Activity Report 2018

Project-Team MIMOVE

Middleware on the Move

RESEARCH CENTER
Paris

THEME
Distributed Systems and middleware

Table of contents

Project-Team MIMOVE

Creation of the Team: 2014 July 01, updated into Project-Team: 2018 February 01

Keywords:

Computer Science and Digital Science:

A1.2.1. - Dynamic reconfiguration
A1.2.3. - Routing
A1.2.4. - QoS, performance evaluation
A1.2.5. - Internet of things
A1.2.6. - Sensor networks
A1.2.7. - Cyber-physical systems
A1.3. - Distributed Systems
A1.4. - Ubiquitous Systems
A1.5. - Complex systems
A1.5.1. - Systems of systems
A1.5.2. - Communicating systems
A2.5. - Software engineering
A2.6.2. - Middleware
A3.1.7. - Open data
A3.1.8. - Big data (production, storage, transfer)
A3.3. - Data and knowledge analysis
A3.5. - Social networks

Other Research Topics and Application Domains:

B6.3. - Network functions
B6.4. - Internet of things
B6.5. - Information systems
B8.2. - Connected city
B8.5.1. - Participative democracy

1. Team, Visitors, External Collaborators

Research Scientists

Nikolaos Georgantas [Team leader, Inria, Researcher, HDR]
Vassilis Christophides [Inria, Advanced Research Position]
Renata Cruz Teixeira [Inria, & Stanford University from Sept 2018, Senior Researcher, HDR]
Valérie Issarny [Inria, & Inria@SiliconValley until Aug 2018, Senior Researcher, HDR]

Post-Doctoral Fellows

Sara Ayoubi [Inria]
Georgios Bouloukakis [Inria, & University of California, Irvine from Feb 2018]
Francesco Bronzino [Inria, until Oct 2018]
Giulio Grassi [Inria, from Oct 2018]
Bruno Lefèvre [Inria, until Mar 2018]

PhD Students

Fethi Dilmi [Inria]
Yifan Du [Inria]
Grigorios Piperagkas [Inria, from Oct 2018]
Sarah Wassermann [Inria]
Patient Ntumba [Inria, from Aug 2018]
Israel Marquez Salinas [Université Pierre et Marie Curie, from Oct 2018]

Technical staff

Rachit Agarwal [Inria]
Ehsan Ahvar [Inria, from Jul 2018]
Shohreh Ahvar [Inria, from Jul 2018]
Giulio Grassi [Inria, from Apr 2018 until Sep 2018]
Bruno Lefèvre [Inria, from Apr 2018]
Patient Ntumba [Inria, until Jul 2018]
Pierre-Guillaume Raverdy [Inria & Inria SED, until Mar 2018]

Interns

Israel Marquez Salinas [Inria, from Mar 2018 until Jul 2018]
Ivan Morandi [Inria, from Feb 2018 until Jul 2018]
Shaan Chopra [Inria, from May 2018 until Jul 2018]
Souha Sboui [Inria, from Apr 2018 until Oct 2018]

Administrative Assistant

Nathalie Gaudechoux [Inria]

Visiting Scientist

Mark Crovella [Boston University, from Aug 2018]

External Collaborators

Françoise Sailhan [Associate Professor, CNAM]
Rafael Angarita Arocha [Associate Professor, ISEP]
Francesco Bronzino [Nokia Bell Labs, from Nov 2018]
Zied Ben Houidi [Nokia Bell Labs]
Diego Neves Da Hora [Google Brazil]

2. Overall Objectives

2.1. Overall Objectives

Given the prevalence of global networking and computing infrastructures (such as the Internet and the Cloud), mobile networking environments, powerful hand-held user devices, and physical-world sensing and actuation devices, the possibilities of new mobile distributed systems have reached unprecedented levels. Such systems are dynamically composed of networked resources in the environment, which may span from the immediate neighborhood of the users – as advocated by pervasive computing – up to the entire globe – as envisioned by the Future Internet and one of its major constituents, the Internet of Things. Hence, we can now talk about truly ubiquitous computing.

The resulting ubiquitous systems have a number of unique – individually or in their combination – features, such as dynamicity due to volatile resources and user mobility, heterogeneity due to constituent resources developed and run independently, and context-dependence due to the highly changing characteristics of the execution environment, whether technical, physical or social. The latter two aspects are particularly manifested through the physical but also social sensing and actuation capabilities of mobile devices and their users. More specifically, leveraging the massive adoption of smart phones and other user-controlled mobile devices, besides physical sensing – where a device's sensor passively reports the sensed phenomena – *social sensing*/*crowd sensing* comes into play, where the user is aware of and indeed aids in the sensing of the environment.

Mobile systems with the above specifics further push certain problems related to the Internet and user experience to their extreme: (i) Technology is too complex. Most Internet users are not tech-savvy and hence cannot fix performance problems and anomalous network behavior by themselves. The complexity of most Internet applications makes it hard even for networking experts to fully diagnose and fix problems. Users can't even know whether they are getting the Internet performance that they are paying their providers for. (ii) There is too much content. The proliferation of user-generated content (produced anywhere with mobile devices and immediately published in social media) along with the vast amount of information produced by traditional media (e.g., newspapers, television, radio) poses new challenges in achieving an effective, near real-time information awareness and personalization. For instance, users need novel filtering and recommendation tools for helping them to decide which articles to read or which movie to watch.

This challenging context raises key research questions:

- How to deal with heterogeneity and dynamicity, which create runtime uncertainty, when developing and running mobile systems in the open and constantly evolving Internet and IoT environment?
- How to enable automated diagnosis and optimization of networks and systems in the Internet and IoT environment for improving the QoE of their users?
- How to raise human centric crowd-sensing to a reliable means of sensing world phenomena?
- How to deal with combination, analysis and privacy aspects of Web/social media and IoT crowd-sensing data streams?

3. Research Program

3.1. Introduction

The research questions identified above call for radically new ways in conceiving, developing and running mobile distributed systems. In response to this challenge, MiMove's research aims at enabling next-generation mobile distributed systems that are the focus of the following research topics.

3.2. Emergent mobile distributed systems

Uncertainty in the execution environment calls for designing mobile distributed systems that are able to run in a beforehand unknown, ever-changing context. Nevertheless, the complexity of such change cannot be tackled at system design-time. Emergent mobile distributed systems are systems which, due to their automated, dynamic, environment-dependent composition and execution, *emerge* in a possibly non-anticipated way and manifest *emergent properties*, i.e., both systems and their properties take their complete form only at runtime and may evolve afterwards. This contrasts with the typical software engineering process, where a system is finalized during its design phase. MiMove's research focuses on enabling the emergence of mobile distributed systems while assuring that their required properties are met. This objective builds upon pioneering research effort in the area of *emergent middleware* initiated by members of the team and collaborators [1], [3].

3.3. Large-scale mobile sensing and actuation

The extremely large scale and dynamicity expected in future mobile sensing and actuation systems lead to the clear need for algorithms and protocols for addressing the resulting challenges. More specifically, since connected devices will have the capability to sense physical phenomena, perform computations to arrive at decisions based on the sensed data, and drive actuation to change the environment, enabling proper coordination among them will be key to unlocking their true potential. Although similar challenges have been addressed in the domain of networked sensing, including by members of the team [7], the specific challenges arising from the *extremely large scale* of mobile devices – a great number of which will be attached to people, with uncontrolled mobility behavior – are expected to require a significant rethink in this domain. MiMove's research investigates techniques for efficient coordination of future mobile sensing and actuation systems with a special focus on their dependability.

3.4. Mobile social crowd-sensing

While mobile social sensing opens up the ability of sensing phenomena that may be costly or impossible to sense using embedded sensors (e.g., subjective crowdedness causing discomfort or joyfulness, as in a bus or in a concert) and leading to a feeling of being more socially involved for the citizens, there are unique consequent challenges. Specifically, MiMove's research focuses on the problems involved in the combination of the physically sensed data, which are quantitative and objective, with the mostly qualitative and subjective data arising from social sensing. Enabling the latter calls for introducing mechanisms for incentivising user participation and ensuring the privacy of user data, as well as running empirical studies for understanding the complex social behaviors involved. These objectives build upon previous research work by members of the team on mobile social ecosystems and privacy, as well as a number of efforts and collaborations in the domain of smart cities and transport that have resulted in novel mobile applications enabling empirical studies of social sensing systems.

3.5. Active and passive probing methods

We are developing methods that actively introduce probes in the network to discover properties of the connected devices and network segments. We are focusing in particular on methods to discover properties of home networks (connected devices and their types) and to distinguish if performance bottlenecks lie within the home network versus in the different network segments outside (e.g., Internet access provider, interconnects, or content provider). Our goal is to develop adaptive methods that can leverage the collaboration of the set of available devices (including end-user devices and the home router, depending on which devices are running the measurement software).

