability of the underlying network. We presented a new clustering algorithm to identify the set of stable nodes once the host coordinates reach their stationary regime. We highlighted the utility of such finding with an application that tracks changes in network delays by continuously monitoring this set of stable nodes.

Applications' traffic measurements One of the most important challenges for network administrators is the identification of applications behind the Internet traffic. We came up with an online iterative probabilistic method that identifies applications quickly and accurately by only using the size of packets, the time between packets and the communication profiles of hosts. Our method, described

[6] http://planete.inria.fr/MonLab/

in [JB09, JCB11b, JCB11a], associates a configurable confidence level to the port number carried in the transport header and is able to consider a variable number of packets at the beginning of a flow. For this we had to propose a model to preprocess the inter-packet time and to use the result as input to the learning process. The profiles of hosts were updated online based on the result of the classification of previous flows. The validation results confirm that our method can provide a better classification accuracy than existing methods because of its capacity to integrate different features. It is also capable to derive profiles for the traffic of Internet hosts and to identify the services they provide [Jab11].

Concerning the characterization of video streaming traffic, we studied in [RLB$^+$11] the traffic characteristics for Youtube and netflix for a large variety of scenarios taking into account the impact of the browser (Firefox, Internet Explorer, Chrome), the impact of the container (Flash, HTLM5), and the impact of the access network (residential, academic, wireless). We identified 3 different strategies used, each one with very fundamental traffic characteristics. In particular, we have shown that the usage of HTML5 (that is presented as the successor of Flash for video streaming) leads to traffic characteristics that are dependent on the Web browser used.

Network topology inference Along this direction, we worked on increasing the robustness and efficiency of solutions for inferring network topology. The first contribution is related to hiding path similarity exchanges among monitors that are exchanging sensitive topology data. This is done for the purpose of making measurements safer for applications relying on such similarity (e.g. content distribution networks). Based on real data we collected, we advocate that path similarity comparisons between different Internet entities can be much simplified and secured using lossy coding techniques, such as Bloom filters, to exchange compressed topology information. In addition, we demonstrated that such technique is scalable as it requires a small amount of active probing and is not targets-dependent [DGK11].

Then we worked on detecting Triangular Inequality Violations (TIV) in network coordinate systems by means of supervised machine learning techniques [KGC$^+$08, LKG$^+$09, KCGL09]. We first showed that path lengths do have an effect on the impact of these TIVs. In particular, the shorter the link between any two nodes is, the less severe TIVs involved in are. In a second step, we focused on the particular case of the Vivaldi coordinate system and we explored how TIVs may impact its accuracy and stability. In particular, we observed correlation between the (in)stability and high effective error of nodes' coordinates with respect to their involvement in TIVs situations. We finally proposed a Two-Tier architecture opposed to a flat structure of Vivaldi that do mitigate the effect of TIVs on the distances predictions.

Finally we worked on the reliability of geolocation databases, the most widely used technique for IP geolocation. We conducted a comparison of several current geolocation databases -both commercial and free- to have an insight of the limitations in their usability. First, the vast majority of entries in the databases refer only to a few popular countries (e.g., U.S.). This creates an imbalance in the representation of countries across the IP blocks of the databases. Second, these entries do not reflect the original allocation of IP blocks, nor BGP announcements. In addition, we quantified the accuracy of geolocation databases on a large European ISP based on ground truth information. This was the first study using a ground truth showing that the overly fine granularity of database entries makes their accuracy worse, not better. Geolocation databases can claim country-level accuracy, but certainly not city-level [PUK$^+$11].

2.6.4 Collaborations

The ANR CMON project supports the research on edge measurements. Within CMON, we collaborate with the other partners, especially Grenouille.com, on the design of a new software architecture for probing the Internet and collecting measurements. The European project ECODE focuses on cognitive routing, we took in charge the monitoring component of the routers and we provided an adaptive component based on NetFlow that adjusts its sampling rates according to application needs. In addition to these intra-project collaborations, we had a long collaboration with Politecnico di Bari on the spectral analysis of sampled traffic. This collaboration resulted in several papers [GB09, GB10, GBM11, VGBB11]. The work on the characterization of video streaming traffic was done in collaboration with UMASS. As for the works on making topology inference methods more robust they were done in collaboration with the University of Louvain La Neuve, University of Liege, UQAM Quebec, University of Lancaster and Deutsche Telekom R&D.

2.6.5 External support

The measurement activities profited from two main fundings, the ANR CMON project for edge measurements and the FP7 ECODE project for traffic measurements in the core. The student who performed the research on traffic classification was granted by the French Ministry for his doctoral studies.

2.6.6 Self assessment

The research on measurements led to several interesting results and new tools that can serve the community and the large public, and can be the subject of a technology transfer in the future. Our platform MonLab is a precious tool for carrying out monitoring experiments in a controlled environment. We added to NetFlow components to adaptively set the sampling rate in routers. Our iterative bayesian approach for traffic classification is novel and our observations on the characteristics of video traffic can have applications for network dimensioning. All our results are promising and motivate us for more research in this area, despite the difficulties we meet in terms of validation. Whereas it is relatively easy to validate a solution by modeling and controlled experiments, it becomes difficult to deploy it in the real world and make people participate to experiments. This is the reason for which we developed MonLab and we count on grenouille.com for the deployment of edge measurement solutions. We also try to profit from platforms as PlanetLab as much as possible in our experiments (the work on coordinates was validated on PlanetLab). We have always managed to validate our work in a solid manner, however we strongly believe that the measurements community is in a serious need of a large scale measurement infrastructure for the obtaining of data and the validation of solution. We hope the Metroscope project led by Inria will fill this gap.

2.7 Work Progress on Objective 4 : Networking Evaluation Platforms

A vast majority of the theoretical work carried out among the Planète project-team deserves experimental activities to accurately assess the relevance of the contribution. The ideal requirements to achieve scientifically viable results are very difficult to meet, as there is an inherently hard to meet contradiction between on the one hand *realism*, i.e. the ability to run against external conditions close to what can be found in the real world, and which in addition often implies scale, and on the other hand *repeatability*, which is key when it comes to comparing several scenarios where only a controlled set of variables are changed.

In the InterNetworking area, one pragmatic answer to this resides in the diversity of evaluation platforms. It is typical to find papers that involve simulation, emulation in controlled environments and/or experimentation "in the wild" in order to touch on both aspects.

Aiming at leveraging the ability of the project-team as a whole to conduct sound and relevant protocol evaluations in a timely fashion, together with an aspiration at contributing the overall simulation and experimetation capability of the community in the large, our team has devoted a substantial effort in the fields of experimental testbeds (with its major implication in *PlanetLab* both from a software development point of view, as well as, even if more marginally, in the operations of such facilities) and realistic simulation (with its major implication in the *ns-3* project).

Evaluating networking protocols on wireless testbeds is not an easy task due to the difficulty to repeat the experiments with the same conditions. Our team also contributed to the domain of protocol *benchmarking* in wireless environments by defining an experimentation methodology enabling fair performance comparison, in the absence of full experience repeatability.

Last, it is important to provide researchers with a tool which can be used to describe, control, analyze, and compare the result of a large variety of networking experiments with varying levels of granularity corresponding to varying degrees of maturity of the networking technologies being studied. We designed and implemented an integrated research environment which is able to cater to the needs of inexpensive coarse-grained and fine-frained simulations, as well as the needs of more expensive experimentation platforms which integrate heterogeneous hardware components. Our goal was to ease the task of network researchers by providing them a *unified evaluation environment* to control and seamlessly mix and match simulation tools and real hardware to easily and cheaply study the behavior of large and complex networking systems.

2.7.1 Personnel

Walid Dabbous (DR1), Thierry Turletti (DR2), Thierry Parmentelat (Senior Engineer), Mathieu Lacage (Research Engineer), Alina Quereilhac (PhD), Shafqat ur-Rehman (PhD).

2.7.2 Project-team positioning

The Planète project-team has been playing a major role in the development and operations of the PlanetLab open testbed, and its underlying software MyPLC that is readily available for deploying private or consortium-wide testbeds. In this context we have a strong and recurrent partnership with **Princeton University**, in particular in terms of the codevelopment of the MyPLC software.

Having pioneered a federation approach to decentralized operations of PlanetLab, we also are very involved in the architecture, design and development of the SFA (Slice-based Federation Architecture) that aims at providing a more general, testbed-neutral federation pardigm, and in this context we have very good working relationship with the **FIRE** (Future Internet Research & Experimentation) Project in Europe, and with the **GENI** (Global Environment for Network Innovations)

Project in the U.S.A. SFA is currently deployed in real scale and allows to browse and provision resources at more than 20 locations worldwide.

Our contribution to the international ecosystem of networking testbeds has been fostered by various projects within FIRE, with a privileged connection to **UPMC** (University Pierre et Marie Curie) in Paris, with whom we co-chair the **PlanetLab Europe** consortium that now operates a wide range of publicly available testbeds whose spectrum is now much wider than the historical PlanetLab platform, and acts as a hub for experimental work in Europe.

In the field of wireless testbeds, our team has a strong cooperation with **NICTA** in Australia, specifically the main development team of OMF, the software that operates the ORBIT platform at Rutgers in New Jersey, and many other deployments in the world, as well as with **University of Thessaloniki** who provide NITOS, such a wireless testbed, that is part of the open testbeds managed by the PlanetLab Europe consortium.

Our major contribution in this field and current focus is an implementation of the SFA layer, that although initially targetting PlanetLab-like systems, has now moved towards a generic approach, and is being used for plugging into the SFA federation platforms as diverse as SensLab (sensor networks), Federica (routing), and Teagle (SOA-oriented platform). Finally, we have frequent interactions with **University of UTAH**, namely the Flux Research Group, who have come up with the other most common implementation of SFA.

On the other hand, ns-3 is an international open source project with contributors from several countries. However, our project-team and the **University of Washington** (Tom Henderson) have been the most active contributors ns-3 during the last four year.

Our overall approach to solve the protocol evaluation problem involves the creation and improvement of realistic simulators (ns-3) and experimentation testbeds (PlanetLab) to make sure that they can play well together as well as the development of integration tools (NEPI) which intend to automate the task of using these tools together. In contrast to most research teams who focus solely on the use and improvement of simulation tools (e.g. INESC Porto, the Centre Tecnologic de Telecomunicacions de Catalunya or Deutsche Telekom Laboratories in Berlin), experimentation platforms (e.g. Princeton University) or integration tools (OpenLab project partners), we take a holistic approach to this problem to ensure the smooth integration of these very different tools.

2.7.3 Scientific achievements

PlanetLab software The Planète project-team has been very active in the codevelopment of the PlanetLab software. The progression of the public platform over the evaluation period has proceeded to reach an overall 1030 nodes on 540 sites as of February 2012; we still marginally contribute to the operations activities of PlanetLab Europe, by providing a level2-like support to users in particularly challenging cases.

Although this growth has admittedly had a less impressive pace as it had in the past, is it still steady. In addition, quite a lot of smaller deployments have been made using the PlanetLab software known as MyPLC around the world, in many cases for prototyping alternative features such as supernodes, or dedicated high speed or optical links.

In this context we have provided quite a few innovative features to the core software, with the goal to provide a very flexible tool for deploying a networking testbed [FFP10].

These include a native emulation layer right into the node kernel, based on dummynet and developped in collaboration with Luigi Rizzo at University of Pisa, implemented as a linux module, and controllable from the hosted virtual machines using a secure mechanism. This feature is in particular suitable to emulate wireless links over a physical wired link, but can also be used at will to tweak the networking environment in a controllable and reproducible fashion. We have also added the ability to create tunnels ending in a PlanetLab slice, which is a very powerful tool to create arbirary topologies that can even span outside the PlanetLab scope as such. Also, it is now possible to operate PlanetLab nodes in reservable mode, which requires specific booking prior to running an experiment, so as to provide maximum guarantee about resource usage and availability. This feature was not primarily targetting the public infrastructure where it is only marginally used, but can come in very handy for operating smaller scale, for example wireless, testbeds. Finally, we have recently started to support linux containers (lxc) as an alternative to the historical linux-vservers technology for hosting virtual machines in the system, as well as a first draft of an OpenFlow-capable virtual switch implementation that sits as a kernel module and provides extensive topology-building capabilities.

These new mechanisms are better suited for current research needs, which tend to lose interest for mere content distribution or peer-to-peer paradigms, for which PlanetLab was historically well suited, and provide more powerful means to address challenging experimentation needs, like for example the ability to run CCN-like networking paradigms on a somewhat representative scale.

Testbeds federation As outlined in the introduction, testbeds diversity is key to properly addressing experimentation needs. Our work on PlanetLab and wireless testbeds aims at providing easy-to-use and flexible testbed management software, but we also acknowledge that it is often simpler to just use some other, already up and running testbed, if it provides the good properties required for the experimentation at hand.

Along those lines, we have been very active in the architecture, design and implementation of the SFA (Slice-based Facility Architecture). There currently are primarily 2 distinct implementations of SFA, one running on the Emulab-based testbeds, the other one being the PlanetLab implementation that we codevelop with Princeton University. The federation is currently up-and-running with all the main PlanetLab and Emulab testbeds reachable for resource discovery and provisioning.

As part of this effort, we have refactored the initial implementation that was totally MyPLC-specific, and have come up with a second release which is much cleaner, and generic enough to be used as a basis for implementing an SFA interface to virtually any testbed. We are currently working in conjunction with several testbed providers in order to achieve this. Namely, within the OpenLab (E.U.) project, FOKUS is developpping an SFA interface for their industry-oriented Teagle testbed; within the NOVI (E.U.) project, we are working with GARR in Italy and with I2CAT in Spain, to make the Federica real-scale routing testbed SFA-complient. Finally, within the (French) F-Lab project, we are working with SensLab project to provide it with the same capacity.

Wireless testbed often exhibit very different constraints than for distributed platforms. In particular, the ORBIT platform at Rutgers, like any other OMF-based testbed, is essentially exclusively based on time reservation. A user can generally hope to get access to the platform for a few hours.

For this reason, OMF provides very efficient tools for describing, and running an experiment, so that people do not waste their precious time at setting up things manually. This is in contrast with PlanetLab for example, where the job of the testbed software goes up to raw provisioning, with little or no help for controlling the experiment, that can be tooled with third-party software.

However in the context of international federation, and for the sake of PlanetLab users too, we have worked with NICTA in Australia, the main developer of the OMF software, and made it possible to control PlanetLab slivers through the OMF tools, which can thus control cross-testbed experiments.

This is a very exciting stage, as FIRE in E.U. and GENI in the U.S.A. have created a substantial momentum for testbeds federation, and the potential outcome, were we to succeed, are huge in terms of lowering the barrier to entry for doing experimental work efficient and cost-effective, in many respects, including learning curve, and cross testbed fertilization, especially in terms of experimenter-level tools.

Benchmarking in Wireless testbeds We have shown that unknown conditions (due e.g. to complex adaptation schemes in wireless hardware) and uncontrolled parameters (especially those characterizing the wireless channel) make performance comparison particularly difficult in wireless environments [URTD11, TAD+11].

In order to enable fair performance comparison in a wireless network we developed a benchmarking methodology that extends classical experimentation approaches to include steps for enhancing repeatability and monitoring and storing the required information to cope with lack of full repeatability.

Our benchmarking methodology [UR12, BVGI+11] defines step by step actions to perform including initial testbed preperation (system under test definition and tools deployment) and "cyclic" experimentation phases to 1) configure the target scenario, 2) calibrate the testbed, 3) run the experiment (potentially a large number of times), 4) undertake measurements and data collection, 5) undertake data cleansing and archiving, 6) Compute metrics, 7) prepare and publish results as full disclosure reports.

To realize all the above, we have developed software tools to automate the overall experimentation process. We have also deployed our own in-house wireless testbed for protocol benchmarking named PLECS that provides around 25 wireless laptop nodes deployed in an office environment[7]. The PLECS testbed and the benchmarking methodology have been set up using the Wireless Experimentation (WEX) Toolbox, developped locally, that aims to set up, run and make easier the analysis of wireless experiments. It is a flexible and scalable open-source set of tools that covers all the experimentation steps, from the definition of the experiment scenario to the storage and analysis of results. For more details see https://twiki-sop.inria.fr/twiki/bin/view/Projets/Planete/WEXToolkit.

ns-3 Through our very early involvement with the open source ns-3 event-driven simulator, we contributed to the architecture and the implementation of its core facilities. This allowed us to steer the project towards tighter integration with other experimentation platforms and protocol implementations, hence making ns-3

[7] There are plans as part of the FIT project to deploy a second wireless testbed, with more heterogeneous and less power-intensive nodes, in a dedicated building that remains to be built at this point.

a simulation platform ideally suited to become the high-impact integration hub it is now.

We did benefit from this very wide impact when we contributed the Direct Code Execution (DCE) simulation module to ns-3 [Lac10]: this new module represents such a considerable improvement in terms of ease of use and performance over previous approaches that, despite its early state, it is now being used by numerous research teams around the world to study the behavior within ns-3 of existing user-space (quagga, CCNx) and kernel-space (Linux) protocol implementations under complex network conditions.

ns-3 is available to the whole research community and is more and more used in research papers. Since a major release in December 2009, the number of monthly downloads of the ns-3 simulator code is between 4000 and 10000. See the ns-3 web site http://www.nsnam.org for more details.

NEPI The Network Experimentation Programming Interface (NEPI) project is complementary to our efforts to improve simulation tools towards more realism and better integration with the real world. The originality of NEPI lies in its new experiment plane which can already be used to perform ns-3 simulations, planetlab and emulation experiments and which has been designed to integrate any experimentation tool used for networking research. Its goal is to make it easier for experimenters to describe the network topology and the configuration parameters, to specify trace collection information, to deploy and monitor experiments, and, finally, collect experiment trace data into a central datastore. NEPI [LFH+10, QLF+11] is a python API (with an implementation of that API) to perform all the above-mentioned tasks and allows users to access these features through a simple yet powerful graphical user interface called NEF.

Both ns-3 and NEPI software development projects were originally started as part of the Ph.D. thesis pursued by Mathieu Lacage [Lac10]. The more mature components are already in wide use within ns-3 while DCE and NEPI are still undergoing rapid development and research, most notably within the context of the Ph.D. thesis of Alina Quereilhac.

