RESEARCH CENTRE

Sophia Antipolis - Méditerranée

2021 ACTIVITY REPORT

Project-Team DIANA

Design, Implementation and Analysis of Networking Architectures

DOMAIN

Networks, Systems and Services, Distributed Computing

THEME

Networks and Telecommunications

Contents

Project-Team DIANA

Creation of the Project-Team: 2015 July 01

Keywords

Computer sciences and digital sciences

A1.1.13. – Virtualization
A1.2.1. – Dynamic reconfiguration
A1.2.2. – Supervision
A1.2.3. – Routing
A1.2.4. – QoS, performance evaluation
A1.2.5. – Internet of things
A1.2.9. – Social Networks
A1.3. – Distributed Systems
A1.3.4. – Peer to peer
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
B9.5.1. – Computer science
B9.5.6. – Data science
B9.8. – Reproducibility
B9.10. – Privacy

1 Team members, visitors, external collaborators

Research Scientists

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

PhD Students

- Othmane Belmoukadam [Univ Côte d'Azur, until Oct 2021]
- Yanis Boussad [Inria, until Oct 2021]
- Giuseppe Di Lena [Orange Labs, CIFRE, until Mar 2021]
- Mamoutou Diarra [Ekinops, CIFRE]
- Houssam Elbouanani [Inria]
- Bernard Tamba Sandouno [YDATA, CIFRE, from May 2021]

Technical Staff

- Thierry Parmentelat [Inria, Engineer]

Interns and Apprentices

- Anass El Boujidi [Inria, from Mar 2021 until Aug 2021]
- Téo Haÿs [Inria, Apprentice, from Sep 2021]

Administrative Assistant

- Christine Foggia [Inria]

External Collaborator

- Mondi Ravi [Self employed, until May 2021]

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 meaningful information 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 Application domains

The DIANA project-team conducts research activities to provide network architectural support for improving citizen rights in the Internet. The main application domains of the teams are:

- Network and quality of experience measurement
- Detection of private information leaks
- Industrial deterministic networks
- Data center networks
- Deployment of future open radio networks
- Realistic simulations and reproducible experiments

5 Social and environmental responsibility

Public health politics and scientist evaluating the impact of EMF radiations on human beings all face the same challenge: How to assess the real exposure of human beings in order to correlate it with observed symptoms and illness. This problem is even harder considering that the period of observation must be long (years), and the number of observed persons must be large.

ElectroSmart is a technological breakthrough that 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.

6 Highlights of the year

6.1 Awards

The PhD thesis of Othmane Belmoukadam, titled “QoE aware content management in the Internet: caching and transport” and defended in September 2021, was awarded the second best PhD thesis award in Computer Science at University Côte d’Azur.

6.2 Electrosmart large-scale deployment

The ElectroSmart App has exceeded 1 million downloads in 2021. This milestone was internally indicated as the target number of downloads at the launch of this app. By end of December 2021, the application has been downloaded 1.1 M times and has 200k active users. The ElectroSmart application made it possible to collect 13 billion measurements in 150 countries.

6.3 SophiaNode initial deployment

We started in collaboration with Eurecom the deployment of an open programmable 5G platform in Sophia-Antipolis in the context of the Eureopean project SLICES-RI. A 600 Gbps fiber links has been installed between Inria and Eurecom and servers and switches are being deployed as the nucleus of the future SophiaNode of the SLICES-RI project.

7 New software and platforms

Several if not all contributions of Diana project-team are validated through experimental software and platform, built for this purpose. However, the scope of these tools extend largely beyond the needs of the project-team and are proposed to the larger research community.

7.1 New software

7.1.1 ACQUAmobile

Name: Application for prediCting Quality of User Experience at Internet Access

Keywords: Android, Internet access, Performance measure, Quality of Experience

Scientific 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 was 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. The project also got the approval of Inria COERLE and French CNIL for the part on experimentation with real users. 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).

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.

Functional Description: 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 a prior work on Skype and YouTube Quality of Experience, or by soliciting the crowd (i.e. crowdsourcing) for combinations (i.e. tuples) of network- and application-level measurements and corresponding user-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.

Assessment: Audience = 3, Software Originality = 4, Software Maturity = 3, Evolution and Maintenance = 3, Software Distribution and Licensing = 5.

URL: http://project.inria.fr/acqua/

Authors: Thierry Spetebroot, Chadi Barakat

Contact: Chadi Barakat

7.1.2 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 changed fundamentally 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.

This project is supported by the UCN@Sophia Labex in 2016/2017/2018 (funding the engineer Mondi Ravi), by an Inria ADT (funding the engineer Abdelhakim Akodadi) 2017/2018, by and Inria ATT (funding the business developer David Migliacci) in 2017/2018, and by the academy 1 of UCAJedi (funding a Ph.D. student Yanis Boussad) 2017/2021.

In August 2016, we released the first stable public release of ElectroSmart. On the 30th December 2021 the app has been downloaded 1.1 million time, we have 200 000 active users and a score of 4,5/5 on Google Play. We collected 13 billions measurements in 150 countries.

Assessment: A-5, SO-4, SM-4, EM-3-up4, SDL-1

URL: https://electrosmart.app

Contact: Arnaud Legout

Participants: Arnaud Legout, Mondi Ravi

7.1.3 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.

- 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

URL: http://nepi-ng.inria.fr

Contact: Thierry Parmentelat

7.1.4 Distrinet

Name: Distrinet

Keywords: SDN, Emulation, Large-scale Emulators, Network simulator

Scientific Description: Networks have become complex systems that combine various concepts, techniques, and technologies. As a consequence, modelling or simulating them now is extremely complicated and researchers massively resort to prototyping techniques. Mininet is the most popular tool when it comes to evaluate SDN propositions. Mininet allows to emulate SDN networks on a single computer but shows its limitations with resource intensive experiments as the emulating host may become overloaded. To tackle this issue, we propose Distrinet, a distributed implementation of Mininet over multiple hosts, based on LXD/LXC, Ansible, and VXLAN tunnels. Distrinet uses the same API than Mininet, meaning that it is compatible with Mininet programs. It is generic and can deploy experiments on Linux clusters (e.g., Grid'5000), as well as on the Amazon EC2 cloud platform.

Assessment: A5, SO3, SM2, EM2-down, SDL4

Functional Description: Distrinet is an extension of Mininet that relies on LXC to be distributed in the cloud, and particularly in Amazon. The extension has been designed to be fully compatible with Mininet. As using Distrinet potentially involves the collaboration of multiple machines we focused on guaranteeing the correctness (in a sense that results are trustworthy) of simulations when running on multiple machines. To speedup deployments, loading and unloading operations have been parallelised with asynchronous calls. The pool of machines used for simulations is automatically provisioned thanks to Ansible.

Release Contributions: First release

URL: https://distrinet-emu.github.io

Publication: hal-03000617v1

Contact: Walid Dabbous

Participants: Damien Saucez, Giuseppe Di Lena, Andrea Tomassilli, Frédéric Giroire, Thierry Turletti

Partner: Orange Labs

7.1.5 OMNET-TSN

Name: OMNET-TSN

Keywords: Real time, Network simulator

Functional Description: Time Sensitive Networking (TSN) aims at providing real time capabilities to Ethernet networks. To achieve this goal, the 2018 revision of the IEEE802.1Q standard (i.e., IEEE802.1Q-2018) provides the ability to define transmission selection algorithms for queues in IEE802.1Q Ethernet switches.

We are developing an IEEE802.1Q-2018 simulation module for OMNeT++. The usual way to implement a simulation module is to abstract the components to be simulated and leverage the simulator functionalities. As a result, simulations can run very efficiently and scale well. The drawback is that the code base is only understandable by well trained developers and hardly understandable by engineers.

With our implementation we decided to make a direct transpose of the normative documents into the simulator such that anyone that reads the standards can read the code and adapt it if needed. This approach makes simulations runs last much longer than if they code were optimized for simulations but the major advantage is that engineers used to read standards and evaluate appliances provided by third parties can use it with minimal training.

The simulator code is implemented in C++ in Omnest (the commercial version of OMNeT++). As in most industrial environments it is extremely complex to install non corporate software, we implemented a front-end in javascript with the REACT framework. As a consequence, it is possible to install and run the simulation tool on a dedicated machine/VM/container and to drive simulations from any workstation solely by using a web browser.

Release Contributions:

- Multiple bug fixes in CBS
- CBS as a transmission selection algorithm instead of a transmission queue
- Complete tutorials for users
- Full code documentation for developper
- Delete message from QeuingFrame::handleMessage after it has been cloned to reduce memory footprint for large experiments
- Includes std algorithm library to compile on Linux std::set_intersection
- Add processing delay simulation
- Automated testing with docker

URL: https://github.com/dsaucez/

Contact: Damien Saucez

7.2 New platforms

Participants: Thierry Parmentelat, Thierry Turletti, Damien Saucez, Walid Dabbous.