We are also developing passive methods that simply observe network traffic to infer the performance of networked applications and the location of performance bottlenecks, as well as to extract patterns of web content consumption. We are working on techniques to collect network traffic both at user's end-devices and at home routers. We also have access to network traffic traces collected on a campus network and on a large European broadband access provider.

3.6. Inferring user online experience

We are developing hybrid measurement methods that combine passive network measurement techniques to infer application performance with techniques from HCI to measure user perception. We will later use the resulting datasets to build models of user perception of network performance based only on data that we can obtain automatically from the user device or from user's traffic observed in the network.

3.7. Real time data analytics

The challenge of deriving insights from the Internet of Things (IoT) has been recognized as one of the most exciting and key opportunities for both academia and industry. The time value of data is crucial for many IoT-based systems requiring *real-time* (or near real-time) *control* and *automation*. Such systems typically collect data continuously produced by "things" (i.e., devices), and analyze them in (sub-) seconds in order to act promptly, e.g., for detecting security breaches of digital systems, for spotting malfunctions of physical assets, for recommending goods and services based on the proximity of potential clients, etc. Hence, they require to both *ingest* and *analyze in real-time* data arriving with different velocity from various IoT data streams.

Existing incremental (online or streaming) techniques for descriptive statistics (e.g., frequency distributions, frequent patterns, etc.) or predictive statistics (e.g., classification, regression) usually assume a good enough quality dataset for mining patterns or training models. However, IoT raw data produced in the wild by sensors embedded in the environment or wearable by users are prone to errors and noise. Effective and efficient algorithms are needed for *detecting* and *repairing data impurities* (for controlling data quality) as well as *understanding data dynamics* (for defining alerts) in real-time, for collections of IoT data streams that might be geographically distributed. Moreover, supervised deep learning and data analytics techniques are challenged

by the presence of sparse ground truth data in real IoT applications. Lightweight and adaptive semi-supervised or unsupervised techniques are needed to power real-time anomaly and novelty detection in IoT data streams. The effectiveness of these techniques should be able to reach a useful level through training on a relatively small amount of (preferably unlabeled) data while they can cope distributional characteristics of data evolving over time.

4. Application Domains

4.1. Mobile urban systems for smarter cities

With the massive scale adoption of mobile devices and further expected significant growth in relation with the Internet of Things, mobile computing is impacting most – if not all – the ICT application domains. However, given the importance of conducting empirical studies to assess and nurture our research, we focus on one application area that is the one of "*smart cities*". The smart city vision anticipates that the whole urban space, including buildings, power lines, gas lines, roadways, transport networks, and cell phones, can all be wired together and monitored. Detailed information about the functioning of the city then becomes available to both city dwellers and businesses, thus enabling better understanding and consequently management of the city's infrastructure and resources. This raises the prospect that cities will become more sustainable environments, ultimately enhancing the citizens' well being. There is the further promise of enabling radically new ways of living in, regulating, operating and managing cities, through the increasing active involvement of citizens by ways of crowd-sourcing/sensing and social networking.

Still, the vision of what smart cities should be about is evolving at a fast pace in close concert with the latest technology trends. It is notably worth highlighting how mobile and social network use has reignited citizen engagement, thereby opening new perspectives for smart cities beyond data analytics that have been initially one of the core foci for smart cities technologies. Similarly, open data programs foster the engagement of citizens in the city operation and overall contribute to make our cities more sustainable. The unprecedented democratization of urban data fueled by open data channels, social networks and crowd sourcing enables not only the monitoring of the activities of the city but also the assessment of their nuisances based on their impact on the citizens, thereby prompting social and political actions. However, the comprehensive integration of urban data sources for the sake of sustainability remains largely unexplored. This is an application domain that we focus on, further leveraging our research on emergent mobile distributed systems, large-scale mobile sensing & actuation, and mobile social crowd-sensing.

In particular, we concentrate on the following specialized applications, which we have investigated in close collaboration with other researchers as part of the dedicated Inria Project Lab *CityLab@Inria*:

- **Democratization of urban data for healthy cities.** We integrate the various urban data sources, especially by way of crowd-Xing, to better understand city nuisances. This goes from raw pollution sensing (e.g., sensing noise) to the sensing of its impact on citizens (e.g., how people react to urban noise and how this affects their health).
- **Social applications.** Mobile applications are being considered by sociologists as a major vehicle to actively involve citizens and thereby prompt them to become activists. We study such a vehicle from the ICT perspective and in particular elicit relevant middleware solutions to ease the development of such "*civic apps*".

More specifically, MiMove led CityLab@Inria [1] from Jan 2014 to Nov 2018. CityLab focused on the study of ICT solutions promoting social sustainability in smart cities, and involved the following Inria project-teams in addition to MiMove: CLIME/ANGE, DICE, FUN, MYRIADS, SMIS/PETRUS, URBANET/AGORA. CityLab further involved strong collaboration with California universities affiliated with CITRIS (Center for Information Technology Research in the Interest of Society) and especially UC Berkeley, in relation with the *Inria@SiliconValley* program.

[1] http://citylab.inria.fr

4.2. Home network diagnosis

With the availability of cheap broadband connectivity, Internet access from the home has become a ubiquity. Modern households host a multitude of networked devices, ranging from personal devices such as laptops and smartphones to printers and media centers. These devices connect among themselves and to the Internet via a local-area network–a home network–that has become an important part of the "Internet experience". In fact, ample anecdotal evidence suggests that the home network can cause a wide array of connectivity impediments, but their nature, prevalence, and significance remain largely unstudied.

Our long-term goal is to assist users with concrete indicators of the quality of their Internet access, causes of potential problems and–ideally–ways to fix them. We intend to develop a set of easy-to-use home network monitoring and diagnosis tools. The development of home network monitoring and diagnosis tools brings a number of challenges. First, home networks are heterogeneous. The set of devices, configurations, and applications in home networks vary significantly from one home to another. We must develop sophisticated techniques that can learn and adapt to any home network as well as to the level of expertise of the user. Second, Internet application and services are also heterogeneous with very diverse network requirements. We must develop methods that can infer application quality solely from the observation of (often encrypted) application network traffic. There are numerous ways in which applications can fail or experience poor performance in home networks. Often there are a number of explanations for a given symptom. We must devise techniques that can identify the most likely cause(s) for a given problem from a set of possible causes. Finally, even if we can identify the cause of the problem, we must then be able to identify a solution. It is important that the output of the diagnosis tools we build is "actionable". Users should understand the output and know what to do.

We are working with Princeton University (associate team HOMENET) to deploy monitoring infrastructure within users' homes. Our goal is to develop a mostly passive measurement system to monitor the performance of user applications, which we call NetMicroscope. We are developing NetMicroscope to run in a box acting as home gateway. Our current deployments use Raspberry Pi and Odroid boxes. We have these boxes deployed in 50 homes in the US and 10 in France. The US deployment is run and financed by the Wall Street Journal. We are collaborating with them to understand the relationship between Internet access speed and video quality. We have been discussing with Internet regulators (in particular, FCC, ACERP, and BEREC) as well as residential access ISP in how NetMicroscope can help overcome the shortcomings of existing Internet quality monitoring systems.

4.3. Mobile Internet quality of experience

Mobile Internet usage has boomed with the advent of ever smarter handheld devices and the spread of fast wireless access. People rely on mobile Internet for everyday tasks such as banking, shopping, or entertainment. The importance of mobile Internet in our lives raises people's expectations. Ensuring good Internet user experience (or Quality of Experience—QoE) is challenging, due to the heavily distributed nature of Internet services. For mobile applications, this goal is even more challenging as access connectivity is less predictable due to user mobility, and the form factor of mobile devices limits the presentation of content. For these reasons, the ability to monitor QoE metrics of mobile applications is essential to determine when the perceived application quality degrades and what causes this degradation in the chain of delivery. Our goal is to improve QoE of mobile applications.