2.7.4 Collaborations

Over the past 5 years, we have pursued a strong collaboration with several teams working on simulation on testbeds:

- Princeton University and University of Pierre and Marie Curie on PlanetLab software
- NICTA (Max Ott team) and University of Thessaloniki on wireless testbeds
- IBBT (Ingrid Moerman team) on Wireless protocol benchmarking
- Tom Henderson, (UWA, Boeing), both through our continued involvement in the development of the ns-3 simulation core and numerous simulation models, and through regular discussions on the respective evolutions of CORE and NEPI. The ns-3 software we built collaboratively is now widely used within academia and the industry.

On the other hand, we have founded recently a consortium between INRIA and University of Washington that provides a point of contact between industrial members and the NS-3 project. Georgia Tech and Bucknell University are joining the consortium as executive members. It is expected that several industrial and academic partners join the consortium in the next period.

2.7.5 External Support

These contributions have been sponsored by the following projects.

- EU-funded Integrated Project OneLab2 (2008-2010)
- EU-funded Integrated Project OpenLab (2011-2014)
- EU-funded STREP NOVI (2010-2013)
- Fr-funded F-Lab (2010-2013)
- Fr-funded FIT (2011-2017)
- Fr-funded PFT (2011-2014).
- Fr-funded CONNECT, ANR project
- Inria-funded GENESIM, Associated Team with Univeristy of Washington.

2.7.6 Self assessment

Our global contribution in the field of networking experimental testbed is a solid and acknowledged one. Our role in the development of PlanetLab, both in terms of the deployed public platform and in related software development, is quite visible. Our involvement in wireless testbeds, although more marginal, has become substantial too with the bridging of PlanetLab with the NITOS and NICTA testbeds through experiment-control tools. We believe these testbeds will still remain in active service and provide valuable capabilities for quite some time.

However, in order to more effectively open innovative tracks, we strongly believe that federation is the way to go in order to make available testbeds easier to find about, and more natural to use effectively and in a timely fashion. There obviously is also a huge potential gain in cross fertilizing user tools between the various testbed management softwares and user-tool communities out there.

We are a visible partner in the Europe effort towards building a large-scale federation, and are very excited to see the running embryo for an international federation of testbeds, that currently spans the U.S.A., the E.U., as well as Korea and Japan even if more marginally, fueled by the international momentum that FIRE and GENI have triggered around these topics.

Ironically, we still see relatively small usage of federation-based tools, as compared with the native tools that come with the individual testbeds. This is understandable as most of the offer up to this point has been focused on the internals of the federation mechanisms, and the user tools are still quite rough anc cannot decently compete with their counterpart, that in most cases have several years of existence, and also need to perform the simpler job of handling a single testbed.

However, with more and more different kinds of testbeds becoming federation-compliant, and with a little more time for federation-based user tools to mature, it seems realistic to expect massive adoption of that paradigm in the field over the next couple of years.

Because we believed that it was important for our work on ns-3 and NEPI to quickly solve the problems we faced daily as experimenters, we chose early on to focus on delivering production-grade software that would embody our ideas. This naturally led to a quick rate of adoption and a large diffuse impact on the networking research community but this detracted us from producing the scientific papers associated with this research effort. We are now attempting to correct this through the submission of papers describing our research results.

3 Knowledge dissemination

3.1 Publications

	year1 (2008)	year2 (2009)	year3 (2010)	year4 (2011)
PhD Thesis		3	4	5
H.D.R (*)	2	1		
Journal	1	7	6	6
Conference proceedings (**)	15	32	22	22
Book chapters			1	1
IETF standards (RFCs)	1	1	1	1
Other IETF drafts	15	10	14	16
General audience papers		1	2	
Technical reports	5	12	5	4
Total	39	67	55	55

(*) HDR Habilitation à diriger des Recherches

(**) Conference or workshop with a program committee

In addition to the above, one HDR, two PhD thesis and three journal papers were published or will appear in 2012 for a total of 222 publications listed in section 6 among them 15 were chosen as reference publications and are listed in subsection 6.6.

Here follow the major journals in the field and, for each, we indicate the number of papers coauthored by members of the project-team that have been accepted during the evaluation period.

1. ACM/IEEE ToN (Transactions on Networking) (1 publication)
2. IEEE JSAC (Journal on Selected Areas in Communications)
3. CCR (Computer Communication Review) (1publication)
4. Computer Networks (3 publications)
5. Computer Communications (1 publication)
6. ACM Transactions on Sensor Networks (2 publications)
7. Wireless Networks (2 publications)
8. IEEE TMC (Transactions on Mobile Computing) (1 publication)

Here follow the **major conferences** in the field and, for each, we indicate the number of papers coauthored by members of the project-team that have been accepted during the evaluation period.

1. SIGCOMM
2. SIGMETRICS
3. IMC (3 publications)
4. CoNext (1 publication)
5. IEEE Security and Privacy (1 publication)
6. ACM CCS (3 publications)
7. INFOCOM (1 publication)
8. NSDI
9. PETS, Privacy Enhancing Technologies Symposium (2 publications)
10. NDSS, Network and Distributed System Security Symposium
11. Information Hiding (1 publication)
12. WiSec, ACM Conference on Wireless Network Security (1 publication)
13. COMSNET, International Conference on COMmunication Systems and NETworkS (2 publications)

Here follows a list of **interesting workshops** in the field.

1. IPTPS, USENIX International workshop on Peer-To-Peer Systems (1 publication)
2. LEET, USENIX Workshop on Large-Scale Exploits and Emergent Threats (2 publications)
3. HotNets, ACM SIGCOMM Hot Topics in Networks
4. WPES, Workshop on Privacy in the Electronic Society (1 publication)
5. ROADS, ACM SIGOPS on Real Overlays and Distributed Systems (2 publications)

3.2 Software

The following software descriptions follow the Inria Evaluation Committee criteria for software self-assessment[8]:

ns-3:

ns-3 is a discrete-event network simulator for Internet systems, targeted primarily for research and educational use. ns-3 includes a solid event-driven simulation core as well as an object framework focused on simulation configuration and event tracing, a set of solid 802.11 MAC and PHY models, an IPv4, UDP, and TCP stack and support for nsc (integration of Linux and BSD TCP/IP network stacks).

[8] See http://www.inria.fr/medias/recrutement-metiers/pdf/criteria-for-software-self-assessment

Audience	Wide-audience software
Software originality	Original software implementing a fair number of original ideas
Software maturity	Major software project
Evolution & Maintenance	Well-defined and implemented plan for future maintenance and evolution
Distribution & Licensing	External packaging and distribution
Web site	http://www.nsnam.org
Users community	ns-3 is free software, licensed under the GNU GPLv2 license, and is publicly available for research, development, and use.
Position wrt. competition	ns-3 is the new network simulator software, and version 2, largely used by the network research community, is now totally obsoleted by ns-3. Its early emphasis on realism and real world integration puts ns-3 in a unique position with regard to other simulators such as omnetpp or opnet which lack such native well integrated support.
Software metrics	C++, python

MyPLC (Portable Installation Software for PlanetLab:

MyPLC is a complete PlanetLab Central (PLC) portable installation. The default installation consists of a web server, an XML-RPC API server, a boot server, and a database server: the core components of PLC. The installation is customized through an easy-to-use graphical interface. All PLC services are started up and shut down through a single script installed on the host system. We have strongly contributed to the development of this installation software. See the PlanetLab page for a reference to our contribution. https://svn.planet-lab.org/wiki/MyPLCUserGuide

Audience	Large-Audience software
Software originality	Original software reusing known ideas and introducing a few new ideas.
Software maturity	Well-developed software
Evolution & Maintenance	Good quality middle-term maintenance
Distribution & Licensing	Public source or binary distribution on the web
Web site	http://svn.planet-lab.org/ https://svn.planet-lab.org/wiki/MyPLCUserGuide https://svn.planet-lab.org/wiki/MyPLCUserGuide
Users community	Network research community
Position wrt. competition	OMF (NICTA) and Emulab (U. Utah) have similar software for their testbeds. However, MyPLC is the only "Ready to . use" package
Software metrics	Python/C for kernel modules, caml for vsys

SFA (Slice-based Federation Architecture) for PlanetLab:

We are codevelopping with Princeton University a reference implementation for the Testbed-Federation architecture known as SFA for Slice-based Federation Architecture. During 2011 we have focused on the maturation of the SFA codebase, with several objectives in mind: better interoperabity between the PlanetLab world and the EmuLab, a more generic shelter that other testbeds can easily leverage in order to come up with their own SFA-compliant wrapper and support for 'reservable' mode, which breaks the usual best-effort PlanetLab model.

Audience	Ambitious software
Software originality	Original software reusing known ideas and introducing a few new ideas.
Software maturity	Basic usage should work, up to Well-developed software
Evolution & Maintenance	Good quality middle-term maintenance
Distribution & Licensing	Public source or binary distribution on the web
Web site	https://svn.planet-lab.org/wiki/SFATutorial http://git.onelab.eu/?p=sfa.git https://www.planet-lab.eu/db/doc/PLCAPI.php
Users community	Network research community
Position wrt. competition	No similar software
Software metrics	Python

NEPI (Network Experimentation Programming Interface):

NEPI implements a new experiment plane used to perform ns-3 simulations, planetlab and emulation experiments, and more generally any experimentation tool used for networking research. Its goal is to make it easier for experimenters to describe the network topology and the configuration parameters, to specify trace collection information, to deploy and monitor experiments, and finally collect experiment trace data into a central datastore. NEPI is a python API to perform all the above-mentioned tasks and it allows users to access these features through a simple yet powerful graphical user interface called NEF.

Audience	Ambitious software
Software originality	Original software implementing a fair number of original ideas.
Software maturity	Basic usage should work, up to Well-developed software
Evolution & Maintenance	Good quality middle-term maintenance
Distribution & Licensing	Public source or binary distribution on the web, GPL licence, APP deposit
Web site	http://nepihome.org
Users community	Network research community.
Position wrt. competition	No similar software
Software metrics	Python

EphPub (Ephemeral Publishing):

EphPub implements a novel key storage mechanism for time-bounded content, that relies on the caching mechanism of the Domain Name System (DNS). EphPub exploits the fact that DNS servers temporarily cache the response to a recursive DNS query for potential further requests. Therefore EphPub comes with high usability as it does not require users to install and execute any extra additional software. EphPub also lets users define data lifetime with high granularity. We provide EphPub as an Android Application to provide ephemeral exchanged SMS or emails, and as a Firefox or Thunderbird extensions so as to support ephemeral publication of any online document.

Audience	Large audience software
Software originality	Original software implementing a fair number of original ideas.
Software maturity	Basic usage should work
Evolution & Maintenance	Basic Maintenance, Work under progress
Distribution & Licensing	Public source or binary distribution on the web, APP deposit
Web site	http://planete.inrialpes.fr/projects/ephemeral-publication/
Users community	Research community for the moment, the goal with the Android version is to enlarge the community to end-users.
Position wrt. competition	EphPub provides higher security than Vanish (it is immune to Sybil attacks). EphPub is also more easily deployable since it does not require any additional infrastructure like DHT.
Software metrics	Python, Java

UsernameTester:

Usernames are ubiquitous on the Internet. Almost every web site uses them to identify its users and, by design, they are unique within each service. The UsernameTester tool estimates how unique and linkable usernames are, with the possibility for any user to check by itself through the web page. For example, according to our tool, ”ladygaga” or ”12345678” only carry 24 and 17 bits of entropy, respectively. They are therefore not likely to be unique on the Internet. On the other hand, usernames such as ”pdjkwerl” or ”yourejerky” carry about 40 bits of entropy and are therefore very good identifiers. This tool can help users to select a username that has low entropy and can’t be used to track them on the Internet.

Audience	Ambitious software
Software originality	Original software implementing a fair number of original ideas.
Software maturity	Well-developed software
Evolution & Maintenance	Basic maintenance to keep the software alive
Distribution & Licensing	None, but the tool is available as an online web service
Web site	http://planete.inrialpes.fr/projects/how-unique-are-your-usernames/
Users community	This tool is destined both to the research and public communities.
Position wrt. competition	No equivalent.
Software metrics	6 person*months

OpenFEC.org: because open, free AL-FEC codes and codecs matter

The OpenFEC.org project aims at providing high performance, ready-to-use, free, C-language software codecs for AL-FEC (Application Level Foward Erasure Correction) codes, all of them being gathered in the same library, under the same API. This project also provides automated performance evaluation tools and demonstration applications. The current AL-FEC codes are provided: Reed-Solomon over $GF(2^4)$ and $GF(2^8)$, 2D parity check codes, LDPC-Staircase codes, LDPC from file.

Audience	Large audience software
Software originality	Original software implementing a fair number of original ideas
Software maturity	Well-developed software
Evolution & Maintenance	Good quality middle-term maintenance
Distribution & Licensing	Public source or binary distribution on the web, CeCiLL-C licence, APP deposit done
Web site	http://openfec.org
Users community	This software is intended to be used by: (1) users who do not want to know the details of AL-FEC schemes but need to use them in the software they are designing; (2) users who want to test new codes or new encoding or decoding techniques; (3) users who need to do extensive tests for certain AL-FEC schemes in a given use-case, with a well defined channel model. This is therefore of interest to several communities: research, industrial, open source community.
Position wrt. competition	This is the first and only initiative of that kind, with this ambition.
Software metrics	26 000 lines of C code, 27 person*months, C and Perl

LDPC-Staircase Advanced codec:

This is a high performance, well-tuned, codec implementing LDPC-Staircase codes. It is a fork of the OpenFEC LDPC-Staircase codec that has been highly optimized for speed and memory, so that it can be used in constrained devices.

Audience	Wide audience software
Software originality	Original software implementing a fair number of original ideas
Software maturity	Major software project
Evolution & Maintenance	Well-defined and implemented plan for future maintenance and evolution, APP deposit done
Distribution & Licensing	External packaging and distribution (commercialized through an external partner to third parties)
Web site	N/A
Users community	This software is commercialized by a private company. Use-cases currently include the ISDB-Tmm (mobile multimedia Japanese standard) broadcast servers and terminals.
Position wrt. competition	This is the most advanced LDPC-Staircase codec. Other codecs exist for competing AL-FEC codes, in particular Qualcomm's Raptor/RaptorQ$^{(tm)}$ codec.
Software metrics	18 000 lines of C code, 38 person*months (includes efforts spent on the early OpenFEC design), C and Perl

MCLv3-P (MultiCast Library Version 3 - Private version):

MCLv3-P is a commercial implementation of the ALC (Asynchronous Layered Coding) content delivery Protocols and of the FLUTE file transfer application. This software is an implementation of the large scale content distribution protocols standardized by the RMT (Reliable Multicast Transport) IETF working group and adopted by several standardization organizations, in particular 3GPP for the MBMS (Multimedia Broadcast/Multicast Service), DVB for the CBMS (Convergence of Broadcast and Mobile Services), OMA BCAST (Mobile Broadcast Services Enabler Suite), and others similar standards for multimedia content broadcast (IPTV, push VOD, DTV). Unlike the MCLv3 open-source version, the MCLv3-P fork is not publically available.

Audience	Wide-audience software
Software originality	Original software reusing known ideas and introducing a few new ideas
Software maturity	Major software project
Evolution & Maintenance	Well-defined and implemented plan for future maintenance and evolution
Distribution & Licensing	External packaging and distribution (commercialized through an external partner to third parties)
Web site	N/A
Users community	This software is commercialized by a private company, both for the server side (broadcaster) and client side (terminal manufacturer). Several use-cases have appeared over the years in the context of digital multimedia distribution.
Position wrt. competition	An open-source FLUTE/ALC protocol stack exists (MAD-FLUTE) but with fewer functionalities (it does not support OpenFEC AL-FEC codes for instance) and it is no longer supported.
Software metrics	51 000 lines of C code.

BitHoc (Tracker-less BitTorrent for Mobile Ad Hoc networks):

BitHoc enables content sharing among spontaneous communities of mobile users using wireless multi-hop connections. BitHoc is the real testbed over which we evaluate our solutions for the support and optimization of file sharing in a mobile wireless environment where the existence of an infrastructure is not needed. The proposed BitHoc architecture includes two principal components: a membership management service and a content sharing service. In its current form it is composed of PDAs and smartphones equipped with WIFI adapters and Windows Mobile 6 operating system.

Audience	Ambitious software
Software originality	Original software implementing a fair number of original ideas.
Software maturity	Basic usage should work
Evolution & Maintenance	No real future plans
Distribution & Licensing	Public source distribution on the web, organized by the development team, GPLv3
Web site	http://planete.inria.fr/bithoc/
Users community	Network research community. This tool has been developped as part of the ITEA ExpeShare European project (http://virtual.vtt.fi/virtual/expeshare/ and has been used in particular by the project partners.
Position wrt. competition	BitHoc is based on BitTorrent. To the best of our knowledge, BitHoc is the only tracker-less version of the highly efficient BitTorrent protocol to support P2P data exchange between ad-hoc communities
Software metrics	C++, C# (.net compact framework), 32494 lines of code

MonLab (Emulation Platform for Network Wide Traffic Sampling and Monitoring):

MonLab is a platform for the emulation and monitoring of traffic in virtual ISP networks that has been developed in the context of the FP7 ECODE project. In its current version, the traffic is sampled at the packet level in each router of the platform, then monitored at the flow level. We put at the disposal of users real traffic emulation facilities coupled to a set of libraries and tools capable of

Cisco NetFlow data export, collection and analysis. Our aim is to enable running and evaluating advanced applications for network wide traffic monitoring and optimization. We believe that this framework can play a significant role in the systematic evaluation and experimentation of these applications algorithms. Among the direct candidates one can mention algorithms for traffic engineering and distributed anomaly detection, as well as methods for placing monitors, sampling traffic, coordinating monitors and inverting sampling traffic.

Audience	Ambitious software
Software originality	Original software implementing a fair number of original ideas.
Software maturity	Basic usage should work
Evolution & Maintenance	Basic maintenance to keep the software alive
Distribution & Licensing	Public source distribution on the web, organized by the development team, GPLv3
Web site	http://planete.inria.fr/MonLab/
Users community	Network research community.
Position wrt. competition	MonLab works with NetFlow, but can work also with other platforms such as Emulab or OneLab.
Software metrics	C++/Linux, 10525 lines of code

TinyRNG:

TinyRNG is a random number generator of cryptographic quality. It uses entropy collection and accumulates entropy into two pools in order to enable forward and backward security. One of the entropy source is the erroneous packet received from the radio, with careful selection between what could be modified by an attacker and what could not be attacker controlled.