7.2.1 Reproducible research Lab - R2lab

Scientific work around network protocols and related software stacks requires experiments, hence experimental conditions, to be reproducible. This is a particularly challenging requirement in the wireless networking area, where characteristics of wireless channels are known to be variable, unpredictable and hardly controllable.

The R2lab wireless testbed was designed with reproducibility as its central characteristics; it is built around an isolated and anechoic chamber, featuring RF absorbers that prevent radio waves reflections, and a Faraday cage blocking external interferences. R2lab thus provides an ideal environment for running reproducible wireless experiments.

R2lab has been operated since December 2015, in the context of the FIT (Future Internet of Things) Equipment of Excellence project, and as such, it is now federated with the other testbeds that are part of the FIT initiative. As of early 2019, it has been also federated within the Fed4Fire initiative.

Available toolsets, both hardware and software, are mostly stable apart from low noise marginal deployment of new kinds of radio devices, that now encompass among 5G and LoRa, among others. Our focus at this point of the project is to leverage our initial technical and financial investment, and to produce scientific work around reproducibility, particularly from a methodological standpoint, as illustrated by various publications listed in the R2lab web site.

Worth being mentioned as well, as part of a partnership with the OpenAirInterface initiative, R2lab is used on a daily basis for system-wide regression tests of the OAI stack, which in return allows us to offer up-to-date images for running OAI-based experiments. Emphasis has been put lately on offering tools that leverage kubernetes as the swiss-knife for orchestrating the deployment of a complete 5G infrastructure as an elastic set of microservices.

In 2021, the management tools on R2lab were extended to support running docker images, in addition to the historic, metal-based image format (.ndz). This allows experimenters to build their images using mainstream tools and base images. Thanks to that, we can now expose the latest OAI code as R2lab-ready docker images, both for EPC and RAN.

Access to R2lab is open 24/7. R2lab is used by more than 150 users half of them from France and the other half from all over the world (Australia, Belgium, Brazil, Canada, Chile, Spain Finland, Germany, India, Indonesia, Italy, Japan, Luxembourg, Netherlands Norway, Tunisia, Turkey, UK, US, Vietnam, etc.) to evaluate a wide range of wireless networking scenarios in realistic and reproducible environment. Examples are deployment of a 4G network in less than 4 minutes (demonstrated at Sigcomm 2017 [6]), Mininet-WiFi calibration (demon paper at Sigcomm 2016 [34]), Orientation Estimation, Joint range extension and localization, Mobile Edge Cloud scenarios, Ad-hoc routing protocols comparison, side channel analysis (CCS), QoE for Internet Video, Mobile video streaming leveraging in-network functions, Reproducible 5G network automation with Kube5G, Fidelity Monitoring of Distributed Network Emulation, Cross-layer Loss Discrimination Algorithms for MEC in 4G networks, Reproducing the OpenRF experiment (Reproducibility workshop, Sigcomm [36]). For more details see R2lab home page. As the future of R2lab, there are wider plans to rebuild the Inria Sophia facility more or less from scratch. As R2lab is still used by the community, the plan is to include R2lab as an addition the Sophia Node that is currently being developed in collaboration with the OpenAirInterface consortium at Eurecom.

7.2.2 SophiaNode: an open programmable 5G platform

Our project-team started in collaboration with Eurecom to deploy an open programmable platform to test 5G services. This year, R2lab was connected to sister site at Eurecom with 600 Gbps fibers forming together the so-called SophiaNode of the ESFRI SLICES-RI project. We plan to enrich R2lab with 5G professional radio units and with Kubernetes clustered server farm, and therefore to provide an experimental environment to test with open source (OAI, SrsLTE) software and some commercially licensed software

(e.g. Amarisoft) for 5G/6G networks supporting for example scenarios with disaggregated 5G networks elements. Initial deployment has already started in 2021 and will continue in 2022.

7.2.3 Network simulator for aircrafts

- Keywords: network, simulation, real-time
- Functional Description: In collaboration with Safran Electrical and Power we produced a network design tool for aircrafts. This tool simulates aircraft networks. The tool is about 10,000 lines of code, out of which we produced 2,000.
- Assessment: A-2up,SO-3,SM-2up,EM-4,SDL-3,OC-DA-CD-TPM
- Licence: confidential
- URL: confidential
- Contact: Damien Saucez

8 New results

8.1 Service Transparency

8.1.1 Leveraging Website Popularity Differences to Identify Performance Anomalies

Participants: Giulio Grassi, Chadi Barakat.

This study was funded by Inria within IPL BetterNet and was performed in collaboration with the EPI MIMOVE of Inria Paris and the Department of Computer Science of Boston University. Its focus is on the detection of performance anomalies across websites with a large span of popularities. In fact, Web performance anomalies (e.g. time periods when metrics like page load time are abnormally high) have significant impact on user experience and revenue of web service providers. Existing methods to automatically detect web performance anomalies focus on popular websites (e.g. with tens of thousands of visits per minute). However, across a more representative set of websites, passive measurement volume varies enormously, and some sites will only have small numbers of measurements per hour. Low rates of measurement creates gaps and noise that prevent the use of existing methods. This study develops WMF, a web performance anomaly detection method applicable across a range of websites with highly variable measurement volume. To demonstrate our method, we leverage data from a website monitoring company partner of the IPL BetterNet, which allows us to leverage cross-site measurements. WMF uses matrix factorization to mine patterns that emerge from a subset of the websites to "fill in" missing data on other websites. Our validation using both a controlled website and synthetic anomalies shows that WMF's F1-score is more than double that of the state-of-the-art method. We then apply WMF to three months of web performance measurements to shed light on performance anomalies across a variety of 125 small to medium websites. The results of this work were presented at the IEEE INFOCOM 2021 conference [24].

8.1.2 Unveiling the end-user viewport resolution from encrypted video traces

Participants: Othmane Belmoukadam, Chadi Barakat.

Video streaming is without doubt the most requested Internet service, and main source of pressure on the Internet infrastructure. At the same time, users are no longer satisfied by the Internet's best effort service, instead, they expect a seamless service of high quality from the side of the network. As result,

Internet Service Providers (ISP) engineer their traffic so as to improve their end-users' experience and avoid economic losses. Content providers from their side, and to enforce customers privacy, have shifted towards end-to-end encryption (e.g., TLS/SSL). Video streaming relies on the dynamic adaptive streaming over HTTP protocol (DASH) which takes into consideration the underlying network conditions (e.g., delay, loss rate, and throughput) and the viewport capacity (e.g., screen resolution) to improve the experience of the end user in the limit of the available network resources. In this work, we propose an experimental framework able to infer fine-grained video flow information such as chunk sizes from encrypted YouTube video traces. We also present a novel technique to separate video and audio chunks from encrypted traces based on Gaussian Mixture Models (GMM). Then, we leverage our dataset to train models able to predict the class of viewport (either SD or HD) per video session with an average 92% accuracy and 85% F1-score. The prediction of the exact viewport resolution is also possible but shows a lower accuracy than the viewport class. This work fits within the PhD thesis of Othmane Belmoukadam [26] and is published in IEEE Transactions on Network and Service Management [11].

8.1.3 Longitudinal Study of Exposure to Radio Frequencies at Population Scale

Participants: Yanis Boussad, Arnaud Legout, Walid Dabbous.

Evaluating population-scale exposure to the radio frequencies (RF) used in wireless telecommunication technologies is important for conducting sound epidemiological studies on the health impacts of these RF. Numerous studies have reported population exposure, but have used very small population samples. In this context, the real exposure of the population to RF remains subject to controversy. Here, to the best of our knowledge, we report the largest crowd-based measurement of population exposure to RF produced by cellular antennas, Wi-Fi access points, and Bluetooth devices for 254,410 unique users in 13 countries from January 2017 to December 2020. All measurements were obtained from the ElectroSmart Android app, and we applied a thorough methodology to clean and consolidate the measurements. We show that total exposure has been multiplied by 2.3 in the four-year period considered, with Wi-Fi as the largest contributor. The cellular exposure levels are orders of magnitude lower than the regulation limits and not significantly impacted by national regulation policies. Therefore, the mere comparison of exposure levels to regulation limits is a poor way to describe the real evolution of exposure. The population tends to be more exposed at home; for half of the study subjects, personal Wi-Fi routers and Bluetooth devices contributed to more than 50% of their total exposure. We make our dataset publicly available to provide a starting point for sound epidemiological studies on the health impacts of RF, and for other types of studies interested in population exposure to RF or the usage of wireless communication technologies. This work is under submission [29].

8.1.4 Evaluating smartphone performance for cellular power measurement

Participants: Yanis Boussad, Mondi Ravi, Arnaud Legout, Walid Dabbous.

