Activity Report 2014

Team INFINE

INFormatIon NEtworks

RESEARCH CENTER
Saclay - Île-de-France

THEME
Networks and Telecommunications

Table of contents

Team INFINE

Keywords: Network Protocols, Network Dynamics, Wireless Sensor Networks, Data Analysis, Social Networks, Modeling

Creation of the Team: 2014 April 01.

1. Members

Research Scientists

Laurent Massoulié [Team leader, Inria, HdR]
Cédric Adjih [Inria, Researcher, from June 2014]
Emmanuel Baccelli [Inria, Researcher, HdR]
Aline Carneiro Viana [Inria, Researcher, HdR]

Engineers

Ichrak Amdouni [Inria, from Oct 2014]
Kuang Xu [Inria, from Sep 2014]

PhD Students

Guangshuo Chen [Inria, from Sep 2014]
Oliver Hahm [Inria, from Apr 2014]
Eduardo Oliveira Mucelli [Inria, from Nov 2011]
Felipe Cunha [Capes-Brazil, from Jan 2013]

Post-Doctoral Fellow

Naveen Kolar Purushothama [Inria, from Apr 2014]

Visiting Scientists

José Ignacio Alvarez-Hamelin [Inria, Oct 2014]
Jorge Brea [Inria, Oct 2014]
Carlos Sarraute [Inria, Oct 2014]
Anelise Munaretto [Inria, Oct 2014]
Artur Ziviani [Inria, Oct 2014]
Don Towsley [Inria, Jun 2014]

Administrative Assistant

Hélène Bessin-Rousseau [Inria, from Apr 2014]

Other

Lavanya Addepalli [Inria, from Jun 2014 until Nov 2014]

2. Overall Objectives

2.1. Overall Objectives

The INFINE proposal aims to design and analyse novel communication paradigms, protocols and architectures based on concepts of ultra distributed, information- and user-centric networking. The project is motivated by the recent and forthcoming evolution of Internet uses. Based on an information- and user-centric perspective, not only does it address issues pertaining to physical communication networks such as traffic routing, regulation and caching, but also issues about online social networks such as content recommendation and privacy protection.

INFINE team is engaged in research along three main themes: *Online social networking*, *Traffic and Resource Management*, and *Spontaneous Wireless Networks*. All these research activities encompass both theoretical research (on elaboration of models, algorithms, protocols and formal characterization of their performance), and applied research (to validate and/or experiment the proposed concepts against real networking scenarios). INFINE fits in the theme "Networks and Telecommunications" of the research field "Networks, Systems and Services, Distributed Computing" at Inria.

2.1.1. New challenging demands

Nowadays, we use networks not only to transport information from where it resides to ourselves but also, with online social networks, to determine what information might be of interest to us. Such a social recommendation functionality holds the promise of allowing us to access more relevant information. At the same time there is ample scope for improving its efficiency. Moreover it creates threats to user privacy.

At the same time, the physical context in which we access communication networks has drastically changed. While in the past, Internet was mostly accessed through fixed desktop computers, users are now mobile about 50% of their time online. In addition, while communicating machines used to be sparse and wired, with the advent of the Internet of Things we now evolve in a dense, interconnected environment of heterogeneous devices communicating via wireless and/or via wires.

This new context of Internet uses challenges several aspects of currently deployed networks. Some aspects pertain to the physical architecture of the Internet. In particular, at the core of the Internet, a drastic increase in volume of data traffic is anticipated due to the emergence of new applications, generalization of cloud services, or the advent of the Internet of Things (IoT) and Machine-to-Machine (M2M) communications. On the other hand, at the edge of the Internet, user mobility and today's pervasiveness of computing devices with increasingly higher capabilities (i.e., processing, storage, sensing) have a fundamental impact on the adequacy of algorithms and communication mechanisms.

Other aspects concern the logical architecture of the network. For instance, currently deployed protocols at layers above IP must now carry massive publish-subscribe traffic, preserve user privacy, be social-aware, and support delay tolerant communications and paradigms for which they were not initially designed. Concerning actual content distribution, the avalanche of data and privacy concerns puts more and more pressure on filter/push mechanisms to provide users with relevant information.