Audience	Ambitious software
Software originality	Original software implementing a fair number of original ideas
Software maturity	Well-developed software
Evolution & Maintenance	No real future plans
Distribution & Licensing	Public source or binary distribution on the web
Web site	http://www.ist-ubisecsens.org/downloads/tinyrng/tinyrng.php
Users community	Network research community.
Position wrt. competition	No similar software.
Software metrics	C

DroidMonitor:

We have developed a novel, private data leakage monitoring tool, DroidMonitor. It aims to serve as an educational tool for regular Android Smartphones users to make them aware of existing privacy threats while they are using Location-Based Services.

Audience	Ambitious software
Software originality	Original software reusing known ideas and introducing a few new ideas
Software maturity	Basic usage should work
Evolution & Maintenance	No real future plans
Distribution & Licensing	Public source or binary distribution on the web
Web site	http://planete.inrialpes.fr/android-privacy/
Users community	Network research community, end-users.
Position wrt. competition	Many tools start to appear. This tool is a first step toward a more ambitious monitoring framework.
Software metrics	Java

3.3 Valorization and technology transfer

3.3.1 Software commercialization

Commercial distribution of MCLv3 (MultiCast Library Version 3) by Expway: A first agreement between Inria and the French company Expway (http://expway.com) was signed in 2007, for a commercial distribution of Inria's FLUTE(RFC3927) / ALC(RFC5775) protocol stack as part of the MCLv3 software. This software is still a leading solution for digital multimedia distribution (Mobile DTV, IPTV and Digital TV) and has been used by several broadcasters, telecom operators and device manufacturers. Customers include companies such as NTT in Japan, TDF in France and Mobile500 Alliance in the US. Thanks to this success, a second version of the Inria/Expway agreement has been signed in 2011 to further extend the collaboration and the distribution of Expway's adaptation of MCLv3.

Impacts: This software is a key component of Expways's commercial offer. For the Planète EPI point of view, this work has significantly increased our credibility at IETF where the protocol stack has been designed (V. Roca is co-author of some of these RFCs).

Commercial distribution of the LDPC-Staircase Advanced codec by Expway: An agreement between Inria and the French company Expway (http://expway.com) has been signed in 2011 in order to commercialize our Advanced LDPC-Staircase codec. This success is the follow up of the recent adoption of our LDPC-Staircase codes (RFC5170) by the ISDB-Tmm Japanese standard for mobile push VOD services.
(press release: http://www.expway.com/telechargement/1314893215.pdf)

Impacts: This software will hopefully become a key component of Expways's commercial offer if the adoption of the LDPC-Staircase codes continues with other use-cases and standards. For the Planète EPI point of view, the development of a high performance LDPC-Staircase codec is a prerequisite for our research activities (the design of high performance AL-FEC codes and codecs have been the main objective of three PhD thesis plus a two-year engineer work). This is also the recognition of the work performed, and it has increased our credibility at IETF where LDPC-Staircase codes have been standardized (V.Roca is the main author of one RFC on the subject, and several related RFCs/Internet Drafts).

3.4 Teaching

2008

Networks and protocols:	Undergraduate course at Ecole Polytechnique, by W. Dabbous (36h).
Understanding Networks:	Course at Master IFI, University of Nice-Sophia Antipolis, by W. Dabbous and C. Barakat(42h).
Internet Measurements and Traffic Analysis:	(i) Networking and Distributed Systems Master at the University of Nice Sophia Antipolis and (ii) Master RIM, ENSI, Tunis, by C. Barakat (15h).
Introduction to Networking:	Undergraduate course at IUT Nice - LPSIL class, by C. Barakat (15h).
Voice over IP:	Graduate Course at (i) Master RTM of the IUP Avignon and (ii) Master TIM UNSA, by C. Barakat (7h).
Network Simulator ns-2:	7 hours, Master RTM of IUP Avignon, 2008.
Wireless Communications:	Undergraduate course at Polytech' Grenoble, on Wireless Communications, by V. Roca (12h).
Networking:	IUT Informatique, University Pierre Mendes France, Grenoble, by V. Roca (28h).
Wireless Security:	Course given to the students of the Ensimag "crypto and security" Master 2, Ensimag, Grenoble by C.Castelluccia (20h).
Wireless Security:	Course given to the students of the Ensimag/INPG "MOSIG" Master 2, Ensimag/INPG, Grenoble by C.Castelluccia (12h).
Networks:	Undergraduate course at University of Nice-Sophia Antipolis, by C. Barakat (6h).
Programming:	Course IUT GTR 2005 (36h), by Arnaud Legout
Programming:	Course IUT GTR 2006 (30h), by Arnaud Legout
Networks:	Course IUT GTR 2006 (30h), by Arnaud Legout
Peer-to-peer networks:	Course master RSD at University of Nice-Sophia Antipolis 2006 (15h), by Arnaud Legout
Programming:	Course IUT GTR 2007 (30h), by Arnaud Legout
Peer-to-peer networks:	Course master RSD at University of Nice-Sophia Antipolis 2007 (15h), by Arnaud Legout

2009

Networks and protocols:	Undergraduate course at Ecole Polytechnique, Palaiseau, by W. Dabbous (36h).
Evolving Internet:	Course at the UbiNet Master program, University of Nice-Sophia Antipolis, by W. Dabbous and C. Barakat(42h).
An introduction to Internet monitoring:	3h, *(i)* Telecom Paris, 2009-2010, and *(ii)* ETH Zurich, 2009, by C. Barakat.
Wireless networking:	7h, Master RTM, IUP Avignon, 2009, by C. Barakat.
Local Aarea Networks:	21 hours course + 10.5 hours practical work, IUT of the University of Nice-Sophia Antipolis, 2007-2010, by C. Barakat.
Wireless Communications:	Undergraduate course at Polytech' Grenoble, on Wireless Communications , by V. Roca (12h).
Wireless Security:	Course given to the students of the Ensimag “crypto and security” Master 2, Ensimag, Grenoble by C.Castelluccia (20h).
Wireless Security:	Course given to the students of the Ensimag/INPG “MOSIG” Master 2, Ensimag/INPG, Grenoble by C.Castelluccia (12h).
Networks:	Undergraduate course at University of Nice-Sophia Antipolis, by C. Barakat (6h).
Peer-to-peer networks:	Course in the UbiNet master at University of Nice-Sophia Antipolis 2009 (15h), by Arnaud Legout
Networks and Telecommunications:	Programming Courses given to the students of the Ensimag, Grenoble by Mohamed Ali Kaafar (18h).
Networks Introduction:	Programming Courses given to the students of the Phelma/INPG , Grenoble by Mohamed Ali Kaafar (12h).

2010	
Networks and protocols:	Undergraduate course at Ecole Polytechnique, Palaiseau, by W. Dabbous (36h).
Evolving Internet:architectural challenges:	Course at the IFI-UbiNet Master program, University of Nice-Sophia Antipolis, by W. Dabbous and C. Barakat(42h).
An introduction to Internet monitoring:	3h, *(i)* Telecom Paris, 2009-2010, and *(ii)* ETH Zurich, 2009, C. Barakat.
Wireless networking:	7h, Master RTM, IUP Avignon, 2009 - 2010, C. Barakat.
Local Aarea Networks:	21 hours course + 10.5 hours practical work, IUT of the University of Nice-Sophia Antipolis, 2007-2010, C. Barakat.
Internet Measurements and Traffic Analysis:	15h, *(i)* Networking and Distributed Systems Master at the University of Nice-Sophia Antipolis, 2004-2007, and *(ii)* Master RIM, ENSI, Tunis, 2003-2007.
Introduction to Networking:	15h, IUT Nice, Licence LPSIL, 2006-2008, C. Barakat.
Voice over IP:	7h, *(i)* Master TIM, UNSA, 2007 - present, *(ii)* Master RTM, IUP Avignon, 2008, C. Barakat.
Network Simulator ns-2:	7 hours, Master RTM of IUP Avignon, 2008, C. Barakat.
Peer-to-peer networks:	Course in the UbiNet master at University of Nice-Sophia Antipolis 2010 (21h course), by Arnaud Legout
Peer-to-peer networks:	Course in the IUP GMI Avignon 2010 (24h course, 21h TP), by Arnaud Legout.
Wireless Security:	Course given to the students of the Ensimag “crypto and security” Master 2, Ensimag, Grenoble by C.Castelluccia (20h).
Wireless Security:	Course given to the students of the Ensimag/INPG “MOSIG” Master 2, Ensimag/INPG, Grenoble by C.Castelluccia (12h).
Wireless Communications:	Undergraduate course at Polytech’ Grenoble, on Wireless Communications , by V. Roca (12h).
Networking:	Undergraduate course at IUT-2 (UPMF University), on network Communications , by V. Roca (24h).
Networks and Telecommunications:	Programming Courses given to the students of the Ensimag, Grenoble by Mohamed Ali Kaafar (18h).
Networks Introduction:	Programming Courses given to the students of the Phelma/INPG , Grenoble by Mohamed Ali Kaafar (12h).

2011	
Evolving Internet:	Master Ubinet, by Walid Dabbous and Chadi Barakat 42 hours, University of Nice-Sophia Antipolis, France.
Internet monitoring:	Master CAR, by Chadi Barakat, 3h, Telecom Paris Tech, France.
Introduction to Networking:	Undergraduate course at IUT, Nice-Sophia Antipolis University, by Chadi Barakat (28h), France.
Networks and protocols:	Undergraduate course, by Walid Dabbous (36h), Ecole Polytechnique, Palaiseau, France.
Networks and Telecommunications:	Undergraduate course, by Mohamed Ali Kaafar (40h), Ensimag Engineering school, France.
Networking:	Undergraduate course at IUT-2 (UPMF University), by V. Roca (24h).
Networking:	Master Phelma, by Mohamed Ali Kaafar (12h), INPG , France.
Peer-to-peer networks:	Undergraduate course at IUP GMI Avignon, by Arnaud Legout (38h), France.
Peer-to-peer networks:	Master Ubinet, by Arnaud Legout (21), University of Nice-Sophia Antipolis.
Peer-to-peer performance and security challenges:	Master FST, by Mohamed Ali Kaafar (21h), Tunisia.
Voice over IP:	Master TSM, by Chadi Barakat, (7h), University of Nice-Sophia Antipolis, France.
Wireless Communications:	Undergraduate course at Polytech Grenoble, V. Roca (12h).
Wireless Security:	Master Crypto and Security, by Claude Castelluccia (20h), Ensimag/University of Grenoble, France.
Wireless Security:	Master MOSIG, by Claude Castelluccia (12h), Ensimag/INPG, France.
Wireless networking:	Master RTM, by Chadi Barakat, 7h, IUP Avignon, France.

3.5 General Audience Actions

Walid Dabbous gave two lectures on “Internet de l’Avenir” in high schools in Draguignan and Marseille in 2011 in the context of “Univsersité de tous les savoirs au lycée”, a french association that promotes scientific dissemination in high schools. These lectures are available on Canal-U (See http://www.canal-u.tv/).

Walid Dabbous wrote an article on Internet Architecture Evolution published in Pour la Science magazine in 2010 [Dab10].

Claude Castelluccia wrote an article on Internet Privacy published in Pour la Science magazine in 2010 [Cas10], and another article in “Science et Vie Junior”.

Mohamed Ali Kaafar gave an interview on Privacy and owner’s Data control for the SVM magazine August 2009.

Tens of articles in the technical and general audience press have covered several of our results[9].

[9] See pointers to the articles in http://planete.inrialpes.fr/˜ccastel/Publications.html and/or http://www-sop.inria.fr/members/Arnaud.Legout/Projects/bluebear.html

3.6 Visibility

Walid Dabbous served in the programme committees of GC'08 CCNS, GC'09 CISS, Mobiquitous 2010, ICC'10 CISS, WNS3'2011, NOMEN'2012, ICC'12 NGN, NOM'12. He was co-chair of the ROADS'09 workshop. He is co-editor or a special issue of the PPNA journal on Experimental Evaluation of Peer-to-Peer Applications to appear in 2012. He is member of the scientific council of the INRIA Bell-Labs laboratory on Self Organizing Networks. He was an affiliate professor at Ecole Polytechnique, Palaiseau until september 2011. He also serves regularly as an expert to the European Commission to evaluate EC funded projects.

Claude Castelluccia served in the program committees of the following international conferences: ACM WiSEC 2008, SecureCom 2008, ACM WiSEC 2009, IEEE SECON 2009, SARSII 2009, WiSEC 2010, SESOC 2010, WWW 2010, WiSEC 2011 SESOC 2011 and ACM CCS 2012. He is the co-founder of the ACM WiSec (Wireless Security) conference. He also served as an expert to the European Commission to evaluate EC funded projects.

Thierry Turletti Senior ACM and IEEE member, served in the program committees of the following international conferences: 3rd Workshop on Performance Analysis and Enhancement of Wireless Networks (PAEWN'08), Packet Video'09, 2nd Workshop on mobile Video Delivery (Movid'09), Packet Video'10, 3rd Workshop on mobile Video Delivery (Movid'10), 1st Conference on Wireless and Ubiquitous Systems (ICWUS'10), 19th Packet Video Workshop (PV'11) and 4th ACM Workshop on mobile Video Delivery (Movid'11). He is member of the Editorial Boards of the Journal of Mobile Communication, Computation and Information (WINET) published by Springer Science and of the Advances in Multimedia Journal published by Hindawi Publishing Corporation. He is associated editor of the Wireless Communications, Mobile Computing (WCMC) Weslay Journal published by John Wiley & Sons. Thierry Turletti has served several times as an expert to the European Commission to evaluate and review EC funded projects and also to review French ANR funded projects.

Chadi Barakat is General Co-Chair of the upcoming ACM CoNEXT 2012 conference on emerging Networking EXperiments and Technologies to be held in Nice on Dec 10-13, 2012. He served on the Technical Program Committee for Broadnets 2008, Algotel 2009, ITC 2009, Infocom 2009, Comsnets 2010, IEEE Infocom 2011 and ACM CoNEXT 2011. He was invited to give talks at the CFIP Conference in Sainte Maxime - 2011. He is currently the scientific referee for international affairs at Inria Sophia Antipolis and member of the Conseil d'Orientation Scientifique et Technologique (COST) at Inria within the working group of international affairs (COST-GTRI).

Vincent Roca is strongly involved at IETF and served as co-chair of the MSEC (Multicast Security) working group in 2010-2011. He is also member of the SecDir group (security directory) of the IETF. He was also member of the program committees of CFIP'09, SPACOMM'09, SPACOMM'10, the Cyber and Physical Security and Privacy symposium at IEEE SmartGridComm'11, CFIP'11, SPACOMM'11.

Arnaud Legout was program committee co-chair of the ICCCN 2009 conference track on P2P networking. He also served in the program committee of the following international conferences and workshops: CoNext'2008 and IPTPS'2009. He was also reviewer of journals (IEEE/ACM Transactions on Networking, IEEE/ACM Transactions on Computers, IEEE Network, Computer Communications, ACM SIGCOMM CCR), and conferences (IEEE Infocom, ACM Sigmetrics). He also served as an expert to the European Commission to evaluate EC funded projects.

Mohamed Ali Kaafar In 2011, he served in the program committees of the following international conferences: Security, Inforensics and Cyber Criminality 2011, CCSEIT-2011, ACC 2011, IWTMP2PS 2011. He is member of the steering commitee of Colloque Jacques Cartier: Security, Inforensics and Cyber Criminality. He is member of the editorial board of IEEE Transactions on Parallel and Distributed Systems TPDS, Computer Communications, IEEE letters of communications, Computer Networks, the International Journal of peer-to-peer networks (IJP2P).

4 External Funding

(k euros)	year1	year2	year3	year4
National initiatives				
ANR ARESA2			21	39
ANR CMON			8	52
ANR F-Lab				60
ANR CONNECT				86
ANR SCATTER				58
ANR RFID-AP	5	23	41	
FUI/SHIVA MInalogic			8	39
ANR CAPRI-FEC	47	43	7	
ANR HIPCAL	48	12	25	
ANR Divine	3			
CPER PLEXUS	40	120	100	80
European projects				
FP7 ECODE	16	100	135	
FP7 NOVI				60
FP7 OPENLAB				10
FP7 OneLab2	102	320	279	
ITEA ExpeShare	112	110		
FP7 WSN4CIP		53	63	164
FP6 UbiSec & Sens	99	31		
Associated teams				
Community		10	10	10
Genesim	20			
Ubisec	8	15	10	
Industrial contracts				
ALU	6	13	40	59
CEA-Leti	12	36	36	36
Scholarships				
Inria PhD *	12pm	12pm	12pm	14pm
CIFRE PhD	12pm	12pm	6pm	
MESR PhD				12 pm
Inria Post Doc*	12pm	3pm	19pm	12 pm
AI+	33pm	22pm	28pm	18pm
EI#				3pm
ICT Labs KIC				
FITTING				9
SDN				30
ICDC				40
Total	518 + 69pm	885 + 49pm	783 + 65pm	832 + 56pm

∗ other than those supported by one of the above projects

+ junior engineer supported by INRIA

# experienced engineer supported by INRIA

4.1 ARCs

Collaborative Action CAPRIS (2011-2014): the Collaborative Action on the Protection of Privacy Rights in the Information Society (CAPRIS), is an Inria national project, which goal is to tackle privacy-related challenges and provide solutions to enhance the privacy protection in the Information Society. His main tasks are the identification of existing and future threats to privacy, and the design of appropriate measures to assess and quantify privacy. Preparatory meetings have taken place in 2011, and the collaborative action will start in 2012.

4.2 National initiatives

ANR FIT (2011-2108): FIT (Future Internet of Things) aims to develop an experimental facility, a federated and competitive infrastructure with international visibility and a broad panel of customers. It will provide this facility with a set of complementary components that enable experimentation on innovative services for academic and industrial users. The project will give French Internet stakeholders a means to experiment on mobile wireless communications at the network and application layers thereby accelerating the design of advanced networking technologies for the Future Internet. FIT is one of 52 winning projects from the first wave of the French Ministry of Higher Education and Research's "Équipements d'Excellence" (Equipex) research grant program. The project will benefit from a 5.8 million euro grant from the French government. Other partners are UPMC, IT, Strasbourg University and CNRS. See also http://fit-equipex.fr/. Total funding for Planète: 626 Keuros.