Smartphones are affordable devices nowadays, capable of embedding a large variety of sensors such as magnetometers or orientation sensors, but also the hardware needed to connect them to most wireless communication technologies such as Wi-Fi, Bluetooth, or cellular networks. Therefore, they are handy devices able to perform received signal strength indicator RSSI measurements for a wide variety of applications such as cellular coverage maps, indoor localization, or proximity tracking. However, to the best of our knowledge, the accuracy of such measurements has never been rigorously assessed. The goals of this article are to assess the accuracy of the RSSI measurements made with a commercial off-the-shelf smartphone in a variety of conditions and how possible inaccuracies can be corrected. We primarily focus on the long-term evolution (LTE) RSSI, but we also extend our results to the Bluetooth RSSI. In this article, we build a controlled experimental setup based on commodity hardware and on open-source software. We evaluate the granularity and limitations of the Android application programming

interface that returns the RSSI. We explore how reliable the measurements in a controlled environment with a mono-polarized antenna are. We show that the orientation of the smartphone, the position or orientation of the source, and the transmission power have a significant impact on the accuracy of the measurements. We introduce several correction techniques based on radiation matrix manipulations and on machine learning in order to improve measurement accuracy to less than 5 dB Root Mean Square Error, when compared to a professional equipment. We also explore the reliability of measurements made in an outdoor realistic environment. We show that although transmission diversity available in LTE base stations significantly improves the measured RSSI regardless of the smartphone orientation, the Bluetooth RSSI remains largely sensitive to the smartphone orientation. The complete work is published in Transactions on Instrumentation and Measurements [13].

8.1.5 Did I delete my cookies? Cookies respawning with browser fingerprinting

Participants: Imane Fouad, Arnaud Legout.

Stateful and stateless web tracking gathered much attention in the last decade, however they were always measured separately. To the best of our knowledge, our study is the first to detect and measure cookie respawning with browser and machine fingerprinting. We develop a detection methodology that allows us to detect cookies dependency on browser and machine features. Our results show that 1,150 out of the top 30, 000 Alexa websites deploy this tracking mechanism. We further uncover how domains collaborate to respawn cookies through fingerprinting. We find out that this technique can be used to track users across websites even when third-party cookies are deprecated. Together with a legal scholar, we conclude that cookie respawning with browser fingerprinting lacks legal interpretation under the GDPR and the ePrivacy directive, but its use in practice may breach them, thus subjecting it to fines up to 20 million. This work is under submission [31]. This study was performed in collaboration with Nataliia Bielova from the EPI PRIVATICS of Inria Grenoble Rhône-Alpes.

8.1.6 In-Depth Technical and Legal Analysis of Tracking on Health Related Websites with ERNIE Extension

Participants: Imane Fouad, Yanis Boussad, Arnaud Legout.

Searching the Web to find doctors and make appointments online is a common practice nowadays. However, simply visiting a doctors website might disclose health related information. As the GDPR only allows processing of health data with explicit user consent, health related websites must ask consent before any data processing, in particular when they embed third party trackers. Admittedly, it is very hard for owners of such websites to both detect the complex tracking practices that exist today and to ensure legal compliance. In this research work, ee proposed Ernie, a browser extension we designed to visualise six state-of-the-art tracking techniques based on cookies. Using Ernie, we analysed 385 health related websites that users would visit when searching for doctors in Germany, Austria, France, Belgium, and Ireland. More specifically, we explored the tracking behavior before any interaction with the consent pop-up and after rejection of cookies on websites of doctors, hospitals, and health related online phone-books. We found that at least one form of tracking occurs on 62% of the websites before interacting with the consent pop-up, and 15% of websites include tracking after rejection. Finally, we performed a detailed technical and legal analysis of three health related websites that demonstrate impactful legal violations. This paper shows that while, from a legal point of view, health related websites are more privacy-sensitive than other kinds of websites, they are exposed to the same technical difficulties to implement a legally compliant website. We believe Ernie, the browser extension we developed, to be an invaluable tool for policy-makers and regulators to improve detection and visualization of the complex tracking techniques used on these websites. The results are published in the Workshop on Privacy in the Electronic Society [25] and the software is described in [33].

8.2 Open Network Architecture

8.2.1 On Accounting for Screen Resolution in Adaptive Video Streaming: QoE driven bandwidth sharing framework

Participants: Othmane Belmoukadam, Chadi Barakat.

Screen resolution along with network conditions are main objective factors impacting the user experience, in particular for video streaming applications. User terminals on their side feature more and more advanced characteristics resulting in different network requirements for good visual experience. Previous studies tried to link MOS (Mean Opinion Score) to video bitrate for different screen types (e.g., Common Intermediate Format (CIF), Quarter Common Intermediate Format (QCIF), and High Definition (HD)). We leverage such studies and formulate a QoE driven resource allocation problem to pinpoint the optimal bandwidth allocation that maximizes the QoE over all users of a network service provider located behind the same bottleneck link, while accounting for the characteristics of the screens they use for video playout. For our optimization problem, QoE functions are built using curve fitting on datasets capturing the relationship between MOS, screen characteristics, and bandwidth requirements. We propose a simple heuristic based on Lagrangian relaxation and KKT (Karush Kuhn Tucker) conditions to efficiently solve the optimization problem. Our numerical simulations show that the proposed heuristic is able to increase overall QoE up to 20% compared to an allocation with a TCP look-alike strategy implementing max-min fairness. Results of this work are published in the Journal of Network Management [12].

8.2.2 QoE-driven Cache Placement for Adaptive Video Streaming: Minding the Viewport

Participants: Othmane Belmoukadam, Chadi Barakat.

To handle the increasing demand for video streaming, ISP's and service providers use edge servers to cache video content to reduce the rush on their servers, balance the load between them and over the network, and smooth out the traffic variability. The dynamic adaptive streaming over HTTP protocol (DASH) makes videos available in multiple representations, and end-users can switch video resolution as a function of their network conditions and terminal display capacity (e.g., bandwidth, screen resolution). In this context, we study a viewport-aware caching optimization problem for dynamic adaptive video streaming that appropriately considers the client viewport size and access speed, the join time, and the characteristics of videos. We formulate and study the proposed optimization problem as an Integer Linear Program (ILP) that balances minimal join time and maximal visual experience, subject to the cache storage capacity. Our framework sheds light on optimal caching performance. Our proposed heuristic provides guidelines on the videos, and the representations of each video, to cache based on the video popularity, its encoding information, and the distribution of end-user display capacity and access speed in a way to maximize the overall end-user QoE. This work fits within the PhD thesis of Othmane Belmoukadam [26] and was first published in [19].

8.2.3 Machine Learning for Next-Generation Intelligent Transportation Systems

Participants: Chadi Barakat, Thierry Turletti.

Intelligent Transportation Systems, or ITS for short, includes a variety of services and applications such as road traffic management, traveler information systems, public transit system management, and autonomous vehicles, to name a few. It is expected that ITS will be an integral part of urban planning and future cities as it will contribute to improved road and traffic safety, transportation and transit efficiency, as well as to increased energy efficiency and reduced environmental pollution. On the other

hand, ITS poses a variety of challenges due to its scalability and diverse quality-of-service needs, as well as the massive amounts of data it will generate. In the Transactions on emerging telecommunications technologies journal, we published a survey that explores the use of Machine Learning (ML), which has recently gained significant traction, to enable ITS. We provided a comprehensive survey of the current state-of-the-art of how ML technology has been applied to a broad range of ITS applications and services, such as cooperative driving and road hazard warning, and identify future directions for how ITS can use and benefit from ML technology [16].

8.2.4 Dynamic Controller Assignment in Software Defined Internet of Vehicles through Multi-Agent Deep Reinforcement Learning

Participants: Chadi Barakat, Thierry Turletti.

We introduced a novel dynamic controller assignment algorithm targeting connected vehicle services and applications, also known as Internet of Vehicles (IoV). The proposed approach considers a hierarchically distributed control plane, decoupled from the data plane, and uses vehicle location and control traffic load to perform controller assignment dynamically. We modeled the dynamic controller assignment problem as a multi-agent Markov game and solve it with cooperative multi-agent deep reinforcement learning. Simulation results using real-world vehicle mobility traces showed that the proposed approach outperforms existing ones by reducing control delay as well as packet loss. This work has been published in the IEEE Transactions on Network and Service Management journal [17].

8.2.5 Harnessing UAVs for Fair 5G Bandwidth Allocation in Vehicular Communication via Deep Reinforcement Learning

Participants: Chadi Barakat, Thierry Turletti.

Terrestrial wireless infrastructure-based networks do not always guarantee that their resources will be shared uniformly by nodes in vehicular networks mostly due to the uneven and dynamic distribution of vehicles in the network as well as path loss effects. In this wor we leveraged multiple fifth-generation (5G) unmanned aerial vehicles (UAVs) to enhance network resource allocation among vehicles by positioning UAVs on-demand as "flying communication infrastructure". We proposed a deep reinforcement learning (DRL) approach to determine the position of UAVs in order to improve the fairness and efficiency of network resource allocation while considering the UAVs' flying range, communication range, and limited energy resources. We used a parametric fairness function for resource allocation that can be tuned to reach different allocation objectives ranging from maximizing the total throughput of vehicles, maximizing minimum throughput, as well as achieving proportional bandwidth allocation. Simulation results showed that the proposed DRL approach to UAV positioning can help improve network resource allocation according to the targeted fairness objective. These results are published in the IEEE Transactions on Network and Service Management journal [18].