To achieve this goal, we are working on three main scientific objectives. First, we are working on novel methods to monitor mobile QoE. Within the IPL BetterNet we are developing the HostView for Android tool that runs directly on mobile devices to monitor network and system performance together with the user perception of performance. Second, we plan to develop models to predict QoE of mobile applications. We will leverage the datasets collected with HostView for Android to build data-driven models. Finally, our goal is to develop methods to optimize QoE for mobile users. We are currently developing optimization methods for interactive video applications. We envision users walking or driving by road-side WiFi access points (APs) with full 3G/LTE coverage and patchy WiFi coverage (i.e., community Wifi or Wifi APs on Lampposts). To

achieve this goal, we plan to leverage multi-path and cross-layer optimizations. We have started conducting experiments in the Paris subway and walking around Inria to measure the quality of FreeWiFi as well as LTE connectivity. We are experimenting with existing multipath protocols (MP-TCP and MP-DASH). We are also analyzing connectivity in datasets from the MONROE project (which measure LTE in Europe) and CarFi (which measures WiFi quality from APs deployed in cars).

5. Highlights of the Year

5.1. Highlights of the Year

5.1.1. Awards

- Renata Teixeira was named ACM Distinguished member for outstanding scientific contributions to computing in 2018.
- Our paper "Narrowing the gap between QoS metrics and Web QoE using Above-the-fold metrics" received the Best Dataset Award at the Passive and Active Measurement Conference 2018.

BEST PAPER AWARD:

[27]
R. GOMES, G. BOULOUKAKIS, F. COSTA, N. GEORGANTAS, R. DA ROCHA. *QoS-Aware Resource Allocation for Mobile IoT Pub/Sub Systems*, in "2018 International Conference on Internet of Things (ICIOT)", Seattle, United States, June 2018, https://hal.inria.fr/hal-01797933

6. New Software and Platforms

6.1. SocialBus

Universal Social Network Bus
KEYWORDS: Middleware - Interoperability - Social networks - Software Oriented Service (SOA)
FUNCTIONAL DESCRIPTION: Online social network services (OSNSs) have become an integral part of our daily lives. At the same time, the aggressive market competition has led to the emergence of multiple competing siloed OSNSs that cannot interoperate. As a consequence, people face the burden of creating and managing multiple OSNS accounts and learning how to use them, to stay connected. The goal of the Universal Social Network Bus (USNB) is to relieve users from such a burden, letting them use their favorite applications to communicate.

- Authors: Rafael Angarita Arocha, Nikolaos Georgantas and Valérie Issarny
- Contact: Valérie Issarny
- URL: https://gitlab.inria.fr/usnb/universal-social-network-bus

6.2. WeBrowse

KEYWORDS: Web Usage Mining - Content analysis - Recommendation systems
FUNCTIONAL DESCRIPTION: The amount of information available on the web today, and the fast rate with which new information appears, overwhelm most users. The goal of our research is to assist Web users in discovering content. One of the most powerful means today to help people discover new web content is sharing between members of online communities. In the case of communities of a place (e.g., people who live, study, or work together) people share common interests, but often fail to actively share content. To address this problem, we have developed WeBrowse, a passive crowdsourced content discovery system for communities of a place.

WeBrowse leverages the passive observation of web-clicks (i.e., the URLs users intentionally visit) as an indication of users' interest in a piece of content. Intuitively, the more users click on a URL, the higher the interest in the content on the corresponding page. Our approach is then to leverage the collective clicks in a community to automatically discover relevant content to promote to users of the community.

To implement passive crowdsourcing, one must be in a position to observe the aggregated web-clicks of the community. Luckily, in many communities of a place, users will connect to the Internet from the same network, such as, e.g., the campus/enterprise network or the network of a residential Internet Service Provider (ISP) in a neighborhood. WeBrowse (i) observes web packets flowing through a network link, (ii) passively extracts HTTP logs (i.e., streams recording the headers of HTTP requests), and (iii) detects and decides on-the-fly the set of URLs to show to users.

- Contact: Renata Cruz Teixeira
- URL: https://team.inria.fr/muse/webrowse-info-page/

6.3. EEE

Experiment Execution Engine
KEYWORD: Iot
FUNCTIONAL DESCRIPTION: Experiment Execution Engine (EEE) eases the development of IoT applications that perform analysis of recent or continuously increasing volumes of IoT data from various data stores. To this end, EEE provides APIs for scheduling queries on federated large-scale semantically-enabled IoT data stores. Queries are expressed in the FIESTA-IoT (http://fiesta-iot.eu) Experiment Description Specification (FEDSpec), which acts as a Domain Specific Language (DSL). EEE (in combination with Experiment Management Console - EMC) further provides features such as: (a) manage scheduled experiments, (b) (un)subscribe to already existing FEDSpec objects, (c) execute an on-demand query outside a described schedule, (d) monitor execution logs, (e) interact with third-party services such as Analytics and Result Storage, (e) handle dynamic attributes, and (f) API sandbox. EEE is integrated within the FIESTA-IoT Platform. It can be customised depending on needs.

- Authors: Rachit Agarwal and Nikolaos Georgantas
- Contact: Nikolaos Georgantas
- URL: https://github.com/fiesta-iot/ExperimentExecutionEngine

6.4. EMC

Experiment Management Console
KEYWORD: Iot
FUNCTIONAL DESCRIPTION: Experiment Management Console (EMC) is a simple easy-to-use user interface that enables experimenters to control their experiments/queries that execute using the Experiment Execution Engine (EEE). EMC provides features such as managing experiment execution, (un)subscribing an existing experiment, and monitoring execution logs.

- Authors: Nikolaos Georgantas and Rachit Agarwal
- Contact: Nikolaos Georgantas
- URL: https://github.com/fiesta-iot/ExperimentManagementConsole

6.5. VSB

eVolution Service Bus
KEYWORDS: Service and Thing choreographies - Middleware protocol interoperability - Enterprise service bus

FUNCTIONAL DESCRIPTION: VSB is a development and runtime environment dedicated to complex distributed applications of the Future Internet. Such applications are open, dynamic choreographies of extremely heterogeneous services and Things, including lightweight embedded systems (e.g., sensors, actuators and networks of them), mobile systems (e.g., smartphone applications), and resource-rich IT systems (e.g., systems hosted on enterprise servers and Cloud infrastructures). VSB's objective is to seamlessly interconnect, inside choreographies, services and Things that employ heterogeneous interaction protocols at the middleware level, e.g., SOAP Web services, REST Web services, Things using CoAP. This is based on runtime conversions between such protocols, with respect to their primitives and data type systems, while properly mapping between their semantics. This also includes mapping between the public interfaces of services/Things, regarding their operations and data, from the viewpoint of the middleware: the latter means that operations and data are converted based on their middleware-level semantics, while their business semantics remains transparent to the conversion. VSB follows the well-known Enterprise Service Bus (ESB) paradigm. We propose a generic interface description, which we call GIDL, for application components that employ VSB. Based on GIDL, we enable automated synthesis of binding components for connecting heterogeneous services and Things onto VSB.

- Participants: Georgios Bouloukakis, Nikolaos Georgantas and Patient Ntumba
- Contact: Nikolaos Georgantas
- URL: https://gitlab.ow2.org/chorevolution/evolution-service-bus

6.6. Service traceroute

KEYWORDS: Network monitoring - Network diagnosis
FUNCTIONAL DESCRIPTION: Traceroute is often used to help diagnose when users experience issues with Internet applications or services. Unfortunately, probes issued by classic traceroute tools differ from application traffic and hence can be treated differently by middleboxes within the network. We propose a new traceroute tool, called Service traceroute. Service traceroute leverages the idea from paratrace, which passively listens to application traffic to then issue traceroute probes that pretend to be part of the application flow. We extend this idea to work for modern Internet services with support for identifying the flows to probe automatically, for tracing of multiple concurrent flows, and for UDP flows. We implement command-line and library versions of Service traceroute, which we release as open source.

- Partner: Princeton University
- Contact: Renata Cruz Teixeira
- URL: https://github.com/wontoniii/service-traceroute

6.7. TA

TA - Traffic Analysis
KEYWORDS: Quality of Experience - Network monitoring - Video analysis
FUNCTIONAL DESCRIPTION: System running at the home getaway that analyzes traffic generated by DASH on-demand and live video streams. The system tracks traffic patterns to infer key video QoE metrics such as average bitrate and re-buffering events. Moreover, the system exploits novel algorithms that use probing techniques, i.e. lightweight pings and traceroutes, to detect possible congestion location.

- Participants: Francesco Bronzino and Renata Cruz Teixeira
- Contact: Francesco Bronzino

6.8. HostView Mobile

KEYWORDS: Quality of Experience - Network monitoring
FUNCTIONAL DESCRIPTION: HostView for mobile runs on Android devices to monitor user system and network performance together with user feedback on Internet experience.

- Contact: Giulio Grassi

7. New Results

7.1. Ontology categorization for IoT semantics

Participants: Rachit Agarwal, Nikolaos Georgantas, Valérie Issarny.

