Activity Report 2019

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, until Jun 2019]
Renata Cruz Teixeira [Inria, Senior Researcher, HDR]
Valérie Issarny [Inria, Senior Researcher, HDR]

Post-Doctoral Fellows

Sara Ayoubi [Inria, Post-Doctoral Fellow, until Apr 2019]
Georgios Bouloukakis [Inria, Post-Doctoral Fellow, until Apr 2019]
Giulio Grassi [Inria, Post-Doctoral Fellow]

PhD Students

William Aboucaya [Sorbonne Université, PhD Student, from Oct 2019]
Fethi Dilmi [Inria, PhD Student, until Nov 2019]

Yifan Du [Inria, PhD Student]
Israel Marquez Salinas [Univ Pierre et Marie Curie, PhD Student, until Mar 2019]
Patient Ntumba [Inria, PhD Student]
Grigorios Piperagkas [Inria, PhD Student]
Sarah Wassermann [Inria, PhD Student, until Apr 2019]

Technical staff

Ehsan Ahvar [Inria, Engineer, until Jun 2019]
Pierre-Guillaume Raverdy [Inria, Engineer, from Mar 2019]

Interns and Apprentices

Lior Diler [Inria, from Nov 2019]
Mohammad Syed [Inria, from Mar 2019 until Aug 2019]

Administrative Assistants

Nathalie Gaudechoux [Inria, Administrative Assistant]
Meriem Guemair [Inria, Administrative Assistant]

Visiting Scientist

Mark Crovella [Boston University, until Jul 2019]

External Collaborators

Shohreh Ahvar [ISEP, until Jun 2019]
Rafael Angarita Arocha [ISEP]
Sara Ayoubi [Nokia Bell Labs, from May 2019]
Zied Ben Houidi [Nokia Bell Labs, until Apr 2019]
Francesco Bronzino [Nokia Bell Labs]
Vassilis Christophides [Univ de Cergy Pontoise, from Jul 2019]
Bruno Lefevre [Univ Paris Nord]
Françoise Sailhan [CNAM]
Rachit Agarwal [IIT Kanpur]
Georgios Bouloukakis [Univ of California, Irvine, from May 2019]

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 as well as methods to directly measure application quality. We 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. One such domain 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 has been and keeps 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:

- **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*".

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.

In our patternship with Princeton University (associate team HOMENET) we have deployed monitoring infrastructure within users' homes. We are developing a mostly passive measurement system to monitor the performance of user applications, which we call Network Microscope. We are developing Network Microscope to run in a box acting as home gateway. We have deployed these boxes in 50 homes in the US and 10 in France. The US deployment was ran and financed by the Wall Street Journal. They were interested in understanding 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 Network Microscope 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) or devices with multiple 3G/LTE links. To achieve this goal, we plan to leverage multi-path and cross-layer optimizations.

4.4. Internet Scanning

Internet-wide scanning has enabled researchers to answer a wealth of new security and measurement questions ranging from "How are authoritarian regimes spying on journalists?" to "Are security notifications effective at prompting operators to patch?" Most of these studies have used tools like ZMap, which operates naıvely, scanning every IPv4 address once. This simplicity enables researchers to easily answer a question once, but the methodology scales poorly when continually scanning to detect changes, as networks change at dramatically

different rates. Service configurations change more frequently on cloud providers like Amazon and Azure than on residential networks. Internet providers in developing regions often have extremely short DHCP windows. Some networks are unstable with host presence varying wildly between different hours and others have distinct periodic patterns, e.g., hosts are only available during regional business hours. A handful of large autonomous systems have not had hosts present in decades. Our work in collaboration with Stanford University is developing more intelligent Internet-wide scanning methods to then implement a system that can scan continuously. Such a system will allow for up-to-date analysis of Internet trends and threats with real-time alerts of important events.

5. Highlights of the Year

5.1. Highlights of the Year

The Wall Street Journal published an article, "The Truth About Faster Internet: It's Not Worth It" (on the front page of the printed version of the paper on August 21, 2019) based on the results of the models we have developed to infer video quality from encrypted network traffic and the Network Microscope system.

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. 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

6.4. 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

6.5. Network Microscope

KEYWORDS: Quality of Experience - Network monitoring - Video analysis

FUNCTIONAL DESCRIPTION: A system that accurately infers video streaming quality metrics in real time, such as startup delay or video resolution, by using just a handful of features extracted from passive traffic measurement. Network Microscope passively collects a corpus of network features about the traffic flows of interest in the network and directs those to a real-time analytics framework that can perform more complex inference tasks. Network Microscope enables network operators to determine degradations in application quality as they happen, even when the traffic is encrypted.