8.2.6 Cross-layer Loss Discrimination Algorithms for MEC in 4G networks

Participants: Mamoutou Diarra, Thierry Turletti, Walid Dabbous.

Traditional loss-based Congestion Control Algorithms (CCAs) suffer from performance issues over wireless networks mostly due to their inability to distinguish wireless random losses from congestion losses. Different loss discrimination algorithms have been proposed to tackle this issue but they are not efficient for 4G networks since they do not consider the impact of various link layer mechanisms such as adaptive modulation and coding and retransmission techniques on congestion in LTE Radio

Access Networks (RANs). We proposed MELD (MEC-based Edge Loss Discrimination), a novel server-side loss discrimination mechanism that leverages recent advancements in Multi-access Edge Computing (MEC) services to discriminate packet losses based on real-time RAN statistics. Our approach collects the relevant radio information via MEC's Radio Network Information Service and uses it to correctly distinguish random losses from congestion losses. Our experimental study made with the QUIC transport protocol shows over 80% higher goodput when MELD is used with NewReno and 8% higher goodput when used with Cubic. This work has been presented at the IEEE International Conference on High Performance Switching and Routing (HPSR) conference in June 2021 at Paris [21].

8.2.7 RAN-aware Proxy-based Flow Control for High Throughput and Low Delay eMBB

Participants: Mamoutou Diarra, Thierry Turletti, Walid Dabbous.

5G enhanced Mobile broadband (eMBB) aims to provide users with a peak data rate of 20 Gbps in the Radio Access Network (RAN). However, since most Congestion Control Algorithms (CCAs) rely on startup and probe phases to discover the bottleneck bandwidth, they cannot quickly utilize the available RAN bandwidth and adapt to fast capacity changes without introducing large delay increase, especially when multiple flows are sharing the same Radio Link Control (RLC) buffer. To tackle this issue, we proposed RAPID, a RAN aware proxy-based flow control mechanism that prevents CCAs from overshooting more than the available RAN capacity while allowing near optimal link utilization. Based on analysis of up-todate radio information using Multi-access Edge Computing (MEC) services and packet arrival rates, RAPID is able to differentiate slow interactive flows from fast download flows and allocate the available bandwidth accordingly. Our experiments with concurrent Cubic and BBR flows show that RAPID can reduce delay increase by a factor of 10 to 50 in both Line-of-Sight (LOS) and Non-LOS (NLOS) conditions while preserving high throughput. To handle the ever growing demand of resource intensive experiments distributed, network emulation tools such as Mininet and Maxinet have been proposed. They automatically allocate experimental resources. In this work, we show that resources are poorly allocated, leading to resource overloading and hence to dubious experimental results. This is why we propose and implement a new placement module for distributed emulation. Our algorithms take into account both link and node resources and minimize the number of physical hosts needed to carry out the emulation. Through extensive numerical evaluations, simulations, and actual experiments , we show that our placement methods outperform existing ones and allowing to re-establish trust in experimental results. This work has been presented at the 24th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM) conference in November 2021 at Alicante, Spain [22].

8.2.8 Robust Programmable Networks

Participants: Giuseppe Di Lena, Damien Saucez, Thierry Turletti.

More than ever, data networks have demonstrated their central role in the world economy, but also in the well-being of humanity that needs fast and reliable networks. In parallel, with the emergence of NFV and SDN, efficient network algorithms considered too hard to be put in practice in the past now have a second chance to be considered again. In this context, as new networks will be deployed and current ones get significant upgrades, it is thus time to rethink the network dimensioning problem with protection against failures. In this work, we considered a path-based protection scheme with the global rerouting strategy in which, for each failure situation, there may be a new routing of all the demands. Our optimization task was to minimize the needed amount of bandwidth. We developed two scalable mathematical models that we handle using both Column Generation and Benders Decomposition techniques. Through extensive simulations on real-world IP network topologies and on randomly generated instances, we showed the effectiveness of our methods: they lead to savings of 40 to 48% of

the bandwidth to be installed in a network to protect against failures compared to traditional schemes. Our implementation in OpenDaylight demonstrates the feasibility of the approach. Its evaluation with Mininet shows that our solution provides sub-second recovery times, but the way it is implemented may greatly impact the amount of signaling traffic exchanged. In our evaluations, the recovery phase requires only a few tens of milliseconds for the fastest implementation, compared to a few hundreds of milliseconds for the slowest one. This work has been done in collaboration with the COATI team and has been published in the Computer Networks journal [15].

8.3 Experimental Evaluation

8.3.1 Time Sensitive Networks Evaluation

Participants: Damien Saucez.

Since he's back from Safran (October 2020), Damien Saucez has been working on developing a Time Sensitive Networking (TSN) simulation module for OMNeT++ [37]. This simulation module is designed for industrial usage and thus closely follows the standards. The idea here is that most potential industrial actors that would use the simulator have no real developper or networking simulation teams. Instead they have large pool of engineers used to read standards and evaluate appliances provided by third parties. Therefore it is essential to provide a solution that is as close as possible to what their day-to-day business is. As a consequence, some choices for the implementation of the simulation would appear inefficient for the ones used to write network simulators but they are not the target of the tool. Our focus on being usable directly by R&D teams in industries explains why exceptionally were are not using ns-3. In addition, as most corporate environments wouldn't allow to install specific development tools, we are working on providing a cloud based support for simulation that we call Simulation-as-a-Service. The backend is an Omnest (the commercial version of OMNeT++) engine while the frontend is purely javascript (REACT) such that simulations can be designed and analysed directly from a web browser. As OMNeT++ relies on a relatively old C++ technology writing a purely reactive javascript simulation environment for OMNeT++ is somehow challenging since nothing as really been thought for such usage.

8.3.2 Distrinet: Distributed Network Emulation

Participants: Giuseppe Di Lena, Damien Saucez, Thierry Turletti.

Networks have become complex systems that combine various concepts, techniques, and technologies. As a consequence, modelling or simulating them now is extremely complicated and researchers massively resort to prototyping techniques. Mininet is the most popular tool when it comes to evaluate SDN propositions. Mininet allows to emulate SDN networks on a single computer but shows its limitations with resource intensive experiments as the emulating host may become overloaded. To tackle this issue, we proposed Distrinet, a distributed implementation of Mininet over multiple hosts, based on LXD/LXC, Ansible, and VXLAN tunnels. Distrinet uses the same API than Mininet, meaning that it is compatible with Mininet programs. It is generic and can deploy experiments on Linux clusters (e.g., Grid'5000), as well as on the Amazon EC2 cloud platform. This tool and an evaluation of its performance has been published in the ACM Computer Communication Review (CCR) journal [14]. This work has been done in collaboration with the COATI team.

8.3.3 A Right Placement Makes a Happy Emulator: a Placement Module for Distributed SDN/NFV Emulation

Participants: Giuseppe Di Lena, Damien Saucez, Thierry Turletti.

To handle the ever growing demand of resource intensive experiments distributed, network emulation tools such as Mininet and Maxinet have been proposed. They automatically allocate experimental resources. In this work, we showed that resources are poorly allocated, leading to resource overloading and hence to dubious experimental results. This is why we proposed and implemented a new placement module for distributed emulation. Our algorithms take into account both link and node resources and minimize the number of physical hosts needed to carry out the emulation. Through extensive numerical evaluations, simulations, and actual experiments , we showed that our placement methods outperform existing ones and allowing to re-establish trust in experimental results. This work is reported in [30] and has been presented at the IEEE International Conference on Communications (ICC) conference in June 2021 at Montréal, Canada [20]. This work has been done in collaboration with the COATI team.

8.3.4 CloudTrace Demo: Tracing Cloud Network Delay

Participants: Giuseppe Di Lena, Damien Saucez, Thierry Turletti.

Many companies and organizations are moving their applications from on-premises data centers to the cloud. The cloud infrastructures can potentially provide an infinite amount of computation (e.g., Elastic Compute) and storage (e.g., Simple Service Storage). In addition, all cloud providers propose different offers: IaaS, PaaS, and SaaS. In June 2021, we presented a demo at the IEEE International Conference on Network Softwarization (NetSoft) virtual conference that focuses on the IaaS services. In a nutshell, we presented a simple tool to measure the network delay in a virtual infrastructure built entirely in the cloud. These measurements are useful for organizations that are moving current applications to, or creating new applications in, the cloud, but have requirements on the maximum, or average, network delay that these applications can tolerate. We showcased CloudTrace, a simple CLI tool that allows regional and multiregional experiments to measure delay, using Amazon AWS [32]. This work has been done in collaboration with the COATI team.

8.3.5 Passive Delay Measurement for Fidelity Monitoring of Distributed Network Emulation

Participants: Houssam Elbouanani, Chadi Barakat, Walid Dabbous, Thierry Turletti.