IoT systems are now being deployed worldwide to sense phenomena of interest. The existing IoT systems are often independent which limits the use of sensor data to only one application. Semantic solutions have been proposed to support reuse of sensor data across IoT systems and applications. This allows integration of IoT systems for increased productivity by solving challenges associated with their interoperability and heterogeneity. Several ontologies have been proposed to handle different aspects of sensor data collection in IoT systems, ranging from sensor discovery to applying reasoning on collected sensor data for drawing inferences. In this work, we study and categorise the existing ontologies based on the fundamental ontological concepts (e.g., sensors, context, location, and more) required for annotating different aspects of data collection and data access in an IoT application. We identify these fundamental concepts by answering the 4Ws (What, When, Who, Where) and 1H (How) identified using the 4W1H methodology.

7.2. Massively-Parallel Feature Selection for Big Data

Participant: Vassilis Christophides.

We present the Parallel, Forward-Backward with Pruning (PFBP) algorithm for feature selection (FS) in Big Data settings (high dimensionality and/or sample size). To tackle the challenges of Big Data FS, PFBP partitions the data matrix both in terms of rows (samples, training examples) as well as columns (features). By employing the concepts of p-values of conditional independence tests and meta-analysis techniques, PFBP manages to rely only on computations local to a partition while minimizing communication costs. Then, it employs powerful and safe (asymptotically sound) heuristics to make early, approximate decisions, such as Early Dropping of features from consideration in subsequent iterations, Early Stopping of consideration of features within the same iteration, or Early Return of the winner in each iteration. PFBP provides asymptotic guarantees of optimality for data distributions faithfully representable by a causal network (Bayesian network or maximal ancestral graph). Our empirical analysis confirms a superlinear speedup of the algorithm with increasing sample size, linear scalability with respect to the number of features and processing cores, while dominating other competitive algorithms in its class.

7.3. Universal Social Network Bus

Participants: Ehsan Ahvar, Shohreh Ahvar, Rafael Angarita, Nikolaos Georgantas, Valérie Issarny, Bruno Lefèvre.

Online social network services (OSNSs) are changing the fabric of our society, impacting almost every aspect of it. Over the last decades, the aggressive market rivalry has led to the emergence of multiple competing, "closed" OSNSs. As a result, users are trapped in the walled gardens of their OSNS, encountering restrictions about what they can do with their personal data, the people they can interact with and the information they get access to. As an alternative to the platform lock-in, "open" OSNSs promote the adoption of open, standardized APIs. However, users still massively adopt closed OSNSs to benefit from the services' advanced functionalities and/or follow their "friends", although the users' virtual social sphere is ultimately limited by the OSNSs they join. Our work aims at overcoming such a limitation by enabling users to meet and interact beyond the boundary of their OSNSs, including reaching out to "friends" of distinct closed OSNSs. We specifically introduce USNB -*Universal Social Network Bus*, which revisits the "service bus" paradigm that enables interoperability across computing systems, to address the requirements of "*social interoperability*". USNB features *synthetic profiles* and *personae* for interaction across the boundaries of –closed and open–, –profile- and non-profile-based– OSNSs through a *reference social interaction service*.

USNB enables users to reach out to their social peers independently of the communication service (and especially underlying platform) each one uses in the virtual world. The success and massive adoption of OSNSs -as magnified by the success of Facebook- shows that online social communication is an essential tool for people. This further paves the way for collective and collaborative actions at the Internet scale. However, existing online collaborative tools come along with their communication platform, which is either a proprietary solution or a third-party OSNS. We argue that USNB contributes to enabling participatory systems at a larger inclusive scale by overcoming the technical boundaries set by existing online communication platforms. In that direction, we investigate the customization of USNB for specific applications and more specifically: participatory systems and massive open online courses.

7.4. Middleware for Mobile Crowdsensing

Participants: Yifan Du, Valérie Issarny, Bruno Lefèvre, Françoise Sailhan.

Mobile Phone Sensing (MPS) offers a great opportunity toward the large scale monitoring of urban phenomena, such as the exposition of the population to environmental pollution. Indeed, mobile crowdsensing empowers ordinary citizens to contribute (whether pro-actively or passively) data sensed or generated from their mobile devices. It allows acquiring hyperlocal knowledge at scale, thanks to the proliferation of mobile devices and the ubiquity of wireless broadband connection. On-demand mobile crowdsensing is in particular a cost-effective service model for smart cities. Numerous sensor types embedded in today's smartphones contribute valuable quantitative observations about the urban environment (e.g., noise, temperature, atmospheric pressure, humidity, light, magnetism). The observations further come along with the related spatial and temporal data, which allows for the analysis of hyper-local environmental knowledge. However, mobile crowdsensing brings valuable knowledge only if a sufficiently large crowd contributes and if we overcome the relatively low accuracy of the gathered data. This is the focus of our research.

We have in particular studied how to reduce the gap between the need for the massive collection of relevant data, and the quantity and accuracy of the measurements that are actually gathered. We specifically carried out an iterative research process to tackle this challenge, which combines technological innovation and social design. We have been developing a number of social tools to study the motivations and usages of MPS-based smart city apps, with the Ambiciti app serving as our use case. Our study has been taking into account the cultural and societal contexts that the usages of Ambiciti could feed, spanning health, environment, education, and urban policies. We carried out an online survey together with interviews with users and local actors in Europe, i.e., France, Belgium, and Finland. The research results contribute to a better understanding of why and how people use mobile phone sensing applications; the results also inform how to best leverage mobile crowd-sensing in the development of smart cities and how it may serve addressing urban challenges related to, e.g., public health or urban planning.

The quality of the contributed measurements challenges the aggregation of relevant knowledge from crowdsensed observations. The measurements quality depends on the *accuracy* of the contributing sensors and the adequacy of the *sensing context*. Addressing the former relies on the sensor calibration for which we study both micro- and macro-level solutions. Addressing the latter requires a supporting inference mechanism, for which we introduce a *personalized hierarchical inference* of all the context elements that are relevant to the phenomenon that is monitored through crowdsensing, and under which the crowdsensor operates. This enables accounting for the specific behavior of the contributing end-user across time, as well as for all the features -and only those- that are relevant and locally available, while reducing the feedback required from the user for the personalization.

7.5. QoS-Aware Resource Allocation for Mobile IoT Pub/Sub Systems

Participants: Georgios Bouloukakis, Nikolaos Georgantas.

IoT applications are usually characterized by large-scale demand and the widespread use of mobile devices. Similarly, performing interaction among application and system components in a decoupled and elastic way, and enforcing Quality of Service (QoS) usually also become issues. Hence, paradigms such as pub/sub on top of cloud resources represent a suitable strategy for application development. However, management of QoS-aware resource allocation for pub/sub systems remains challenging, especially when system peers connect in an intermittent way. In this work, we propose a new approach for resource allocation focusing on end-to-end performance in face of peers' disconnections. We evaluate and demonstrate the benefits of our approach using simulations. QoS enforcement was achieved in almost all scenarios, and we have shown that our approach can help reasoning about efficient resource allocation.

7.6. Queueing Network Modeling Patterns for Reliable & Unreliable Pub/Sub Protocols

Participants: Georgios Bouloukakis, Nikolaos Georgantas, Patient Ntumba, Valérie Issarny.

Mobile Internet of Things (IoT) applications are typically deployed on resource-constrained devices with intermittent network connectivity. To support the deployment of such applications, the Publish/Subscribe (pub/sub) interaction paradigm is often employed, as it decouples mobile peers in time and space. Furthermore, pub/sub middleware protocols and APIs consider the Things' hardware limitations and support the development of effective applications by providing Quality of Service (QoS) features. These features aim to enable developers to tune an application by switching different levels of response times and delivery success rates. However, the profusion of pub/sub middleware protocols coupled with intermittent network connectivity result in non-trivial application tuning. In this work, we model the performance of middleware protocols found in IoT, which are classified within the pub/sub interaction paradigm – both reliable and unreliable underlying network layers are considered. We model reliable and unreliable protocols, by considering QoS semantics for data validity, buffer capacities, as well as the intermittent availability of peers. To this end, we rely on queueing network models, which offer a simple modeling environment that can be used to represent IoT interactions by combining multiple queueing model types. Based on these models, we perform statistical analysis by varying the QoS semantics, demonstrating their significant effect on response times and on the rate of successful interactions. We showcase the application of our analysis in concrete scenarios relating to Traffic Information Management systems, that integrate both reliable and unreliable participants. The consequent PerfMP performance modeling pattern may be tailored for a variety of deployments, in order to control fine-grained QoS policies.