- Participants: Francesco Bronzino and Renata Cruz Teixeira
- Contact: Renata Cruz Teixeira
- URL: https://netmicroscope.com/

6.6. 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. Automated Synthesis of Mediators for Middleware-Layer Protocol Interoperability in the IoT

Participants: Georgios Bouloukakis, Nikolaos Georgantas, Patient Ntumba, Valérie Issarny (MiMove)
To enable direct Internet connectivity of Things, complete protocol stacks need to be deployed on resource-constrained devices. Such protocol stacks typically build on lightweight IPv6 adaptations and may even include a middleware layer supporting high-level application development. However, the profusion of IoT middleware-layer interaction protocols has introduced technology diversity and high fragmentation in the IoT systems landscape with siloed vertical solutions. To enable the interconnection of heterogeneous Things across these barriers, advanced interoperability solutions at the middleware layer are required. In this paper, we introduce a solution for the automated synthesis of protocol mediators that support the interconnection of heterogeneous Things. Our systematic approach relies on the Data eXchange (DeX) connector model, which comprehensively abstracts and represents existing and potentially future IoT middleware protocols. Thanks to DeX, Things seamlessly interconnect through lightweight mediators. We validate our solution with respect to: (i) the support to developers when developing heterogeneous IoT applications; (ii) the runtime performance of the synthesized mediators.

7.2. Probabilistic Event Dropping for Intermittently Connected Subscribers over Pub/Sub Systems

Participants: Georgios Bouloukakis, Nikolaos Georgantas (MiMove), Ioannis Moscholios (Univ of Peloponnese)
Internet of Things (IoT) aim to leverage data from multiple sensors, actuators and devices for improving peoples' daily life and safety. Multiple data sources must be integrated, analyzed from the corresponding application and notify interested stakeholders. To support the data exchange between data sources and stakeholders, the publish/subscribe (pub/sub) middleware is often employed. Pub/sub provides additional mechanisms such as reliable messaging, event dropping, prioritization, etc. The event dropping mechanism is often used to satisfy Quality of Service (QoS) requirements and ensure system stability. To enable event dropping, basic approaches apply finite buffers or data validity periods and more sophisticated ones are information-aware. In this paper, we introduce a pub/sub mechanism for probabilistic event dropping by considering the stakeholders' intermittent connectivity and QoS requirements. We model the pub/sub middleware as a network of queues which includes a novel ON/OFF queueing model that enables the definition of join probabilities. We validate our analytical model via simulation and compare our mechanism with existing ones. Experimental results can be used as insights for developing hybrid dropping mechanisms.

7.3. Adaptive Mediation for Data Exchange in IoT Systems

Participants: Georgios Bouloukakis (MiMove & Univ of California, Irvine), Andrew Chio, Sharad Mehrotra, Nalini Venkatasubramanian (Univ of California, Irvine), Cheng-Hsin Hsu (National Tsing Hua Univ)
Messaging and communication is a critical aspect of next generation Internet-of-Things (IoT) systems where interactions among devices, software systems/services and end-users is the expected mode of operation. Given the diverse and changing communication needs of entities, the data exchange interactions may assume different protocols (MQTT, CoAP, HTTP) and interaction paradigms (point to point, multicast, unicast). In this paper, we address the issue of supporting adaptive communications in IoT systems through a mediation-based architecture for data exchange. Here, components called mediators support protocol translation to bridge the heterogeneity gap. Aiming to provide a placement of mediators to nodes, we introduce an integer linear programming solution that takes as input: a set of Edge nodes, IoT devices, and networking semantics. Our proposed solution achieves adaptive placement resulting in timely interactions between IoT devices for larger topologies of IoT spaces.

7.4. Universal Social Network Bus: Toward the Federation of Heterogeneous Online Social Network Services

Participants: Valérie Issarny, Nikolaos Georgantas, Ehsan Ahvar, Bruno Lefèvre, Shohreh Ahvar (MiMove), Rafael Angarita (ISEP Paris)
Online Social Network Services (OSNSs) are changing the fabric of our society, impacting almost every aspect of it. Over the past few decades, an 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 *Universal Social Network Bus* (USNB), 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 and profile- and non-profile-based OSNSs through a reference social interaction service. We ran a 1-day workshop with a panel of users who experimented with the USNB prototype to assess the potential benefits of social interoperability for social network users. Results show the positive evaluation of users for USNB, especially as an enabler of applications for civic participation. This further opens up new perspectives for future work, among which includes enforcing security and privacy guarantees.

