Activity Report 2017

Project-Team DIANA

Design, Implementation and Analysis of Networking Architectures

Table of contents

Project-Team DIANA

Creation of the Team: 2013 January 01, updated into Project-Team: 2015 July 01

Keywords:

Computer Science and Digital Science:

       A1.1.7. - Peer to peer
       A1.1.13. - Virtualization
       A1.2.1. - Dynamic reconfiguration
       A1.2.2. - Supervision
       A1.2.3. - Routing
       A1.2.4. - QoS, performance evaluation
       A1.2.8. - Network security
       A1.2.9. - Social Networks
       A1.3. - Distributed Systems
       A1.4. - Ubiquitous Systems

Other Research Topics and Application Domains:

       B6.2. - Network technologies
       B6.2.1. - Wired technologies
       B6.2.2. - Radio technology
       B6.2.3. - Satellite technology
       B6.3.2. - Network protocols
       B6.3.3. - Network Management
       B6.3.4. - Social Networks
       B8.5.2. - Crowd sourcing
       B9.1.1. - E-learning, MOOC
       B9.4.1. - Computer science
       B9.4.5. - Data science
       B9.6. - Reproducibility
       B9.8. - Privacy

1. Personnel

Research Scientists
    Walid Dabbous [Team leader, Inria, Senior Researcher, HDR]
    Chadi Barakat [Inria, Researcher, HDR]
    Arnaud Legout [Inria, Researcher, HDR]
    Damien Saucez [Inria, Researcher]
    Thierry Turletti [Inria, Senior Researcher, HDR]

Post-Doctoral Fellows
    Osama Arouk [Inria]
    Natasa Sarafijanovic-Djukic [Inria, until Nov 2017]

PhD Students
    Yanis Boussad [Université de Nice - Sophia Antipolis, from Oct 2017]

Thibaut Ehlinger [Inria, from Nov 2017]
Karyna Gogunska [Inria]
Muhammad Jawad Khokhar [Inria]
Mohamed Naoufal Mahfoudi [Inria]
Ghada Moualla [Inria]
Vitalii Poliakov [Université de Nice - Sophia Antipolis]
Hardik Soni [Inria]
Mathieu Thiery [Inria, from April 2017]
Imane Taibi [Inria, from Nov 2017]
Thibaud Trolliet [Inria, from Oct 2017]
Luigi Vigneri [Eurecom, until Jul 2017]

Technical staff
Abdelhakim Akodadi [Inria, from May 2017]
Thierry Parmentelat [Inria]
Mondi Ravi [Inria]
Thierry Spetebroot [Inria]
Mario Antonio Zancanaro [Inria, until Mar 2017]

Interns
Neha Agarwal [Inria, from Apr 2017 until Sep 2017]
Yanis Boussad [Inria, from Mar 2017 until Aug 2017]
Pretesh Chauhan [Inria, from May 2017 until Jul 2017]
Giuseppe Di Lena [Inria, from Mar 2017 until Aug 2017]
David Migliacci [Inria, from Jul 2017]
Yassir Mrabet [Inria, from Mar 2017 until Aug 2017]
Imane Taibi [Inria, from Mar 2017 until Aug 2017]

Administrative Assistant
Christine Foggia [Inria]

Visiting Scientists
Salma Ait Lhaj [Université Côte d'Azur, from Nov 2017]
Boubaker Idir [Inria, from Nov 2017]
Claudio Laterza [Université Côte d'Azur, from Nov 2017]
Nauroz Khan Nausherwani [Université Côte d'Azur, from Nov 2017]
Katia Obraczka [University of California Santa Cruz, Jul 2017]
Gayatri Sivadoss [Université Côte d'Azur, from Nov 2017]
Alberto Zirondelli [Université Côte d'Azur, from Nov 2017]

2. Overall Objectives

2.1. Presentation of the team

The overall objective of the DIANA project-team is to provide network architectural support for improving citizen rights in the Internet. To do so, we work to provide service transparency and user data control in the context of hundreds of billions of both wired and mobile devices. Our methodology includes advanced measurement techniques, design and implementation of architectural solutions, and their validation in adequate experimental facilities.

The high complexity of the Internet architecture, protocols and services, and the economic interests of the big stakeholders result in a lack of transparency concerning information of high interest to the connected "citizen" such as possible privacy leaks, root cause of service degradation or lock-in behavior. It is therefore important to enhance the network to provide service transparency to citizens.

On the other hand, the ossification of the Internet architecture around the IP protocol makes introduction of new functionalities in the network quite difficult. Users currently have no control on their contents and depend on big companies (e.g., Google drive, iCloud, dropbox, Microsoft OneDrive) to easily access and share data at the expense of their privacy. However, the recent development of software-defined network and network functions virtualization concepts open the perspective of faster deployment of network functionalities, as it abstracts the whole network as a single piece of software, instead of a large number of heterogeneous and dedicated devices to be configured one-by-one.

In the DIANA project-team, we have two main research directions:

- designing and deploying a measurement plane providing network service transparency,
- defining and deploying an open network architecture for user control.

Our research program is presented briefly in the next section.

3. Research Program

3.1. Service Transparency

Transparency is to provide network users and application developers with reliable information about the current or predicted quality of their communication services, and about potential leakages of personal information, or of other information related to societal interests of the user as a "connected citizen" (e.g. possible violation of network neutrality, opinion manipulation). Service transparency therefore means to provide information meaningful to users and application developers, such as quality of experience, privacy leakages, or opinion manipulation, etc. rather than network-level metrics such as available bandwidth, loss rate, delay or jitter.

The Internet is built around a best effort routing service that does not provide any guarantee to end users in terms of quality of service (QoS). The simplicity of the Internet routing service is at the root of its huge success. Unfortunately, a simple service means unpredicted quality at the access. Even though a considerable effort is done by operators and content providers to optimise the Internet content delivery chain, mainly by over-provisioning and sophisticated engineering techniques, service degradation is still part of the Internet. The proliferation of wireless and mobile access technologies, and the versatile nature of Internet traffic, make end users quality of experience (QoE) forecast even harder. As a matter of fact, the Internet is missing a dedicated measurement plane that informs the end users on the quality they obtain and in case of substantial service degradation, on the origin of this degradation. Current state of the art activities are devoted to building a distributed measurement infrastructure to perform active, passive and hybrid measurements in the wired Internet. However, the problem is exacerbated with modern terminals such as smartphones or tablets that do not facilitate the task for end users (they even make it harder) as they focus on simplifying the interface and limiting the control on the network, whereas the Internet behind is still the same in terms of the quality it provides. Interestingly, this same observation explains the existing difficulty to detect and prevent privacy leaks. We argue that the lack of transparency for diagnosing QoE and for detecting privacy leaks have the same root causes and can be solved using common primitives. For instance, in both cases, it is important to be able to link data packets to an application. Indeed, as the network can only access data packets, there must be a way to bind these packets to an application (to understand users QoE for this application or to associate a privacy leak to an application). This is however a complex task as the traffic might be obfuscated or encrypted. Our objectives in the research direction are the following:

- Design and develop measurement tools providing transparency, in spite of current complexity
- Deploy those measurement tools at the Internet's edge and make them useful for end users
- Propose measurements plane as an overlay or by exploiting in-network functionalities
- Adapt measurements techniques to network architectural change
- Provide measurements as native functionality in future network architecture

3.2. Open network architecture

We are surrounded by personal content of all types: photos, videos, documents, etc. The volume of such content is increasing at a fast rate, and at the same time, the spread of such content among all our connected devices (mobiles, storage devices, set-top boxes, etc) is also increasing. All this complicates the control of personal content by the user both in terms of access and sharing with other users. The access of the personal content in a seamless way independently of its location is a key challenge for the future of networks. Proprietary solutions exist, but apart from fully depending on one of them, there is no standard plane in the Internet for a seamless access to personal content. Therefore, providing network architectural support to design and develop content access and sharing mechanisms is crucial to allow users control their own data over heterogeneous underlying network or cloud services.

On the other hand, privacy is a growing concern for states, administrations, and companies. Indeed, for instance the French CNIL (entity in charge of citizens privacy in computer systems) puts privacy at the core of its activities by defining rules on any stored and collected private data. Also, companies start to use privacy preserving solutions as a competitive advantage. Therefore, understanding privacy leaks and preventing them is a problem that can already find support. However, all end-users do not *currently* put privacy as their first concern. Indeed, in face of two services with one of higher quality, they usually prefer the highest quality one whatever the privacy implication. This was, for instance, the case concerning the Web search service of Google that is more accurate but less privacy preserving than Bing or Qwant. This is also the case for cloud services such as iCloud or Dropbox that are much more convenient than open source solutions, but very bad in terms of privacy. Therefore, to reach end-users, any privacy preserving solutions must offer a service equivalent to the best existing services.

We consider that it will be highly desirable for Internet users to be able to *easily* move their content from a provider to another and therefore not to depend on a content provider or a social network monopoly. This requires that the network provides built-in architectural support for content networking.