7.7. Lightweight, General Inference of Streaming Video Quality from Encrypted Traffic

Participants: Francesco Bronzino, Sara Ayoubi, Renata Teixeira, Sarah Wasserman.

Accurately monitoring application performance is becoming more important for Internet Service Providers (ISPs), as users increasingly expect their networks to consistently deliver acceptable application quality. At the same time, the rise of end-to-end encryption makes it difficult for network operators to determine video stream quality–including metrics such as startup delay, resolution, rebuffering, and resolution changes–directly from the traffic stream. This work develops general methods to infer streaming video quality metrics from encrypted traffic using lightweight features. Our evaluation shows that our models are not only as accurate as previous approaches, but they also generalize across multiple popular video services, including Netflix, YouTube, Amazon Instant Video, and Twitch. The ability of our models to rely on lightweight features points to promising future possibilities for implementing such models at a variety of network locations along the end-to-end network path, from the edge to the core.

7.8. Service traceroute: Tracing Paths of Application Flows

Participants: Ivan Morandi, Francesco Bronzino, Renata Teixeira.

Traceroute is often used to help diagnose when users experience issues with Internet applications or services. Unfortunately, probes issued by classic traceroute tools differ from application traffic and hence can be treated differently by middleboxes within the network. This work proposes a new traceroute tool, called Service traceroute. Service traceroute leverages the idea from paratrace, which passively listens to application traffic to then issue traceroute probes that pretend to be part of the application flow. We extend this idea to work for modern Internet services with support for identifying the flows to probe automatically, for tracing of multiple concurrent flows, and for UDP flows. We implement command-line and library versions of Service traceroute, which we release as open source. This paper also presents an evaluation of Service traceroute when tracing paths traversed by Web downloads from the top-1000 Alexa websites and by video sessions from Twitch and Youtube. Our evaluation shows that Service traceroute has no negative effect on application flows. Our comparison with Paris traceroute shows that a typical traceroute tool that launches a new flow to the same destination discovers different paths than when embedding probes in the application flow in a significant fraction of experiments (from 40% to 50% of our experiments in PlanetLab Europe).

8. Bilateral Contracts and Grants with Industry

8.1. Bilateral Grants with Industry

"Application Performance Bottleneck Detection", Comcast Gift to R. Teixeira 2018.

9. Partnerships and Cooperations

9.1. National Initiatives

"BottleNet: Understanding and Diagnosing End-to-end Communication Bottlenecks of the Internet", project funded by the French research agency (ANR), from Feb 2016 to Sep 2020.

9.1.1. Inria Support

9.1.1.1. Inria IPL CityLab@Inria

Participants: Valérie Issarny, Bruno Lefèvre.

- **Name:** CityLab@Inria – *Overcoming the Smart City Challenge – Toward Environmental and Social Sustainability*
- **Period:** [January 2014 – November 2018]
- **Inria teams:** CLIME/ANGE, DICE, FUN, MIMOVE, MYRIADS, SMIS/PETRUS, URBANET/AGORA
- **URL:** http://citylab.inria.fr

The Inria Project Lab (IPL) CityLab@Inria studies ICT solutions toward smart cities that promote both social and environmental sustainability. A strong emphasis of the Lab is on the undertaking of a multi-disciplinary research program through the integration of relevant scientific and technology studies, from sensing up to analytics and advanced applications, so as to actually enact the foreseen smart city Systems of Systems. Obviously, running experiments is a central concern of the Lab, so that we are able to confront proposed approaches to actual settings.

9.1.1.2. Inria IPL BetterNet

Participants: Renata Teixeira, Vassilis Christophides, Francesco Bronzino.

- **Name:** BetterNet – *An observatory to measure and improve Internet service access from user experience*
- **Period:** [2016 – 2019]
- **Inria teams:** Diana, Dionysos, Inria Chile, Madynes, MiMove, Spirals
- **URL:** https://project.inria.fr/betternet/

BetterNet aims at building and delivering a scientific and technical collaborative observatory to measure and improve the Internet service access as perceived by users. In this Inria Project Lab, we will propose new original user-centered measurement methods, which will associate social sciences to better understand Internet usage and the quality of services and networks. Our observatory can be defined as a vantage point, where:

1. tools, models and algorithms/heuristics will be provided to collect data,
2. acquired data will be analyzed, and shared appropriately with scientists, stakeholders and civil society,
3. and new value-added services will be proposed to end-users.

9.1.1.3. Inria ADT MOSQUITO

Participants: Renata Teixeira, Francesco Bronzino.

- **Name:** MOSQUITO – *A mobile platform to measure the quality of Internet connectivity*
- **Period:** [November 2016 – October 2018]
- **Partners:** Inria MiMove, Inria SPIRALS.

The ADT MOSQUITO is part of the Inria Project Lab (IPL) initiative BetterNet. This ADT project focuses on the design and the development of a measurement platform for the quality of mobile Internet access by federating the existing mobile platforms identified in the BetterNet IPL. Beyond the priceless value of such a measurement platform for the research community, this ADT also aims to publish live reports on the quality of mobile Internet access through the BetterNet initiative.

9.2. European Initiatives

9.2.1. FP7 & H2020 Projects

9.2.1.1. H2020 ICT FIESTA-IoT

Participants: Valérie Issarny, Nikolaos Georgantas, Rachit Agarwal.

Name: FIESTA-IoT – *Federated Interoperable Semantic IoT/cloud Testbeds and Applications*

URL: http://fiesta-iot.eu

Type: Research & Innovation Action (ICT)

Topic: FIRE+ (Future Internet Research & Experimentation)

Period: [February 2015 - January 2018]

Partners: Fraunhofer FOKUS (Germany) [**coordinator**], INSIGHT @ National University of Galway (Ireland) [**co-coordinator**], University of Southampton IT Innovation Centre (UK), Inria MiMove, University of Surrey (UK), Unparallel Innovation Lda (Portugal), Easy Global Market (France), NEC Europe Ltd (UK), University of Cantabria (Spain), Com4innov (France), Athens Information Technology (Greece), SOCIEDAD PARA EL DESARROLLO REGIONAL DE CANTABRIA (Spain), Ayuntamiento de Santander (Spain), Korea Electronics Technology Institute (Korea).

Despite the proliferation of IoT and smart cities testbeds, there is still no easy way to conduct large scale experiments that leverage data and resources from multiple geographically and administratively distributed IoT platforms. Recent advances in IoT semantic interoperability provide a sound basis for implementing novel cloud-based infrastructures that could allow testbed-agnostic access to IoT data and resources. FIESTA will open new horizons in IoT experimentation at a global scale, based on the interconnection and interoperability of diverse IoT testbeds. FIESTA will produce a first-of-a-kind blueprint experimental infrastructure (tools, techniques and best practices) enabling testbed operators to interconnect their facilities in an interoperable way, while at the same time facilitating researchers in deploying integrated experiments, which seamlessly transcend the boundaries of multiple IoT platforms. FIESTA will be validated and evaluated based on the interconnection of four testbeds (in Spain, UK, France and Korea), as well as based on the execution of novel experiments in the areas of mobile crowd-sensing, IoT applications portability, and dynamic intelligent

discovery of IoT resources. In order to achieve global outreach and maximum impact, FIESTA will integrate an additional testbed and experiments from Korea, while it will also collaborate with IoT experts from USA. The participation of a Korean partner (based its own funding) will maximize FIESTA's value for EC money. Moreover, the project will take advantage of open calls processes towards attracting third-parties that will engage in the integration of their platforms within FIESTA or in the conduction of added-value experiments. As part of its sustainability strategy, FIESTA will establish a global market confidence programme for IoT interoperability, which will enable innovative platform providers and solution integrators to ensure/certify the openness and interoperability of their developments.

9.3. International Initiatives

9.3.1. Inria International Labs

Inria@SiliconValley
Associate Team involved in the International Lab:

9.3.1.1. MINES

Title: Adaptive Communication Middleware for Resilient Sensing & Actuation IN Emergency Response Scenarios

International Partner:

University of California, Irvine (United States) - Information and Computer Science - Nalini Venkatasubramanian

Start year: 2018

See also: http://mimove-apps.paris.inria.fr/mines/index.html

Emerging smart-city and smart-community efforts will require a massive deployment of connected entities (Things) to create focused smartspaces. Related applications will enhance citizen quality of life and public safety (e.g., providing safe evacuation routes in fires). However, supporting IoT deployments are heterogeneous and can be volatile and failure-prone as they are often built upon low-powered, mobile and inexpensive devices - the presence of faulty components and intermittent network connectivity, especially in emergency scenarios, tend to deliver inaccurate/delayed information. The MINES associate team addresses the resulting challenge of enabling interoperability and resilience in large-scale IoT systems through the design and development of a dedicated middleware. More specifically, focusing on emergency situations, the MINES middleware will: (i) enable the dynamic composition of IoT systems from any and all available heterogeneous devices; (ii) support the timely and reliable exchange of critical data within and across IoT in the enabled large-scale and dynamic system over heterogeneous networks. Finally, the team will evaluate the proposed solution in the context of emergency response scenario use cases.