7.5. Social Middleware for Civic Engagement

Participants: Valérie Issarny, Nikolas Georgantas, Grigoris Piperagkas (MiMove), Rafael Angarita (ISEP Paris)
Civic engagement refers to any collective action towards the identification and solving of public issues. Current civic technologies are traditional Web- or mobile-based platforms that make difficult, or just impossible, the participation of citizens via different communication technologies. Moreover, connected objects sensing physical-world data can nourish participatory processes by providing physical evidence to citizens; however, leveraging these data is not direct and still a time-consuming process for civic technologies developers. We introduce the concept of *social middleware* for civic engagement. Social middleware allows citizens to engage in participatory processes -supported by civic technologies- via their favorite communication tools, and to interact not only with other citizens but also with relevant connected objects and software platforms. The mission of social middleware goes beyond the connection of all these heterogeneous entities. It aims at easing the implementation of distributed applications oriented toward civic engagement by featuring dedicated built-in services.

7.6. Mobile Crowd-Sensing as a Resource for Contextualized Urban Public Policies

Participants: Valfie Issarny, Bruno Lefèvre, Rachit Agarwal (MiMove), Vivien Mallet (Inria Ange)
Environmental noise is a major pollutant in contemporary cities and calls for the active monitoring of noise levels to spot the locations where it most affects the people's health and well-being. However, due to the complex relationship between environmental noise and its perception by the citizens, it is not sufficient to quantitatively measure environmental noise. We need to collect and aggregate contextualized –both quantitative and qualitative– data about the urban environmental noise so as to be able to study the objective and subjective relationships between sound and living beings. This complex knowledge is a prerequisite for making efficient territorial public policies for soundscapes that are inclined towards living beings welfare. In this paper, we investigate how Mobile Phone Sensing (MPS) –*aka* crowdsensing– enables the gathering of such knowledge, provided the implementation of sensing protocols that are customized according to the context of use and the intended exploitation of the data. Through three case studies that we carried out in France and Finland, we show that MPS is not solely a tool that contributes to sensitizing citizens and decision-makers about noise pollution; it also contributes to increasing our knowledge about the impact of the environmental noise on people's health and well-being in relation to its physical and subjective perception.

7.7. Multi-Sensor Calibration Planning in IoT-Enabled Smart Spaces

Participants: Valérie Issarny (MiMove), Françoise Sailhan (CNAM), Qiuxi Zhu, Md Yusuf Sarwar Uddin, Nalini Venkatasubramanian (University of California, Irvine)
Emerging applications in smart cities and communities require massive IoT deployments using sensors/actuators (things) that can enhance citizens' quality of life and public safety. However, budget constraints often lead to limited instrumentation and/or the use of low-cost sensors that are subject to drift and bias. This raises concerns of robustness and accuracy of the decisions made on uncertain data. To enable effective decision making while fully exploiting the potential of low-cost sensors, we propose to send mobile units (e.g., trained personnel) equipped with high-quality (more expensive) and freshly-calibrated reference sensors so as to carry out calibration in the field. We design and implement an efficient cooperative approach to solve the calibration planning problem, which aims at minimizing the cost of the recurring calibration of multiple sensor types in the long-term operation. We propose a two-phase solution that consists of a sensor selection phase that minimizes the average cost of recurring calibration, and a path planning phase that minimizes the travel cost of multiple calibrators which have load constraints. We provide fast and effective heuristics for both phases. We further build a prototype that facilitates the mapping of the deployment field and provides navigation guidance to mobile calibrators. Extensive use-case-driven simulations show that our proposed approach significantly reduces the average cost compared to naive approaches: up to 30% in a moderate-sized indoor case, and higher in outdoor cases depending on scale

7.8. User-Centric Context Inference for Mobile Crowdsensing

Participants: Yifan Du, Valérie Issarny (MiMove), Françoise Sailhan (CNAM)
Mobile crowdsensing is a powerful mechanism to aggregate hyperlocal knowledge about the environment. Indeed, users may contribute valuable observations across time and space using the sensors embedded in their smartphones. However, the relevance of the provided measurements depends on the adequacy of the sensing context with respect to the phenomena that are analyzed. Our research concentrates more specifically on assessing the sensing context when gathering observations about the physical environment beyond its geographical position in the Euclidean space, i.e., whether the phone is in-/out-pocket, in-/out-door and on-/under-ground. We introduce an online learning approach to the local inference of the sensing context so as to overcome the disparity of the classification performance due to the heterogeneity of the sensing devices as well as the diversity of user behavior and novel usage scenarios. Our approach specifically features a hierarchical algorithm for inference that requires few opportunistic feedbacks from the user, while increasing the accuracy of the context inference per user.