While considering physical and logical aspects of networks, the INFINE team will pursue research activities combining theoretical and experimental approaches.

2.1.2. Research agenda

Our general goal is to develop distributed mechanisms for optimizing the operation of networks both at the mentioned logical and physical levels of the architecture. Taking an information- and user-centric perspective, we envision networks as means to convey relevant information to users, while adapting to customary practices (in terms of context, interests, or content demands) of such users. At the logical level, online social networks (OSNs) allow users to choose what information to access. At the physical level, communication, computation, and memory resources allow users to retrieve some content eventually selected on the basis of the online social network.

The two setups feature scarce resources: for instance, in OSNs, these are the users' budget of attention, which must be used sparingly by recommending only relatively few potential content items. At the physical level this is typically the channels' capacity or networking resources, which cannot be oversubscribed.

Beyond a formal resemblance between the optimizations that one must carry at these two levels, there is a strong commonality in the methods adequate for conducting optimizations in the two setups. To illustrate this point, consider *contact recommendation*, that is a key objective in our agenda on online social networks. This entails automatically proposing to users potential contacts for optimizing the subsequent efficiency of social content filtering. We envision addressing contact recommendation by first performing some community detection, i.e. identification of similarly behaving users. Similarly, at the physical level, user-centric approaches, sometimes also related to community detection, have guided routing decisions in

challenged network environments, where delay-tolerant networking is used. Still, associated with dynamic centrality metrics, community detection can guide the replication of a specific content in well-selected users, while exploiting the advantages of distributed decentralized storage and opportunistic communications.

As an additional example at the logical level, we consider *content recommendation*, whereby a list of potential contents is filtered before being presented to a user, with the aim of maximizing the chance this user finds an item of interest therein. This has an exact analogue at the physical level, where by taking an information- and user-centric approach, we intend to off-load communication resources via pre-loaded content replicas at various storage points in the network. The problem of determining which content to cache so as to maximize the chance of it being accessed in the vicinity of the corresponding cache memory corresponds precisely to the aforementioned content recommendation problem.

We now detail further our agenda along three main specific axes, namely Online Social Networks, Traffic and Resource Management, and Spontaneous Wireless Networks/Internet of Things, bearing in mind that we will develop generic solutions relevant to several of these axes wherever possible.

3. Research Program

3.1. Online Social Networks (OSN)

Large-scale online social networks such as Twitter or FaceBook provide a powerful means of selecting information. They rely on "social filtering", whereby pieces of information are collectively evaluated and sorted by users. This gives rise to information cascades when one item reaches a large population after spreading much like an epidemics from user to user in a viral manner. Nevertheless, such OSNs expose their users to a large amount of content of no interest to them, a sign of poor "precision" according to the terminology of information retrieval. At the same time, many more relevant content items never reach those users most interested in them. In other words, OSNs also suffer from poor "recall" performance.

This leads to a first challenge: *what determines the optimal trade-off between precision and recall in OSNs? And what mechanisms should be deployed in order to approach such an optimal trade-off?* We intend to study this question at a theoretical level, by elaborating models and analyses of social filtering, and to validate the resulting hypotheses and designs through experimentation and processing of data traces. More specifically, we envision to reach this general objective by solving the following problems.

3.1.1. Community Detection

Identification of implicit communities of like-minded users and contact recommendation for helping users "rewire" the information network for better performance. Potential schemes may include variants of spectral clustering and belief propagation-style message passing. Limitations / relative merits of candidate schemes, their robustness to noise in the input data, will be investigated.

3.1.2. Incentivization

Design of incentive mechanisms to limit the impact of users' selfishness on system behavior: efficiency should be maintained even when users are gaming the system to try and increase their estimated expertise. By offering rewards to users on the basis of their involvement in filtering and propagation of content, one might encourage them to adjust their action and contribute to increase the overall efficiency of the OSN as a content access platform.