In this research direction, we will define a new *service abstraction layer* (SAL) that could become the new waist of the network architecture with network functionalities below (IP, SDN, cloud) and applications on top. SAL will define different services that are of use to all Internet users for accessing and sharing data (seamless content localisation and retrieval, privacy leakage protection, transparent vertical and horizontal handover, etc.). The biggest challenge here is to cope in the same time with large number of content applications requirements and high underlying networks heterogeneity while still providing efficient applications performance. This requires careful definition of the services primitives and the parameters to be exchanged through the service abstraction layer.

Two concurring factors make the concept behind SAL feasible and relevant today. First, the notion of scalable network virtualization that is a required feature to deploy SAL in real networks today has been discussed recently only. Second, the need for new services abstraction is recent. Indeed, more than fifteen years ago the Internet for the end-users was mostly the Web. Only ten years ago smartphones came into the picture of the Internet boosting the number of applications with new functionalities and risks. Since a few years, many discussions in the network communities took place around the actual complexity of the Internet and the difficulty to develop applications. Many different approaches have been discussed (such as CCN, SDN) that intend to solve only part of the complexity. SAL takes a broader architectural look at the problem and considers solutions such as CCN as mere use cases. Our objectives in this research direction include the following:

- Identify common key networking services required for content access and sharing
- Detect and prevent privacy leaks for content communication
- Enhance software defined networks for large scale heterogeneous environments
- Design and develop open Content Networking architecture
- Define a service abstraction layer as the thin waist for the future content network architecture
- Test and deploy different applications using SAL primitives on heterogeneous network technologies

3.3. Methodology

We follow an experimental approach that can be described in the following techniques:

- Measurements: the aim is to get a better view of a problem in quantifiable terms. Depending on the field of interest, this may involve large scale distributed systems crawling tools; active probing techniques to infer the status and properties of a complex and non controllable system as the Internet; or even crowdsourcing-based deployments for gathering data on real-users environments or behaviours.

- Experimental evaluation: once a new idea has been designed and implemented, it is of course very desirable to assess and quantify how effective it can be, before being able to deploy it on any realistic scale. This is why a wide range of techniques can be considered for getting early, yet as significant as possible, feedback on a given paradigm or implementation. The spectrum for such techniques span from simulations to real deployments in protected and/or controlled environments.

4. Highlights of the Year

4.1. Highlights of the Year

4.1.1. *Reproducibility'17 workshop*

Our team was strongly involved in the Reproducibility'17@SIGCOMM workshop. Damien Saucez served as a co-chair of the workshop. Chadi Barakat and Mohamed Naoufal Mahfoudi participated to the workshop discussions. Mohamed Naoufal also presented our paper Lessons Learned while Trying to Reproduce the OpenRF Experiment [21]. See section 6.3.1 for more details about the workshop results.

4.1.2. *R2lab demonstration at SIGCOMM*

We have demonstrated the deployment of a standalone 5G network in less than 5 minutes in the R2lab tesbed. All the network components (base station, subscriber management, serving and packet gateways, network trafic analyzers) were run automatically using the nepi-ng experiment orchestration tool. Download and upload performance to the Internet from a commercial phone located in the anechoic chamber was also performed. This demo has been presented at the ACM SIGCOMM conference in August 2017 [33].

4.1.3. *MOOC Python 3*

Arnaud Legout and Thierry Parmentelat are co-authors of the MOOC: "Python 3 : des fondamentaux aux concepts avancés du langage" that lasts 9 weeks on FUN, UCA. For the first session there were 11677 registered persons. This MOOC is a brand new version of the past MOOC on Python 2, and has been funded by UCA.

5. New Software and Platforms

5.1. ACQUAmobile

KEYWORDS: Android - Internet access - Performance measure - Quality of Experience
FUNCTIONAL DESCRIPTION: ACQUA is an Application for prediCting QUality of Experience (QoE) at Internet Access. It is developed by the Diana team at Inria Sophia Antipolis – Méditerranée and is supported by Inria under the ADT ACQUA grant. The scientific project around ACQUA is supported by Inria Project Lab BetterNet and the French National Project ANR BottleNet. ACQUA presents a new way for the evaluation of the performance of Internet access. Starting from network-level measurements as the ones we often do today (bandwidth, delay, loss rates, jitter, etc), ACQUA targets the estimated Quality of Experience (QoE) related to the different applications of interest to the user without the need to run them (e.g., estimated Skype quality, estimated video streaming quality).

An application in ACQUA is a function, or a model, that links the network-level and device-level measurements to the expected Quality of Experience. Supervised machine learning techniques are used to establish such link between measurements both at the network level and the device level, and estimations of the Quality of Experience for different Internet applications. The required data for such learning can be obtained either by controlled experiments as we did in [19] on YouTube Quality of Experience, or by soliciting the crowd (i.e. crowdsourcing) for combinations (i.e. tuples) of measurements and corresponding application-level Quality of Experience. Our current work is concentrating on using the ACQUA principle in the estimation and prediction of the Quality of Experience for main user's applications. We refer to the web site of the project for further details.

The ACQUA Android application is supposed to be on one hand the reference application for QoE forecasting and troubleshooting for end users at their Internet access, and on the other hand, the feedback channel that allows end users to report to us (if they are willing) on their experience together with the corresponding network measurements so as to help us calibrating better and more realistic models. For this calibration, we are currently performing extensive, efficient and automatic measurements in the laboratory, we will count on end users to help us completing this dataset with further applications and more realistic network and user conditions.

ACQUA is mainly meant for end users, but it is also of interest to (mobile) network operators and to content providers to estimate the QoE of their customers and their networks without each time having to run expensive application-level traffic and to involve real users.

- Authors: Thierry Spetebroot and Chadi Barakat
- Contact: Chadi Barakat
- URL: http://project.inria.fr/acqua/

5.2. ElectroSmart

KEYWORDS: Crowd-sourcing - UMTS - GSM - Bluetooth - Wi-Fi - 4G - 3G - 2G - Electromagnetic waves - Android - LTE

FUNCTIONAL DESCRIPTION: The Internet and new devices such as smartphones have fundamentally changed the way people communicate, but this technological revolution comes at the price of a higher exposition of the general population to microwave electromagnetic fields (EMF). This exposition is a concern for health agencies and epidemiologists who want to understand the impact of such an exposition on health, for the general public who wants a higher transparency on its exposition and the health hazard it might represent, but also for cellular operators and regulation authorities who want to improve the cellular coverage while limiting the exposition, and for computer scientists who want to better understand the network connectivity in order to optimize communication protocols. Despite the fundamental importance to understand the exposition of the general public to EMF, it is poorly understood because of the formidable difficulty to measure, model, and analyze this exposition.

The goal of the ElectroSmart project is to develop the instrument, methods, and models to compute the exposition of the general public to microwave electromagnetic fields used by wireless protocols and infrastructures such as Wi-Fi, Bluetooth, or cellular. Using a pluri-disciplinary approach combining crowd-based measurements, in-lab experiments, and modeling using sparse and noisy data, we address challenges such as designing and implementing a measuring instrument leveraging on crowd-based measurements from mobile devices such as smartphones, modeling the exposition of the general public to EMF to compute the most accurate estimation of the exposition, and analyzing the evolution of the exposition to EMF with time. This technological breakthrough will have scientific, technical, and societal applications, notably on public health politics, by providing the scientific community and potential users with a unique measuring instrument, methods, and models to exploit the invaluable data gathered by the instrument.

In August 2016, we released the first stable public release of ElectroSmart. On the 22th December 2017 we have 35 836 downloads in Google Play, an average score of 4,6/5, 10 538 active users, 350 millions measured signals.

We are in a process of creating a startup to commercialize the exposition maps we can build with the data we are collecting.

- Participants: Arnaud Legout, Hackob Melconian, Inderjeet Singh and Mondi Ravi
- Contact: Arnaud Legout
- URL: https://es.inria.fr/home/index?path_prefix=en

5.3. OpenLISP

KEYWORDS: LISP - Routing - Control-plane

FUNCTIONAL DESCRIPTION: Among many options tackling the scalability issues of the current Internet routing architecture, the Locator/Identifier Separation Protocol (LISP) appears as a viable solution. LISP improves a network's scalability, flexibility, and traffic engineering, enabling mobility with limited overhead. As for any new technology, implementation and deployment are essential to gather and master the real benefits that it provides. We propose a complete open source implementation of the LISP control plane. Our implementation is deployed in the worldwide LISP Beta Network and the French LISP-Lab testbed, and includes the key standardized control plane features. Our control plane software is the companion of the existing OpenLISP dataplane implementation, allowing the deployment of a fully functional open source LISP network compatible with any implementation respecting the standards.

- Contact: Damien Saucez
- URL: http://www.openlisp.org/downloads

5.4. nepi-ng

KEYWORDS: Wireless network - Experimentation

FUNCTIONAL DESCRIPTION: In the specific context of R2lab, we have created a tool suite for orchestrating network experiments, that for historical reasons we refer to collectively as `nepi-ng`, for NEPI new generation. An umbrella website is available at https://nepi-ng.inria.fr/.