7.9. Let Opportunistic Crowdsensors Work Together for Resource-efficient, Quality-aware Observations

Participants: Yifan Du, Valérie Issarny (MiMove), Françoise Sailhan (CNAM)
Opportunistic crowdsensing empowers citizens carrying hand-held devices to sense physical phenomena of common interest at a large and fine-grained scale without requiring the citizens' active involvement. However, the resulting uncontrolled collection and upload of the massive amount of contributed raw data incur significant resource consumption, from the end device to the server, as well as challenge the quality of the collected observations. Our research tackles both challenges raised by opportunistic crowdsensing, that is, enabling the resource-efficient gathering of relevant observations. To achieve so, we introduce the *BeTogether* middleware fostering context-aware, collaborative crowdsensing at the edge so that co-located crowdsensors operating in the same context, group together to share the work load in a cost- and quality-effective way. Our implementation-driven evaluation of the proposed solution, which leverages a dataset embedding nearly one million entries contributed by 550 crowdsensors over a year, shows that *BeTogether* increases the quality of the collected data while reducing the overall resource cost compared to the cloud-centric approach.

7.10. Detecting Mobile Crowdsensing Context in the Wild

Participants: Rachit Agarwal, Shaan Chopra, Vassilis Christophides, Nikolaos Georgantas, Valérie Issarny (MiMove)
Understanding the sensing context of raw data is crucial for assessing the quality of large crowdsourced spatio-temporal datasets and supporting context-augmented personal trajectories. Detecting sensing contexts in the wild is a challenging task and requires features from smartphone sensors that are not always available. In this paper, we propose three heuristic algorithms for detecting sensing contexts such as in/out-pocket, under/over-ground, and in/out-door for crowdsourced spatio-temporal datasets. These are unsupervised binary classifiers with a small memory footprint and execution time. Using a segment of the Ambiciti real dataset-a feature-limited crowdsourced dataset-we report that our algorithms perform equally well in terms of balanced accuracy (within 4.3%) when compared to machine learning (ML) models reported by an AutoML tool.

7.11. Inferring Streaming Video Quality from Encrypted Traffic: Practical Models and Deployment Experience

Participants: Francesco Bronzino, Sara Ayouibi, Renata Teixeira (MiMove), Paul Schmitt (Princeton), Guilherme Martins, Nick Feamster (University of Chicago)
Inferring the quality of streaming video applications is important for Internet service providers, but the fact that most video streams are encrypted makes it difficult to do so. We develop models that infer quality metrics (i.e., startup delay and resolution) for encrypted streaming video services. Our paper builds on previous work, but extends it in several ways. First, the models work in deployment settings where the video sessions and segments must be identified from a mix of traffic and the time precision of the collected traffic statistics is more coarse (e.g., due to aggregation). Second, we develop a single composite model that works for a range of different services (i.e., Netflix, YouTube, Amazon, and Twitch), as opposed to just a single service. Third, unlike many previous models, our models perform predictions at finer granularity (e.g., the precise startup delay instead of just detecting short versus long delays) allowing to draw better conclusions on the ongoing streaming quality. Fourth, we demonstrate the models are practical through a 16-month deployment in 66 homes and provide new insights about the relationships between Internet "speed" and the quality of the corresponding video streams, for a variety of services; we find that higher speeds provide only minimal improvements to startup delay and resolution. This work was accepted for publication at the ACM SIGMETRICS conference. The models we developed in this work and the findings were the basis for a first-page story published on The Wall Street Journal ("The Truth About Faster Internet: It's Not Worth It"). [1]

[1] The article is available online at: https://www.wsj.com/graphics/faster-internet-not-worth-it/?mod=article_inline&mod=hp_lead_pos5.

7.12. Implications of User Perceived Page Load Time Multi-Modality on Web QoE Measurement

Participants: Renata Teixeira, Vassilis Christophides (MiMove), Flavia Salutari, Diego Da Hora (Telecom Paris Tech), Matteo Varvello (Brave Software), Dario Rossi (Huawei)
Web browsing is one of the most popular applications for both desktop and mobile users. A lot of effort has been devoted to speedup the Web, as well as in designing metrics that can accurately tell whether a webpage loaded fast or not. An often implicit assumption made by industrial and academic research communities is that a *single* metric is sufficient to assess whether a webpage loaded fast. In this work we collect and make publicly available a unique dataset which contains webpage features (e.g., number and type of embedded objects) along with both *objective* and *subjective* Web quality metrics. This dataset was collected by crawling over 100 websites—representative of the top 1 M websites in the Web—while crowdsourcing 6,000 user opinions on *user perceived page load time* (uPLT). In contrast to related work, we show that the uPLT distribution is often multimodal and that, in practice, no more than three modes are present. The main conclusion drawn from our analysis is that, for complex webpages, each of the different objective QoE metrics proposed in the literature (such as AFT, TTI, PLT, etc.) is suited to approximate one of the different uPLT modes.