DGCIS PFT (2011-2014) : DGCIS funded project, in the context of the competitivity cluster SCS, whose aim is to provide to PACA region industrials wishing to develop or validate new products related to future mobile networks and services and M2M application, a networking infrastructure and tools helpful for development, test and validation of those products. Other partners : 3Roam, Audilog Groupe Ericsson, Ericsson, Eurecom, Inria, iQsim, MobiSmart, Newsteo, OneAccess, Orange Labs, Pôle SCS, ST Ericsson, Telecom Valley. Our contribution is centred around providing a test methodology and tools for wireless networks experimentation. Total funding for Planète: 228 Keuros.

ANR ARESA2 (2009-2012): The Planète team is involved in the ARESA2 project which aims at advancing the state of the art in Secure, Self-Organizing, Internet-Connected, Wireless Sensor and Actuator Networks (WSANs). These challenges are to be addressed in an energy-efficient way while sticking to memory-usage constraints. The partners are INRIA, CEA-LETI, France Telecom R&D, Coronis Systems, LIG/Drakkar, Verimag and TELECOM Bretagne. Total funding for Planète: 149.5 Keuros.

ANR pFlower (2010-2013): Parallel Flow Recognition with Multi-Core Processor. The main objective of this project is to take advantage of powerful parallelism of multi-thread, multi-core processors, to explore the parallel architecture of pipelined-based flow recognition, parallel signature matching algorithms. The project involves INRIA (Planète), Université de Savoie, and ICT/CAS (China).

Inria Mobilitics (2011-2012): as a joint national project with CNIL (the French national committee of Information freedom). Platform for mobile devices privacy evaluation. This project strives to deploy an experimental mobile platform for studying and analyzing the weaknesses of current online (smartphone) applications and operating systems and the privacy implications for end-users. For instance, one of the objectives is to understand trends and patterns collected when they are aimed at obtaining general knowledge that does not pertain to any specific individual. Examples of such tasks include learning of commuting patterns, inference of recommendation rules, and creation of advertising segments. Inria will provide a funding for an Associate Engineer on this project in 2012.

ANR CMON (2009-2012) : This project involves, in addition to INRIA, Technicolor Paris Lab, LIP6, ENS and the Grenouille.com association. CMON stands for collaborative monitoring. It is an industrial research project that develops the technology needed to allow end-users to collaborate in order to identify the origin and cause of Internet service degradation. The main differentiating assumptions made in this project are that *(i)* ISPs do not cooperate together, and *(ii)* one cannot rely on any information they provide in order to diagnose service problems. Even more, CMON considers that these ISP will try to masquerade the user observations in order to make their service look better. The software designed in this project will be added to the toolbox currently provided by the Grenouille architecture. The hope is that such a project will encourage ISPs to improve their quality of service and will contribute to improve customer satisfaction. See also http://wiki.grenouille.com/index.php/CMON. Total funding for Planète: 78.5 Keuros.

ANR F-Lab (2011-2013): ANR funded project on the federation of computation, storage and network resources, belonging to autonomous organizations operating heterogeneous testbeds (e.g. PlanetLab testbeds and Sensors testbeds). This includes defining terminology, establishing universal design principles, and identifying candidate federation strategies. Other partners : UPMC, A-LBLF and Thales. Total funding for Planète: 185 Keuros.

ANR Connect (2011-2012): ANR funded project on content centric Networking architecture. The aim is to propose adequate naming, routing, cache management and transmission control schemes for CCN based networks. Our contribution is centered on network traffic characterization video streaming and on the integration of the CCNx code in the ns-3 simulator. Other partners: UPMC, Alcatel Lucent, Orange R&D, IT. Total funding for Planète: 173 Keuros.

ANR SCATTER (2011-2012): ANR funded project on Scalable Naming in Information Centric Networks. The goal of this activity is to evaluate the scalability of state of the art naming schemes both from the name resolution and routing points of view. The four main approaches that will be considered are: Content Centric Networking (CCN), Publish-Subscribe Internet Routing Paradigm (PSIRP), Network of Information (NetInf) and Data-Oriented Network Architecture (DONA). Other French partners: UPMC. International KIC partner: SICS. Total funding for Planète: 58 Keuros.

RNRT OSCAR (2006-2008):

The Planète group was a member of the OSCAR RNRT project. This project aimed at studying the attacks against P2P overlays and their impact on the underlying network infrastructure. Planète was responsible of studying the attacks against BitTorrent and the danger those attacks could have on the underlying network infrastructure. In particular, we evaluated how by manipulating the tracker, the BitTorrent traffic could be redirected in such a way to increase significantly the load on some specific links and by how much this could harm the underling network. Extensive simulations and experimentations over Grid 5000 were run for this purpose. The project started in April 2006 and ended in March 2008. It involved teams from both academy and industry, such as LAAS, LIP6, France Telecom, Mitsubishi, ENS Lyon and ENST Bretagne.

RNRT RFID-AP (2008-2010): The Planète group is involved in the RFIDAP RNRT project which aims at designing and prototyping cryptographic algorithms and secure protocols for RFID deployment. Such algorithms and protocols could be used individually, or in combination, and will provide a practical and useful framework within which to apply innovative but practical techniques for device authentication and user privacy. Total funding for Planète: 70 Keuros.

FUI/SHIVA Minalogic (2009-2012):
The goal of the SHIVA (Secured Hardware Immune Versatile Architecture) project is to design a multi-Gbps security platform, compatible with high assurance environments where a clear separation between red and black flows is a must. The partners are CS (leader), EASII IC, INRIA, NETHEOS, UJF-IF, UJF-LJK, UJF-Verimag, TIMA, and IWALL. Total funding for Planète: 132 Keuros.

ANR/VERSO ARSSO (2010-2013):
The goal of this project is to design adaptable robust streaming solutions. The partners include ALU-BL (leader), INRIA, CEA-LETI, ENSICA and Thales. Total funding for Planète: 132 Keuros.

ANR/RNRT CAPRI-FEC (2007-2009):
The goal of this project is to design and analyze Application-Level FEC (AL-FEC) codes for the erasure channel, and their adequacy to wireless applications. The partners are INRIA (leader), CEA-LETI, ENSICA, STMicroelectronics, and Eutelsat). Total funding for Planète: 126 KEuros.

ANR/CIS HIPCAL (2007-2009):
The goal of this project is to design a middle-ware that provides secure communications and assured performances to grids. This middle-ware relies on the HIP (Host Identity Protocol) subsystem, and on host virtualization techniques to dynamically define virtual, confined clusters. The middle-ware will be tested with several biomedical and bio-informatics applications. The partners are INRIA Reso (leader), INRIA Grand Large, INRIA Planète, CNRS IBCP, CNRS I3S. Total funding for Planète: 107 KEuros.

ANR Divine project (2006-2008): The DIVINE ANR project proposes the study and the development of a simple yet realistic system of video and image transmission towards heterogeneous mobile terminals (for instance PDA or digital

TV receiver) through heterogeneous wireless and wired IP links. The application aimed by the DIVINE project is the interactive access to multimedia data in a museum. This is a typical environment where various techniques of wireless (WLAN, WiMAX) and wired transmission can be jointly exploited. The DIVINE project aims to study innovative solutions (scalable video and still image coding, unequal error protection, multicast links, multiple description coding) allowing to perform an end-to-end optimization, based on the detection, optimal management and adaptive processing of this heterogeneity. The project has started in July 2006 and involves teams from both industry and academy as Thales, France Télécom R&D, ETIS, ENST Paris, L2S, LIP6 and the research center of French Museums C2RMF-UMR171. At Planète, we focus in particular on the design and the evaluation of multicast multimedia transmission mechanism for IEEE 802.11 WLANs. Total funding for Planète: 100KEuros.

CPER Plexus (2007-2010) :
This project aims to build an experimental wireless networking platform in several sites in Sophia Antipolis. This platform will be interconnected with the European OneLab platform through INRIA and will integrate Eurecom's radio platform. The goal is to study the performance in terms of bandwidth and radio resources utilization in a heterogeneous radio environment. Total funding for Planète: 340 KEuros.

4.3 European projects

ECODE

- Title: Experimental COgnitive Distributed Engine
- Type: COOPERATION (ICT)
- Defi: New paradigms and experimental facilities
- Instrument: Specific Targeted Research Project (STREP)
- Duration: September 2008 - August 2011
- Coordinator: Alcatel Lucent (Belgium)
- Others partners: UCL (Belgium), ULg (Belgium), IBBT (Belgium), ULANC (UK), CNRS (France).
- See also: http://www.ecode-project.eu/
- Abstract: The goal of the ECODE project is to develop, implement, and validate experimentally a cognitive routing system that can meet the challenges experienced by the Internet in terms of manageability and security, availability and accountability, as well as routing system scalability and quality. By combining both networking and machine learning research fields, the resulting cognitive routing system fundamentally revisits the capabilities of the Internet networking layer so as to address these challenges altogether. For this purpose, the project investigates and elaborates novel semi-supervised, on line, and distributed machine learning techniques kernel of the cognitive routing system. During the building phase, the cognitive routing system is both designed and

prototyped. In the second phase, three sets of use cases are experimented to evaluate the benefits of the developed machine learning techniques. The experimentation and the validation of these techniques are carried out on physical (iLAB) and virtual (e.g.,OneLab) experimental facilities.

- Total funding for Planète: 250 Keuros.

NOVI

- Title: Networking innovations Over Virtualized Infrastructures
- Type: COOPERATION (ICT)
- Defi: CAPACITIES programme.
- Instrument: Specific Targeted Research Project (STREP)
- Duration: September 2010 - February 2013
- Coordinator: NTUA (Greece)
- Others partners: 13 european partners including GARR, ELTE, Cisco, etc.
- See also: http://www.fp7-novi.eu/
- Abstract: NOVI (Networking innovations Over Virtualized Infrastructures) research concentrates on efficient approaches to compose virtualized e-Infrastructures towards a holistic Future Internet (FI) cloud service. Resources belonging to various levels, i.e. networking, storage and processing are in principle managed by separate yet interworking providers. NOVI will concentrate on methods, information systems and algorithms that will enable users with composite isolated slices, baskets of resources and services provided by federated infrastructures.
- Total funding for Planète: 122 Keuros.

OPENLAB

- Title: OpenLab: extending FIRE testbeds and tools
- Type: COOPERATION (ICT)
- Defi: ICT 2011.1.6 Future Internet Research and Experimentation (FIRE)
- Instrument: Integrated Project (IP)
- Duration: September 2011 - January 2014
- Coordinator: Université Pierre et Marie Curie (France)
- Others partners: 18 European partners (including ETH Zurich, Fraunhofer, IBBT, TUB, UAM, etc.) and Nicta from Australia.
- See also: http://www.ict-openlab.eu/

- Abstract: OpenLab brings together the essential ingredients for an open, general purpose and sustainable large scale shared experimental facility, providing advances to the early and successful prototypes serving the demands of Future Internet Research and Experimentation. OpenLab partners are deploying the software and tools that allow these advanced testbeds to support a diverse set of applications and protocols in more efficient and flexible ways. OpenLab's contribution to a portfolio that includes: PlanetLab Europe (PLE), with its over 200 partner/user institutions across Europe; the NITOS and w-iLab.t wireless testbeds; two IMS telco testbeds that can connect to the public PSTN, to IP phone services, and can explore merged media distribution; an LTE cellular wireless testbed; the ETOMIC high precision network measurement testbed; the HEN emulation testbed; and the ns-3 simulation environment. Potential experiments that can be performed over the available infrastructure go beyond what can be tested on the current internet. OpenLab extends the facilities with advanced capabilities in the area of mobility, wireless, monitoring, domain interconnections and introduces new technologies such as OpenFlow. These enhancements are transparent to existing users of each facility. Finally, OpenLab will finance and work with users who propose innovative experiments using its technologies and testbeds, via the open call mechanism developed for FIRE facilities.

- Total funding for Planète: 301 Keuros.

OneLab1 & 2

- Title: OneLab: Future Internet testbeds

- Type: COOPERATION (ICT)

- Defi: ICT 2011.1.6 Future Internet Research and Experimentation (FIRE)

- Instrument: Integrated Project (IP)

- Duration: September 2006 - November 2010

- Coordinator: Université Pierre et Marie Curie (France)

- Others partners: 29 partners (including ALU, BT, Thomson, ETH Zurich, Fraunhofer, UAM, U. Pisa, Thales, etc.) and Nicta from Australia.

- See also: http://www.onelab.eu/

- OneLab has been financed by grants from the European Framework Programmes FP6 and FP7. We refer to each successive round of funding as a different phase of the project. To date there were two phases: The first phase of the project (OneLab1) ran from September 2006 to August 2008. The central aim was to establish an autonomous European testbed for research on the future Internet, which was achieved through the creation of the PlanetLab Europe testbed. Additional aims were to extend, deepen, and federate this testbed: extension to new technologies, notably wireless; deepening through adding monitoring capabilities; and federating with the global PlanetLab system. The second phase and more extensive phase (OneLab2) for 27 months from September 2008 to November 2010) aims to build on the foundations laid under OneLab1, and continue the project of extension, deepening, and federating. Extension now includes "customer" testbeds including SAC testbeds, wireless testbeds, and content-based testbeds. Deepening continues with the incorporation of some major European measurement infrastructures: DIMES and ETOMIC. Federation continues with federation between PlanetLabs, extending to PlanetLab Japan, as well as federation with the "customer" testbeds, such as the Haggle and ANA testbeds.

- Total funding for Planète: 700 Keuros (OneLab2).

WSN4CIP

- Title: Wireless Sensor Networks for critical infrastructures Protection
- Type: COOPERATION (ICT)
- Defi: FP7 Security area, Objective 1.7 Critical Infrastructure Protection
- Instrument: Specific Targeted Research Project (STREP)
- Duration: 2009 - 2011
- Coordinator: Eurescom (Germany)
- Others partners: 11 European partners (including IHP, NEC, BUTE, etc.)
- See also: http://www.wsan4cip.eu/home.html
- Abstract: The goal of WSAN4CIP is to advance the technology of Wireless Sensor and Actuator Networks (WSANs) beyond the current state of the art, in order to improve the protection of Critical Infrastructures (CIs) By advancing WSAN technology, the project contributes to networked information and process control systems which are more secure and resilient. The distributed nature of WSANs enables them to survive malicious attacks as well as accidents and operational failures. It makes them dependable in critical situations, when information is needed to prevent further damage to CIs.
- Total funding for Planète: 364 KEuros

UbiSec&Sens

- Title: Ubiquitous Sensing and Security in the European Homeland
- Type: COOPERATION (ICT)
- Challenge: FP6 Security area, Towards a global dependability and security framework
- Instrument: Specific Targeted Research Project (STREP)
- Duration: 2006 - 2009
- Coordinator: Eurescom (Germany).
- Others partners: RWTH, IHP, NEC, RUB, LTU, INOV and BUTE
- See also: http://www.ist-ubisecsens.org/

- Abstract: The goal of this project is to develop new security protocols for wireless sensor networks. The UbiSec&Sens approach was to use three representative WSN scenarios to iteratively determine solutions for the key WSN issues of scalability, security, reliability, self-healing and robustness. This gave a clearer understanding of the real-world WSN requirements and limitations as well as identifying how to achieve a successful rollout of WSNs. UbiSec&Sens provided a comprehensive architecture for medium and large scale wireless sensor networks with the full level of security that will make them trusted and secure for all applications.
- Total funding for Planète: 235 KEuros.

Expeshare

- Title: Experience Sharing in Mobile Peer Communities
- Type: Collaborative
- Instrument: Eureka ITEA program
- Duration: 2007 - 2009
- Coordinator: VTT (Finland)
- Others partners: Over 25 European partners are involved mainly Philips, Nokia, Telefonica, the GET-INT and the university of Evry.
- See also: http://www.expeshare.org
- Abstract: Expeshare is an ITEA project to enable virtual communities to share media experiences in their personal devices legally and securely. The final aim is to develop and implement an architecture for a wireless peer-to-peer network that links personal devices and realizes DRM and mobile payment functionality and allows for legal and secure sharing of multimedia content and experiences. The role of INRIA in this project was to participate to the design and evaluation of protocols for the network and Peer-to-Peer layer in order to support the sharing of media in a wireless network.
- Total funding for Planète: 222 Keuros.

4.4 Associated teams and other international projects

COMMUNITY Associated team (2009-2011): Planète is an associated team with the UC Santa Cruz's Jack Baskin School of Engineering. The collaborative project is about communication in heterogeneous networks prone to episodic connectivity, see http://inrg.cse.ucsc.edu/community/. Our initial scientific objective throughout the project was to design efficient message delivery mechanisms for challenged and heterogeneous networks, and targeted: (1) The design of a unifying solution to enable message delivery over heterogeneous networks with varying degrees of connectivity. (2) The design of error- and congestion control techniques in episodically connected networks. (3) The exploration of different mechanisms for quality-of-service (QoS) support in such environments. We have re-oriented some of the initial proposed research. In particular, rather than investigating error and congestion control techniques for DTNs, we focused on the development of efficient routing strategies that take into account the utility of nodes to relay messages. Furthermore, we developed a naming scheme that supports message delivery over heterogeneous networks prone to connectivity disruptions.

UBISEC Associated team (2004-2010): is an associated team between UC Irvine (Prof. G.Tsudik) and INRIA Planète project-team. Rapid advances in microelectronics are making it possible to mass-produce tiny inexpensive devices, such as processors, RF-IDs, sensors, and actuators. These devices are already, or soon will be, deployed in many different settings for a variety of purposes, which typically involve tracking (e.g., of hospital patients, military/rescue personnel, wildlife/livestock and inventory in stores/warehouses) or monitoring (e.g., of seismic activity, border/perimeter control, atmospheric or oceanic conditions). In fact, it is widely believed that, in the future, sensors will permeate the environment and will be truly ubiquitous in clothing, cars, tickets, food packaging and other goods. These new highly networked environments create many new exciting security and privacy challenges. The objectives of the UbiSec associated team is to understand and tackle some of them. More specifically, the proposed project will consider the following three topics: infrastructure-less security, nano-security and anonymous association/routing. The team was prolongated for 3 years in November 2007.