At this point, nepi-ng has a much smaller scope than its NEPI ancestor used to have, in that it only supports remote control of network experiments over `ssh`. As a matter of fact, in practice, this is the only access mechanism that we need to have for running experiments on both R2lab, and PlanetLab Europe.

The design of `nepi-ng` of course is modular, so that it will be perfectly possible to add other control mechanisms to this core if and when it becomes necessary.

`nepi-ng` is currently made of 2 separate python libraries :

- asynciojobs:
    - URL: http://asynciojobs.readthedocs.io/en/latest/
    - Version: asynciojobs v0.5.4
    - Keywords: networking experimentation, orchestration
    - License: CC BY-SA 4.0
    - Type of human computer interaction: python library
    - OS/Middleware: Linux
    - Required library or software: python-3.5 / asyncio
    - Programming language: python3
- apssh:
    - URL: http://apssh.readthedocs.io/en/latest/

      – Version: apssh v0.7.1

      – Keywords: networking experimentation, orchestration

      – License: CC BY-SA 4.0

      – Type of human computer interaction: python library

      – OS/Middleware: Linux

      – Required library or software: python-3.5 / asyncio

      – Programming language: python3

- Contact: Thierry Parmentelat
- URL: http://nepi-ng.inria.fr

5.5. Platforms

5.5.1. *Reproducible research laboratory (*R^2lab*)*

Scientific evaluation of network protocols requires for experiments to be reproducible before they can be deemed valid. This is particularly difficult to obtain in the wireless networking area, where characteristics of wireless channels are known to be variable, unpredictable and hardly controllable.

The R^2lab wireless testbed is built around an isolated and anechoic chamber, featuring RF absorbers preventing radio waves reflections and a Faraday cage blocking external interferences. This lab, named R^2lab, represents an ideal environment for experiments reproducibility.

It represents a perfect facility for making wireless experiments reproducible. It has been operated for 3 years now, in the context of the FIT Equipment of Excellence project, and as such, it is now federated with the other testbeds that are part of the FIT initiative. This testbed is for the long-haul, and is scheduled to remain operational until at least 2020.

During 2017, our focus regarding R^2lab has been set on deploying more, and more diverse USRPs (Universal Software Radio Peripherals). The chamber now offers more Ettus-based devices, as well as a couple of lime-sdr devices, and a couple of E3372 LTE dongles. An apple iphone will soon be available, in addition to the first nexus phone. All these additions aim at widening even further the spectrum of experiments that the testbed can support.

For more details see http://r2lab.inria.fr.

6. New Results

6.1. Service Transparency

6.1.1. *On active sampling of controlled experiments for QoE modeling*

Participants: Muhammad Jawad Khokhar, Nawfal Abbasi Saber, Thierry Spetebroot, Chadi Barakat.
For internet applications, measuring, modeling and predicting the quality experienced by end users as a function of network conditions is challenging. A common approach for building application specific Quality of Experience (QoE) models is to rely on controlled experimentation. For accurate QoE modeling, this approach can result in a large number of experiments to carry out because of the multiplicity of the network features, their large span (e.g., band-width, delay) and the time needed to setup the experiments themselves. However, most often, the space of network features in which experimentations are carried out shows a high degree of uniformity in the training labels of QoE. This uniformity, difficult to predict beforehand, amplifies the training cost with little or no improvement in QoE modeling accuracy. So, in this work, we aim to exploit this uniformity, and propose a methodology based on active learning, to sample the experimental space intelligently, so that the training cost of experimentation is reduced. We prove the feasibility of our methodology by validating it over a particular case of YouTube streaming, where QoE is modeled both in terms of interruptions and stalling duration. This first validation has appeared in [19]. In another paper which is currently under submission, we propose an online version of this methodology together with a set of criterion to stop the experiments when the learner is confident enough.

6.1.2. *On the Cost of Measuring Traffic in a Virtualized Environment*

Participants: Karyna Gogunska, Chadi Barakat, Guillaume Urvoy-Keller, Dino Lopez Pacheco.
The current trend in application development and deployment is to package applications and services within containers or virtual machines. This results in a blend of virtual and physical resources with complex interconnection network schemas mixing virtual and physical switches along with specific protocols to build virtual networks spanning over several servers. While the complexity of this set-up is hidden by private/public cloud management solutions, e.g. OpenStack, this constitutes a challenge when it comes to monitor and debug performance related issues. In this work, carried out in collaboration with the Signet team of I3S with the support of the UCN@SOPHIA Labex, we introduce the problem of measuring traffic in a virtualized environment and focus on one typical scenario, namely virtual servers interconnected with a virtual switch. For this scenario, we assess the cost of continuously measuring the network traffic activity of the machines. Specifically, we seek to estimate the competition that exists to access the physical resources (CPU, memory, etc.) of the physical substrate between the measurement task and the legacy application activity. The results of this first study are currently under submission.

6.1.3. *LISP measurements*

Participant: Damien Saucez.
The Locator/Identifier Separation Protocol (LISP) separates classical IP addresses into two categories: one for identifying terminals, the other for routing. To associate identifiers and locators LISP needs a specific mechanism, called mapping system. This technology is still at an early stage but two experimental platforms have already been deployed in the Internet: LISP Beta Network and LISP-Lab. However, only the LISP Beta Network is monitored with LISPmon that partially monitors the mapping system once a day. To accompany the growth of LISP, a dynamic and complete monitoring system is required. Therefore, we propose LISP-Views, a dynamic versatile large scale LISP monitoring architecture. LISP-Views allows to automatically conduct comprehensive and objective measurements. After running LISP-Views in the wild for several months and comparing the monitoring results with LISPmon, we confirm that LISP-Views provides more detailed and accurate information. We observe the different behaviours between every network entity within mapping system, and also explore the current LISP performance for further improvements. A paper on "LISP-Views Monitoring LISP at Large Scale" was published in ITC this year.

6.2. Open Network Architecture

6.2.1. *Controller load in SDN networks*

Participant: Damien Saucez.
In OpenFlow, a centralized programmable controller installs forwarding rules into switches to implement policies. However, this flexibility comes at the expense of extra overhead in signalling and number of rules to install. The community considered that it was essential to install all rules and strictly respect routing requirements, hence working on making extra fast and large memory switches and controllers. Instead we took an opposite direction and came with a new vision that leverages the SDN concept and considers the network as a black box where tailored rules should be used only for network traffic that really matters while for the rest a good-enough (sub-optimal but cheap) default behaviour should be enough. In the past, we applied this vision to limit the needed memory on network switches in [5]. Lately, we proposed solutions to limit the number of exchanged messages between the switches and the controller. More precisely, in [31], [16] we developed a distributed sampling adaptive algorithm that allows switches to locally decide if they can contact the controller or if instead they should make their own decision locally. Numerical evaluation and emulation in Mininet demonstrate the benefit of the approach. The results were published in the PGMO (Gaspard Monge Program for Optimisation) days, Nov 2017, Paris, France.

6.2.2. *Traceroute facility for Content-Centric Network*

Participant: Thierry Turletti.

In the context of the UHD-on-5G associated team with our colleagues at NICT, Japan, we have proposed the Contrace tool for Measuring and Tracing Content-Centric Networks (CCNs). CCNs are fundamental evolutionary technologies that promise to form the cornerstone of the future Internet. The information flow in these networks is based on named data requesting, in-network caching, and forwarding – which are unique and can be independent of IP routing. As a result, common IP-based network tools such as ping and traceroute can neither trace a forwarding path in CCNs nor feasibly evaluate CCN performance. We designed Contrace, a network tool for CCNs (particularly, CCNx implementation running on top of IP) that can be used to investigate 1) the Round-Trip Time (RTT) between content forwarder and consumer, 2) the states of in-network cache per name prefix, and 3) the forwarding path information per name prefix. This tool can estimate the content popularity and design more effective cache control mechanisms in experimental networks. We have published an Internet-Draft [30] describing the specification of Contrace.

6.2.3. *Message Dissemination in Intelligent Transport Systems*

Participant: Thierry Turletti.

We proposed D2-ITS, a flexible and extensible framework to dynamically distribute network control to enable message dissemination in Intelligent Transport Systems (ITS). By decoupling the control from the data plane, D2-ITS leverages network programmability to address ITS scalability, delay intolerance and decentralization. It uses a distributed control plane based on a hierarchy of controllers that can dynamically adjust to environment and network conditions in order to satisfy ITS application requirements. We demonstrate the benefits of D2-ITS through a proof-of-concept prototype using the ns-3 simulation platform. Results indicate lower message delivery latency with minimal additional overhead. This work has been presented at the IEEE/ACM Symposium on Distributed Simulation and Real Time Applications (DS-RT) in October 2017 [18].

6.2.4. *Peer-assisted Information-Centric Network*

Participant: Thierry Turletti.