One promising direction will be to leverage the general class of Vickrey-Clarke-Groves incentive-compatible mechanisms of economic theory to design so-called marginal utility reward mechanisms for OSN users.

3.1.3. Social Recommendation and Privacy

So far we have only alluded to the potential benefits of OSNs in terms of better information access. We now turn to the risks they create. Privacy breaches constitute the greatest of these risks: OSN users disclose a wealth of personal information and thereby expose themselves to discrimination by potential employers, insurers, lenders, government agencies...Such privacy concerns are not specific to OSNs: internauts' online activity is discretely tracked by companies such as Bluekai, and subsequently monetized to advertisers seeking better ad targeting. While disclosure of personal data creates a privacy risk, on the other hand it fuels personalized services and thereby potentially benefits everyone.

One line of research will be to focus on the specific application scenario of content categorization, and to characterize analytically the trade-off between user privacy protection (captured by differential privacy), accuracy of content categorization, and sample complexity (measured in number of probed users).

3.2. Traffic and resource management

Despite the massive increases in transmission capacity of the last few years, one has every reason to believe that networks will remain durably congested, driven among other factors by the steadily increasing demand for video content, the proliferation of smart devices (i.e., smartphones or laptops with mobile data cards), and the forecasted additional traffic due to machine-to-machine (M2M) communications. Despite this rapid traffic growth, there is still a rather limited understanding of the features protocols have to support, the characteristics of the traffic being carried and the context where it is generated. There is thus a strong need for smart protocols that transport requested information at the cheapest possible cost on the network as well as provide good quality of service to network subscribers. One particularly new aspect of up-and-coming networks is that networks are now used to not only (i) access information, but also (ii) distributively process information, en-route.

We intend to study these issues at the theoretical and protocol design levels, by elaborating models and analysis of content demands and/or mobility of network subscribers. The resulting hypothesis and designs will be validated through experimentation, simulation, or data trace processing. It is also worth mentioning the provided solutions may bring benefits to different entities in the network: to content owners (if applied at the core of Internet) or to subscribers or network operators (if applied at the edge of the Internet).

3.2.1. At the Internet Core

One important optimization variable consists in content replication: users can access the closest replica of the content they are interested in. Thus the memory resource can be used to create more replicas and reduce the usage of the bandwidth resource. Another interesting arbitrage between resources arises because content is no longer static but rather dynamic. Here are two simple examples: i) a video could be encoded at several resolutions. There is then a choice between pre-recording all possible resolutions, or alternatively synthesizing a lower-resolution version on the fly from a higher resolution version when a request arises. ii) A user requests the result of a calculation, say the average temperature in a building; this can either be kept in memory, or recomputed each time such a query arises. Optimizing the joint use of all three resources, namely bandwidth, memory, computation, is a complex task. Content Delivery Network companies such as Akamai or Limelight have worked on the memory/bandwidth trade-off for some years, but as we will explain more can be done on this. On the other hand optimizing the memory/computation trade-off has received far less attention. We aim to characterize the best possible content replication strategies by leveraging fine-grained prediction of i) users' future requests, and ii) wireless channels' future bandwidth fluctuations. In the past these two determining inputs have only been considered at a coarse-grained, aggregate level. It is important to assess how much bandwidth saving can be had by conducting finer-grained prediction. We are developing light-weight protocols for conducting these predictions and automatically instantiating the corresponding optimal replication policies. We are also investigating generic protocols for automatically trading replication for computation, focusing initially on the above video transcoding scenario.

3.2.2. At the Internet Edge

Cellular and wireless data networks are increasingly relied upon to provide users with Internet access on devices such as smartphones, laptops or tablets. In particular, the proliferation of handheld devices equipped with multiple advanced capabilities (e.g., significant CPU and memory capacities, cameras, voice to text, text to voice, GPS, sensors, wireless communication) has catalyzed a fundamental change in the way people are connected, communicate, generate and exchange data. In this evolving network environment, users' social relations, opportunistic resource availability, and proximity between users' devices are significantly shaping the use and design of future networking protocols.