7.13. The News We Like Are Not the News We Visit: News Categories Popularity in Usage Data

Participants: Renata Teixeira (MiMove), Giuseppe Scavo (MiMove, Nokia Bell Labs), Zied Ben-Houidi (Nokia Bell-Labs), Stefano Traverso, Marco Mellia (Politecnico di Torino)
Most of our knowledge about online news consumption comes from survey-based news market reports, partial usage data from a single editor, or what people publicly share on social networks. Our work published on the 13th International AAAI Conference on Web and Social Media (ICWSM-2019) complements these sources by presenting the first holistic study of visits across online news outlets that a population uses to read news. We monitored the entire network traffic generated by Internet users in four locations in Italy. Together these users generated 80 million visits to 5.4 million news articles in about one year and a half. This unique view allowed us to evaluate how usage data complements existing data sources. We find for instance that only 16% of news visits in our datasets came from online social networks. In addition, the popularity of news categories when considering all visits is quite different from the one when considering only news discovered on social media, or visits to a single major news outlet. Interestingly, a substantial mismatch emerges between self-reported news-category preferences (as measured by Reuters Institute in the same year and same country) and their actual popularity in terms of visits in our datasets. In particular, unlike self-reported preferences expressed by users in surveys that put “Politics”, “Science” and “International” as the most appreciated categories, “Tragedies and Weird news" and “Sport” are by far the most visited. Our paper discusses two possible causes of this mismatch and conjecture that the most plausible reason is the disassociation that may occur between individuals’ cognitive values and their cue-triggered attraction.

7.14. Classification of Load Balancing in the Internet

Participants: Renata Teixeira (MiMove), Rafael Almeida, Ítalo Cunha (Universidade Federal de Minas Gerais), Darryl Veitch (University of Technology Sydney), Christophe Diot (Google)

Recent advances in programmable data planes, software-defined networking, and the adoption of IPv6, support novel, more complex load balancing strategies. We introduce the Multipath Classification Algorithm (MCA), a probing algorithm that extends traceroute to identify and classify load balancing in Internet routes. MCA extends existing formalism and techniques to consider that load balancers may use arbitrary combinations of bits in the packet header for load balancing. We propose optimizations to reduce probing cost that are applicable to MCA and existing load balancing measurement techniques. Through large-scale measurement campaigns, we characterize and study the evolution of load balancing on the IPv4 and IPv6 Internet with multiple transport protocols. Our results that will appear in the IEEE INFOCOM 2020 conference show that load balancing is more prevalent and that load balancing strategies are more mature than previous characterizations have found.

7.15. MinoanER: Schema-Agnostic, Non-Iterative, Massively Parallel Resolution of Web Entities

Participants: Vassilis Christophides (MiMove), Vasilis Efthymiou (IBM Almaden Research Center), George Papadakis (Univ of Athens), Kostas Stefanidis (Univ of Tampere)
Entity Resolution (ER) aims to identify different descriptions in various Knowledge Bases (KBs) that refer to the same entity. ER is challenged by the Variety, Volume and Veracity of entity descriptions published in the Web of Data. To address them, we propose the MinoanER framework that simultaneously fulfills full automation, support of highly heterogeneous entities, and massive parallelization of the ER process. MinoanER leverages a token-based similarity of entities to define a new metric that derives the similarity of neighboring entities from the most important relations, as they are indicated only by statistics. A composite blocking method is employed to capture different sources of matching evidence from the content, neighbors, or names of entities. The search space of candidate pairs for comparison is compactly abstracted by a novel disjunctive blocking graph and processed by a non-iterative, massively parallel matching algorithm that consists of four generic, schema-agnostic matching rules that are quite robust with respect to their internal configuration. We demonstrate that the effectiveness of MinoanER is comparable to existing ER tools over real KBs exhibiting low Variety, but it outperforms them significantly when matching KBs with high Variety.

8. Bilateral Contracts and Grants with Industry

8.1. Bilateral Grants with Industry

"Application Performance Bottleneck Detection", Comcast Gift to R. Teixeira 2018-2019.

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 BetterNet

Participants: Renata Teixeira, Vassilis Christophides, Giulio Grassi.