Genesim Associated team: (2007-2010) is an associated team between University of Washington (Prof. S. Roy) and INRIA Planète project-team. Evaluation of new network protocols and architectures is at the core of networking research. This evaluation is usually performed using simulations, emulations, or experimental platforms. Each of these evaluation techniques has strengths and weaknesses and therefore they complement one another. However, there is currently no way to combine them in a scientific experimental workflow. On the other hand, wireless network protocols are challenging to evaluate mainly due to the high variability of the channel characteristics and their sensitivity to interference. Indeed, as the wireless environment is very difficult to control, repeatable experiments are complex to perform. In addition, a large number of parameters impact the results of an experiment. It is therefore difficult to find the subset of key parameters to be taken into account to characterise a wireless experiment. The objective of this Associated Team is to contribute toward this area by providing a prototype evaluation environment for wireless experiments. This evaluation environment is based on a common programming interface between ns-3, Orbit and OneLab. This prototype will allow running basic wireless networking scenarios on these three environments and to compare the simulations and experiments' results. Based on University of Washington competence on Orbit and ns-3 and on INRIA's competence on OneLab and ns-3 we expect this common project to have a high impact on both European and International consortiums.

Roseate (STIC AmSud) (2008-2012): this project aims to design realistic models of the physical layer in order to be used in both simulations and experimentation of wireless protocols. In addition to the Planète Project-Team, the partners are Universidad de Valparaiso, Chile, Universidad de Córdoba, Argentina and Universidad Diego Portales, Chile. The collaboration was prolongated for two years in November 2010.

STIC Tunisia : (2007-2010) Collaboration with Sonia Gammar from Professor Farouk Kamoun's team at ENSI (Tunis) in the context of a STIC Tunisia project on Security and Monitoring of Hybrid Wireless Mesh Networks. In this project, we co-supervise a PhD student from University of Tunis (Amine Elabidi) working on Data oriented Networking Architecture.

4.5 Industrial contracts

Selfnet ADR, ALU Bell Labs: (2008-2012)
The goal of this study is the use of AL-FEC techniques in broadcasting systems and in particular on the optimization of FEC strategies for wireless communications. Three persons have been supported: Ferdaouss Mattoussi received a PhD grant, while Amira Alloum and Rodrigue Imad received a post-doc grant. Total funding for Planète: 213 KEuros. *Impacts:* This collaboration enabled us to have more manpower on the AL-FEC subject (in particular GLDPC-Staircase codes and UEP techniques). This collaboration also enabled us to obtain the ANR/VERSO-09 ARSSO (Adaptable Robust Streaming SOlutions) project.

UDcast, Sophia Antipolis: (2007-2010)
UDcast is providing a PhD grant (CIFRE contract) to support the activity on DVB-SH FEC Scheme for Efficient Erasure Recovery. This grant supported Amine Ismail (who defended in 2010).

CEA LETI, Grenoble: (2008-2011)
CEA LETI is providing a PhD grant to support the activity on wireless sensor network security. This grant supports Sana Ben Hamida. Total funding for Planète: 120 KEuros.

4.6 ICT Labs

FITTING

- Title: Future InterneT (of ThINGs) facility
- Activity Number: 10340
- Duration: 2011-2012
- Coordinator: UPMC (France)
- Others partners: Alcatel Lucent, Fraunhofer FOKUS, BME, IT, U. Paris XI.
- Abstract: FITTING develops a testbed federation architecture that combines wireless and wired networks. Through FITTING, components and solutions developed in the projects OneLab2, PII and SensLAB are brought together to facilitate access. These components and devices complement each other: for instance SensLAB enhances the testbed federation by adding wireless sensors. FITTING addresses issues related to usability and accessibility of federated experimentation resources from multiple autonomous organizations. FITTING is a process of federating elements from various European and national initiatives into a global shared resource pool with a standardized interface to access them. Further, FITTING will adopt a user-driven (researchers, developers,

students) approach with its running testbeds allowing experimentation with different technologies to meet the variety of needs of a broad customer base.

- The FITTING activity is mentioned as a "success story" by the EIT ICT LabsKIC[10]. In fact, after an initial funding in 2010, the french partners succeded to get the FIT Equipment of Excellence project accepted with a total budget of 5.8 MEuros to develop a testbed federation in France.
- Total funding for Planète: 25 Keuros

SDN

- Title: Software-Defined Networking
- Activity Number: 11643
- Duration: 2011-2012
- Coordinator: Joerg Ott, Aalto University
- Others partners: SICS, KTH, TU-Berlin, Deutsche Telekom, U Helsinki, Fraunhofer Gesellschaft and TU Munchen.
- Abstract: The objective of this activity is to explore software-defined networking at different positions on the axis between basic flow-level processing (using OpenFlow for end-to-end flows) in controlled fixed networks and cooperation between mobile end nodes in the open wireless Internet (using opportunistic networking for resources communicated hop-by-hop). These complementary elements allow adding value to networking at different increments and support the gradual introduction of information-centric networking concepts from today's perspective. The activity joins researchers with carrier projects in related, yet typically independently investigated areas. By combining 1) the "traditional" notion of SDN with enhancements toward 2) mobile access and 3) mobile ad-hoc environment, this activity makes SDNs more comprehensive.
- Total funding for Planète: 30 Keuros

ICDC

- Title: Information-centric and device clouds
- Activity Number: 11901
- Duration: 2011-2012
- Coordinator: Bengt Ahlgren, SICS
- Others partners: Fraunhofer FOKUS, DFKI, Aalto, Ericsson and UPMC.
- Abstract: Study Information Centric Networking Architectures. Our contribution in this project focuses on the performance of the Content Centric Networking architecture in DTN environment.
- Total funding for Planète: 40 Keuros

[10] See http://eit.europa.eu/kics1/stories-archiv/stories-single-view/article/fitting-from-eit-ict-labs-the-next-generation-testbeds.html

5 Objectives for the next four years

The Planète project-team was created in January 2001. It will reach the 12 years age limit by the end of 2012. The two branches at Sophia-Antipolis and Grenoble will split and two project-team proposals are currently being prepared: one Privacy issues (Privastics) and the other on Networking Architecture. Here follows preliminary description of the two proposals.

5.1 Privatics

Privatics will bring together research scientists coming from different groups in Grenoble and Lyon (4 researchers from the Grenoble branch of the *Planète* project-team, one researcher from *Licit* and one assistant professor from the *Swing* project-team) to work on privacy along the following research lines:

- **Understanding privacy:** Privacy is hard to define and to grasp because it is a multi-faceted concept that actually varies from one context to another or one country to another. In order to understand it, it is important to study the factors that can contribute to strengthen privacy or to threaten it. These factors are not only technical, but also economical, legal, social, etc. This research line will thus be conducted in interaction with other disciplines and all stakeholders (including data protection authorities, legislators, industrial actors, standardization bodies, etc.). Understanding these non technical aspects is critical in order to be able to provide privacy-preserving solutions that have a chance to get deployed. A solution that would not lead to viable business models or that would not be legally or socially acceptable would just be doomed to failure. On the technical side, recent studies in privacy have shown that so-called anonymous datasets can sometimes be de-anonymized to a great extent, thus enabling an adversary to learn personal information about individuals. However, de-anonymization is not the only threat to privacy. For instance, an attacker may access some personal data (possibly sanitized) through legitimate means and then perform an inference attack that combines his *a priori* knowledge with the data obtained legally. Possibly, this new information was thought initially to be out of reach of an adversary by the creator of the privacy policy (or even unknown by himself), mainly because it was difficult for him to assess in advance the inference power and the background knowledge of this adversary. As a result, users are more and more tracked and profiled when they are online. This profiling will still increase with the development of ubiquitous advertising and personalized services. It is therefore important to study how the tracking and profiling techniques evolve. This will be performed via real-time experiments, using, for example, the testbed developed by the members of the group in collaboration with the CNIL.

- **Defining privacy foundations:** One of the key challenges in this area is the need to take into account the necessary trade-off between privacy and other requirements such as utility, security or accountability[11]. One of our objectives will be to provide formal foundations for privacy that encompass, as much as

[11] The requirements for accountability of certain actions (e.g. transactions, access to information, etc.) may themselves introduce new privacy risks because they lead to recording further information in audit logs. Therefore accountability, depending on the actions it applies to, can be either a tool to enhance privacy or a source of risk.

possible, the general principles of privacy (minimization of data, control on personal data, transparency of the treatments, accountability, etc.) as well as other, possibly conflicting, requirements. Defining realistic and formally grounded measures of privacy, adapted and appropriate for specific contexts, is another challenging task; it is also a prerequisite for both evaluating the risks and assessing potential solutions. For instance, being able to define privacy with respect to a particular application is a fundamental issue as it can be used to measure the privacy gained by using protection mechanisms (such as sanitization algorithms). In computer science, several research communities, such as cryptography, databases, statistics and data mining, have tackled the problem of defining and analyzing privacy, each coming with their own approaches. However, it is sometimes difficult to compare these approaches, as their goal as well as the framework on which they are based or the assumption they make can be very different. Promising notions such as differential privacy have recently emerged that attempt to give an ad-omnia definition of privacy and to provide privacy guarantees independently of the adversary background knowledge. Considering the above limitations, one of our objectives will be to compare and contrast how privacy is defined and addressed by the computer science and law communities in different settings. We intend to extract the gist of these different approaches by proposing a more generic perspective to reason about privacy as well as to design new qualitative and quantitative methods to evaluate the privacy protection provided by a particular system.

- **Building Privacy-Preserving Systems:** the third research line of *Privatics* will be the proposal of new solutions to enhance privacy. This research line will be based on the understanding of privacy and the foundations provided by the two other research lines. We will consider both the architectural level (privacy by design) and the component level (privacy enhancing technologies). As far as privacy by design is concerned, as suggested in Section 2.5.3, we intend to build on our formal models for privacy to define systematic methods to explore the design space or check architectures w.r.t. privacy requirements. We also intend to study new ways to enhance transparency in order to provide means for individuals to understand how their personal data (and, ideally, any data that can be used in a processing with potential effects on them) is collected, generated, managed, transferred, etc. These *TETs* (Transparency Enhancing Technologies) can play a critical role in a context where information flows are growing dramatically and the data mining and inference techniques become more and more powerful. We also intend to pursue our effort on *PETs* (Privacy Enhancing Technologies) building blocks typically using cryptography and sanitization techniques (aggregation, perturbation, etc.) and to assess these tools on the basis of the properties and measures developed in the second research line.

The permanent members of *Privatics* are the following:

- Claude Castelluccia (leader, Senior Research Scientist, Inria): networking, protocols, applied cryptography, anonymisation/sanitization algorithms, security, privacy.

- Daniel Le Metayer (Senior Research Scientist, Inria): formal methods, legal aspects, privacy.

- Mohamed-Ali Kaafar (Research Scientist, Inria): networking, protocols, metrology, security, privacy.

- Vincent Roca (Research Scientist, Inria): networking, protocols, coding theory, standardization, security, privacy.

- Cedric Lauradoux (Research Scientist, Inria): applied cryptography, security, privacy.

- Marine Minier (Assistant Professor, Insa): cryptography, privacy.

They have a strong background on privacy and bring the different and complementary expertises which are required to address the objectives of PRIVASTIC. They also have multiple collaborations both in the academic community (in Europe through several European projects and participations in two new action lines on privacy in the EIT ICT Labs, in the US with associated teams at the university of Berkeley and the university of California at Irvine, and in France through the *cappris* action (see below) and national or international bodies such as CNIL or ENISA.

5.2 DIANA

The Internet has been designed almost 50 years ago to interconnect networks. The problem in the 60s was to access remotely dedicated computer resources. The stunning design of the Internet architecture enabled to seamlessly increase its scale by 5 orders of magnitude in less than 40 years.

The Internet has been so successful that it fundamentally changed how people communicate, the worldwide economy and even societies. Everybody can use the Internet from a few years old child to a 90 years old person who was born when computers didn't even exist.

Therefore, it is legitimate to ask whether the Internet has achieved a maturity that makes research on it closer to maintenance engineering than to ground breaking scientific research. As surprising as it might seem, the Internet does not fit current needs, and its usage has been adapted to work around fundamental flaws. These flaws are not minor issues that can be patched, but very complex problems that require to rethink the whole Internet architecture.

The main problem with the Internet architecture of today is the difficulty to find and access contents in an efficient and easy way. It is misleading to believe that Google is solving the issue. All Web indexing services only address publicly available Web contents, but they do not address at all the tremendous amount of personal contents generated by users. These contents are spread on numerous devices (laptop, smartphone, etc.) and private online services (facebook, flick, google docs, dropbox, etc.) Consolidating personal contents, that is locating and accessing them is today a very complex task that leads to data loss, privacy exposure, and users frustration.

Furthermore, we envision that the lack of privacy and security will become the number one concern for the next 10 to 15 years. By design, the Internet exposes to anybody all people's activities. This is a major issue as a growing number of daily activities use the Internet, thus a more invasive privacy leak.

However, these flaws are only the visible part of a more fundamental issue that appeared in early 2000 with the increasing number of Internet devices and services. Whereas the Internet enables a interconnection of devices once devices are properly configured, it fails to transparently give access to features and information. We define a feature as any functionality (or peripheral) available on a device; it might be a speaker or a microphone, a projector, a TV screen, a CPU, a hard drive, etc. We define an information as a generalization of the notion of content. In addition to classical content, an information can be a flow of content updated in real time, a description of a peripheral, a Web service, etc.

Enabling a transparent access to features and information will truly revolutionize the way the Internet is used. To understand this breakthrough long-term objective, let's take two simple examples. Today, it is not possible to send on a TV screen pictures taken from a mobile phone, unless in case of proprietary solutions and at the expense of complex configuration. Now, imagine that the mobile phone is able, without any manual configuration, to transparently and securely find all features available in its vicinity. It will then become possible to send any information from the smartphone to any available feature like, for instance, a TV screen, a projector, or an audio equipment, and this, even in the absence of the infrastructure network.

The second example is related to social networking, which generates lot of activities nowadays and introduces a new way of communication. With social networking, the most important is to find friends and share information with them wherever they are. It is also important to post to friends content generated by all personal devices. Today these networks are operated by proprietary firms that detain lot of information on their subscribers. It will be very useful if the network can localize by itself the content of interest and leverage for that all communication opportunities, whether fixed or mobile.

This long term objective leads to three grand challenges. First, how to enable an automatic, seamless, and secure connection to the Internet? Second, how to seamlessly publish and access information at any scale and in an efficient way, while preserving privacy. The third challenge is how to incrementally deploy these mechanisms and ensure their adoption by end users, content providers and network operators. We will tackle these challenges with an experimental approach by measuring and observing the current behavior of Internet users and available technology and come up with models for the ways information are exchanged. Then we will design and evaluate protocols and system solutions that allow this seamless, automatic, efficient, and secure access to information and features. We will in particular focus on whether we follow a clean-slate approach or leverage on the existing technologies towards the solutions of the above three challenges.

The permanent members of *DIANA* are the following:

- Walid Dabbous (leader, Senior Research Scientist, Inria): networking, protocols, architectures.
- Thierry Turletti (Senior Research Scientist, Inria): networking, protocols, architectures, wireless networks, ad-hoc networks.
- Chadi Barakat (Research Scientist, Inria): networking, protocols, architectures, network tomography, performance evaluation.
- Arnaud Legout (Research Scientist, Inria): networking, protocols, architectures, peer-to-peer systems, large scale measurements, security and privacy.

They have a strong background in Internet protocols and architecture, with a focus an experimental evaluation and measurements. The team members have several national, european, and international collaborations.

6 Bibliography of the project-team

6.1 Doctoral dissertations and "Habilitation" theses

[Bar09] Chadi Barakat. *Solutions efficaces pour la métrologie de l'Internet.* Habilitation à diriger des recherches, Université de Nice Sophia Antipolis, 01 2009.

[BR11] Rao Naveed Bin Rais. *Communication Mechanisms for Message Delivery in Heterogeneous Networks Prone to Episodic Connectivity.* PhD thesis, Université de Nice Sophia-Antipolis, February 2011.

[Cas08] C. Castelluccia. *Sécurité des systèmes sans fil embarqués.* Habilitation à diriger des recherches, 2008.

[Cun10] Mathieu Cunche. *Codes AL-FEC hautes performances pour las canaux á effacements: variations autour des codes LDPC.* PhD Thesis in Computer Science, Joseph Fourier University, Grenoble, May 2010.

[Dab08] W. Dabbous. *Quelle architecture pour l'Internet du futur?* Habilitation à diriger des recherches, 2008.

[Duj09] Diego Dujovne. *Amélioration des expérimentations sur réseaux sans fil.* PhD thesis, Université de Nice Sophia Antipolis, 05 2009.

[Fra09] Aurélien Francillon. *Attacking and Protecting Constrained Embedded Systems from Control Flow Attacks.* Phd thesis, INPG, October 2009.

[Ism10] Mohamed Amine Ismail. *Study and Optimization of Data Protection, Bandwidth Usage and Simulation Tools for Wireless Networks.* PhD Thesis in Computer Science, University of Nice-Sophia Antipolis, June 2010.

[Jab11] Mohamad Jaber. *Internet Traffic Profiling and Identification.* PhD thesis, Université de Nice Sophia Antipolis, October 2011.

[Lac10] Mathieu Lacage. *Outils d'Expérimentation pour la Recherche en Réseaux.* PhD Thesis in Computer Science, University of Nice-Sophia Antipolis, November 2010.

[Las11] Imed Lassoued. *Adaptive Monitoring and Management of Internet Traffic.* PhD thesis, Université de Nice Sophia Antipolis, December 2011.

[LB11] Stevens Le Blond. *Is Privacy Dead (in P2P Networks)?* PhD thesis, Université de Nice Sophia Antipolis, April 2011.

[Leg12] A. Legout. *Efficacité et vie privée : de BitTorrent à Skype.* Habilitation à diriger des recherches, 2012.

[Per11] Daniele Perito. *Exécution sécurisée de code sur systèmes embarqués.* PhD thesis, Université de Grenoble, October 2011.

[Sba10] Mohamed Karim Sbai. *Architecture for Content Sharing in Wireless Networks.* PhD Thesis in Computer Science, University of Nice-Sophia Antipolis, October 2010.

[Soo09] Mate Soos. *Privacy-preserving Security Protocols for RFIDs.* Phd thesis, INPG, October 2009.

[UR12] Shafqat Ur-Rehman. *Enabling Benchmarking in Wireless Networks.* PhD thesis, Université de Nice Sophia Antipolis, January 2012.