Emulation has become a popular approach for the validation and evaluation of network research. It provides researchers with a contained, customizable, and scalable testing environment, which can be easily packaged and published for potential readers to reproduce their results. However, as the network components are only virtual, emulation lacks the inherent realism of physical testbeds. In light of this, monitoring specific metrics of the emulated network has been proposed as a solution to mitigate to some degree inaccuracies caused by emulation. While this is not difficult to implement in a single-machine setting (e.g. with Mininet), monitoring is limited by the lack of time synchronization in scenarios where the emulation is distributed over multiple physical machines (e.g., Distrinet). In this work which fits within the PhD thesis of Houssam ElBouanani, we tackle the case of packet delay monitoring, to which we propose a methodology for passively measuring one-way delays with underlying assumptions about time synchronization, and round-trip delays otherwise. For an efficient implementation of our methodology, we propose an eBPF-based packet measurement tool that performs better than current packet sniffers under emulation-specific assumptions. We implement and evaluate our system in an open testbed and show that it can reach results within few microseconds of perfect accuracy and precision. This work was published in [23] and is currently under further development for its usage for network experimentation fidelity monitoring.

9 Bilateral contracts and grants with industry

Participants: Chadi Barakat, Walid Dabbous, Mamoutou Diarra, Arnaud Legout, Giuseppe Di Lena, Bernard Tamba Sandouno, Damien Saucez, Thierry Turletti.

9.1 Bilateral contracts with industry

Collaboration with Ekinops
Participants: Walid Dabbous, Mamoutou Diarra, Thierry Turletti.

We have started a collaboration with Ekinops on the topic of Multi-access Edge Computing. The activity started with a CIFRE thesis. The PhD student Mamoutou Diarra started his PhD on this topic on November 2019. Currently, he is working on efficient congestion control mechanisms for 5G scenarios in mutli-access edge environments.

Collaboration with Orange Labs
Participant: Giuseppe Di Lena, Damien Saucez, Thierry Turletti.

We have a collaboration with Orange Labs on the topic of Network Function Virtualization. The activity includes the CIFRE PhD thesis of Giuseppe Di Lena that started his PhD on resilient NFV/SDN environments on April 2018 co-supervised by Thierry Turletti and Frédéric Giroire from the COATI project-team.

Collaboration with YDATA
Participant: Chadi Barakat, Walid Dabbous Bernard Tamba Sandouno, Thierry Turletti.

We have started a collaboration with YDATA on the topic of geolocation assessment of mobile network performance. The activity started with a CIFRE thesis. The PhD student Bernard Tamba Sandouno started his PhD on this topic on May 2021. Currently, he is working on developing a tool to predict the geolocation performance of mobile broadband networks using a ray tracing approach coupled with wireless channel propagation models.

9.2 Bilateral grants with industry

QWANT
Participant: Arnaud Legout.
The PIA ANSWER project is led 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 QWANT 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 in 2018, a funding for a 3 years Ph.D. working on Web tracking technologies and privacy protection.

10 Partnerships and cooperations

Participants: Chadi Barakat, Walid Dabbous, Thierry Parmentelat, Damien Saucez, Thierry Turletti.

10.1 International initiatives

10.1.1 Associate Teams in the framework of an Inria International Lab or in the framework of an Inria International Program

DrIVE

Title: DrIVE: Distributed Intelligent Vehicular Environment - Enabling ITS through programmable networks

Duration: 2018 -> 2021

Coordinator: Christian Esteve Rothenberg (chesteve@dca.fee.unicamp.br) , Mateus Augusto Silva Santos (mateus.santos@ericsson.com) , Katia Obraczka (katia@soe.ucsc.edu)

Partners:

- Universidade Estadual de Campinas , Ericsson Research, Indaiatuba-SP, BRAZIL , University of California Santa Cruz

Inria contact: Thierry Turletti

Summary: Transportation systems are part of our society's critical infrastructure and are expected to experience transformative changes as the Internet revolution unfolds. The automotive industry is a notable example: it has been undergoing disruptive transformations as vehicles transition from traditional unassisted driving to fully automated driving, and eventually to the self-driving model. Communication technology advancements such as 5G and support for vehicle-to-infrastructure (V2I) and vehicle-to-vehicle (V2V) communication have been one of the key enablers of next-generation transportation services, also known as Intelligent Transport Systems (ITS). However, ITS services and applications pose significant challenges to the underlying communication and network infrastructure due to their stringent low latency, reliability, scalability, and geographic decentralization requirements. The DrIVE associated team proposal aims at addressing such challenges by:

- developing a programmable network control plane that will dynamically adjust to current environment conditions and network characteristics to support ITS' scalability, quality of service (QoS), and decentralization requirements, and
- applying the proposed distributed network control plane framework to ITS services and applications, such as road hazard warning, autonomous- and self-driving vehicles, and passenger-centric services (e.g., infotainment and video streaming).

10.2 International research visitors

10.2.1 Visits of international scientists

Inria International Chair With our great pleasure, Professor Katia Obraczka has been awarded an Inria International Chair for the 2021-2025 period. The research program entitled: "Smart Networks: When Computational Intelligence and Networking Meet" will explore synergies between machine learning (ML) and networking domains by addressing two complementary goals: (1) As users of ML, explore ML techniques as they apply to networks, their protocols, and their services towards intelligent network systems that are simple, automatically adapt to current conditions, do not require human in the loop, and deliver adequate end-to-end performance; and (2) Investigate how networking can benefit and empower ML. More specifically, as computing and storage shift to the edge of the network, we are interested in exploring how ML can leverage the network as a distributed, decentralized computing resource to handle the enormous amounts of data generated at the edge by providing full decentralization without compromising accuracy and preserving end user privacy.

All the administrative aspects were settled and the first visit of Prof Obraczka is expected in June 2022.

10.3 European initiatives

10.3.1 FP7 & H2020 projects

Fed4FIRE+

Title: Federation for FIRE Plus

Duration: Jan 2017 - Jun 2022

Coordinator: iMinds - Belgium

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

Inria contact: Thierry Parmentelat

Summary: 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 project's 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.

10.3.2 Other European programs/initiatives

ESFRI, the European Strategy Forum on Research Infrastructures, is a strategic instrument to develop the scientific integration of Europe and to strengthen its international outreach. Our project team hasbeen involved in the last year in the preparation of a large-scale infrastructure projet called SLICES-RI (Scientific LargeScale Infrastructure for Computing/Communication Experimental Studies-Research Infrastructure) that was finally accepted on the ESFRI roadmap for funding. Consortium members of this long term project are eligible to specific calls to design, bootstrap, deploy and operate a european wide flexible platform designed to support large-scale, experimental research focused on networking protocols, radio technologies, services, data collection, parallel and distributed computing and in particular cloud and edge-based computing architectures and services. The DIANA project-team is involved in the main SLICES-RI and two companion projects: SLICES-DS focused on the Design Study of the Research Infrastructure) and SLICES-SC focused on bootstrapping the users community for the RI. Here follows more detail on SLICES-SC.

SLICES-SC

Title: Scientific Large-scale Infrastructure for Computing/Communication Experimental Studies– Starting Community

Duration: Mar 2021 - Feb 2024

Coordinator: Sorbonne Université - France

Partners: UTH, Mandat International, PSNC, IMDEA, CNR, EURECOM, COSM, IoTLab, University of Oulu, INRIA, Imec, SZTAKI, TUM.

Inria contact: Walid Dabbous

Summary: Today we are experiencing the digital transformation happening with an unprecedented pace, with the community constantly researching on new solutions to support this transformation with ample computational power and connectivity. Towards addressing such research efforts, Research Infrastructure (RI) specific to addressing Digital Sciences research efforts have been deployed worldwide, towards trying to address key aspects contrary to off-the-shelf commercial infrastructure: 1) Full control over the parameters of an experiment, 2) Repeatable experiments regardless of the physical infrastructure, 3) Valid experimental results, which are easy to cross-reference and replicate. As such, several RIs have emerged, offering experimentation services with bleeding edge resources, that otherwise are only offered only in industrial R&D laboratories, with limited functionality. Towards combating these issues, SLICES Research Infrastructure is about to be deployed, aiming to provide high quality experimentation services with emerging technologies around the area of digital sciences (5G/6G, NFV, IoT and Cloud Computing), in an Internet-scale setup. With SLICES-SC, we aspire to foster the community of researchers around this ecosystem, create and strengthen necessary links with relevant industrial stakeholders for the exploitation of the infrastructure, advance existing methods for research reproducibility and experiment repeatability, and design and deploy the necessary solutions for providing SLICES-RI with an easy to access scheme for users from different disciplines. A set of detailed research activities has been designed to materialize these efforts in tools for providing transnational (remote and physical) access to the facility, as well as virtual access to the data produced over the facilities. The respective networking activities of the project aspire in fostering the community around these infrastructures, as well as open up to new disciplines and industrial stakeholders.

10.4 National initiatives

- **ANR FIT** (2011-2021): 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. The project ended in 2020. See also the FIT Web site.