Information-Centric Networking (ICN) is a promising solution for most of Internet applications where the content represents the core of the application. However, the proposed solutions for the ICN architecture are associated with many complexities including pervasive caching in the Internet and incompatibility with legacy IP networks, so the deployment of ICN in real networks is still an open problem. In this work, we proposed a backward compatible ICN architecture to address the caching issue in particular. The key idea is implementing edge caching in ICN, using a coalition of end clients and edge servers. Our solution can be deployed in IP networks with HTTP requests. We performed a trace-driven simulation for analyzing PICN benefits using IRCache and Berkeley trace files. The results showed that in average, PICN decreases the latency for 78% and increases the content retrieval speed for 69% compared to a direct download from the original web servers. When comparing PICN with a solution based on central proxy servers, we showed that the hit ratio obtained using a small cache size in each PICN client is almost 14% higher than the hit ratio obtained with a central proxy server using an unlimited cache storage. This work has been published in the IEEE Access journal [14].

6.2.5. *Streaming using In-Network Coding and Caching*

Participant: Thierry Turletti.

With the rapid growth in high-quality video streaming over the Internet, preserving high-level robustness against data loss and low latency, while maintaining higher data transmission rates, is becoming an increasingly important issue for high-quality real-time delay-sensitive streaming. We have proposed a low latency, low loss streaming mechanism, L4C2, specialized for high-quality delay-sensitive streaming. Using L4C2, nodes in a network estimate the acceptable delay and packet loss probability in their uplinks, aiming at retrieving lost data packets from in-network cache and/or coded data packets using in-network coding within an acceptable delay, by extending the Content-Centric Networking (CCN) approach. Further, L4C2 naturally provides multiple path and multicast technologies to efficiently utilize network resources while sharing network resources fairly with competing data flows by adjusting the video quality as necessary. We validate through comprehensive simulations that L4C2 achieves a high success probability of data transmission considering the acceptable

one-way delay and outperforms the existing solution. This work has been presented at the IEEE Infocom conference in May 2017 [23].

6.2.6. *Scalable Multicast Service in Software Defined ISP networks*

Participants: Hardik Soni, Thierry Turletti, Walid Dabbous.

In the context of the SDN-based multicast mechanisms activity, we designed an architectural solution to provide scalable multicast service in ISP networks. In fact, new applications where anyone can broadcast video are becoming very popular on smartphones. With the advent of high definition video, ISP providers may take the opportunity to propose new high quality broadcast services to their clients. Because of its centralized control plane, Software Defined Networking (SDN) seems an ideal way to deploy such a service in a flexible and bandwidth-efficient way. But deploying large scale multicast services on SDN requires smart group membership management and a bandwidth reservation mechanism to support QoS guarantees that should neither waste bandwidth nor impact too severely best effort traffic. We have proposed a Network Function Virtualization based solution for Software Defined ISP networks to implement scalable multicast group management. We also proposed a routing algorithm called Lazy Load balancing Multicast (L2BM) for sharing the network capacity in a friendly way between guaranteed-bandwidth multicast traffic and best-effort traffic. Our implementation of the framework made on Floodlight controllers and Open vSwitches has been used to study the performance of L2BM. This work has been presented at the IEEE ICC conference [24] in May 2017 and an extended version has been published to the IEEE TNSM journal [13].

6.2.7. *Placement of Virtual Network Function Chains in 5G*

Participants: Osama Arouk, Thierry Turletti.

We proposed a novel algorithm, namely Multi-Objective Placement (MOP), for the efficient placement of Virtualized Network Function (VNF) chains in future 5G systems. Real datasets are used to evaluate the performance of MOP in terms of acceptance ratio and embedding time when placing the time critical radio access network (RAN) functions as a chain. In addition, we rely on a realistic infrastructure topology to assess the performance of MOP with two main objectives: maximizing the number of base stations that could be embedded in the Cloud and load balancing. The results reveal that the acceptance ratio of embedding RAN functions is only 5% less than the one obtained with the optimal solution for the majority of considered scenarios, with a speedup factor of up to 2000 times. This work has been presented at the IEEE CloudNet conference in September 2017 [17].

6.2.8. *P4Bricks: Enabling multiprocessing using Linker-based network data plane architecture*

Participants: Hardik Soni, Thierry Turletti, Walid Dabbous.

In order to realize NFV-based multicast service as an add-on network capability without having knowledge of implementation level details of other network functions, we proposed a novel data plane architecture, P4Bricks, for modularized control and packet processing in the network... We propose P4Bricks, a system which aims to deploy and execute multiple independently developed and compiled P4 programs on the same reconfigurable hardware device. P4Bricks is based on a Linker component that merges the programmable parsers/deparsers and restructures the logical pipeline of P4 programs by refactoring, decomposing and scheduling the pipelines' tables. It merges P4 programs according to packet processing semantics (parallel or sequential) specified by the network operator and runs the programs on the stages of the same hardware pipeline, thereby enabling multiprocessing. A paper presenting the initial design of our system with an ongoing implementation and studies P4 language's fundamental constructs facilitating merging of independently written programs was submitted to SOSR and publisehd as a research report [37].

6.2.9. *Vehicles as a Mobile Cloud: Modelling, Optimization and Performance Analysis*

Participants: Luigi Vigneri, Thrasyvoulos Spyropoulos, Chadi Barakat.

The large diffusion of handheld devices is leading to an exponential growth of the mobile traffic demand which is already overloading the core network. To deal with such a problem, several works suggest to store content (files or videos) in small cells or user equipments. In this work, done in collaboration with Eurecom with the support of the UCN@SOPHIA Labex, we push the idea of caching at the edge a step further, and we propose to use public or private transportation as mobile small cells and caches. In fact, vehicles are widespread in modern cities, and the majority of them could be readily equipped with network connectivity and storage. The adoption of such a mobile cloud, which does not suffer from energy constraints (compared to user equipments), reduces installation and maintenance costs (compared to small cells). In our work, a user can opportunistically download chunks of a requested content from nearby vehicles, and be redirected to the cellular network after a deadline (imposed by the operator) or when her playout buffer empties. The main goal of the work is to suggest to an operator how to optimally replicate content to minimize the load on the core network. Our main contributions are: (i) Modelling: We model the above scenario considering heterogeneous content size, generic mobility and a number of other system parameters. (ii) Optimization: We formulate some optimization problems to calculate allocation policies under different models and constraints. (iii) Performance analysis. We build a MATLAB simulator to validate the theoretical findings through real trace-based simulations. We show that, even with low technology penetration, the proposed caching policies are able to offload more than 50 percent of the mobile traffic demand. The results of this work has been published in several papers, and are currently the subject of two submissions to journals. In particular, in [25] we consider the case of per-chunk caching for video streaming, whereas in [26] we consider the case of Quality of Experience-aware caching of files where the Quality of Experience is modeled as the slowdown in the file download time. A thorough presentation of this work and of our contributions can be found in the PhD thesis of Luigi Vigneri defended in July 2017 and available at [12].

6.3. Experimental Evaluation

6.3.1. *The Reproducibility'17 workshop*

Participant: Damien Saucez.
Recently, the ACM highlighted that the lack of reproducibility tended to be general in computer science and proposed normalised artifact reviewing and badging definitions [1] with the hope that the various ACM communities would perform artifact reviews based on these definition. We organized a special workshop on reproducibility in conjunction with the ACM SIGCOMM 2017 conference to produce a set of recommendations on how to assess the reproducibility of research published in ACM SIGCOMM'related conferences and journals and ways to promote reproducibility. The proceedings of the workshop is available in [27] and we have produced a set of recommendations to the community in [34] with the following conclusions:

The workshop pointed out that there are several hurdles concerning reproducibility, namely the absence of incentives and the bad habit that our community has grown accustomed to. This is evident in the current typical review process which is not adapted to handle reproducibility. Furthermore, there is no general way to share and preserve artifacts (and related documentation), every author does it in their own way.

The workshop focused on the two most important points to be tackled, namely, i) how to provide incentives for reproducible papers and ii) how to share artifacts.

For the first, a promising approach is to put in place a Reproducibility Committee, which will run in parallel with the normal Technical Program Committee of conferences and workshops, which will assess the level of reproducibility of papers accepted for publication by the TPC. Such approach will solve some of the privacy and anonymity issues while reducing the volume of work for the reviewers that volunteer in assessing the reproducibility level.

---

[1] Artifact Review and Badging, https://www.acm.org/publications/policies/artifact-review-badging, December 2017

For the second, a gradual approach has been suggested. The ACM digital library has been suggested as place to start sharing artifacts, which will be also identified via a DOI number. Beside the artifact itself it is important to share all of the meta-information necessary to actually reproduce prior work, as well as a way to provide feedback in order to make the community learn which meta-information is actually important and build guidelines on how to provide such information.

6.3.2. *Towards Realistic Software-Defined Wireless Networking Experiments*

Participants: Mohamed Naoufal Mahfoudi, Walid Dabbous, Thierry Turletti.