6.2 Articles in referred journals and book chapters

[BRTO11] Rao Naveed Bin Rais, Thierry Turletti, and Katia Obraczka. Message Delivery in Heterogeneous Networks prone to Episodic Connectivity. *ACM/Kluwer Wireless Networks*, 17(8):1775–1794, 2011.

[BUC10] Md. Borhan Uddin and Claude Castelluccia. Toward clock skew based services in wireless sensor networks. *International Journal of Sensor Networks (IJSNet)*, 2010.

[CC11] C. Castelluccia and A. Chan. A security framework for privacy-preserving data aggregation in wireless sensor networks. *ACM ToSN (Transaction on Sensor Networks)*, 7, 2011.

[CCMT09] C. Castelluccia, A. Chan, E. Meykletun, and G. Tsudik. Efficient and provably secure aggregation of encrypted data in wireless sensor networks. *ACM ToSN (Transaction on Sensor Networks)*, 2009.

[DGK10] Benoit Donnet, Bamba Gueye, and Mohamed Ali Kaafar. A survey on network coordinates systems, design, and security. *IEEE Communication Surveys & Tutorials*, 12(4), October 2010.

[DGK11] Benoit Donnet, Bamba Gueye, and Mohamed Ali Kaafar. Path similarity evaluation using bloom filters. *Computer Networks*, 2011.

[DMS$^+$09] R. Di Pietro, L. V. Mancini, C. Soriente, A. Spognardi, and G. Tsudik. Data security in unattended wireless sensor networks. *Autonomic Network Computing, IEEE Transaction on Computers*, 2009.

[DTF10] Diego Dujovne, Thierry Turletti, and Fethi Filali. A Taxonomy of IEEE 802.11 Wireless Parameters and Open Source Measurement Tools. *IEEE Surveys and Tutorials*, 12(2):249–262, Second Quarter 2010.

[FFP10] Serge Fdida, Timur Friedman, and Thierry Parmentelat. *OneLab: An Open Federated Facility for Experimentally Driven Future Internet Research.* Springer, 2010.

[GBM11] Luigi Alfredo Grieco, Chadi Barakat, and Michele Marzulli. Spectral models for bitrate measurement from packet sampled traffic. *IEEE Transactions on Network and Service Management*, 8(2), June 2011.

[KBS12] Amir Krifa, Chadi Barakat, and Thrasyvoulos Spyropoulos. Message drop and scheduling in dtns: Theory and practice. *IEEE Transactions on Mobile Computing*, 2012. to appear.

[LBLD10] Stevens Le Blond, Arnaud Legout, and Walid Dabbous. Pushing BitTorrent Locality to the Limit. *Computer Networks*, doi:10.1016/j.comnet.2010.09.014, 2010.

[LFH$^+$10] Mathieu Lacage, Martin Ferrari, Mads Hansen, Thierry Turletti, and Walid Dabbous. NEPI: Using Independent Simulators, Emulators, and Testbeds for Easy Experimentation. *ACM Operating Systems Review (OSR)*, 43(4), January 2010.

[LNM+09] Tianji Li, Qiang Ni, David Malone, Douglas Leith, Yang Xiao, and Thierry Turletti. Aggregation with Fragment Retransmission for Very High-Speed WLANs. *IEEE/ACM Transactions on Networking*, 17(2), April 2009.

[MBD] Mohammad Malli, Chadi Barakat, and W. Dabbous. Chess: An application-aware space for enhanced scalable services in overlay networks. *Computer Communications Journal*, 31(6):1239–1253, April.

[MLHT09] Mohamed Hossein Manshaei, Mathieu Lacage, Ceilidh Hoffmann, and Thierry Turletti. On Selecting the Best Transmission Mode for WiFi Devices. *Wireless Communications and Mobile Computing*, 9(7), July 2009.

[PUK+11] Ingmar Poese, Steve Uhlig, Mohamed Ali Kaafar, Benoit Donnet, and Bamba Gueye. IP Geolocation Databases: Unreliable? *ACM SIGCOMM Computer Communication Review*, 41(2), April 2011.

[PWC09] S. Peter, D. Westhoff, and C. Castelluccia. A survey on the encryption of convergecast-traffic with in-network processing. *IEEE Transactions on Dependable and Secure Computing (TDSC)*, 2009.

[SB09] Karim Sbai and Chadi Barakat. Experiences on enhancing data collection in large networks. *Computer Networks*, 53(7), May 2009.

[SBRT+10] Thrasyvoulos Spyropoulos, Rao Naveed Bin Rais, Thierry Turletti, Katia Obraczka, and Athanasios Vasilakos. Routing for disruption tolerant networks: taxonomy and design. *Wireless Networks*, 16:2349–2370, November 2010.

[SBRT+11] Thrasyvoulos Spyropoulos, Rao Naveed Bin Rais, Thierry Turletti, Katia Obraczka, and Athanasios Vasilakos. *DTN Routing: Taxonomy & Design*, chapter 2. CRC Press, 2011.

[STO09] T. Spyropoulos, T. Turletti, and K. Obraczka. Routing in Delay Tolerant Networks Comprising Heterogeneous Node Populations. *IEEE Transaction on Mobile Computing*, 8(8), August 2009.

[TLL+11] Pierre-Ugo Tournoux, Emmanuel Lochin, Jérôme Lacan, Amine Bouabdallah, and Vincent Roca. On-the-fly erasure coding for real-time video applications. *IEEE Transactions on Multimedia*, 13(4), August 2011.

6.3 Publications in Conferences and Workshops

[AHLLB09] Anwar Al-Hamra, Nikitas Liogkas, Arnaud Legout, and Chadi Barakat. Swarming overlay construction strategies. In *Proceedings of ICCCN 2009*, San Francisco, CA, USA, 2009.

[ALX11] Eitan Altman, Arnaud Legout, and Yuedong Xu. Network Non-Neutrality Debate: An Economic Analysis. In *IFIP Networking 2011*, page 12p, Valencia, Spain, May 2011. IFIP Technical Committee on Communication Systems (TC 6), Springer.

[ASR09] Amira Alloum, Bessem Sayadi, and Vincent Roca. Reed solomon codes on graph for dvb-sh streaming services. In *22nd Wireless World Research Forum (WWRF'09)*, Paris, France, 2009.

[BBH$^+$11] Elie Bursztein, Romain Beauxis, Paskov Hristo, Daniele Perito, Celine Fabry, and John Mitchell. The failure of noise-based noncontinuous audio captchas. In *Proceedings of the 37th IEEE Symposium on Security & Privacy, Oakland California*, 2011.

[BDT09] Bilel Ben Romdhanne, Diego Dujovne, and Thierry Turletti. Efficient and Scalable Merging Algorithms for Wireless Traces. In *Proceedings of the 4th ROADS Workshop*, Montana, USA, October 2009.

[BHPC10] Sana Ben Hamida, Jean Benoit Pierrot, and Claude Castelluccia. Empirical analysis of uwb channel characteristics for secret key generation in indoor environments. In *proceedings of The 21st IEEE International Symposium on Personal, Indoor and Mobile Radio Communications PIMRC'10*, Istambul, Trukey, September 2010.

[BKM10] Chérifa Boucetta, Mohamed Ali Kaafar, and Marine Minier. How secure are secure localizations protocols in wsns. In *The ICST Conference on Wireless Sensor Network (WSN) Systems and Software*, Miami, December 2010. ICST.

[BPC09] S. BenHamida, J.B. Pierrot, and C. Castelluccia. An adaptive quantization algorithm for secret key generation using radio channel measurements. In *International Conference on New Technologies, Mobility and Security (NTMS)*, Dec. 2009.

[BRATO11] Rao Naveed Bin Rais, Mariem Abdelmoula, Thierry Turletti, and Katia Obraczka. Naming for Heterogeneous Networks Prone to Episodic Connectivity. In *IEEE WCNC*, Cancun, Mexico, March 2011.

[BRMTO10] Rao Naveed Bin Rais, Marc Mendonca, Thierry Turletti, and Katia Obraczka. "message delivery in heterogeneous disruption-prone networks". Demo description in Proc. of ACM Mobicom, September 2010.

[BRMTO11] Rao Naveed Bin Rais, Marc Mendonca, Thierry Turletti, and Katia Obraczka. Towards Truly Heterogeneous Internets: Bridging Infrastructure-based and Infrastructure -less Networks. In *The third International Conference on COMmunication Systems and NETworkS (COMSNETS)*, Bangalore, India, January 2011.

[CA11] Claude Castelluccia and Gergely Acs. I have a dream! (differentially private smart metering). In *Proceedings of the The 13th Information Hiding Conference (IH)*, 2011.

[CAK12] Abdelberi Chaabane, Gergely Acs, and Mohamed Ali Kaafar. You Are What You Like! Information leakage through users' Interests. In *proceedings of the The Network & Distributed System Security Symposium (NDSS)*, San Diego, February 2012.

[CAL11] Claude Castelluccia, Gergely Acs, and William Lecat. Protecting against physical resource monitoring. In *Proceedings of the 10th ACM Workshop on Privacy in the Electronic Society (WPES)*, 2011.

[CB10] Roberto Cascella and Chadi Barakat. Estimating the access link quality by active measurements. In *proceedings of the 22nd International Teletraffic Congress (ITC 22)*, Amsterdam, Netherlands, September 2010.

[CC08] A.C.-F. Chan and C. Castelluccia. On the (im)possibility of aggregate message authentication codes. In *Information Theory, 2008. ISIT 2008. IEEE International Symposium on*, pages 235–239, July 2008.

[CDCFAK11] Claude Castelluccia, Emiliano De Cristofaro, Aurélien Francillon, and Mohamed Ali Kaafar. EphPub: Toward Robust Ephemeral Publishing. In *proceedings of the 19th IEEE International Conference on Network Protocols (ICNP)*, Vancouver, October 2011.

[CDCP10] Claude Castelluccia, Emiliano De Cristofaro, and Daniele Perito. Private information disclosure from web searches. In *proceedings of the 2010 Privacy Enhancing Technologies Symposium (PETS)*, Berlin, Germany, 2010.

[CET10] Claude Castelluccia, Karim ElDefrawy, and Gene Tsudik. Link-layer encryption effect on the capacity of network coding in wireless networks. In *proceedings of IEEE Infocom MiniConference*, San Diego, CA, March 2010.

[CF08] C. Castelluccia and A. Francillon. Proteger les reseaux de capteurs sans fil. In *SSTIC2008*, 2008.

[CFSP09] C. Castelluccia, A. Francillon, C. Soriente, and D. Perito. On the difficulty of software-based attestation of embedded devices. In *ACM CCS'09: Proceedings of the 16th ACM conference on Computer and Communications Security*, Nov. 2009.

[CK09a] Claude Castelluccia and Mohamed Ali Kaafar. Owner-centric networking: A new architecture for a pollution-free internet. *ERCIM News, Special Theme on Future Internet Technology*, 77, 2009.

[CK09b] Claude Castelluccia and Mohamed Ali Kaafar. Owner-centric networking (ocn): Toward a data pollution-free internet. In *SAINT Workshop on Trust and Security in the Future Internet, FIST*, Seattle, USA, July 2009. IEEE Communications Society.

[CK09c] Abdelberi Chaabane and Mohamed Ali Kaafar. Revisiting unstructured overlay network security. In *Foundations and Practice Of Security workshop*, Grenoble, June 2009.

[CKMP09] C. Castelluccia, M. A. Kaafar, P. Manils, and D. Perito. Geolocalization of proxied services and its application to fast-flux hidden servers. In *ACM/Usenix Internet Measurement Conference IMC 2009*, Chicago, USA, November 2009. ACM.

[CMK10] Abdelberi Chaabane, Pere Manils, and Mohamed Ali Kaafar. Digging into anonymous traffic: a deep analysis of the tor anonymizing network. In *IEEE International Conference in Network and System Security (NSS)*, Melbourne, September 2010. IEEE.

[CR08] M. Cunche and V. Roca. Optimizing the error recovery capabilities of ldpc-staircase codes featuring a gaussian elimination decoding scheme. In *10th IEEE International Workshop on Signal Processing for Space Communications (SPSC'08), Rhodes Island, Greece*, October 2008.

[CR09a] M. Cunche and V. Roca. Adding integrity verification capabilities to the ldpc-staircase erasure correction codes. In *IEEE Global Communications Conference (GLOBECOM 2009)*, November 2009.

[CR09b] Mathieu Cunche and Vincent Roca. Le RFC 5170 en pratique : conception et évaluation d'un codec AL-FEC LDPC-staircase hautes performances. In *Colloque Francophone sur l'Ingénierie des Protocoles (CFIP)*, October 2009.

[CS08] C. Castelluccia and C. Soriente. Abba: A balls and bins approach to secure aggregation in wsns. In *Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks and Workshops, 2008. WiOPT 2008. 6th International Symposium on*, pages 185–191, April 2008.

[CSR$^+$08] M. Cunche, V. Savin, V. Roca, G. Kraidy, A. Soro, and J. Lacan. Low-rate coding using incremental redundancy for gldpc codes. In *IEEE International Workshop on Satellite and Space Communications 2008 (IWSSC'08)*, October 2008.

[CSR10] Mathieu Cunche, Valentin Savin, and Vincent Roca. Analysis of Quasi-Cyclic LDPC codes under ML decoding over the erasure channel. In *proceedings of IEEE International Symposium on Information Theory and its Applications (ISITA'10) (http://arxiv.org/abs/1004.5217)*, Taichung, Taiwan, April 2010.

[DMS$^+$08] R. Di Pietro, L. V. Mancini, C. Soriente, A. Spognardi, and G. Tsudik. Catch me (if you can): Data survival in unattended sensor networks. In *IEEE PerCom'08*, 2008.

[DMS$^+$09] R. Di Pietro, L. V. Mancini, C. Soriente, A. Spognardi, and G. Tsudik. Data security in unattended wireless sensor networks. In *Autonomic Network Computing, IEEE Transaction on Computers*, 2009.

[DSST09] R. Di Pietro, C. Soriente, A. Spognardi, and G. Tsudik. Intrusion-resilience via collaborative authentication in unattended wsns. In *ACM WiSec '09*, 2009.

[FC08] A. Francillon and C. Castelluccia. Code injection attacks on harvard-architecture devices. In *CCS '08: Proceedings of the 15th ACM conference on Computer and communications security*, pages 15–26, New York, NY, USA, 2008. ACM.

[FPC09] A. Francillon, D. Perito, and C. Castelluccia. Defending embedded systems against control flow attacks. In *SECUCODE'09: 1st ACM wokshop on secure code execution*, Nov. 2009.

[FT09] J. Farooq and T. Turletti. An IEEE 802.16 WiMAX Module for the NS-3 Simulator. In *ICST Simutools*, Roma, Italy, March 2009.

[GB09] Luigi Alfredo Grieco and Chadi Barakat. An analysis of packet sampling in the frequency domain. In *proceedings of the ACM Internet Measurement Conference (IMC)*, Chicago, November 2009.

[GB10] Luigi Alfredo Grieco and Chadi Barakat. A frequency domain model to predict the estimation accuracy of packet sampling. In *proceedings of the IEEE Infocom MiniConference*, San Diego (CA), March 2010.

[IPGT10] Mohamed Amine Ismail, Giuseppe Piro, Luigi Alfredo Grieco, and Thierry Turletti. An Improved IEEE 802.16 WiMAX Module for the NS-3 Simulator; Best Student Award Paper. In *proceedings of ICST Simutools'2010*, Torremolinos, Malaga, Spain, March 2010.

[ITD09] Mohamed Amine Ismail, Thierry Turletti, and Walid Dabbous. Optimizing the DVB-SH FEC Scheme for Efficient Erasure Recovery. In *Mobile Video Delivery (MOVID) Workshop at Infocom*, Rio de Janeiro, Brazil, April 2009.

[JB09] Mohamad Jaber and Chadi Barakat. Enhancing application identification by means of sequential testing. In *proceedings of IFIP/TC6 Networking Conference*, Aachen, Germany, May 2009.

[JCB11a] Mohamad Jaber, Roberto Cascella, and Chadi Barakat. Boosting statistical application identification by flow correlation. In *proceedings of EuroNF-TCCFI (International Workshop on Traffic and Congestion Control for the Future Internet)*, Greece, April 2011.

[JCB11b] Mohamad Jaber, Roberto Cascella, and Chadi Barakat. Can we trust the inter-packet time for traffic classification? In *proceedings of IEEE International Conference on Communications (ICC)*, Kyoto, Japan, June 2011.

[JCB12] Mohamad Jaber, Roberto Cascella, and Chadi Barakat. Using host profiling to refine statistical application identification. In *proceedings of IEEE INFOCOM Mini-Conference*, Orlando, FL, March 2012.

[JNB10] Mohamad Jaber, Cao-Cuong Ngo, and Chadi Barakat. A view from inside a distributed internet coordinate system. In *proceedings of the Global Internet Symposium at IEEE Infocom*, San Diego (CA), March 2010.

[JRRAA10] Ludovic Jacquin, Vincent Roca, Jean-Louis Roch, and Mohamed Al Ali. Parallel arithmetic encryption for high-bandwidth communications on multicore/gpgpu platforms. In *ACM workshop on Parallel Symbolic Computation (PASCO 2010)*, July 2010.

[KBS08a] Amir Krifa, Chadi Barakat, and T. Spyropoulos. An Optimal Joint Scheduling and Drop Policy for Delay Tolerant Networks. In *in proceedings of the WoWMoM Workshop on Autonomic and Opportunistic Communications*, Newport Beach (CA), June 2008.

[KBS08b] Amir Krifa, Chadi Barakat, and T. Spyropoulos. Optimal Buffer Management Policies for Delay Tolerant Networks. In *in proceedings of the 5th IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON 2008)* - **Best Paper**, San Francisco, June 2008.

[KBS11] Amir Krifa, Chadi Barakat, and Thrasyvoulos Spyropoulos. Mobitrade: Trading content in disruption tolerant networks. In *proceedings of ACM Mobicom Workshop on Challenged Networks (CHANTS)*, Las Vegas, September 2011.

[KCGL09] Mohamed Ali Kaafar, Francois Cantin, Bamba Gueye, and Guy Leduc. Detecting triangle inequality violations for internet coordinate systems. In *Proceedings of International Workshop on the Network of the Future*, June 2009.