- **ANR BOTTLENET** (2016-2021): 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. All this is meant to assist users, network and service operators as well as regulators in understanding Internet QoE and the sources of performance bottleneck. This national project got extended until March 2021 because of the COVID pandemic, and was coordinated, starting from November 2020, by Chadi Barakat from the Diana team.

10.5 Regional initiatives

During 2021, in the context of the CPER PUV project in collaboration with Eurecom and the LEAT, we have started to deploy the first servers and switches of the SophiaNode. Our goal is to provide an open prgrammable 5G infrastructure and will serve as blueprint for advanced next generation wireless nodes in the context of the SLICES-RI project. The deployment continues in 2022 with funding from the SLICES projects.

11 Dissemination

Participants: Chadi Barakat, Walid Dabbous, Arnaud Legout, Thierry Parmentelat, Damien Saucez, Thierry Turletti.

11.1 Promoting scientific activities

Chadi Barakat is on the editorial board of the Computer Networks journal, and was/is on the Technical Program Committee for the Passive and Active Measurement conference (PAM 2022), the International Teletraffic Congress (ITC 2021), the Network Traffic Measurement and Analysis Conference (TMA 2021), the International Conference on Network and Service Management (CNSM 2021), the first International Workshop on Machine Learning in Networking (MaLeNe 2021), and the Mediterranean Communication and Computer Networking Conference (MedComNet 2021 and 2022). He is currently in charge of international affairs at Inria Sophia Antipolis and is member of the COST-GTRI of Inria. He gave a keynote at the Inria Workshop on Systems (WOS'21) titled "Bridging the gap between network measurements and quality of experience: the video streaming case, Inria Rennes, October 2021.

Walid Dabbous is Director of the Academy of Excellence NIDS (Networks, Information and Digital Society). He is also member of the scientific committee of the DS4H Graduate school and member of the Ubinet International Master program steering committee. In 2021, he served in an High Council for Evaluation of Research and Higher Education (HCERES) review committee.

Arnaud Legout is member of the scienfic board of the Regalia project whose aim is to build a software environment for testing and regulation support to deal with the risks of bias and disloyalty generated by the algorithms of digital platforms. He is also in the technical program committee and editorial board of the Proceedings on Privacy Enhancing Technologies (PoPETs) in 2021 and 2022.

Thierry Turletti was in the program committees of the following workshops and conferences: 41st IEEE ICDCS 2021 on Next-Generation Mobile Networking and Computing (NG Mobile 2021), July 7-10, 2021, fully virtual; 12th Workshop on ns-3, June 23rd 2021, fully virtual; 23rd Algotel Conference, May 31 – June 4 2021, La Rochelle, France, 2021; IEEE International Conference on Communications (ICC'21), June 14-18 2021, Montreal, Canada; and IFIP/IEEE International Symposium on Integrated Network Management (IM'21), 17-21 May 2021, Bordeaux, France. He is still on the editorial boards of the "Wireless Networks" journal published by Springer Science and of the "Advances in Multimedia" Journal published by Hindawi Publishing Corporation. Since June 2020, he became a member of the Comité de Suivi Doctoral (CSD) at INRIA Sophia Antipolis Méditerranée.

11.2 Teaching - Supervision - Juries

11.2.1 Teaching

- Master 2 Ubinet: Chadi Barakat and Walid Dabbous, Evolving Internet, 31.5 hours, M2, Université Côte d'Azur, France.
- Master 2 Ubinet: Chadi Barakat and Walid Dabbous, Internet Measurements and New Architectures, 31.5 hours, M2, Université Côte d'Azur, France.
- Master 1 in Computer Science: Chadi Barakat, Computer Networks, 30 hours, M1, Université Côte d'Azur, France.

- Master 2 Estel: Chadi Barakat, Voice over IP, 9 hours, Université Côte d'Azur, France.
- Master Ubinet: Arnaud Legout, From BitTorrent to Privacy, 22.5 hours, M2, Université Côte d'Azur, France.
- Master module AWARE (Awarness-Raising to research) : Arnaud Legout, lecture 2 hours, Eurecom, France.
- Thierry Parmentelat helps coordinating the CS courses for the 1st year students of École des Mines de Paris, covering general topics like Numerical Programming with Python, Advanced Python Programming, and Introduction to Web technologies.
- **E-learning**
  - Python: Arnaud Legout and Thierry Parmentelat are co-authors of the MOOC: "Python 3 : des fondamentaux aux concepts avancés du langage". Since its creation in 2014, this MOOC has reached a consolidated number of 111 907 registered students and 10948 attestations. This MOOC is in French, as we felt like there was a gap in the previous offering, with the vast majority of teaching resources available in English only.

11.2.2 Supervision

PhD students

- PhD defended: Othmane Belmoukadam finished his PhD on "QoE aware content management in the Internet caching and transport" in September 2021 [26]. His PhD was supervised by Chadi Barakat and funded by the doctoral school EDSTIC of Université Côte d'Azur.
- PhD defended: Yanis Boussad finished his PhD on "Large scale characterization of the exposition to microwaves" in September 2021 [27]. His PhD was co-supervised by Arnaud Legout and Leonardo Lizzi from LEAT and funded by the DS4H Graduate school of Université Côte d'Azur.
- PhD defended: Giuseppe Di Lena finished his PhD on "Building a resilience methodology for NFV/SDN" in March 2021 [28]. His PhD was co-supervised by Thierry Turletti, Damien Saucez and Frédéric Giroire from the Coati project-team and Chidung Lac from Orange Labs. and funded by a CIFRE grant in collaboration with Orange Labs.
- PhD defended: Imane Fouad finished her PhD on "Web tracking technologies and privacy protection" in June 2021 [35]. Her PhD was co-supervised by Arnaud Legout and and Nataliia Bielova from the PRIVATICS project-team and funded by a grant from QWANT.
- PhD in progress: Bernard Tamba Sandouno works on a "Geolocation assessment model of mobile network performance". His PhD is co-supervised by Chadi Barakat, Thierry Turletti and Walid Dabbous from the Diana team, and by Yamen Alsaba from YDATA. His thesis is funded by a CIFRE grant in collaboration with YDATA.
- PhD in progress: Mamoutou Diarra works on "Multi-access Edge Computing". His PhD is co-supervised by Thierry Turletti and Walid Dabbous from the Diana team, and by Amine Ismail from Ekinops. His thesis is funded by a CIFRE grant in collaboration with Ekinops.
- PhD in progress: Houssam Elbouanani works on "Experiment control for reproducible research". His PhD is co-supervised by Chadi Barakat, Thierry Turletti and Walid Dabbous from the Diana team and funded by the Fed4Fire+ H2020 project.
- PhD in progress: Imane Taibi works on "Big data analysis for network monitoring and troubleshooting". She is co-supervised by Gerardo Rubino, Yassine Hadjadj-Aoul from the Dionysos project-team and Chadi Barakat.

Interns

- Anas El Bouzidi:
  - Master: Ubinet
  - Institution: Université Côte D'Azur
  - Date: Mar 2021 - Aug 2021
  - Subject: QoE-aware bandwidth sharing for video streaming traffic
  - Supervisor: Chadi Barakat
- Rostyslav Khudov:
  - Master: Ubinet
  - Institution: Université Côte D'Azur
  - Date: November 2021 - December 2021
  - Subject: 7- How Wi-Fi SSIDs and Bluetooth device names reveal political opinions
  - Supervisor: Arnaud Legout
- Stéphane Tolly:
  - Master: Ubinet
  - Institution: Université Côte D'Azur
  - Date: November 2021 - December 2021
  - Subject: Could we get rid of 5G with wifi ?
  - Supervisor: Arnaud Legout, Damien Saucez
- Mariella Al Jreidy:
  - Master: Ubinet
  - Institution: Université Côte D'Azur
  - Date: November 2021 - December 2021
  - Subject: Leveraging the wealth of data available in the browser for network monitoring and troubleshooting
  - Supervisor: Chadi Barakat
- Viacheslav Herasymov:
  - Master: Ubinet
  - Institution: Université Côte D'Azur
  - Date: November 2021 - December 2021
  - Subject: Efficient Monitoring Service for 5G Wireless Access Networks
  - Supervisor: Walid Dabbous, Thierry Turletti

Apprentices

- Téo Haÿs
  - Master 1 Informatique, Parcours Informatique et Interactions
  - Institution: Université Côte D'Azur
  - Date: September 2021 - September 2022
  - Subject: Tools for reproducible experimentation on R2Lab
  - Supervisor: Thierry Parmentelat