Software-Defined Wireless Networking (SDWN) is an emerging approach based on decoupling radio control functions from the radio data plane through programmatic interfaces. Despite diverse ongoing efforts to realize the vision of SDWN, many questions remain open from multiple perspectives such as means to rapid prototype and experiment candidate software solutions applicable to real world deployments. To this end, emulation of SDWN has the potential to boost research and development efforts by re-using existing protocol and application stacks while mimicking the behavior of real wireless networks. In this work, we provided an in-depth discussion on that matter focusing on the Mininet-WiFi emulator design to fill a gap in the experimental platform space. We showcased the applicability of our emulator in an SDN wireless context by illustrating the support of a number of use cases aiming to address the question on how far we can go in realistic SDWN experiments, including comparisons to the results obtained in a wireless testbed. Finally, we discussed the ability to replay packet-level and radio signal traces captured in the real testbed towards a virtual yet realistic emulation environment in support of SDWN research. This works has been published in a Special Issue on Software Defined Wireless Networks of the Computer Journal [15].

6.3.3. *ORION: Orientation Estimation Using Commodity Wi-Fi*

Participants: Mohamed Naoufal Mahfoudi, Thierry Turletti, Thierry Parmentelat, Walid Dabbous.

With MIMO, Wi-Fi led the way to the adoption of antenna array signal processing techniques for finegrained localization using commodity hardware. These techniques, previously exclusive to specific domains of applications, open the road to reach beyond localization, and now allow to consider estimating the device's orientation in space, that once required other sources of information. Wi-Fi's popularity and the availability of metrics related to channel propagation (CSI), makes it a candidate readily available for experimentation. We have recently proposed the ORION system to estimate the orientation (heading and yaw) of a MIMO Wi-Fi equipped object, relying on a joint estimation of the angle of arrival and the angle of departure. Although the CSI's phase data is plagued by several phase inconsistencies, we demonstrate that an appropriate phase compensation strategy significantly improves estimation accuracy. By feeding the estimation to a Kalman filter, we further improve the overall system accuracy, and lay the ground for an efficient tracking. Our technique allows estimating orientations within high precision. The results of the study were presented at a IEEE specialized workshop on Network Localization on Navigation [22].

6.3.4. *Lessons Learned while Trying to Reproduce the OpenRF Experiment*

Participants: Mohamed Naoufal Mahfoudi, Thierry Turletti, Thierry Parmentelat, Walid Dabbous.

Evaluating and comparing performance of wireless systems, like for any other scientific area, requires the ability to reproduce experimental results. In this work, we described the specific issues that we encountered when focusing on reproducing the experiments described in a paper related to wireless systems. We selected the OpenRF paper published in SIGCOMM 2013, a very interesting research work allowing to perform beamforming on commodity WiFi devices. We illustrated how reproducibility is strongly dependent on the used hardware, and why an extensive knowledge of the used hardware and its design is necessary. On the basis of this experience, we proposed some recommendations and lessons for the design of reproducible wireless experiments. This work has been presented at the ACM SIGCOMM 2017 Reproducibility Workshop in August 2017 [21].

6.3.5. *Deploying a 5G network in less than 5 minutes*

Participants: Mohamed Naoufal Mahfoudi, Thierry Parmentelat, Thierry Turletti, Walid Dabbous.

We proposed a demonstration run on R2lab, an anechoic chamber located at Inria Sophia Antipolis, France. This demonstration consists in deploying a standalone 5G network in less than 5 minutes. All the network components (base station, subscriber management, serving and packet gateways, network traac analyzers) were run automatically using the nepi-ng experiment orchestration tool. Download and upload performance to the Internet from a commercial phone located in the anechoic chamber are shown. This demo has been presented at the ACM SIGCOMM conference in August 2017 [33].

7. Bilateral Contracts and Grants with Industry

7.1. Bilateral Contracts with Industry

We have signed a bilateral contract for one year with Safran in order to build a network simulator specialised for aeronautical networks.

We are involved in an ADR "Rethinking the Network: Virtualizing Network Functions, from Middleboxes to Applications" with Nokia Bell Labs. The idea is to work on Unified control plane for fast NFV deployment.

7.2. Bilateral Grants with Industry

The ANSWER project is leaded by the QWANT search engine and the Inria Sophia Antipolis Méditerranée research center. This proposal is the winner of the "Grand Challenges du Numérique" (BPI) and aims to develop the new version of the search engine http://www.qwant.com with radical innovations in terms of search criteria, indexed content and privacy of users. In the context of this project, we got with Nataliia Bielova from the INDES project-team a funding for a 3 years Ph.D. student to work on Web tracking technologies and privacy protection.

8. Partnerships and Cooperations

8.1. Inria internal funding

ADT ACQUA: In the context of the Inria ADT call, we have a funding for a two year engineering position on the ACQUA project for the 2015-2017 period. Thierry Spetebroot is hired on this position. In 2017, this ADT got extended by six months beyond the two years period to therefore end on March 2018.

IPL BetterNet: The Diana team is part of the Inria Project Lab BetterNet (http://project.inria.fr/betternet/). Within this lab, Inria has funded two PhD students in 2017 co-supervised by Chadi Barakat from the Diana project-team. The first PhD student is Thibaut Ehlinger hosted within the Diana team and co-supervised by Vassilis Christophides from the MiMove team in Paris. The second PhD student is Imane Taibi hosted by the Dionysos team in Rennes and co-supervied by Gerardo Rubino and Yassine Hadjadj-Aoul. Both PhDs started on the 1st of November 2017.

8.2. Regional Initiatives

ElectroSmart: This project benefits form the following fundings:
- a three year engineering position from the UCN@Sophia Labex for the 2016-2018 period (Ravi Mondi is hired on this position)
- 30KEuros from academy 1 of UCAJedi
- a two years engineering position from an Inria ADT for 2017/2018 (Abdelhakim Akodadi)
- a one year business developer from Inria ATT for june 2017-june 2018(David Migliacci)
- a 3 years 2017/2020 Ph.D. thesis from academy 1 of UCAJedi (Yanis Boussad)

D2D Indoor: This project is joint with the NFCOM startup in Nice, specialized in the development of new services for mobile phones. The project aims at leveraging mobile to mobile communications for offloading the cellular infrastructure, and will target a solution based on algorithms previously developed in the Diana project-team (BitHoc and HBSD). The project got a funding for one year engineer from the Labex. A position is open.

8.3. National Initiatives

8.3.1. ANR

- **ANR FIT** (2011-2018): FIT (Future Internet of Things) aims at developing 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 Equipements of Excellence (Equipex) research grant programme. 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/.

- **ANR DISCO** (2014-2017): DISCO (DIstributed SDN COntrollers for rich and elastic network services) aims at exploring the way how Software Defined Networking changes network monitoring, control, urbanisation and abstract description of network resources for the optimisation of services. The project works throughout experimentations and application use cases on the next generation of Software-Defined Networking solutions for large and critical distributed systems. The project studied the distribution of the current SDN control plane and the optimization of network operations that the integrated system view of cloud computing-based architectures allows. See also http://anr-disco.ens-lyon.fr/.

- **ANR REFLEXION** (2015-2017): REFLEXION (REsilient and FLEXible Infrastructure for Open Networking) research project will study the robustness and scalability of the current SDN architectures and the flexibility leveraged by SDN for provisioning resources and virtualized network functions (VNF). The project will address four main scientific objectives: (1) Fault and disruption management for virtualized services, (2) Robust and scalable control plane for next generation SDN, (3) Dynamic performance management of low level resources in SDN/NFV environments and (4) Distribution and optimization of virtual network functions in SDN environments. Our contribution in this project was focused on fault and disruption management for virtualized services. See also http://anr-reflexion.telecom-paristech.fr/.

- **ANR BottleNet** (2016-2019): BottleNet aims to deliver methods, algorithms, and software systems to measure Internet Quality of Experience (QoE) and diagnose the root cause of poor Internet QoE. This goal calls for tools that run directly at users' devices. The plan is to collect network and application performance metrics directly at users' devices and correlate it with user perception to model Internet QoE, and to correlate measurements across users and devices to diagnose poor Internet QoE. This data-driven approach is essential to address the challenging problem of modeling user perception and of diagnosing sources of bottlenecks in complex Internet services. ANR BottleNet will lead to new solutions to assist users, network and service operators as well as regulators in understanding Internet QoE and the sources of performance bottleneck.

8.4. European Initiatives

8.4.1. FP7 & H2020 Projects

Program: FP7 FIRE programme

Project acronym: Fed4Fire+

Project title: Federation for FIRE Plus

Duration: January 2017 - December 2021

Coordinator: iMinds (Belgium)

Other partners: 20 european partners including IMEC (Belgium), UPMC (Fr), Fraunhofer (Germany), TUB (Germany), etc.