- **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.

9.1.1.2. Inria ADT SocialBus

Participants: Valérie Issarny , Rafael Angarita, Nikolaos Georgantas, Ehsan Ahvar , Lior Diler .

- **Name:** SocialBus – *Contributing to the development of SocialBus - A Universal Social Network Bus*
- **Period:** [July 2018 – June 2019 ; November 2019 – October 2020]
- **Partners:** Inria MiMove.

Computer-mediated communication can be defined as any form of human communication achieved through computer technology. From its beginnings, it has been shaping the way humans interact with each other, and it has influenced many areas of society. There exist a plethora of social interaction services enabling computer-mediated social communication (e.g., Skype, Facebook Messenger, Telegram, WhatsApp, Twitter, Slack, etc.). Based on personal preferences, users may prefer a social interaction services rather than another. As a result, users sharing same interests may not be able to interact since they are using incompatible technologies.

To tackle the above interoperability barrier, we propose SocialBus, a middleware solution targeted to enable the interaction via heterogeneous social interaction services. The ADT specifically supports the related implementation through the funding an engineer, toward technology transfer in the mid-term.

The SocialBus software is available under the AGPL open source license at https://gitlab.inria.fr/usnb/universal-social-network-bus.

9.2. International Initiatives

9.2.1. Inria International Labs

Inria@EastCoast

Associate Team involved in the International Lab:

9.2.1.1. HOMENET

Title: Home network diagnosis and security

International Partner (Institution - Laboratory - Researcher):

Princeton (United States) - Computer Science - 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 Muse brings an existing Web-based measurement platform, and expertise in traffic-based profiling and anomaly detection.

Inria@SiliconValley

Associate Team involved in the International Lab:

9.2.1.2. MINES

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

International Partner (Institution - Laboratory - Researcher):

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.

9.2.2. Inria International Partners

9.2.2.1. Informal International Partners

- Northeastern University (Prof. David Choffnes): We are working on methods based on active probing to diagnose poor video quality.
- Universidade Federal do Rio de Janeiro, Brazil (Prof. Edmundo Souza e Silva): We are working on characterizing Internet bottlenecks.
- Universidade Federal de Goias, Brazil (Prof. Fabio Costa): We are working on service selection and cloud resource allocation for QoS-aware enactment of service choreographies.

9.3. International Research Visitors

9.3.1. Visits of International Scientists

Mark Crovella from Boston University was visiting professor at Inria.

9.3.2. Visits to International Teams

9.3.2.1. Sabbatical programme

Renata Teixeira is visiting scholar at the Computer Science Department at Stanford University.

9.3.2.2. Research Stays Abroad

Georgios Bouloukakis was 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: Organisation

10.1.1.1. Member of the Organizing Committees

- R. Teixeira, member of the selection committee of the Heidelberg Laureate Forum 2019–2021.

10.1.2. Scientific Events: Selection

10.1.2.1. Chair of Conference Program Committees

- V. Issarny, TPC co-chair of: The International Conference on the IoT (ICIOT) at SCF'19 (June 2019, San Diego, CA, USA); and The 5th 2CM/IEEE Conference on Internet of Things Design & Implementation (IoTDI) at CPS-IoT week (April 2020, Sydney, Australia).

10.1.2.2. Member of the Conference Program Committees

- V. Issarny, member of the TPC of the following international conferences: IEEE CIC'19, ACM EUROSYS'20 (Heavy PC), ACM ESEC/FSE'19, IEEE ICDCS'20, ICSE'20, ICSE-SEIS'19 & 20, ACM/IEEE IoTDI'19, and ACM/IFIP Middleware'19.
- R. Teixeira, member of the TPC of: USENIX Symposium on Networked Systems Design and Implementation (NSDI) 2020, Passive and Active Measurement Conference (PAM) 2020, and the Buffer Sizing Workshop 2019.
- N. Georgantas, member of the TPC of the following international conferences: SAC'19&'20, SOSE'19&'20, The Web Conference'19&'20.
- N. Georgantas, member of the TPC of the following international workshops: MRT'19, IoT-ASAP'19, ARM'19, ASYDE'19.
- V. Christophides, member of the TPC of the EDBT/ICDT 2020 Joint Conference, IEEE ICDE 2020, and the 28th ACM International Conference on Information and Knowledge Management (CIKM).

10.1.3. Journal

10.1.3.1. Member of the Editorial Boards

- V. Issarny, member of the following editorial boards: The ACM Transaction on Autonomous and Adaptive Systems (TAAS), The ACM Transactions on the Internet of Things (TIOT), The IEEE Transactions on Services Computing (TSC), and The 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) and MDPI Open Access Journal of Future Internet.