11.2.3 Juries

- Chadi Barakat served as examiner of Duncan Deveaux's PhD thesis, "On the Networking of Knowledge in Vehicular Networks", defended in December 2021, at Eurecom.
- Chadi Barakat served as examiner and president of Ishani Sarkar's PhD thesis, "Data mining and optimisation of a multi-content hybrid P2P/Web distribution network", defended in November 2021, at I3S and Université Côte d'Azur.
- Chadi Barakat served as examiner of Flavia Salutari's PhD thesis, "Longitudinal, large-scale and unbiased Internet measurements", defended in September 2021, at Telecom Paris.
- Chadi Barakat served as reviewer of Loick Bonniot's PhD thesis, "Computer Network Modeling and Root Cause Analysis with Statistical Learning", defended in June 2021, at Université de Rennes 1.
- Chadi Barakat served as jury member for the mid-term review of Raza Ul Mustafa's PhD thesis, "Machine Learning Assisted DASH QoE Inference Through Network Level QoS Feature", University of Campinas, Brazil, held in October 2021.
- Chadi Barakat served as jury member for the mid-term review of Zhejiayu Ma's PhD thesis, "Machine learning in the context of WebRTC-based multimedia diffusion", EasyBroadcast and Université Côte d'Azur, held in May 2021.
- Walid Dabbous served as mentor and jury member for the HDR of Nataliia Bielova on "Protecting Privacy of Web Users" defended on June 7, 2021 at Université Côte d'Azur.
- Walid Dabbous served as reviewer and jury member for the HDR of Hicham Lakhlef on "Towards scalable and secure control in wireless communicating things" defended on December 3, 2021 at University of Technology, Compiègne.
- Walid Dabbous served jury president for the PhD of Luca Santamaria on "Integrated Reconfigurable Antenna Systems for Dynamic Wireless IoT Networks" defended on December 13, 2021 at Université Côte d'Azur.
- Thierry Turletti served as reviewer of Anouar Rkhami's PhD thesis "From services placement to services monitoring in 5G and post-5G networks", defended on December 14, 2021 at Université de Rennes 1.
- Thierry Turletti served as examiner of Adrien Thibaud's PhD thesis "Répartition de flux dans les réseaux de contenu, application à un contexte satellite", defended on September 2, 2021 at INP Toulouse.
- Thierry Turletti served as reviewer of Adrien Wion's PhD thesis "Control Plane in Dynamic Software Networks, defended on March 5, 2021 at Telecom Paris.
- Thierry Turletti served as examiner and president of El Fadel Bonfoh's PhD thesis "VTL : Une Architecture Stable pour la Conception, l'Implémentation, et le Déploiement de Protocoles de Communication d'Internet", defended on January 26, 2021 at INSA Toulouse.
- Thierry Turletti served as reviewer of Salma Matoussi's PhD thesis "User-Centric Slicing with Functional Splits in 5G Cloud-RAN", defended on January 22, 2021 at Sorbonne Université.

11.3 Popularization

11.3.1 Internal or external Inria responsibilities

Damien Saucez is chargé de mission médiation scientifique interne. The role is to promote and favour scientific exchanges between Sophia's center researchers but also to popularise sciences within the center, not only for the researchers but for all personnel of the center.

The two main activities that are put in place are the so-called In'tro and Café In.

The concept of In'Tro has been proposed by Fabien Gandon and we implemented it successfully in March 2021. In'Tro aim to promote recently hired researcher in order to foster new collaborations. Every month, a researcher is allocated a slot of 30 minutes during lunch time to presents her/his work during 15 minutes and to answer questions during 15 minutes. Due to the pandemics the 8 events of 2021 were purely virtual meetings. Presentations are made publicly available on VoD. On average we counted about 30 live attendees. We also observe that the videos are watched on demand after the presentation, not only by Inria people but also by anonymous viewers that received YouTube suggestions.

The Café-In events are different in their objective and organised since 2012. Once a month a one hour slot is dedicated to popularise sciences to Inria personel, regardless of whether they are scientific or not. After lunch, a researcher, or a panel of researchers, is invited to talk about a subject of his/her choice around a coffee. The main objective is not to foster new collaborations but to allow everyone at Inria to understand important research subjects that are work on by Inria researchers and to open their curiosity. Because of the constrains imposed by the pandemic it was not possible in 2021 to really talk around a coffee and even though the 6 talks that has been organised were of high quality, we noticed very low attendance levels compared to "normal years". Nevertheless, similarly to the In'Tros, talk have been recorded and are publicly available and views are increasing from month to month.

In parallel, we promoted the use of the center's zimbra calendar to advertise internal team seminars. As of now, we lack of indicators to determine if this calendar has boosted inter-team or cross disciplines collaborations.

For more information, check the MASTIC website.

11.3.2 Interventions

Damien Saucez made 4 Chiche! sessions this year. See the Chiche Web site.

Walid Dabbous presented R2lab and an introduction to electromagntic waves to young interns from different schools on december 16, 2021.

Walid Dabbous gave a presentaton to "Lycée" students in the context of "La fête de la Science" on October 4, 2021. The video was published on YouTube.

Bernard Tamba Sandounou presented his work on predict the performance of mobile networks using geometric methods and wireless channel propagation models in the "Nuit Européenne des chercheur.es" at Université Côte d'Azur on September 24, 2021.

12 Scientific production

12.1 Major publications

[1] O. Belmoukadam and C. Barakat. 'Unveiling the end-user viewport resolution from encrypted video traces'. In: *IEEE Transactions on Network and Service Management* 18.3 (Sept. 2021), pp. 3324–3335. DOI: 10.1109/TNSM.2021.3083070. URL: https://hal.inria.fr/hal-03230168.

[2] F. De Pellegrini, L. Maggi, A. Massaro, D. Saucez, J. Leguay and E. Altman. 'Blind, Adaptive and Robust Flow Segmentation in Datacenters'. In: *INFOCOM 2018 - IEEE International Conference on Computer Communications*. Honolulu, United States, Apr. 2018. URL: https://hal.inria.fr/hal-01666905.

[3] M. Flittner, M. N. Mahfoudi, D. Saucez, M. Wählisch, L. Iannone, V. Bajpai and A. Afanasyev. 'A Survey on Artifacts from CoNEXT, ICN, IMC, and SIGCOMM Conferences in 2017'. In: *Computer Communication Review* 48.1 (Apr. 2018), pp. 75–80. URL: https://hal.inria.fr/hal-01968401.

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

[5] M. J. Khokhar, N. A. Saber, T. Spetebroot and C. Barakat. 'An Intelligent Sampling Framework for Controlled Experimentation and QoE Modeling'. In: *Computer Networks* 147 (Dec. 2018), pp. 246–261. DOI: 10.1016/j.comnet.2018.10.011. URL: https://hal.inria.fr/hal-01906145.

[6] M. N. Mahfoudi, T. Parmentelat, T. Turletti, W. Dabbous and R. Knopp. *Deploy a 5G network in less than 5 minutes: Demo Abstract*. ACM SIGCOMM Posters and Demos. Poster. Aug. 2017. URL: https://hal.inria.fr/hal-01580065.

[7] M. N. Mahfoudi, G. Sivadoss, O. Bensouda Korachi, T. Turletti and W. Dabbous. 'Joint range extension and localization for LPWAN'. In: *Internet Technology Letters* (June 2019). DOI: 10.1002/itl2.120. URL: https://hal.archives-ouvertes.fr/hal-02170466.

[8] D. Saucez, L. Iannone, C. Albert and F. Coras. *Locator/ID Separation Protocol (LISP) Impact*. Internet Engineering Task Force (IETF), Request for Comments: 7834. Apr. 2016. URL: https://hal.inria.fr/hal-01423163.

[9] H. Soni, W. Dabbous, T. Turletti and H. Asaeda. 'NFV-based Scalable Guaranteed-Bandwidth Multicast Service for Software Defined ISP networks'. In: *IEEE Transactions on Network and Service Management* 14.4 (Dec. 2017), p. 14. DOI: 10.1109/TNSM.2017.2759167. URL: https://hal.inria.fr/hal-01596488.

[10] L. Vigneri, T. Spyropoulos and C. Barakat. 'Low Cost Video Streaming through Mobile Edge Caching: Modelling and Optimization'. In: *IEEE Transactions on Mobile Computing* (2018). DOI: 10.1109/TMC.2018.2861005. URL: https://hal.inria.fr/hal-01855304.

12.2 Publications of the year

International journals

[11] O. Belmoukadam and C. Barakat. 'Unveiling the end-user viewport resolution from encrypted video traces'. In: *IEEE Transactions on Network and Service Management* 18.3 (Sept. 2021), pp. 3324–3335. DOI: 10.1109/TNSM.2021.3083070. URL: https://hal.inria.fr/hal-03230168.

[12] O. Belmoukadam, M. J. Khokhar and C. Barakat. 'On Accounting for Screen Resolution in Adaptive Video Streaming: QoE driven bandwidth sharing framework'. In: *International Journal of Network Management* 31.1 (7th Jan. 2021), e2128. DOI: 10.1002/nem.2128. URL: https://hal.inria.fr/hal-02615576.

[13] Y. Boussad, M. N. Mahfoudi, A. Legout, L. Lizzi, F. Ferrero and W. Dabbous. 'Evaluating Smartphone Accuracy for RSSI Measurements'. In: *IEEE Transactions on Instrumentation and Measurement* 70 (5th Jan. 2021), pp. 1–12. DOI: 10.1109/tim.2020.3048776. URL: https://hal.inria.fr/hal-03063997.