Web site: http://www.fed4fire.eu/

Abstract: The Fed4FIRE+ project has the objective to run and further improve Fed4FIRE as best-in-town federation of experimentation facilities for the Future Internet Research and Experimentation initiative. Federating a heterogeneous set of facilities covering technologies ranging from wireless, wired, cloud services and open flow, and making them accessible through common frameworks and tools suddenly opens new possibilities, supporting a broad range of experimenter communities covering a wide variety of Internet infrastructures, services and applications. Fed4FIRE+ will continuously upgrade and improve the facilities and include technical innovations, focused towards increased user satisfaction (user-friendly tools, privacy-oriented data management, testbed SLA and reputation, experiment reproducibility, service-level experiment orchestration, federation ontologies, etc.). It will open this federation to the whole FIRE community and beyond, for experimentation by industry and research organisations, through the organization of Open Calls and Open Access mechanisms. The project will also establish a flexible, demand-driven framework which allows test facilities to join during the course of its lifetime by defining a set of entry requirements for new facilities to join and to comply with the federation. FIRE Experimental Facilities generate an ever increasing amount of research data that provides the foundation for new knowledge and insight into the behaviour of FI systems. Fed4FIRE+ will participate in the Pilot on Open Research Data in Horizon 2020 to offer open access to its scientific results, to the relevant scientific data and to data generated throughout the projectâs lifetime. Fed4FIRE+ will finally build on the existing community of experimenters, testbeds and tool developers and bring them together regularly (two times a year) in engineering conferences to have maximal interaction between the different stakeholders involved.

8.5. International Initiatives

8.5.1. *Inria Associate Teams Not Involved in an Inria International Labs*

8.5.1.1. *UHD-on-5G*

Title: Ultra High Definition video streaming on future 5G networks

International Partner (Institution - Laboratory - Researcher):

National Institute of Information and Communications Technology (NICT) (Japan) ICN project-team - Hitoshi Asaeda

Start year: 2016

See also: https://team.inria.fr/diana/uhd-on-5g/

The aim of this collaboration is to design and develop efficient mechanisms for streaming UHD video on 5G networks and to evaluate them in a realistic and reproducible way by using novel experimental testbeds.

Our approach leverages and extends when necessary ICN and SDN technologies to allow very high quality video streaming at large scale. We also plan to use Virtual Network Functions (VNF) in order to place easily and dynamically different functions (e.g. transcoding, caching) at strategic locations within the network. Specifically, the placement of these functions will be decided by SDN controllers to optimize the quality of experience (QoE) of users. Moreover, we plan to integrate ICN functionalities (e.g., name-based forwarding and multipath transport using in-network caching) with SDN/NFV to provide better QoE and mobility services support to users than traditional IP

architectures. Monitoring mechanisms such as the Contrace tool we developed in the SIMULBED associated team will be helpful to provide an accurate view of the network at the SDN controllers side. In addition, we will build a large-scale testbed to evaluate our solutions through reproducible experimentations based on two testbeds: the ICN wired CUTEi tesbed developed by NICT and the wireless R2lab testbed developed by Inria.

8.6. International Research Visitors

8.6.1. Visits of International Scientists

Katia Obraczka is Professor of Computer Engineering and Graduate Director at Department of Computer Engineering, UC Santa Cruz where she leads the Internetworking Research Group (i-NRG). She has visited us for four weeks in July 2017. The Labex UCN@Sophia has supported two one-month visits at the DIANA project-team, in July 2017 and during summer 2018 to work in particular on the decentralization of the SDN control plane applied to Intelligent Transport Systems (ITS).

8.6.1.1. Internships

Neha Agarwal

Date: from Apr 2017 until Sep 2017

Institution: Ubinet Master intern, University of Nice Sophia Antipolis

Supervisor: Arnaud Legout

Subject: Automated Tests for ElectroSmart in Android Studio

Yanis Boussad

Date: from Mar 2017 until Aug 2017

Institution: Ubinet Master intern, University of Nice Sophia Antipolis

Supervisor: Arnaud Legout

Subject: Exploration of Electromagnetic Fields Metrics

Pretesh Chauhan

Date: from May 2017 until Jul 2017

Institution: Third year intern, National Institute of Technology, Hamirpur, INDIA.

Supervisor: Arnaud Legout

Subject: User Exposure Profiles in ElectroSmart

Giuseppe Di Lena

Date: from Mar 2017 until Aug 2017

Institution: Ubinet Master intern, University of Nice Sophia Antipolis

Supervisor: Damien Saucez and Thierry Turletti

Subject: Robust Virtualized services in OpenStack

David Migliacci

Date: from Jul 2017

Institution: Intern, Skema Business School

Supervisor: Arnaud Legout

Subject: Business Developement for ElectroSmart

Yassir Mrabet

Date: from Mar 2017 until Aug 2017

Institution: Ubinet Master intern, University of Nice Sophia Antipolis

Supervisor: Walid Dabbous and Thierry Turletti

Subject: Anechoic Chamber Characterization for Trustful Evaluation of Wireless Protocols.

Imane Taibi

Date: from Mar 2017 until Aug 2017

Institution: Ubinet Master intern, University of Nice Sophia Antipolis

Supervisor: Chadi Barakat

Subject: Experimenting and modeling Web Quality of Experience

8.6.2. *Visits to International Teams*

Thierry Turletti visited NICT in Tokyo Japan in the context of the UHD-on-5G associated team in December 2017.

9. Dissemination

9.1. Promoting Scientific Activities

Chadi Barakat is on the editorial board of the Computer Networks journal, and was on the Technical Program Committee for the ACM Internet Measurement Conference (IMC) 2017, the International Teletraffic Congress (ITC) 2017 and the IEEE Measurement and Networking (M&N) Workshop 2017. He co-chaired the poster session at the ACM Internet Measurement Conference (IMC) 2017 and was session chair at the same conference. He is currently the scientific referee for international affairs at Inria Sophia Antipolis and member of the Conseil d'Orientation Scientifique et Technologique at Inria within the working group of international affairs (COST-GTRI).

Walid Dabbous is member of the scientific council of the Inria Bell-Labs laboratory on Communication networks of the future. He also servers as a chair of the scientific committee of the User Centric Networking (UCN@Sophia) Laroratory of Excellence, and member of the scientific committee of the Academy 1 of the UCAJedi Idex. He is also member of the Ubinet International Master program steering committee. He is member of the Inria Sophia Antipolis project committee's bureau (Bureau du CP).

Arnaud Legout is on the editorial board of the Computer Networks journal. Arnaud Legout is the president of the Commission of the users of IT resources of Sophia Antipolis Inria research center.

Damien Saucez was TPC member of IEEE ICC 2017, IEEE CCNC 2017. In 2017, Damien Saucez have co-organised the ACM SIGCOMM Reproducilbity workshop, as detailed above. He also co-organized the RESCOM 2017 summer school that aimed at gathering French communities working around networking. We were in charge of organizing the student poster sessions [28] and to give labs on computer-network simulations.

Thierry Turletti, Senior ACM and IEEE member, served in the program committees of the following international workshops and conferences: the 8th Workshop on ns-3 at Porto, Portugal in June 13-14, 2nd CoRes Workshop at Quiberon, France in May 29-30 2017. 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. Thierry Turletti is president of the Committee for Technological Development (CDT) and member of the committee NICE that studies postdoc and visiting researcher applications at Inria Sophia Antipolis.

9.2. Teaching - Supervision - Juries

9.2.1. Teaching

Master Ubinet: Chadi Barakat and Walid Dabbous, Evolving Internet, 31.5 hours, M2, University of Nice-Sophia Antipolis, France.

Master Ubinet: Chadi Barakat and Walid Dabbous, Internet Measurements and New Architectures, 31.5 hours, M2, University of Nice-Sophia Antipolis, France.

Master Ubinet: Arnaud Legout, From BitTorrent to Privacy, 36 hours, M2, University of Nice-Sophia Antipolis, France.

Master 1 International: Chadi Barakat, Algorithms for Networking, 22.5 hours, M1, University of Nice-Sophia Antipolis, France.

Master 1 International: Arnaud Legout, Oral and written communications, 18 hours, M1, University of Nice-Sophia Antipolis, France.

Master Estel: Chadi Barakat, Voice over IP, 9 hours, University of Nice-Sophia Antipolis, France.

Master IUP GMI: Damien Saucez, Security and privacy in networks, 38h, M2, University of Avignon, France.

IUT : Damien Saucez, Advanced Network Services and Operator Network Technologies, 27h, L1 and L2, University of Nice-Sophia Antipolis, France.

E-learning

Python: Arnaud Legout and Thierry Parmentelat are co-authors of the MOOC Python 3 : "Python 3 : des fondamentaux aux concepts avancés du langage" that lasts 9 weeks on FUN (https://www.france-universite-numerique-mooc.fr/), UCA. For the first session there were 11677 registered persons.

This MOOC is a brand new version of the past MOOC "Python: des fondamentaux à l'utilisation du langage" in Python 2. This new MOOC has been funded by UCA.

9.2.2. Supervision

PhD in progress: Yanis Boussad started his PhD on "Large scale characterization of the exposition to microwaves" in October 2017. He is co-supervised with Leonardo Lizzi, LEAT.