9.3.2. Inria Associate Teams Not Involved in an Inria International Lab

9.3.2.1. HOMENET

Title: Home network diagnosis and security

International Partner:

Princeton University (United States) - Computer Science Department - Nick Feamster

Start year: 2017

See also: https://team.inria.fr/homenet/

Modern households connect a multitude of networked devices (ranging from laptops and smartphones to a number of Internet of Things devices) via a home network. Most home networks, however, do not have a technically skilled network administrator for managing the network, for example to identify faulty equipment or take steps to secure end hosts such as applying security patches. Home networks represent a particularly challenging environment due to the diversity of devices, applications, and services users may connect. The goal of HOMENET is to assist users in diagnosing and securing their home networks. Our approach is based on developing new algorithms and mechanisms that will run on the home router (or in-collaboration with the router). The router connects the home network to the rest of the Internet; it is hence the ideal place to secure home devices and to distinguish problems that happen in the home from those happening elsewhere. We will address a number of research challenges for example in device discovery and fingerprinting, anomaly detection in the Internet of Things, home network diagnosis (including wireless diagnosis). HOMENET will bring together two leading research teams in the network measurement arena with successful prior collaboration. Moreover, Princeton brings an existing home router platform and expertise in security, wireless, and software-defined networks; and MiMove brings an existing Web-based measurement platform, and expertise in traffic-based profiling and anomaly detection.

9.3.2.2. ACHOR

Title: Adaptive enactment of service choreographies
International Partner:

Universidade Federal de Goiás (Brazil) - Computer Science Department - Fabio Costa

Start year: 2016
See also: http://www.inf.ufg.br/projects/achor
Service choreographies are distributed compositions of services (e.g., Web services) that coordinate their execution and interactions without centralized control. Due to this decentralized coordination and the ability to compose third-party services, choreographies have shown great potential as an approach to automate the construction of large-scale, on-demand, distributed applications. Technologies to enable this approach are reaching maturity level, such as modeling languages for choreography specification and engines that operate the deployment of services and enactment of choreographies at Future Internet scales. Nevertheless, a number of problems remain open on the way to fully realize the approach, among them: (i) Deployment of multiple choreographies on top of a collection of shared services (considering service sharing as an effective way to increase the utilization of resources); (ii) Dynamic adaptation of functional and non-functional properties due to runtime changes in the environment and user requirements (adapting the set of services and/or the resources used to run the services in order to add/remove/change functions and maintain QoS properties, respectively); and (iii) Seamless and dynamic integration of mobile services (e.g., smartphone apps, sensors and actuators on handhelds and wearables) and cloud- based services (including the need to consider: mobility of both devices and services, resource constraints of mobile devices, temporary disconnection, interoperability between different interaction paradigms (message-passing, event-based, data-sharing) at the middleware layer, and effect of these paradigms on end-to-end QoS). The overall goal of the project is to design an architecture for adaptive middleware to support service choreographies in large-scale scenarios that involve dynamicity and diversity in terms of application requirements, service interaction protocols, and the use of shared local, mobile and cloud resources.

9.3.3. Inria International Partners

9.3.3.1. Informal International Partners

- Northeastern University (Prof. David Choffnes): We are working on methods based on active probing to diagnose poor video quality.

9.4. International Research Visitors

9.4.1. Visits of International Scientists

Mark Crovella from Boston University is visiting professor at Inria.

9.4.2. Visits to International Teams

9.4.2.1. Research Stays Abroad

- Valérie Issarny was visiting scholar at the EECS Department at UC Berkeley till August 2018. She was hosted by CITRIS in the context of which she was carrying out collaborative research in the area of smart cities and acting as scientific coordinator of the Inria@SiliconValley program.
- Renata Teixeira is visiting scholar at the Computer Science department at Stanford University.
- Georgios Bouloukakis is Inria postdoctoral fellow at University of California, Irvine, in the context of the Inria@SiliconValley program.

10. Dissemination

10.1. Promoting Scientific Activities

10.1.1. Scientific Events Selection

10.1.1.1. Chair of Conference Program Committees

- V. Issarny, TPC co-chair of ICSE-SEIS'2018 - Software Engineering in Society Track of the ICSE'18 conference. Gotheborg, SE, May 2018.
- V. Issarny, TPC co-chair of CIC'2018 - 4th IEEE International Conference on Collaboration and Internet Computing. Philadelphia, USA, October 2018
- V. Issarny, TPC co-chair of SCF-ICIOT - 2019 International Conference on Internet of Things Services at the Services Conference Federation. San Diego, USA, June 2019.

10.1.1.2. Member of the Conference Program Committees

- V. Christophides, PC member of IEEE Data Engineering Conference 2018.
- V. Issarny, PC member of COOPIS'18, FASE'18 & 20, IC2E'19, ICDCS'18, ICSE'18 & 20, ICSE-SEIS-19, IOTDI'19, Middleware'18, OPODIS'18, WWW'19.
- R. Teixeira, PC member of ACM SIGCOMM 2018.
- R. Teixeira, PC member of ACM HotNets workshop 2018, ACM SIGCOMM Workshop on Self-Driving Networks 2018.
- N. Georgantas is PC member of the following international conferences: SAC'18&'19, AmI'18, SOSE'18&'19, WWW'19.
- N. Georgantas is PC member of the following international workshops: MRT'18, SERENE'18, IoT-ASAP'18&'19, ARM'18, SESOS-ICSE'18.

10.1.2. Journal

10.1.2.1. Member of the Editorial Boards

- V. Issarny, Associate editor, ACM Transactions on Internet of Things (TIOT)
- V. Issarny, Associate editor, IEEE Transactions on Services Computing (TSC)
- V. Issarny, Associate editor, IEEE Transactions on Software Engineering (TSE)
- N. Georgantas, Associate editor, International Journal of Ambient Computing and Intelligence (IJACI)
- V. Christophides, Associate editor, MDPI Open Access Journal of Internet of Things (IoT)
- V. Christophides, Associate editor, MDPI Open Access Journal of Future Internet

10.1.2.2. Reviewer - Reviewing Activities

- V. Christophides, Reviewer for the ACM Journal Transactions on Internet of Things (TIOT)
- V. Christophides, Reviewer for the Springer Journal Knowledge and Information Systems (KAIS)

10.1.3. Invited Talks

- V. Christophides "IoT Data Analytics", Invited Tutorial at the French Database Community (BDA) Summer School on Scalable Data Analytics Aussois, Modane June 17-22, 2018.
- V. Issarny, "USNB - Enabling Universal Online Social Interactions", Keynote, ECSCW - The 16th European Conference on Computer-Supported Cooperative Work: The International venue on Practice-centred computing and the Design of cooperation technologies. Nancy, FR, June 2018.
- V. Issarny, "When service-oriented and cloud computing meets the IoT: A use case in the context of urban mobile crowd-sensing", Keynote, ESOCC Conference. COMO, IT, September 2018.
- R. Teixeira, “Diagnosis of Internet Quality of Experience in Home Networks”, Keynote, 6th International Conference, NETYS 2018, Essaouira, Morocco, May 9–11, 2018.
- N. Georgantas, "Challenges in Enabling Effective Smart Cities and Communities: Global Perspectives", Plenary panel, IEEE SMARTCOMP 2018, Taormina, Sicily, June 18-20, 2018.

10.1.4. Leadership within the Scientific Community

- V. Christophides, Member of the EDBT Association (since 2014).
- V. Issarny, Council member, ACM Europe (since 2017).

10.1.5. Scientific Expertise

- V. Issarny, Elected member: *Commission d’Evaluation Inria*.
- V. Issarny, Scientific council member: ARCEP (since September 2018).
- V. Issarny, Committee member: *Advanced Data science Alliance Expert Panel of the Canadian Networks of Centres of Excellence* (NCE) (Ottawa, Canada, 08/18); *JPI Urban Europe and NSFC China pilot call on Sustainable and Liveable Cities and Urban Areas* (Paris, 10/18); *FWO Expert panel for PhD fellowships strategic basic research* (Bruxelles, 11/18)
- R. Teixeira, evaluator for the H2020-ICT-2018-2 call, topic ICT-21-2018 “EU-US Collaboration for advanced wireless platforms” organised by the European Commission.
- R. Teixeira, monitor for the H2020 projects CogNet, SUPERFLUIDITY.
- R. Teixeira, Member of the acceptability committee for the competitive selection of young researchers (CR), 2018.
- R. Teixeira, Member of the admission committee for the competitive selection of young researchers (CR), 2018.
- V. Christophides, Member of the appointment committee for faculty position on «Machine Learning with focus on Bioinformatics» University of Crete, December 2017.
- N. Georgantas, monitor for the ANR project INTEROP.