[14] G. Di Lena, A. Tomassilli, D. Saucez, F. Giroire, T. Turletti and C. Lac. 'Distrinet: a Mininet Implementation for the Cloud'. In: *Computer Communication Review* 51.1 (31st Jan. 2021), pp. 2–9. DOI: 10.1145/3457175.3457177. URL: https://hal.inria.fr/hal-03000617.

[15] A. Tomassilli, G. Di Lena, F. Giroire, I. Tahiri, D. Saucez, S. Pérennes, T. Turletti, R. Sadykov, F. Vanderbeck and C. Lac. 'Design of Robust Programmable Networks with Bandwidth-optimal Failure Recovery Scheme'. In: *Computer Networks* 192.108043 (June 2021). DOI: 10.1016/j.comnet.2021.108043. URL: https://hal.inria.fr/hal-03441630.

[16] T. Yuan, W. B. Da Rocha Neto, C. E. Rothenberg, K. Obraczka, C. Barakat and T. Turletti. 'Machine Learning for Next-Generation Intelligent Transportation Systems: A Survey'. In: *Transactions on emerging telecommunications technologies* (2021). URL: https://hal.inria.fr/hal-02284820.

[17] T. Yuan, W. D. Rocha Neto, C. E. Rothenberg, K. Obraczka, C. Barakat and T. Turletti. 'Dynamic Controller Assignment in Software Defined Internet of Vehicles through Multi-Agent Deep Reinforcement Learning'. In: *IEEE Transactions on Network and Service Management* 18.1 (Mar. 2021), p. 12. DOI: 10.1109/TNSM.2020.3047765. URL: https://hal.inria.fr/hal-03000911.

[18] T. Yuan, C. E. Rothenberg, K. Obraczka, C. Barakat and T. Turletti. 'Harnessing UAVs for Fair 5G Bandwidth Allocation in Vehicular Communication via Deep Reinforcement Learning'. In: *IEEE Transactions on Network and Service Management* (2021). URL: https://hal.inria.fr/hal-03001383.

International peer-reviewed conferences

[19] O. Belmoukadam and C. Barakat. 'QoE-driven Cache Placement for Adaptive Video Streaming: Minding the Viewport'. In: MeditCom 2021 - IEEE International Mediterranean Conference on Communications and Networking. Athènes, Greece, 7th Sept. 2021. DOI: 10.1109/MeditCom49071.2021.9647613. URL: https://hal.inria.fr/hal-03320414.

[20] G. Di Lena, A. Tomassilli, F. Giroire, D. Saucez, T. Turletti and C. Lac. 'A Right Placement Makes a Happy Emulator: a Placement Module for Distributed SDN/NFV Emulation'. In: ICC 2021 - IEEE International Conference on Communications. Montréal, Canada: IEEE, 14th June 2021. DOI: 10.1109/ICC42927.2021.9500359. URL: https://hal.inria.fr/hal-03001913.

[21] M. Diarra, W. Dabbous, A. Ismail and T. Turletti. 'Cross-layer Loss Discrimination Algorithms for MEC in 4G networks'. In: IEEE International Conference on High Performance Switching and Routing (HPSR). Paris, France, 7th June 2021. URL: https://hal.inria.fr/hal-03363851.

[22] M. Diarra, W. Dabbous, A. Ismail and T. Turletti. 'RAN-aware Proxy-based Flow Control for High Throughput and Low Delay eMBB'. In: 24th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM). Alicante, Spain, 22nd Nov. 2021. URL: https://hal.inria.fr/hal-03363827.

[23] H. Elbouanani, C. Barakat, W. Dabbous and T. Turletti. 'Passive Delay Measurement for Fidelity Monitoring of Distributed Network Emulation'. In: MedComNet 2021 - 19th Mediterranean Communication and Computer Networking Conference. Virtual, France, 15th June 2021. DOI: 10.1109/MedComNet52149.2021.9501246. URL: https://hal.inria.fr/hal-03001876.

[24] G. Grassi, R. Teixeira, C. Barakat and M. Crovella. 'Leveraging Website Popularity Differences to Identify Performance Anomalies'. In: INFOCOM 2021 - IEEE International Conference on Computer Communications. Vancouver / Virtual, Canada, 10th May 2021. DOI: 10.1109/INFOCOM42981.2021.9488832. URL: https://hal.inria.fr/hal-03109717.

Conferences without proceedings

[25] V. Wesselkamp, I. Fouad, C. Santos, Y. Boussad, N. Bielova and A. Legout. 'In-Depth Technical and Legal Analysis of Tracking on Health Related Websites with ERNIE Extension'. In: 20th Workshop on Privacy in the Electronic Society. Seoul, South Korea, 15th Nov. 2021. URL: https://hal.archives-ouvertes.fr/hal-03241333.

Doctoral dissertations and habilitation theses

[26] O. Belmoukadam. 'QoE-aware content management in the Internet : caching and transport'. Université Côte d'Azur, 27th Sept. 2021. URL: https://tel.archives-ouvertes.fr/tel-03482558.

[27] Y. Boussad. 'Large-scale characterization of exposure to downlink radiofrequencies'. DIANA; Laboratoire d'Electronique, Antennes et Télécommunications; Université Côté d'Azur, 30th Sept. 2021. URL: https://hal.inria.fr/tel-03528597.

[28] G. Di Lena. 'Distributed and trustable SDN-NFV-enabled network emulation on testbeds and cloud infrastructures'. Université Côte d'Azur, 22nd Mar. 2021. URL: https://tel.archives-ouvertes.fr/tel-03343533.

Reports & preprints

[29] Y. Boussad, X. L. Chen, A. Legout, A. Chaintreau and W. Dabbous. *Longitudinal Study of Exposure to Radio Frequencies at Population Scale*. 1st Oct. 2021. URL: https://hal.inria.fr/hal-03361556.

[30] G. Di Lena, A. Tomassilli, F. Giroire, D. Saucez, T. Turletti and C. Lac. *Placement Module for Distributed SDN/NFV Network Emulation*. RR-9391. Inria Sophia Antipolis - Méditerranée; I3S, Université Côte d'Azur; Orange Labs R&D [Lannion] (France Télécom), 5th Feb. 2021, p. 32. URL: https://hal.inria.fr/hal-03132873.

[31] I. Fouad, C. Santos, A. Legout and N. Bielova. *Did I delete my cookies? Cookies respawning with browser fingerprinting*. 5th May 2021. URL: https://hal.archives-ouvertes.fr/hal-03218403.

Other scientific publications

[32] G. Di Lena, F. Giroire, T. Turletti and C. Lac. 'CloudTrace Demo: Tracing Cloud Network Delay'. In: IEEE International Conference on Network Softwarization (NetSoft). Fully Virtual, France: IEEE, 28th June 2021. DOI: 10.1109/NetSoft51509.2021.9492583. URL: https://hal.inria.fr/hal-03364025.

12.3 Other

Softwares

[33] [SW] V. Wesselkamp, I. Fouad, A. Legout and N. Bielova, *Ernie Extension*, 16th Nov. 2021. LIC: GNU General Public License v3.0 only. HAL: ⟨hal-03413044⟩, URL: https://hal.archives-ouvertes.fr/hal-03413044, VCS: https://github.com/vwesselkamp/ernie-extension, SWHID: ⟨swh:1:dir:6cf9d69e8467428628b93fd5e91b2c5f1684e7e7;origin=https://hal.archives-ouvertes.fr/hal-03413044;visit=swh:1:snp:69b5f385b8f53f754ddbc6174b9d6cefc869e8fa;anchor=swh:1:rev:55f6308a857dd97100e4b06a5d6647d5330eadaf;path=/⟩.

12.4 Cited publications

[34] R. d. R. Fontes and C. E. Rothenberg. 'Mininet-WiFi: A Platform for Hybrid Physical-Virtual Software-Defined Wireless Networking Research'. In: *Proceedings of the 2016 ACM SIGCOMM Conference*. Florianopolis, Brazil, 2016, pp. 607–608. URL: https://doi.org/10.1145/2934872.2959070.

[35] I. Fouad. 'Detection and measurement of web tracking'. Theses. Université Côte d'Azur, June 2021. URL: https://tel.archives-ouvertes.fr/tel-03278529.

[36] M. N. Mahfoudi, T. Turletti, T. Parmentelat and W. Dabbous. 'Lessons Learned while Trying to Reproduce the OpenRF Experiment'. In: *Reproducibility'17 - ACM SIGCOMM 2017 Reproducibility Workshop*. Vol. 41. 1. Los Angeles, United States, Aug. 2017, pp. 21–23. DOI: 10.1145/3097766.3097772. URL: https://hal.archives-ouvertes.fr/hal-01615398.

[37] A. Varga. 'OMNeT++'. In: *Modeling and tools for network simulation*. Springer, 2010, pp. 35–59.