10.1.4. Invited Talks

- R. Teixeira "Passive Measurements of Residential Internet Performance", Google Networking Summit, 2019.
- R. Teixeira, "Lightweight, General Inference of Streaming Video Quality from Encrypted Traffic", Workshop em ciencia de dados LNCC, Brazil, 2019.
- N. Georgantas, "Enabling Emergent Mobile IoT Systems: Middleware Interoperability & QoS Analysis", Eclipse IoT Day Paris Saclay 2019.
- V. Christophides, "Detecting Outliers in IoT Data Streams", IRT SystemX, 2019.

10.1.5. Leadership within the Scientific Community

- V. Christophides, Member of the Executive Board of the EDBT Association.

10.1.6. Scientific Expertise

- V. Issarny, member of: the FWP expert panel for PhD fellowships on strategic basic research, and the expert panel of the 2019 FET open call.
- V. Issarny, member of: the ARCEP Scientific Council (appointed), and The ACM Europe Council (Elected).
- N. Georgantas, member of: the EDITE Doctoral School selection committe for PhD fellowships, 2018 & 2019 .

10.2. Teaching - Supervision - Juries

10.2.1. Supervision

PhD: Sara el Aouad, “Personalized, Aspect-based Summarization of Movie Reviews”, Sorbonne Université, April 26, 2019, R. Teixeira, V. Christophides, C. Diot.

PhDs in progress:

- Yifan Du (Since October 2017): "In-network collaborative crowd-Xing", Sorbonne University, V. Issarny and F. Sailhan (CNAM)
- Grigoris Piperagkas (Since October 2018): "Leveraging universal social networking and the IoT for urban-scale participatory systems", Sorbonne University, V. Issarny and R. Angarita (ISEP).
- William Aboucaya (Since October 2019): "Version control for urban participatory systems", Sorbonne University, V. Issarny and R. Angarita (ISEP).
- Dimitris Tsolovos (Since March 2017): "A privacy-by-design middleware for urban-scale mobile crowdsensing", UVSQ, N. Anciaux (Inria PETRUS @ Saclay) and V. Issarny.
- Patient Ntumba (Since August 2018): “Dynamic management of IoT data stream analytics in the edge-fog-cloud continuum”, Sorbonne University, N. Georgantas and Vassilis Christophides.

10.2.2. Juries

- V. Issarny, member of the following juries: Christian Cabrera (11/19, PhD, Trinity College Dublin, Ireland), Gilles Tredan (06/19, HDR, University of Toulous), and Nic Volanschi (11/19, HDR, University of Bordeaux).

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^{o} 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] Í. 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] S. EL AOUAD. *Personalized, Aspect-based Summarization of Movie Reviews*, Sorbonne universite, April 2019, https://hal.inria.fr/tel-02444980

Articles in International Peer-Reviewed Journals

[11] R. ANGARITA, B. LEFÈVRE, S. AHVAR, E. AHVAR, N. GEORGANTAS, V. ISSARNY. *Universal Social Network Bus: Towards the Federation of Heterogeneous Online Social Network Services*, in "ACM Transactions on Internet Technology", 2019, forthcoming [*DOI :* 10.1145/3323333], https://hal.inria.fr/hal-02072544

[12] G. BOULOUKAKIS, N. GEORGANTAS, P. NTUMBA, V. ISSARNY. *Automated synthesis of mediators for middleware-layer protocol interoperability in the IoT*, in "Future Generation Computer Systems", December 2019, vol. 101, pp. 1271-1294 [*DOI :* 10.1016/J.FUTURE.2019.05.064], https://hal.inria.fr/hal-02304074

[13] F. BRONZINO, P. SCHMITT, S. AYOUBI, G. MARTINS, R. TEIXEIRA, N. FEAMSTER. *Inferring Streaming Video Quality from Encrypted Traffic: Practical Models and Deployment Experience*, in "Proceedings of the ACM on Measurement and Analysis of Computing Systems ", December 2019, vol. 3, n° 3, https://hal.inria.fr/hal-02396405

[14] Y. DU, V. ISSARNY, F. SAILHAN. *When the Power of the Crowd Meets the Intelligence of the Middleware : The Mobile Phone Sensing Case*, in "Operating Systems Review", July 2019, vol. 53, n° 1, pp. 85-90 [*DOI :* 10.1145/3352020.3352033], https://hal.inria.fr/hal-02238256