[KGC$^+$08] M. A. Kaafar, B. Gueye, F. Cantin, G. Leduc, and L. Mathy. Towards a two-tier internet coordinate system to mitigate the impact of triangle inequality violations. In *In Proc. of Networking'2008*, Singapore, May 2008.

[KLB10] Amir Krifa, Imed Lassoued, and Chadi Barakat. Emulation platform for network wide traffic sampling and monitoring. In *proceedings of the 1st International Workshop on TRaffic Analysis and Classification (TRAC)*, Caen, France, June 2010.

[KM10] Mohamed Ali Kaafar and Pere Manils. Why spammers should thank google. In *ACM EUROSYS on Social Network Systems (SNS 2010)*, Paris, April 2010. ACM.

[KMB$^+$09] Mohamed Ali Kaafar, Laurent Mathy, Chadi Barakat, Kave Salamatian, Thierry Turletti, and Walid Dabbous. Certified internet coordinates. In *proceedings of IEEE ICCCN conference*, San Francisco, August 2009.

[KMR$^+$11] Amir Krifa, Marc Mendonca, Rao Naveed Bin Rais, Chadi Barakat, Thierry Turletti, and Katia Obraczka. "efficient content dissemination in heterogeneous networks prone to episodic connectivity". Demo at ACM Sigcomm, August 2011.

[KSBT09a] Amir Krifa, Karim Sbai, Chadi Barakat, and Thierry Turletti. Bithoc: A content sharing application for wireless ad hoc networks. In *demo description in proceedings of the IEEE Percom conference*, Galveston, Texas, March 2009.

[KSBT09b] Amir Krifa, Karim Sbai, Chadi Barakat, and Thierry Turletti. A standalone content sharing application for spontaneous communities of mobile handhelds. In *demo description in proceedings of the ACM SIGCOMM MobiHeld Workshop*, Barcelona, August 2009.

[LB10] Imed Lassoued and Chadi Barakat. Adaptive Multi-task Monitoring System Based on Overhead Prediction. In *proceedings of the ACM CoNext PRESTO workshop on Programmable Routers for Extensible Services of Tomorrow*, Philadelphia (PA), November 2010.

[LB11] Imed Lassoued and Chadi Barakat. A multi-task adaptive monitoring system combining different sampling primitives. In *proceedings of the 23rd International Teletrafic Congress (ITC)*, San Francisco, September 2011.

[LBA11] Imed Lassoued, Chadi Barakat, and Konstantin Avrachenkov. Network-wide monitoring through self-configuring adaptive system. In *proceedings of IEEE INFOCOM*, Shanghai, China, April 2011.

[LBLFLM09a] Stevens Le Blond, Fabrice Le Fessant, and Erwan Le Merrer. Choix de partenaires en p2p suivant des critères de disponibilité. In *conférence francaise sur les systèmes d'exploitation*, Toulouse, France, 2009.

[LBLFLM09b] Stevens Le Blond, Fabrice Le Fessant, and Erwan Le Merrer. Finding good partners in availability-aware p2p networks. In *International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS'09)*, Lyon, France, 2009.

[LBLL$^+$10] Stevens Le Blond, Arnaud Legout, Fabrice Lefessant, Walid Dabbous, and Mohamed Ali Kaafar. Spying the World from your Laptop - Identifying and Profiling Content Providers and Big Downloaders in BitTorrent. In *proceedings of LEET'10*, San Jose, CA, USA, April 2010.

[LBMC$^+$10] Stevens Le Blond, Pere Manils, Abdelberi Chaabane, Mohamed Ali Kaafar, Arnaud Legout, Claude Castelluccia, and Walid Dabbous. De-anonymizing BitTorrent Users on Tor. In *poster session at the 7th USENIX Symposium on Network Design and Implementation (NSDI '10)*, San Jose, CA, USA, April 2010.

[LBMC$^+$11] Stevens Le Blond, Pere Manils, Abdelberi Chaabane, Mohamed Ali Kaafar, Dali, Claude Castelluccia, Arnaud Legout, and Walid Dabbous. One Bad Apple Spoils the Bunch: Exploiting P2P Applications to Trace and Profile Tor Users. In *4th USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET '11)*, Boston, United States, March 2011. USENIX.

[LBZL$^+$11] Stevens Le Blond, Chao Zhang, Arnaud Legout, Keith Ross, W., and Walid Dabbous. I Know Where You are and What You are Sharing: Exploiting P2P Communications to Invade Users' Privacy. In *Internet Measurement Conference (ACM/USENIX IMC)*. ACM/USENIX, November 2011.

[LFHT09] Mathieu Lacage, Martin Ferrari, Mads Hansen, and Thierry Turletti. NEPI: Using Independent Simulators, Emulators, and Testbeds for Easy Experimentation. In *Proceedings of the 4th ROADS Workshop*, Montana, USA, October 2009.

[LKG+09] Yongjun Liao, Mohamed Ali Kaafar, Bamba Gueye, Francois Cantin, Pierre Geurts, and Guy Leduc. Detecting triangle inequality violations in internet coordinate systems by supervised learning. In *Proceedings of the IFIP Networking Conference 2009*, May 2009.

[LMGK+11] Amine Labidi, Sonia Mettali Gammar, Farouk Kamoun, Walid Dabbous, Thierry Turletti, and Arnaud Legout. Hybrid approach for experimental networking research. In *13th International Conference on Distributed Computing and Networking, (ICDCN)*, Hong Kong, China, January 2011.

[MCLB+10] Pere Manils, Abdelberi Chaabane, Stevens Le Blond, Mohamed Ali Kaafar, Arnaud Legout, Claude Castelluccia, and Walid Dabbous. Compromising tor anonymity exploiting p2p information leakage. In *ACM Hot Topics in Privacy Enhancing Technologies*, Berlin, July 2010. ACM.

[MDC+10] Kazuhisa Matsuzono, Jonathan Detchart, Mathieu Cunche, Vincent Roca, and Hitoshi Asaeda. Performance analysis of a high-performance real-time application with several al-fec schemes. In *35th IEEE Conference on Local Computer Network (LCN'10)*, Denver, Colorado, U.S.A., October 2010.

[MLLK08] P. Marciniak, N. Liogkas, A. Legout, and E. Kohler. Small is not always beautiful. In *In Proc. of IPTPS'2008*, Tampa Bay, FL, USA, February 2008.

[MLT08] F. Maguolo, M. Lacage, and T. Turletti. Efficient Collision Detection for Auto Rate Fallback Algorithm. In *3rd Workshop on multiMedia Applications over Wireless Networks*, Marrakech, Morocco, July 2008.

[PCKM11] Daniele Perito, Claude Castelluccia, Mohamed Ali Kaafar, and Pere Manils. How Unique and Traceable are Usernames. In *proceedings of the 11th Privacy Enhancing Technologies Symposium(PETS)*, Waterloo, July 2011.

[QLF+11] Alina Quereilhac, Mathieu Lacage, Claudio Freire, Thierry Turletti, and Walid Dabbous. NEPI: An Integration Framework for Network Experimentation. In *19th International Conference on Software Telecommunications and Computer Networks (SoftCOM)*, Dubrovnik, Croatia, September 2011.

[RCHBC09] K. Rasmussen, C. Castelluccia, T. Heydt-Benjamin, and S. Capkun. Proximity-based access control for implantable medical devices. In *ACM CCS'09: Proceedings of the 16th ACM conference on Computer and Communications Security*, Nov. 2009.

[RLB+11] Ashwin Rao, Yeon-Sup Lim, Chadi Barakat, Arnaud Legout, Don Towsley, and Walid Dabbous. Network Characteristics of Video Streaming Traffic. In *proceedings of ACM CoNEXT'11*, Tokyo, Japan, December 2011.

[RLD10a] Ashwin Satish Rao, Arnaud Legout, and Walid Dabbous. BitTorrent Experiments on Testbeds: A Study of the Impact of Network Latencies. In *JDIR'10*, Sophia Antipolis, France, March 2010.

[RLD10b] Ashwin Satish Rao, Arnaud Legout, and Walid Dabbous. Can Realistic BitTorrent Experiments Be Performed on Clusters? In *proceedings of P2P'2010*, Delft, Netherlands, August 2010.

[RTO08] R.N. Bin Rais, T. Turletti, and K. Obraczka. Coping with Episodic Connectivity in Heterogeneous Networks. In *ACM MSWiM*, Vancouver, Canada, October 2008.

[SB09] Karim Sbai and Chadi Barakat. Revisiting content sharing in wireless ad hoc networks. In *proceedings of the fourth workshop on self-organizing systems (IWSOS)*, Zurich, December 2009.

[SBC$^+$08] Karim Sbai, Chadi Barakat, Jaeyoung Choi, Anwar Al Hamra, and T. Turletti. Adapting BitTorrent to wireless ad hoc networks. In *in proceedings of the AdHoc-Now Networks and Wireless conference*, Sophia Antipolis, September 2008.

[SCLR09] A. Soro, M. Cunche, J. Lacan, and V. Roca. Erasure codes with a banded structure for hybrid iterative-ml decoding. In *IEEE Global Communications Conference (GLOBECOM 2009)*, November 2009.

[SNC09] Mate Soos, Karsten Nohl, and Claude Castelluccia. Extending sat solvers to cryptographic problem. In *International Conference on Theory and Applications of Satisfiability Testing (SAT)*, jul 2009.

[Soo08] Mate Soos. Analysing the molva and di pietro private rfid authentication scheme. In *RFIDSec 00*, 2008.

[SSB09a] Emna Salhi, Mohamed Karim Sbai, and Chadi Barakat. Neighborhood selection in mobile p2p networks. In Augustin Chaintreau and Clemence Magnien, editors, *Algotel*, Carry-Le-Rouet, France, 2009.

[SSB09b] Karim Sbai, Emna Salhi, and Chadi Barakat. A membership management protocol for mobile p2p networks. In *proceedings of the ACM Mobility Conference*, Nice, September 2009.

[SSB10] Mohamed Karim Sbai, Emna Salhi, and Chadi Barakat. P2p content sharing in spontaneous multi-hop wireless networks. In *proceedings of the second International Conference on COMmunication Systems and NETworkS (COMSNETS)*, Bengalore, India, January 2010.

[TAD$^+$11] Cristian Tala, Luciano Ahumada, Diego Dujovne, Shafqat Ur-Rehman, Thierry Turletti, and Walid Dabbous. Guidelines for the accurate design of empirical studies in wireless networks. In *7th International ICST Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TridentCom)*, Shangai, China, April 2011.

[URTD11] Shafqat Ur-Rehman, Thierry Turletti, and Walid Dabbous. Multicast Video Streaming over WiFi Networks: Impact of Multipath

Fading and Interference. In *IEEE Workshop on multiMedia Applications over Wireless Networks (MediaWiN)*, Corfu, Greece, June 2011.

[VGBB11] Rosa Vilardi, Luigi Alfredo Grieco, Gennaro Boggia, and Chadi Barakat. Adaptation of real-time temporal resolution for bitrate estimates in ipfix systems. In *proceedings of the 2nd International Workshop on TRaffic Analysis and Classification (TRAC)*, Istanbul, July 2011.

[VRM$^+$09] Pascale Vicat-Blanc Primet, Vincent Roca, Johan Montagnat, Jean-Patrick Gelas, Olivier Mornard, Lionel Giraud, Guilherme Koslovski, and Tram Truong Huu. A scalable security model for enabling dynamic virtual private execution infrastructures on the internet. In *9th IEEE International Symposium on Cluster Computing and the Grid (CCGrid'09), Shanghai, China*, May 2009.

6.4 Standardization Documents

[AMR08a] H. Asaeda, K. Mishima, and V. Roca. *Requirements for IP Multicast Session Announcement in the Internet*, October 2008. IETF RMT Working Group (individual document), Work in Progress: <draft-ietf-mboned-session-announcement-req-00>.

[AMR08b] H. Asaeda, K. Mishima, and V. Roca. *Requirements for IP Multicast Session Announcement in the Internet*, July 2008. IETF RMT Working Group (individual document), Work in Progress: <draft-asaeda-mboned-session-announcement-req-00>.

[AR08a] B. Adamson and V. Roca. *Security and Reliable Multicast Transport Protocols: Discussions and Guidelines*, July 2008. IETF RMT Working Group, Work in Progress: <draft-ietf-rmt-sec-discussion-02.txt>.

[AR08b] B. Adamson and V. Roca. *Security and Reliable Multicast Transport Protocols: Discussions and Guidelines*, February 2008. IETF RMT Working Group, Work in Progress: <draft-ietf-rmt-sec-discussion-01.txt>.

[ARA08] B. Adamson, V. Roca, and H. Asaeda. *Security and Reliable Multicast Transport Protocols: Discussions and Guidelines*, November 2008. IETF RMT Working Group, Work in Progress: <draft-ietf-rmt-sec-discussion-03.txt>.

[ARA10] Brian Adamson, Vincent Roca, and Hitoshi Asaeda. *Security and Reliable Multicast Transport Protocols: Discussions and Guidelines*, May 2010. IETF RMT Working Group, Work in Progress: <draft-ietf-rmt-sec-discussion-05.txt>.

[GPR10a] Sarit Galanos, Orly Peck, and Vincent Roca. *RTP Payload Format for Reed Solomon FEC*, August 2010. IETF FECFRAME Working Group, Work in Progress: <draft-galanos-fecframe-rtp-reedsolomon-02>.

[GPR10b] Sarit Galanos, Orly Peck, and Vincent Roca. *RTP Payload Format for Reed Solomon FEC*, March 2010. IETF FECFRAME Working Group, Work in Progress: <draft-galanos-fecframe-rtp-reedsolomon-01>.

[LRPP09] J. Lacan, V. Roca, J. Peltotalo, and S. Peltotalo. *Reed-Solomon Forward Error Correction (FEC) Schemes*, April 2009. IETF Request for Comments, RFC 5510 (Standards Track/Proposed Standard).

[PWL$^+$08] T. Paila, R. Walsh, M. Luby, R. Lehtonen, and V. Roca. *FLUTE - File Delivery over Unidirectional Transport (revised)*, September 2008. IETF RMT Working Group, Work in Progress: <draft-ietf-rmt-flute-revised-06.txt>.

[PWL$^+$10a] Toni Paila, Rod Walsh, Michael Luby, Vincent Roca, and Rami Lehtonen. *FLUTE - File Delivery over Unidirectional Transport (revised)*, March 2010. IETF RMT Working Group, Work in Progress: <draft-ietf-rmt-flute-revised-11.txt>.

[PWL$^+$10b] Toni Paila, Rod Walsh, Michael Luby, Vincent Roca, and Rami Lehtonen. *FLUTE - File Delivery over Unidirectional Transport (revised)*, January 2010. IETF RMT Working Group, Work in Progress: <draft-ietf-rmt-flute-revised-10.txt>.

[PWL$^+$10c] Toni Paila, Rod Walsh, Michael Luby, Vincent Roca, and Rami Lehtonen. *FLUTE - File Delivery over Unidirectional Transport (revised)*, January 2010. IETF RMT Working Group, Work in Progress: <draft-ietf-rmt-flute-revised-09.txt>.

[PWL$^+$11] Toni Paila, Rod Walsh, Michael Luby, Vincent Roca, and Rami Lehtonen. *FLUTE - File Delivery over Unidirectional Transport (revised)*, February 2011. IETF RMT Working Group, Work in Progress: <draft-ietf-rmt-flute-revised-12.txt>.

[RA08a] V. Roca and B. Adamson. *FCAST: Scalable Object Delivery for the ALC and NORM Protocols*, September 2008. IETF RMT Working Group (individual document), Work in Progress: <draft-roca-rmt-newfcast-03.txt>.

[RA08b] V. Roca and B. Adamson. *FCAST: Scalable Object Delivery for the ALC and NORM Protocols*, July 2008. IETF RMT Working Group (individual document), Work in Progress: <draft-roca-rmt-newfcast-02.txt>.

[RA09a] V. Roca and B. Adamson. *FCAST: Scalable Object Delivery for the ALC and NORM Protocols*, October 2009. IETF RMT Working Group (individual document), Work in Progress: <draft-roca-rmt-newfcast-05.txt>.

[RA09b] V. Roca and B. Adamson. *FCAST: Scalable Object Delivery for the ALC and NORM Protocols*, July 2009. IETF RMT Working Group (individual document), Work in Progress: <draft-roca-rmt-newfcast-05.txt>.

[RA09c] V. Roca and B. Adamson. *FCAST: Scalable Object Delivery for the ALC and NORM Protocols*, March 2009. IETF RMT Working Group (individual document), Work in Progress: <draft-roca-rmt-newfcast-04.txt>.

[RA10a] Vincent Roca and Brian Adamson. *FCAST: Scalable Object Delivery for the ALC and NORM Protocols*, October 2010. IETF RMT Working Group (individual document), Work in Progress: <draft-ietf-rmt-newfcast-02.txt>.

[RA10b] Vincent Roca and Brian Adamson. *FCAST: Scalable Object Delivery for the ALC and NORM Protocols*, July 2010. IETF RMT Working Group (individual document), Work in Progress: <draft-ietf-rmt-newfcast-01.txt>.

[RA11a] Vincent Roca and Brian Adamson. *FCAST: Scalable Object Delivery for the ALC and NORM Protocols*, October 2011. IETF RMT Working Group (individual document), Work in Progress: <draft-ietf-rmt-newfcast-05.txt>.

[RA11b] Vincent Roca and Brian Adamson. *FCAST: Scalable Object Delivery for the ALC and NORM Protocols*, July 2011. IETF RMT Working Group (individual document), Work in Progress: <draft-ietf-rmt-newfcast-04.txt>.

[RA11c] Vincent Roca and Brian Adamson. *FCAST: Scalable Object Delivery for the ALC and NORM Protocols*, February 2011. IETF RMT Working Group (individual document), Work in Progress: <draft-ietf-rmt-newfcast-03.txt>.

[RCL09a] V. Roca, M. Cunche, and J. Lacan. *LDPC-Staircase Forward Error Correction (FEC) Schemes for FECFRAME*, July 2009. IETF FECFRAME Working Group, Work in Progress: <draft-roca-fecframe-ldpc-00.txt>.

[RCL$^+$09b] V. Roca, M. Cunche, J. Lacan, A. Bouabdallah, and K. Matsuzono. *Reed-Solomon Forward Error Correction (FEC) Schemes for FECFRAME*, July 2009. IETF FECFRAME Working Group, Work in Progress: <draft-roca-fecframe-rs-01.txt>.