PhD in progress: Thibaut Ehlinger: started his PhD on "Mapping Quality of Service metrics to user Quality of Experience in the Internet" in November 2017. He is co-supervised by Chadi Barakat and Vassilis Christophides (EPI MiMove, Inria Paris).

PhD in progress: Karyna Gogunska works on "Empowering Virtualized Networks with Measurement As a Service (MaaS)" since October 2016. Her thesis is co-supervised by Chadi Barakat and Guillaume Urvoy-Keller (I3S).

PhD in progress: Muhammad Jawad Khokhar works on "From Network Level Measurements to Expected Quality of User Experience" since September 2016. His PhD is supervised by Chadi Barakat.

PhD in progress: Mohamed Naoufal Mahfoudi works on cross-layer optimization techniques for next generation MIMO-based networks since November 2015. His thesis is co-supervised by Walid Dabbous and Robert Staraj (LEAT).

PhD in progress: Ghada Moualla works on "the problem of network faults and how to circumvent them by the means of Software Defined Networking, virtualization, and service function chaining" since November 2015. Her thesis is co-supervised by Thierry Turletti and Damien Saucez.

PhD in progress: Vitalii Poliakov works on "the application of Software Defined Networking on 5G networks in order to optimise the Quality of Experience of network services" since November 2015. His thesis is co-supervised by Damien Saucez and Lucile Sassatelli (I3S).

PhD in progress: Hardik Soni works on "Software Defined Networking in challenged environments" since September 2014. His thesis is co-supervised by Thierry Turletti and Walid Dabbous.

PhD in progress: Iman Taibi started his PhD in the Dionysos project-team on "Big data analysis for network monitoring and troubleshooting" in November 2017. She is co-supervised by Gerardo Rubino, Yassine Hadjadj-Aoul and Chadi Barakat.

PhD in progress: Mathieu Thiery started his PhD in the Privatics project-team on "Data protection of connected objects and smartphones" in April 2017. He is co-supervised by Vincent Roca and Arnaud Legout.

PhD in progress: Thibaud Trolliet started his PhD on "Exploring trust on Twitter" in octobre 2017. He is co-supervised with Frederic Giroire.

PhD: Luigi Vigneri defended his PhD on "Vehicles as a Mobile Cloud: Leveraging mobility for content storage and dissemination" in July 2017. His thesis was co-supervised by Chadi Barakat and Thrasyvoulos Spyropoulos (Eurecom).

9.2.3. *Juries*

Chadi Barakat served as reviewer of Wenqin Shao PhD thesis, "Measurement-based inter-domain traffic engineering: Scalability, data interpretation and network event visibility" defended in November at Telecom ParisTech.

Chadi Barakat served as reviewer of Danilo Giordano PhD thesis, "Machine Learning and Bigdata Methodologies for Network Traffic Monitoring" defended in March at Politecnico di Torino.

Walid Dabbous serves as reviewer of Cédric Baudoin HDR thesis, "Optimisation et intégration des réseaux de télécommunication par satellite", that will be defended on January 2018 Institut National Polytechnique de Toulouse (INP Toulouse).

Walid Dabbous serves as reviewer of Mathieu Bouet HDR thesis, "Software networks : Orchestration, resilience, and programmability concerns", defended in July 2017 at UPMC Sorbonne Universités.

Walid Dabbous serves as jury member of Giovanni Neglia HDR thesis, "Delay Tolerant Networks", defended in February 2017 at University of Nice Sophia Antipolis.

Walid Dabbous served as jury member of Nicolas Huin PhD thesis, "Energy Efficient Software Defined Networks", defended on September 28, 2017 at the University of Nice Sophia Antipolis.

Walid Dabbous served as a jury member for the mid-term reivew of the Ph.D. thesis of Sumit Kymar (Eurecom).

Arnaud Legout serves as a jury member for the intermediate validation of the Ph.D. thesis of Athanasios Andreou (Eurecom, UPMC)

Thierry Turletti served as reviewer of Mediwaththe Gedara Chathurika Prasadini Mediwaththe PhD thesis, "Game-theoretic Methods for Small-scale Demand-side Management in Smart Grid", defended on January 2017 at the Faculty of Engineering, University of New South Wales, Australia.

Thierry Turletti served as reviewer of Pascal Thubert PhD thesis, "Converging over Deterministic Networks: For an Industrial Internet", defended on March 2017 at the Université Bretagne Loire, France.

Thierry Turletti served as jury member of Myriana Rifai PhD thesis, "Next-Generation SDN Based Virtualized Networks", defended on September 25 2017 at the University of Nice-Sophia Antipolis, France.

Thierry Turletti served as reviewer of Gentian Jaklari HDR thesis, "Thinking Outside the Networking Box for Realizing the Potential of Wireless Networks", defended on November 14 2017 at INP Toulouse, France.

9.3. Popularization

Chadi Barakat keeps participating to the organization of the Mediterranean Students Days @ Campus SophiaTech. The fifth edition took place on Feb 28 - March 2, 2017, and the sixth edition will take place on February 12 - 14, 2018. All details on this event can be found at http://univ-cotedazur.fr/events/meddays.

Arnaud Legout and Damien Saucez participated in the "fête de la science" event. They presented the ElectroSmart project and the R2lab Testbed.

Walid Dabbous published a popularization article on Network Architecture for Big Data in the book "Les Big Data à découvert" edited by CNRS editions [29].

Damien Saucez is part of the MASTIC [2] group at Inria. MASTIC groups all the activities for scientific dissemination for Inria Sophia Antipolis. Within this activity we visit schools to make conferences on Internet and privacy related matters and Arnaud Legout made a talk during MathC2+. Within the team, we gave a 3h game-based class to Collège interns that made an internship at Inria in November to introduce them to the fundamental algorithms used in the Internet (shortest-path, longest prefix matching, distance-vector, AIMD). We also served on the Inria stand for "fête de la science" to welcome kids and families and present mathematical games, the R2LAB testbed, and Electrosmart. In December 2017 Damien Saucez made a Café-in in Sophia Antipolis which is a 1h informal talk on the presenter's research with all Inria collaborators (administrative, technical, and scientific). In addition to these regional activities, we also took part in the ICN Fun MOOC [3] where Damien Saucez made a 10 minutes explanations of the fundamentals behind Operating Systems. As of today, the video [35] has been watched more than 9k times on YouTube.

10. Bibliography

Major publications by the team in recent years

[1] M. Gabielkov, A. Ramachandran, A. Chaintreau, A. Legout. *Social Clicks: What and Who Gets Read on Twitter?*, in "ACM SIGMETRICS / IFIP Performance 2016", Antibes Juan-les-Pins, France, June 2016, https://hal.inria.fr/hal-01281190

[2] M. Gabielkov, A. Rao, A. Legout. *Studying Social Networks at Scale: Macroscopic Anatomy of the Twitter Social Graph*, in "ACM Sigmetrics 2014", Austin, United States, June 2014, https://hal.inria.fr/hal-00948889

[3] Y.-H. Kim, A. Quereilhac, M. A. Larabi, J. Tribino, T. Parmentelat, T. Turletti, W. Dabbous. *Enabling Iterative Development and Reproducible Evaluation of Network Protocols*, in "Computer Networks", January 2014, vol. 63, pp. 238-250, https://hal.inria.fr/hal-00861002

[4] D. Kutscher, S. Eum, K. Pentikousis, I. Psaras, D. Corujo, D. Saucez, T. C. Schmidt, M. Waehlisch. *Information-Centric Networking (ICN) Research Challenges*, July 2016, Internet Research Task Force (IRTF), Request for Comments: 7927, https://hal.inria.fr/hal-01423154

[5] X.-N. Nguyen, D. Saucez, C. Barakat, T. Turletti. *OFFICER: A general Optimization Framework for OpenFlow Rule Allocation and Endpoint Policy Enforcement*, in "The 34th Annual IEEE International Conference on Computer Communications (INFOCOM 2015)", Hongkong, China, IEEE, April 2015, https://hal.inria.fr/hal-01104519

[6] R. Ravaioli, G. Urvoy-Keller, C. Barakat. *Towards a General Solution for Detecting Traffic Differentiation At the Internet Access*, in "27th International Teletraffic Congress (ITC-27)", Ghent, Belgium, September 2015, https://hal.inria.fr/hal-01161795

---

[2] https://project.inria.fr/mastic
[3] https://www.fun-mooc.fr/courses/inria/41014/session01/about

[7] J. REN, A. RAO, M. LINDORFER, A. LEGOUT, D. CHOFFNES. *ReCon: Revealing and Controlling PII Leaks in Mobile Network Traffic*, in "Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services", Singapore, Singapore, ACM, June 2016, pp. 361 - 374 [*DOI :* 10.1145/2906388.2906392], https://hal.inria.fr/hal-01386899

[8] D. SAUCEZ, L. IANNONE, C. ALBERT, F. CORAS. *Locator/ID Separation Protocol (LISP) Impact*, April 2016, Internet Engineering Task Force (IETF), Request for Comments: 7834, https://hal.inria.fr/hal-01423163