[15] B. LEFÈVRE, R. AGARWAL, V. ISSARNY, V. MALLET. *Mobile Crowd-Sensing as a Resource for Contextualized Urban Public Policies: A Study using Three Use Cases on Noise and Soundscape Monitoring*, in "Cities & Health", May 2019 [*DOI :* 10.1080/23748834.2019.1617656], https://hal.inria.fr/hal-02127052

International Conferences with Proceedings

[16] R. AGARWAL, S. CHOPRA, V. CHRISTOPHIDES, N. GEORGANTAS, V. ISSARNY. *Detecting Mobile Crowdsensing Context in the Wild*, in "20th IEEE International Conference on Mobile Data Management (MDM) 2019", Hong Kong, Hong Kong SAR China, June 2019, https://hal.inria.fr/hal-02151434

[17] R. ALMEIDA, Í. CUNHA, R. TEIXEIRA, D. VEITCH, C. DIOT. *Classification of Load Balancing in the Internet*, in "IEEE INFOCOM 2020 - International Conference on Computer Communications", Beijing, China, April 2020, https://hal.inria.fr/hal-02396406

[18] R. ANGARITA, N. GEORGANTAS, V. ISSARNY. *Social Middleware for Civic Engagement*, in "ICDCS 2019 - 39th IEEE International Conference on Distributed Computing Systems - Vision Track", Dallas, United States, July 2019, https://hal.inria.fr/hal-02162736

[19] G. BOULOUKAKIS, I. MOSCHOLIOS, N. GEORGANTAS. *Probabilistic Event Dropping for Intermittently Connected Subscribers over Pub/Sub Systems*, in "ICC 2019 - IEEE International Conference on Communications", Shanghai, China, May 2019, https://hal.inria.fr/hal-02058417

[20] A. CHIO, G. BOULOUKAKIS, C.-H. HSU, S. MEHROTRA, N. VENKATASUBRAMANIAN. *Adaptive Mediation for Data Exchange in IoT Systems*, in "ARM 2019 - 18th Workshop on Adaptive and Reflexive Middleware", Davis, United States, December 2019 [*DOI :* 10.1145/3366612.3368122], https://hal.inria.fr/hal-02319447

[21] Y. DU, V. ISSARNY, F. SAILHAN. *User-centric context inference for mobile crowdsensing*, in "IoTDI 2019: ACM/IEEE International Conference on Internet of Things Design and Implementation", Montreal, Canada, ACM Press, April 2019, pp. 261-266 [*DOI :* 10.1145/3302505.3310088], https://hal.inria.fr/hal-02082034

[22] 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

[23] V. ISSARNY, B. BILLET, G. BOULOUKAKIS, D. FLORESCU, C. TOMA. *LATTICE: A Framework for Optimizing IoT System Configurations at the Edge*, in "ICDCS 2019 - 39th IEEE International Conference on Distributed Computing Systems", Dallas, Texas, United States, July 2019, https://hal.inria.fr/hal-02161795

[24] I. MORANDI, F. BRONZINO, R. TEIXEIRA, S. SUNDARESAN. *Service traceroute: Tracing Paths of Application Flows*, in "PAM 2019 - Passive and Active Measurement Conference", Puerto Varas, Chile, March 2019, https://hal.inria.fr/hal-02058307

[25] G. TEXIER, V. ISSARNY. *La puissance du crowdsensing au service de la durée de vie des réseaux IoT urbains*, in "ALGOTEL 2019 - 21èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications", Saint Laurent de la Cabrerisse, France, June 2019, https://hal.archives-ouvertes.fr/hal-02118074

[26] Q. ZHU, F. SAILHAN, Y. S. UDDIN, V. ISSARNY, N. VENKATASUBRAMANIAN. *Multi-Sensor Calibration Planning in IoT-Enabled Smart Spaces*, in "39th International Conference on Distributed Computing Systems (ICDCS)", Dallas, United States, July 2019, https://hal.inria.fr/hal-02136852

[27] B. H. ZIED, G. SCAVO, S. TRAVERSO, R. TEIXEIRA, M. MELLIA, S. GANGULY. *The News We Like Are Not the News We Visit: News Categories Popularity in Usage Data*, in "ICWSM-2019 - 13th International AAAI Conference on Web and Social Media", Munich, Germany, June 2019, https://hal.inria.fr/hal-02268409

Other Publications

[28] R. AGARWAL, S. CHOPRA, V. CHRISTOPHIDES, N. GEORGANTAS, V. ISSARNY. *Inferring Context of Mobile Data Crowdsensed in the Wild*, July 2019, NetMob 2019 - Conference on the scientific analysis of mobile phone datasets, Poster, https://hal.inria.fr/hal-02132194