10.1.6. Research Administration

- V. Issarny, Scientific coordinator: IIL Inria@SiliconValley (till 08/18) and IPL CityLab@Inria (till 11/18).
- N. Georgantas, member of the PhD monitoring committee at Inria Paris (till 10/18).

10.2. Teaching - Supervision - Juries

10.2.1. Teaching

Master: V. Christophides, "Big Data Processing and Analytics", 25h M2, Ecole CentraleSupelec, France. Created this new master’s level class in Spring 2016.

Master : R. Teixeira, “Methodology for research in networking”, 4h CM, M2, Sorbonne University, France

Master : R. Teixeira, “Network Metrology”, 4h CM, M2, Sorbonne University, France

E-learning

V. Issarny, MOOC & SPOC *Implementing Successful Processes for Citizen Participation: Supporting Methods and Civic Tech*, 5 weeks, 1,600 registered students in the session of November 2018.

R. Teixeira, MOOC “Internet Measurements: a Hands-on Introduction”, 5 weeks in FUN platform, session May 2018.

10.2.2. Supervision

PhDs in progress:

- Yifan Du (Since October 2017): "In-network collaborative crowd-Xing", Sorbonne University, V. Issarny and F. Sailhan (CNAM)
- Thibaut Ehlinger (Since November 2017): «Causal Modelling of User Quality Experience (QoE) from Network Quality of Service (QoS)» University of Côte d’Azur. advisors: V. Christophides and C. Barakat (Inria)
- Grigoris Piperagkas (Since October 2018): "Leveraging universal social networking and the IoT for urban-scale participatory systems", Sorbone University, V. Issarny and R. Angarita (ISEP).
- Israel Marquez Salinas (Since October 2018): “Optimization of Internet Quality of Experience in Home Networks”, S. Ayoubi, R. Teixeira.
- Dimitris Tsolovos (Since March 2017): "A privacy-by-design middleware for urban-scale mobile crowdsensing", UVSQ, N. Anciaux (Inria PETRUS @ Saclay) and V. Issarny.
- Fethi Dilmi (Since October 2017): “End-to-end monitoring and diagnosis of video Quality of Experience”, F. Bronzino and R. Teixeira.
- Sarah Wassermann (Since October 2017): “Passive analysis and opitimization of Internet Quality of Experience”, R. Teixeira.
- Sara El Aouad (Since May 2013): “Building a personalized summary from movie reviews”, V. Christophides, R. Teixeira, C. Diot (Google).
- Patient Ntumba (Since August 2018): “Dynamic adaptation of middleware-layer protocols for emergent mobile systems”, Sorbonne University, N. Georgantas.

10.2.3. Juries

V. Christophides: External evaluator of the Ph.D. Thesis of Giorgos Borboudakis entitled "Efficient and Accurate Feature Selection, with Extensions for Multiple Solutions and to Big Data", University of Crete, Greece, November 2018

V. Issarny: President of the PhD defense committee of Saad El Jaouhari on the topic "A secure design of WoT services for smart cities", IMT-Atlantique, Rennes, FR, December 2018

R. Teixeira, external evaluator of the Ph.D dissertation of Danilo Cicalese on “Anycast nowadays” (Telecom ParisTech, 2018).

R. Teixeira, external evaluator of the Ph.D dissertation of Luca Vassio on “Data Analysis and Modelling of Users’ Behaviour on the Web” (Politecnico di Torino, 2018).

10.3. Popularization

10.3.1. Internal or external Inria responsibilities

- V. Issarny, Appointed member: CODEV, *Métropole du Grand Paris* (Since December 2018).

11. Bibliography

Major publications by the team in recent years

[1] A. BENNACEUR, V. ISSARNY. *Automated Synthesis of Mediators to Support Component Interoperability*, in "IEEE Transactions on Software Engineering", 2015, 22 p. , https://hal.inria.fr/hal-01076176

[2] B. BILLET, V. ISSARNY. *Spinel: An Opportunistic Proxy for Connecting Sensors to the Internet of Things*, in "ACM Transactions on Internet Technology", March 2017, vol. 17, n° 2, pp. 1 - 21 [*DOI :* 10.1145/3041025], https://hal.inria.fr/hal-01505879

[3] G. BLAIR, A. BENNACEUR, N. GEORGANTAS, P. GRACE, V. ISSARNY, V. NUNDLOLL, M. PAOLUCCI. *The Role of Ontologies in Emergent Middleware: Supporting Interoperability in Complex Distributed Systems*, in "Big Ideas track of ACM/IFIP/USENIX 12th International Middleware Conference", Lisbon, Portugal, 2011, http://hal.inria.fr/inria-00629059/en

[4] M. CAPORUSCIO, P.-G. RAVERDY, V. ISSARNY. *ubiSOAP: A Service Oriented Middleware for Ubiquitous Networking*, in "IEEE Transactions on Services Computing", 2012, vol. 99 [*DOI :* 10.1109/TSC.2010.60], http://hal.inria.fr/inria-00519577

[5] I. CUNHA, R. TEIXEIRA, D. VEITCH, C. DIOT. *DTrack: A System to Predict and Track Internet Path Changes*, in "IEEE/ACM Transactions on Networking", August 2014, vol. 22, n° 4, pp. 1025–1038, https://hal.inria.fr/hal-01097439

[6] O. GOGA, P. LOISEAU, R. SOMMER, R. TEIXEIRA, K. P. GUMMADI. *On the Reliability of Profile Matching Across Large Online Social Networks*, in "KDD'15: ACM SIGDD Conference on Knowledge Discovery and Data Mining", Sydeny, Australia, August 2015 [*DOI :* 10.1145/2783258.2788601], https://hal.inria.fr/hal-01162402

[7] S. HACHEM, A. PATHAK, V. ISSARNY. *Service-Oriented Middleware for Large-Scale Mobile Participatory Sensing*, in "Pervasive and Mobile Computing", 2014, http://hal.inria.fr/hal-00872407

[8] V. ISSARNY, N. GEORGANTAS, S. HACHEM, A. ZARRAS, P. VASSILIADIS, M. AUTILI, M. A. GEROSA, A. BEN HAMIDA. *Service-Oriented Middleware for the Future Internet: State of the Art and Research Directions*, in "Journal of Internet Services and Applications", May 2011, vol. 2, n° 1, pp. 23-45 [*DOI :* 10.1007/s13174-011-0021-3], http://hal.inria.fr/inria-00588753/en

[9] K. MIRYLENKA, V. CHRISTOPHIDES, T. PALPANAS, I. PEFKIANAKIS, M. MAY. *Characterizing Home Device Usage From Wireless Traffic Time Series*, in "19th International Conference on Extending Database Technology (EDBT)", Bordeaux, France, March 2016, https://hal.inria.fr/hal-01249778

Publications of the year

Doctoral Dissertations and Habilitation Theses

[10] N. GEORGANTAS. *Service Oriented Computing in Mobile Environments: Abstractions and Mechanisms for Interoperability and Composition*, Sorbonne Université, February 2018, Habilitation à diriger des recherches, https://hal.inria.fr/tel-01740629

Articles in International Peer-Reviewed Journals

[11] G. BAJAJ, R. AGARWAL, P. SINGH, N. GEORGANTAS, V. ISSARNY. *4W1H in IoT semantics*, in "IEEE Access", October 2018, vol. 6, pp. 65488-65506 [*DOI :* 10.1109/ACCESS.2018.2878100], https://hal.inria.fr/hal-01898506

[12] J. LANZA, L. SANCHEZ, J. R. SANTANA, R. AGARWAL, N. KEFALAKIS, P. GRACE, T. ELSALEH, M. ZHAO, E. TRAGOS, H. NGUYEN, F. CIRILLO, R. STEINKE, J. SOLDATOS. *Experimentation as a Service*

over Semantically Interoperable Internet of Things Testbeds, in "IEEE Access", September 2018, vol. 18, n^{o} 1, pp. 51607 - 51625 [*DOI :* 10.1109/ACCESS.2018.2867452], https://hal.inria.fr/hal-01876881