[9] D. SAUCEZ, L. IANNONE, O. BONAVENTURE. *Locator/ID Separation Protocol (LISP) Threat Analysis*, April 2016, Internet Engineering Task Force (IETF), Request for Comments: 7835, https://hal.inria.fr/hal-01423162

[10] T. SPETEBROOT, S. AFRA, N. AGUILERA, D. SAUCEZ, C. BARAKAT. *From network-level measurements to expected Quality of Experience: the Skype use case*, in "IEEE International Workshop on Measurement and Networking (M&N)", Coimbra, Portugal, October 2015, https://hal.inria.fr/hal-01071373

[11] L. VIGNERI, T. SPYROPOULOS, C. BARAKAT. *Storage on Wheels: Offloading Popular Contents Through a Vehicular Cloud*, in " IEEE 17th International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM)", Coimbra, Portugal, June 2016 [*DOI :* 10.1109/WoWMoM.2016.7523506], https://hal.inria.fr/hal-01315370

Publications of the year

Doctoral Dissertations and Habilitation Theses

[12] L. VIGNERI. *Vehicles as a mobile cloud : modelling, optimization and performance analysis*, Université Côte d'Azur, July 2017, https://tel.archives-ouvertes.fr/tel-01614566

Articles in International Peer-Reviewed Journals

[13] H. SONI, W. DABBOUS, T. TURLETTI, H. ASAEDA. *NFV-based Scalable Guaranteed-Bandwidth Multicast Service for Software Defined ISP networks*, in "IEEE Transactions on Network and Service Management", December 2017, vol. 14, n⁰ 4, 14 p. [*DOI :* 10.1109/TNSM.2017.2759167], https://hal.inria.fr/hal-01596488

[14] Z. ZALI, E. ASLANIAN, M. H. MANSHAEI, M. R. HASHEMI, T. TURLETTI. *Peer-assisted Information-Centric Network (PICN): A Backward Compatible Solution*, in "IEEE Access", November 2017, vol. 5, pp. 25005 - 25020 [*DOI :* 10.1109/ACCESS.2017.2762697], https://hal.inria.fr/hal-01661617

[15] R. DOS REIS FONTES, M. MAHFOUDI, W. DABBOUS, T. TURLETTI, C. ROTHENBERG. *How far can we go? Towards Realistic Software-Defined Wireless Networking Experiments*, in "The Computer Journal (Oxford)", October 2017, vol. 60, n⁰ 10, 13 p. , https://hal.inria.fr/hal-01480973

International Conferences with Proceedings

[16] E. ALTMAN, F. DE PELLEGRINI, A. MASSARO, J. LEGUAY, D. SAUCEZ, L. MAGGI, J. LEGUAY. *Blind, Adaptive and Robust Flow Segmentation in Datacenters*, in "INFOCOM 2018 - IEEE International Conference on Computer Communications", Honolulu, United States, April 2018, https://hal.inria.fr/hal-01666905

[17] O. AROUK, N. NIKAEIN, T. TURLETTI. *Multi-Objective Placement of Virtual Network Function Chains in 5G*, in "IEEE International Conference on Cloud Networking (CloudNet)", Prague, Czech Republic, September 2017, https://hal.inria.fr/hal-01577933

[18] A. KAUL, K. OBRACZKA, M. A. S. SANTOS, C. E. ROTHENBERG, T. TURLETTI. *Dynamically Distributed Network Control for Message Dissemination in ITS*, in "IEEE/ACM DS-RT 2017 - 21st International Symposium on Distributed Simulation and Real Time Applications", Rome, Italy, October 2017, https://hal.inria.fr/hal-01579425

[19] M. J. KHOKHAR, N. A. SABER, T. SPETEBROOT, C. BARAKAT. *On active sampling of controlled experiments for QoE modeling*, in "ACM SIGCOMM 2017 2nd Workshop on QoE-based Analysis and Management of Data Communication Networks (Internet-QoE 2017)", Los Angeles, United States, August 2017, https://hal.inria.fr/hal-01525723

[20] Y. LI, A. ABOUSEIF, L. IANNONE, D. SAUCEZ. *LISP-Views: Monitoring LISP at Large Scale*, in "29th International Teletraffic Congress", Genova, Italy, September 2017, https://hal.inria.fr/hal-01676534

[21] M. N. MAHFOUDI, T. TURLETTI, T. PARMENTELAT, W. DABBOUS. *Lessons Learned while Trying to Reproduce the OpenRF Experiment*, in "Reproducibility'17 - ACM SIGCOMM 2017 Reproducibility Workshop", Los Angeles, United States, August 2017, vol. 41, n^o 1, pp. 21 - 23 [*DOI : 10.1145/3097766.3097772*], https://hal.archives-ouvertes.fr/hal-01615398

[22] M. N. MAHFOUDI, T. TURLETTI, T. PARMENTELAT, F. FERRERO, L. LIZZI, R. STARAJ, W. DABBOUS. *ORION: Orientation Estimation Using Commodity Wi-Fi*, in "Workshop on Advances in Network Localization and Navigation (ANLN)", Paris, France, May 2017, pp. 1033-1038, https://hal.archives-ouvertes.fr/hal-01424239

[23] K. MATSUZONO, H. ASAEDA, T. TURLETTI. *Low Latency Low Loss Streaming using In-Network Coding and Caching*, in "IEEE INFOCOM", Atlanta, United States, May 2017, https://hal.inria.fr/hal-01437074

[24] H. SONI, T. TURLETTI, W. DABBOUS, H. ASAEDA. *Scalable Guaranteed-Bandwidth Multicast Service in Software Defined ISP networks*, in "IEEE International Conference on Communications (ICC)", Paris, France, May 2017, https://hal.inria.fr/hal-01400688

[25] L. VIGNERI, S. PECORARO, T. SPYROPOULOS, C. BARAKAT. *Per Chunk Caching for Video Streaming from a Vehicular Cloud*, in "CHANTS 2017- ACM MobiCom Workshop on Challenged Networks", Snowbird, Utah, USA, United States, October 2017, https://hal.inria.fr/hal-01567671

[26] L. VIGNERI, T. SPYROPOULOS, C. BARAKAT. *Quality of Experience-Aware Mobile Edge Caching through a Vehicular Cloud*, in "MSWiM 2017 - ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems", Miami, United States, November 2017, https://hal.inria.fr/hal-01567678

Books or Proceedings Editing

[27] D. SAUCEZ, L. IANNONE (editors). *Proceedings of the Reproducibility Workshop*, August 2017, https://hal.inria.fr/hal-01666893

[28] D. SAUCEZ, A. LÈBRE, S. SECCI (editors). *RESCOM 2017 Summer school*, CNRS, June 2017, https://hal.inria.fr/hal-01558074

Scientific Popularization

[29] W. DABBOUS. *Quels réseaux pour les Big Data?*, in "Les Big Data à découvert", M. BOUZEGHOUB, R. MOSSERI (editors), CNRS Editions, 2017, https://hal.archives-ouvertes.fr/hal-01423096

Other Publications

[30] H. ASAEDA, X. SHAO, T. TURLETTI. *Contrace: Traceroute Facility for Content-Centric Network*, October 2017, Internet-Draft, ICN Research Group, draft-asaeda-icnrg-contrace-04, https://hal.inria.fr/hal-01661663

[31] F. DE PELLEGRINI, L. MAGGI, A. MASSARO, D. SAUCEZ, J. LEGUAY, E. ALTMAN, J. LEGUAY. *Learning how to segment flows in the dark*, November 2017, PGMO Days 2017 - Gaspard Monge Program for Optimisation, Poster, https://hal.inria.fr/hal-01666903

[32] G. D. LENA. *Robust Service Function Chains in OpenStack*, Université Côte d'Azur, August 2017, pp. 1-37, https://hal.inria.fr/hal-01651440

[33] M. N. MAHFOUDI, T. PARMENTELAT, T. TURLETTI, W. DABBOUS, R. KNOPP. *Deploy a 5G network in less than 5 minutes: Demo Abstract*, August 2017, ACM SIGCOMM Posters and Demos, Poster, https://hal.inria.fr/hal-01580065

[34] D. SAUCEZ, I. LUIGI, L. IANNONE. *Thoughts and Recommendations from the ACM SIGCOMM 2017 Reproducibility Workshop*, December 2017, working paper or preprint, https://hal.inria.fr/hal-01666895

[35] D. SAUCEZ. *Le système d'exploitation en trois idées clé*, May 2017, MOOC sur l'Informatique et la Création Numérique, https://hal.inria.fr/hal-01666909

[36] D. SAUCEZ. *Vers la fin programmée des réseaux?*, December 2017, working paper or preprint, https://hal.inria.fr/hal-01671752

[37] H. SONI, T. TURLETTI, W. DABBOUS. *P4Bricks: Enabling multiprocessing using Linker-based network data plane architecture*, November 2017, working paper or preprint, https://hal.inria.fr/hal-01632431