[RCL$^+$09c] V. Roca, M. Cunche, J. Lacan, A. Bouabdallah, and K. Matsuzono. *Reed-Solomon Forward Error Correction (FEC) Schemes for FECFRAME*, March 2009. IETF FECFRAME Working Group, Work in Progress: <draft-roca-fecframe-rs-00.txt>.

[RCL10a] Vincent Roca, Mathieu Cunche, and Jerome Lacan. *LDPC-Staircase Forward Error Correction (FEC) Schemes for FECFRAME*, October 2010. IETF FECFRAME Working Group, Work in Progress: <draft-roca-fecframe-ldpc-01.txt>.

[RCL$^+$10b] Vincent Roca, Mathieu Cunche, Jerome Lacan, Amine Bouabdallah, and Kazuhisa Matsuzono. *Reed-Solomon Forward Error Correction (FEC) Schemes for FECFRAME*, March 2010. IETF FECFRAME Working Group, Work in Progress: <draft-roca-fecframe-rs-02.txt>.

[RCL$^+$10c] Vincent Roca, Mathieu Cunche, Jerome Lacan, Amine Bouabdallah, and Kazuhisa Matsuzono. *Simple Reed-Solomon Forward Error Correction (FEC) Schemes for FECFRAME*, July 2010. IETF FECFRAME Working Group, Work in Progress: <draft-roca-fecframe-rs-03.txt>.

[RCL$^+$11a] Vincent Roca, Mathieu Cunche, Jerome Lacan, Amine Bouabdallah, and Kazuhisa Matsuzono. *Simple Reed-Solomon Forward Error Correction (FEC) Scheme for FECFRAME*, November 2011. IETF FECFRAME Working Group, Work in Progress: <draft-ietf-fecframe-simple-rs-02>.

[RCL$^+$11b] Vincent Roca, Mathieu Cunche, Jerome Lacan, Amine Bouabdallah, and Kazuhisa Matsuzono. *Simple Reed-Solomon Forward Error Correction (FEC) Scheme for FECFRAME*, September 2011. IETF FECFRAME Working Group, Work in Progress: <draft-ietf-fecframe-simple-rs-01>.

[RCL$^+$11c] Vincent Roca, Mathieu Cunche, Jerome Lacan, Amine Bouabdallah, and Kazuhisa Matsuzono. *Simple Reed-Solomon Forward Error Correction (FEC) Scheme for FECFRAME*, February 2011. IETF FECFRAME Working Group, Work in Progress: <draft-ietf-fecframe-simple-rs-01>.

[RFF08a] V. Roca, A. Francillon, and S. Faurite. *TESLA source authentication in the ALC and NORM protocols*, December 2008. IETF RMT Working Group, Work in Progress: <draft-msec-tesla-for-alc-norm-07.txt>.

[RFF08b] V. Roca, A. Francillon, and S. Faurite. *TESLA source authentication in the ALC and NORM protocols*, October 2008. IETF RMT Working Group, Work in Progress: <draft-msec-tesla-for-alc-norm-06.txt>.

[RFF08c] V. Roca, A. Francillon, and S. Faurite. *TESLA source authentication in the ALC and NORM protocols*, July 2008. IETF RMT Working Group, Work in Progress: <draft-msec-tesla-for-alc-norm-05.txt>.

[RFF08d] V. Roca, A. Francillon, and S. Faurite. *TESLA source authentication in the ALC and NORM protocols*, February 2008. IETF RMT Working Group, Work in Progress: <draft-msec-tesla-for-alc-norm-04.txt>.

[RFF09a] V. Roca, A. Francillon, and S. Faurite. *TESLA source authentication in the ALC and NORM protocols*, October 2009. IETF RMT Working Group, Work in Progress: <draft-msec-tesla-for-alc-norm-10.txt>.

[RFF09b] V. Roca, A. Francillon, and S. Faurite. *TESLA source authentication in the ALC and NORM protocols*, October 2009. IETF RMT Working Group, Work in Progress: <draft-msec-tesla-for-alc-norm-09.txt>.

[RFF09c] V. Roca, A. Francillon, and S. Faurite. *TESLA source authentication in the ALC and NORM protocols*, September 2009. IETF RMT Working Group, Work in Progress: <draft-msec-tesla-for-alc-norm-08.txt>.

[RFF10] Vincent Roca, Aurélien Francillon, and Sébastien Faurite. *TESLA source authentication in the ALC and NORM protocols*, April 2010. IETF Request for Comments, RFC 5776 (Experimental).

[RNF08a] V. Roca, C. Neumann, and D. Furodet. *Low Density Parity Check (LDPC) Staircase and Triangle Forward Error Correction (FEC) Schemes*, June 2008. IETF Request for Comments, RFC 5170.

[RNF08b] V. Roca, C. Neumann, and D. Furodet. *Low Density Parity Check (LDPC) Staircase and Triangle Forward Error Correction (FEC) Schemes*, January 2008. IETF RMT Working Group, Work in Progress: <draft-ietf-rmt-fec-bb-ldpc-08.txt>.

[Roc08a] V. Roca. *FCAST: Scalable Object Delivery on top of the ALC Protocol*, February 2008. IETF RMT Working Group (individual document), Work in Progress: <draft-roca-rmt-newfcast-01.txt>.

[Roc08b] V. Roca. *Simple Authentication Schemes for the ALC and NORM Protocols*, October 2008. IETF RMT Working Group, Work in Progress: <draft-ietf-rmt-simple-auth-for-alc-norm-00.txt>.

[Roc09] V. Roca. *Simple Authentication Schemes for the ALC and NORM Protocols*, March 2009. IETF RMT Working Group, Work in Progress: <draft-ietf-rmt-simple-auth-for-alc-norm-01.txt>.

[Roc10] Vincent Roca. *Simple Authentication Schemes for the ALC and NORM Protocols*, July 2010. IETF RMT Working Group, Work in Progress: <draft-ietf-rmt-simple-auth-for-alc-norm-03.txt>.

[Roc11a] Vincent Roca. *Simple Authentication Schemes for the ALC and NORM Protocols*, December 2011. IETF RMT Working Group, Work in Progress: <draft-ietf-rmt-simple-auth-for-alc-norm-06.txt>.

[Roc11b] Vincent Roca. *Simple Authentication Schemes for the ALC and NORM Protocols*, September 2011. IETF RMT Working Group, Work in Progress: <draft-ietf-rmt-simple-auth-for-alc-norm-05.txt>.

[Roc11c] Vincent Roca. *Simple Authentication Schemes for the ALC and NORM Protocols*, July 2011. IETF RMT Working Group, Work in Progress: <draft-ietf-rmt-simple-auth-for-alc-norm-04.txt>.

[RRS11a] Vincent Roca, Aline Roumy, and Bessem Sayadi. *The Generalized Object Encoding (GOE) Approach for the Forward Erasure Correction (FEC) Protection of Objects and its Application to Reed- Solomon Codes over GF(2x)*, July 2011. IETF RMT Working Group, Work in Progress: <draft-roca-rmt-goe-fec-00.txt>.

[RRS11b] Vincent Roca, Aline Roumy, and Bessem Sayadi. *The Generalized Object Encoding (GOE) LDPC-Staircase FEC Scheme*, October 2011. IETF RMT Working Group, Work in Progress: <draft-roca-rmt-goe-ldpc-00.txt>.

[WBR11a] Mark Watson, Ali Begen, and Vincent Roca. *Forward Error Correction (FEC) Framework*, June 2011. IETF FECFRAME Working Group, Work in Progress: <draft-ietf-fecframe-framework-15.txt>.

[WBR11b] Mark Watson, Ali Begen, and Vincent Roca. *Forward Error Correction (FEC) Framework*, March 2011. IETF FECFRAME Working Group, Work in Progress: <draft-ietf-fecframe-framework-14.txt>.

[WBR11c] Mark Watson, Ali Begen, and Vincent Roca. *Forward Error Correction (FEC) Framework*, February 2011. IETF FECFRAME Working Group, Work in Progress: <draft-ietf-fecframe-framework-13.txt>.

[WBR11d] Mark Watson, Ali Begen, and Vincent Roca. *Forward Error Correction (FEC) Framework*, January 2011. IETF FECFRAME Working Group, Work in Progress: <draft-ietf-fecframe-framework-12.txt>.

[WBR11e] Mark Watson, Ali Begen, and Vincent Roca. *Forward Error Correction (FEC) Framework*, June 2011. IETF Request for Comments, RFC 6363 (Standards Track/Proposed Standard).

6.5 Technical and Internal Reports

[BRDTD09] Bilel Ben Romdhanne, Diego Dujovne, Thierry Turletti, and Walid Dabbous. Efficient and scalable merging algorithms for wireless traces. Rapport de recherche, INRIA, 2009. RR-6969.

[BVGI$^+$11] Stefan Bouckaert, Vanhie Van Gerwen, Moerman Ingrid, Stephen C. Phillips, Jerker Wilander, Shafqat Ur-Rehman, Thierry Turletti, and Walid Dabbous. "benchmarking computers and computer networks". Whitepaper, IST Fire projects, September 2011.

[Cas10] Claude Castelluccia. Internet et vie privée, des frères ennemis ? Pour La Science Magazine, Number 66, January 2010.

[CR08] M. Cunche and V. Roca. Improving the decoding of ldpc codes for the packet erasure channel with a hybrid zyablov iterative decoding/gaussian elimination scheme. Research report, INRIA, March 2008.

[CR09a] Mathieu Cunche and Vincent Roca. Adding integrity verification capabilities to the ldpc-staircase erasure correction codes. Technical report, INRIA, 2009.

[CR09b] Mathieu Cunche and Vincent Roca. Le rfc 5170 en pratique : conception et évaluation d'un codec al-fec ldpc-staircase hautes performances. Rapport de recherche, INRIA, 2009.

[Dab10] Walid Dabbous. L'architecture d'internet à l'ère du mouvement. Pour La Science Magazine, Number 66, January 2010.

[DMS$^+$08] R. Di Pietro, L. V. Mancini, C. Soriente, A. Spognardi, and G. Tsudik. Maximizing data survival in unattended wireless sensor networks against a focused mobile adversary. Cryptology ePrint Archive, Report 2008/293, 2008.

[DTD08] D. Dujovne, T. Turletti, and W. Dabbous. "experimental methodology for wireless networks". INRIA Research Report, RR-6667, October 2008.

[FT08] J. Farooq and T. Turletti. "an ieee 802.16 wimax module for the ns-3 simulator". Technical Report, inria-00336858, November 2008.

[KBS11] Amir Krifa, Chadi Barakat, and Thrasyvoulos Spyropoulos. MobiTrade: Interest driven content dissemination architecture for Disruption Tolerant Networks. Technical report, INRIA, February 2011.

[LBLD09] Stevens Le Blond, Arnaud Legout, and Walid Dabbous. Pushing bittorrent locality to the limit. Technical report, INRIA, 2009.

[LBLFLM09] Stevens Le Blond, Fabrice Le Fessant, and Erwan Le Merrer. Finding good partners in availability-aware p2p networks. Rapport de recherche, INRIA, 2009. RR-6795.

[LBLLD10] Stevens Le Blond, Arnaud Legout, Fabrice Lefessant, and Walid Dabbous. *Angling for Big Fish in BitTorrent*, January 2010.

[LFH$^+$09] Mathieu Lacage, Martin Ferrari, Mads Hansen, Thierry Turletti, and Walid Dabbous. Nepi: Using independent simulators, emulators, and testbeds for easy experimentation. Rapport de recherche, INRIA, 2009. RR-6967.

[LLT$^+$09] Jerome Lacan, Emmanuel Lochin, Pierre-Ugo Tournoux, Amine Bouabdallah, and Vincent Roca. On-the-fly coding for time-constrained applications. Research report, Computing Research Repository, April 2009.

[LLT$^+$10a] Jerome Lacan, Emmanuel Lochin, Pierre-Ugo Tournoux, Amine Bouabdallah, and Vincent Roca. On-the-fly coding for time-constrained applications, September 2010.

[LLT$^+$10b] Jerome Lacan, Emmanuel Lochin, Pierre-Ugo Tournoux, Amine Bouabdallah, and Vincent Roca. On-the-fly coding for time-constrained applications, February 2010.

[LMC$^+$09] Stevens Le Blond, Pere Manils, Abdelberi Chaabane, Mohamed Ali Kaafar, Arnaud Legout, and Claude Castelluccia. De-anonymizing bittorrent users on tor. Technical report, INRIA, December 2009.

[MCLB$^+$10] Pere Manils, Abdelberi Chaabane, Stevens Le Blond, Mohamed Ali Kaafar, Arnaud Legout, Claude Castelluccia, and Walid Dabbous. *Compromising Tor Anonymity Exploiting P2P Information Leakage*, April 2010.

[MDD09] Juan-Carlos Maureira, Diego Dujovne, and Olivier Dalle. Network provisioning for high speed vehicles moving along predictable routes - part 1: Spiderman handover. Rapport de recherche, INRIA, 2009. RR-6850.

[RL09] Ashwin Rao and Arnaud Legout. Impact of network latencies on the outcome of bittorrent experiments performed on testbeds. Technical report, INRIA, December 2009.

[RRSI11] Aline Roumy, Vincent Roca, Bessem Sayadi, and Rodrigue Imad. Unequal Erasure Protection and Object Bundle Protection with the Generalized Object Encoding Approach. Research Report RR-7699, INRIA, July 2011. submitted to Infocom 2012.

[RTD11] Shafqat Rehman, Thierry Turletti, and Walid Dabbous. A Roadmap for Benchmarking in Wireless Networks. Research Report RR-7707, INRIA, August 2011.

[SB09] Mohamed Karim Sbai and Chadi Barakat. Revisiting content sharing in wireless ad hoc networks. Rapport de recherche, INRIA, 2009.

[SCLR09] Alexandre Soro, Mathieu Cunche, Jerome Lacan, and Vincent Roca. Erasure codes with a banded structure for hybrid iterative-ml decoding. Research report, Computing Research Repository, January 2009.

[SSB08] Karim Sbai, Emna Salhi, and Chadi Barakat. Adaptive overlay for P2P membership management in MANET. Technical Report inria-00342691, INRIA, November 2008.

[URTD10] Shafqat Ur-Rehman, Thierry Turletti, and Walid Dabbous. *"Benchmarking of Wireless Experimentations"*, October 2010.

6.6 Reference Publications of the Planète project-team

[CDCP10] Claude Castelluccia, Emiliano De Cristofaro, and Daniele Perito. Private information disclosure from web searches. In *Proceedings of the 2010 Privacy Enhancing Technologies Symposium (PETS)*, 2010.

[CFSP09] Claude Castelluccia, Aurélien Francillon, Claudio Soriente, and Daniele Perito. On the difficulty of software-based attestation of embedded devices. In *CCS '09: Proceedings of the 16th ACM conference on Computer and communications security*, 2009.

[CKMP09] Claude Castelluccia, Mohamed Ali Kaafar, Pere Manils, and Daniele Perito. Geolocalization of proxied services and its application to fast-flux hidden servers. In *ACM/Usenix Internet Measurement Conference (IMC 2009)*, Chicago, USA, November 2009. ACM.

[CSR10] Mathieu Cunche, Valentin Savin, and Vincent Roca. Analysis of quasi-cyclic ldpc codes under ml decoding over the erasure channel. In *IEEE International Symposium on Information Theory and its Applications (ISITA'10) (http://arxiv.org/abs/1004.5217)*, April 2010.

[LBA11] Imed Lassoued, Chadi Barakat, and Konstantin Avrachenkov. Network-wide monitoring through self-configuring adaptive system. In *proceedings of IEEE INFOCOM*, Shanghai, China, April 2011.

[LBZL+11] Stevens Le Blond, Chao Zhang, Arnaud Legout, Keith Ross, and Walid Dabbous. I Know Where You are and What You are Sharing: Exploiting P2P Communications to Invade Users' Privacy. In *proceedings of ACM SIGCOM/USENIX IMC'11*, Berlin, Germany, November 2011.

[LNM+09] Tianji Li, Qiang Ni, David Malone, Douglas Leight, Yang Xiao, and Thierry Turletti. Aggregation with fragment retransmission for very high-speed wlans. *IEEE/ACM Transactions on Networking Journal*, 17(2), 2009.

[PCKM11] Daniele Perito, Claude Castelluccia, Mohamed Ali Kaafar, and Pere Manils. How Unique and Traceable are Usernames. In *proceedings of the 11th Privacy Enhancing Technologies Symposium (PETS)*, Waterloo, July 2011.

[RCHBC09] Kasper B. Rasmussen, Claude Castelluccia, Thomas Heydt-Benjamin, and Srdjan Capkun. Proximity-based access control for implantable medical devices. In *CCS '09: Proceedings of the 16th ACM conference on Computer and communications security*, 2009.

[RLB$^+$11] Ashwin Rao, Yeon-Sup Lim, Chadi Barakat, Arnaud Legout, Don Towsley, and Walid Dabbous. Network Characteristics of Video Streaming Traffic. In *proceedings of ACM CoNEXT'11*, Tokyo, Japan, December 2011.

[RNF08] Vincent Roca, Christoph Neumann, and David Furodet. *Low Density Parity Check (LDPC) Staircase and Triangle Forward Error Correction (FEC) Schemes*, June 2008. IETF Request for Comments, RFC 5170 (Standards Track/Proposed Standard).

[SB09] Karim Sbai and Chadi Barakat. Experiences on enhancing data collection in large networks. *Computer Networks*, 53(7):1073–1086, May 2009.

[SBRT$^+$10] Thrasyvoulos Spyropoulos, Rao Naveed Bin Rais, Thierry Turletti, Katia Obraczka, and Athanasios Vasilakos. Routing for disruption tolerant networks: Taxonomy and design. *ACM/Springer Wireless Networks (WINET)*, 16(8), 2010.

[STO09] Thrasyvoulos Spyropoulos, Thierry Turletti, and Katia Obraczka. Routing in delay tolerant networks comprising heterogeneous node populations. *IEEE Transactions on Mobile Computing (TMC)*, 8(8), 2009.

[WBR11] Mark Watson, Ali Begen, and Vincent Roca. *Forward Error Correction (FEC) Framework*, June 2011. IETF Request for Comments, RFC 6363 (Standards Track/Proposed Standard).