[13] L. SANCHEZ, J. LANZA, J. R. SANTANA, R. AGARWAL, P. G. RAVERDY, T. ELSALEH, Y. FATHY, S. JEONG, A. DADOUKIS, T. KORAKIS, S. KERANIDIS, P. O'BRIEN, J. HORGAN, A. SACCHETTI, G. MASTANDREA, A. FRAGKIADAKIS, P. CHARALAMPIDIS, N. SEYDOUX, C. ECREPONT, M. ZHAO. *Federation of Internet of Things Testbeds for the Realization of a Semantically-Enabled Multi-Domain Data Marketplace*, in "Sensors", October 2018, vol. 18, n^{o} 10, 34 p. [*DOI :* 10.3390/s18103375], https://hal.inria.fr/hal-01889896

[14] I. TSAMARDINOS, G. BORBOUDAKIS, P. KATSOGRIDAKIS, P. PRATIKAKIS, V. CHRISTOPHIDES. *A greedy feature selection algorithm for Big Data of high dimensionality*, in "Machine Learning Journal", August 2018, https://hal.inria.fr/hal-01960888

[15] R. VENTURA, V. MALLET, V. ISSARNY. *Assimilation of mobile phone measurements for noise mapping of a neighborhood*, in "Journal of the Acoustical Society of America", September 2018, vol. 144, n^{o} 3, pp. 1279 - 1292 [*DOI :* 10.1121/1.5052173], https://hal.inria.fr/hal-01909933

Articles in Non Peer-Reviewed Journals

[16] V. MALLET, R. VENTURA, V. ISSARNY, P. G. RAVERDY. *Assimilation d'observations participatives issues de l'application mobile Ambiciti*, in "Acoustique et Techniques : trimestriel d'information des professionnels de l'acoustique", January 2018, https://hal.inria.fr/hal-01909880

Invited Conferences

[17] V. ISSARNY, G. BOULOUKAKIS, N. GEORGANTAS, F. SAILHAN, G. TEXIER. *When Service-oriented Computing Meets the IoT: A Use Case in the Context of Urban Mobile Crowdsensing*, in "ESOCC 2018 - 7th European Conference on Service-oriented and Cloud Computing", Como, Italy, September 2018, https://hal.inria.fr/hal-01871515

[18] B. LEFÈVRE, V. ISSARNY. *Matching Technological & Societal Innovations: The Social Design of a Mobile Collaborative App for Urban Noise Monitoring*, in "4th IEEE International Conference on Smart Computing", Taormina, Italy, June 2018, https://hal.inria.fr/hal-01801314

International Conferences with Proceedings

[19] K. E. BENSON, G. BOULOUKAKIS, C. GRANT, V. ISSARNY, S. MEHROTRA, I. MOSCHOLIOS, N. VENKATASUBRAMANIAN. *FireDeX: a Prioritized IoT Data Exchange Middleware for Emergency Response*, in "ACM/IFIP/USENIX Middleware conference", Rennes, France, December 2018, vol. 18 [*DOI :* 10.1145/3274808.3274830], https://hal.inria.fr/hal-01877555

[20] G. BOULOUKAKIS, A. KATTEPUR, N. GEORGANTAS, V. ISSARNY. *Queueing Network Modeling Patterns for Reliable and Unreliable Publish/Subscribe Protocols*, in "MobiQuitous 2018 - 15th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services", New York, United States, November 2018, https://hal.inria.fr/hal-01893926

[21] G. BOULOUKAKIS, I. MOSCHOLIOS, N. GEORGANTAS, V. ISSARNY. *Simulation-based Queueing Models for Performance Analysis of IoT Applications*, in "11th International Symposium on Communication Systems,

Networks, and Digital Signal Processing (CSNDSP)", Budapest, Hungary, July 2018, https://hal.inria.fr/hal-01797930

[22] F. M. COSTA, N. GEORGANTAS, R. GOMES, R. DA ROCHA, G. BOULOUKAKIS. *Cross-layer QoS-Aware Resource Allocation for IoT-Enabled Service Choreographies*, in "The 5th International Workshop on Middleware and Applications for the Internet of Things (M4IoT)", Rennes, France, December 2018, https://hal.inria.fr/hal-01943541

[23] D. DA HORA, A. S. ASRESE, V. CHRISTOPHIDES, R. TEIXEIRA, D. ROSSI. *Narrowing the gap between QoS metrics and Web QoE using Above-the-fold metrics*, in "PAM 2018 - International Conference on Passive and Active Network Measurement", Berlin, Germany, March 2018, pp. 1-13, https://hal.inria.fr/hal-01677260

[24] D. DA HORA, K. VAN DOORSELAER, K. VAN OOST, R. TEIXEIRA. *Predicting the effect of home Wi-Fi quality on QoE*, in "INFOCOM 2018 - IEEE International Conference on Computer Communications", Honolulu, United States, April 2018, pp. 1-10, https://hal.inria.fr/hal-01677214

[25] V. EFTHYMIOU, G. PAPADAKIS, K. STEFANIDIS, V. CHRISTOPHIDES. *Simplifying Entity Resolution on Web Data with Schema-agnostic, Non-iterative Matching*, in "ICDE 2018 - 34th IEEE International Conference on Data Engineering", Paris, France, April 2018, pp. 1-4, https://hal.inria.fr/hal-01718040

[26] V. EFTHYMIOU, G. PAPADAKIS, K. STEFANIDIS, V. CHRISTOPHIDES. *MinoanER: Schema-Agnostic, Non-Iterative, Massively Parallel Resolution of Web Entities*, in "EDBT/ICDT 2019 Joint Conference", Lisbon, Portugal, March 2019, https://hal.inria.fr/hal-01960933

[27] *Best Paper*
R. GOMES, G. BOULOUKAKIS, F. COSTA, N. GEORGANTAS, R. DA ROCHA. *QoS-Aware Resource Allocation for Mobile IoT Pub/Sub Systems*, in "2018 International Conference on Internet of Things (ICIOT)", Seattle, United States, June 2018, https://hal.inria.fr/hal-01797933.

[28] P. NTUMBA, G. BOULOUKAKIS, N. GEORGANTAS. *Interconnecting and Monitoring Heterogeneous Things in IoT Applications*, in "International Conference on Web Engineering (ICWE)", Caceres, Spain, June 2018, https://hal.inria.fr/hal-01771857

[29] G. TEXIER, V. ISSARNY. *Leveraging the Power of the Crowd and Offloading Urban IoT Networks to Extend their Lifetime*, in "LANMAN 2018 - IEEE International Symposium on Local and Metropolitan Area Networks", Washington, DC, United States, June 2018, pp. 1-6, https://hal-imt-atlantique.archives-ouvertes.fr/hal-01813313

National Conferences with Proceedings

[30] D. TSOLOVOS, N. ANCIAUX, V. ISSARNY. *A Privacy Aware Approach for Participatory Sensing Systems*, in "34ème Conférence sur la Gestion de Données – Principes, Technologies et Applications", Bucharest, Romania, October 2018, https://hal.inria.fr/hal-01947863

Research Reports

[31] D. DA HORA, K. VAN DOORSELAER, K. VAN OOST, R. TEIXEIRA. *Predicting the effect of home Wi-Fi quality on QoE: Extended Technical Report*, Inria ; Technicolor ; Telecom ParisTech, January 2018, https://hal.inria.fr/hal-01676921

Other Publications

[32] R. AGARWAL, N. GEORGANTAS, V. ISSARNY. *Impact of Spatial Scale on Mobility Parameters*, June 2018, NetSci 2018 - International Conference on Network Science, Poster, https://hal.inria.fr/hal-01923856

[33] D. DA HORA, D. ROSSI, V. CHRISTOPHIDES, R. TEIXEIRA. *A practical method for measuring Web above-the-fold time*, August 2018, ACM SIGCOMM 2018 Conference, Poster, https://hal.inria.fr/hal-01960984

[34] I. MORANDI. *Service traceroute: Tracing Paths of Application Flows*, UPMC - Paris 6 Sorbonne Universités, September 2018, https://hal.inria.fr/hal-01888618

[35] M.-I. SALINAS. *Home Wi-Fi Impairments*, UPMC (Paris 6) - Sorbonne Université, September 2018, https://hal.inria.fr/hal-01889806

[36] L. SCALZOTTO, K. E. BENSON, G. BOULOUKAKIS, P. BELLAVISTA, V. ISSARNY, S. MEHROTRA, N. VENKATASUBRAMANIAN. *An Implementation Experience with SDN-enabled IoT Data Exchange Middleware*, December 2018, Middleware 2018 - ACM/IFIP/USENIX Middleware conference, Poster, https://hal.inria.fr/hal-01895274