One consequence of these changes is that mobile data traffic has recently experienced a staggering growth in volume: Cisco has recently foreseen that the mobile data traffic will increase 18-fold within 2016, in front of a mere 9-fold increase in connection speeds. Hence, one can observe today that the inherently centralized and terminal-centric communication paradigm of currently deployed cellular networks cannot cope with the increased traffic demand generated by smartphone users. This mismatch is likely to last because (1) forecasted mobile data traffic demand outgrows the capabilities of planned cellular technological advances such as 4G or LTE, and (2) there is strong skepticism about possible further improvements brought by 5G technology.

Congestion at the Internet's edge is thus here to stay. Solutions to this problem relates to: densify the infrastructure, opportunistically forward data among neighbors wireless devices, to offload data to alternate networks, or to bring content from the Internet closer to the subscribers. Our recent work on leveraging user mobility patterns, contact and inter-contact patterns, or content demand patterns constitute a starting point to these challenges. The projected increase of mobile data traffic demand pushes towards additional complementary offloading methods. Novel mechanisms are thus needed, which must fit both the new context that Internet users experience now, and their forecasted demands. In this realm, we will focus on new approaches leveraging ultra-distributed, user-centric approaches over IP.

3.3. Spontaneous Wireless Networks (SWN) and Internet of Things (IoT)

The unavailability of end-to-end connectivity in emergent wireless mobile networks is extremely disruptive for IP protocols. In fact, even in simpler cases of spontaneous wireless networks where end-to-end connectivity exists, such networks are still disruptive for the standard IP protocol stack, as many protocols rely on atomic link-local services (such as link-local multicast/broadcast), while these services are inherently unavailable in such networks due to their opportunistic, wireless multi hop nature. In this domain, we will aim to characterize the achievable performance in such IP-disruptive networks and to actively contribute to the design of new, deployable IP protocols that can tolerate these disruptions, while performing well enough compared to what is achievable and remaining interoperable with the rest of the Internet.

Spontaneous wireless networking is also a key aspect of the Internet of Things (IoT). The IoT is indeed expected to massively use this networking paradigm to gradually connect billions of new devices to the Internet, and drastically increase communication without human source or destination – to the point where the amount of such communications will dwarf communications involving humans. Large scale user environment automation require communication protocols optimized to efficiently leverage the heterogeneous and unreliable wireless vicinity (the scope of which may vary according to the application). In fact, extreme constraints in terms of cost, CPU, battery and memory capacities are typically experienced on a substantial fraction of IoT devices. We expect that such constraints will not vanish any time soon for two reasons. On one hand the progress made over the last decade concerning the cost/performance ratio for such small devices is quite disappointing. On the other hand, the ultimate goal of the IoT is ubiquitous Internet connectivity between devices as tiny as dust particles. These constraints actually require to redesign not only the network protocol stack running on these devices, but also the software platform powering these machines. In this context, we will aim at contributing to the design of novel network protocols and software platforms optimized to fit these constraints while remaining compatible with legacy Internet.

3.3.1. Design & Development of Open Experimental IoT Platforms

Based initially on "Demonstration abstract: Simply RIOT — Teaching and experimental research in the Internet of Things" Manufacturers announce on a regular basis the availability of novel tiny devices, most of them featuring network interfaces: the Internet of Things (IoT) is already here, from the hardware perspective, and it is expected in the near future that we will see a massive increase of the number of muti-purpose smart objects (from tiny sensors in industrial automation to devices like smart watches and tablets). Thus, one of the challenges is to be able to test architectures, protocols and applications, in realistic conditions and at large scale.

One necessity for research in this domain is to establish and improve IoT hardware platforms and testbeds, that integrate representative scenarios (such as Smart Energy, Home Automation etc.) and follow the evolution of technology, including radio technologies, and associated experimentation tools. For that, we plan to build upon the IoT-LAB federated testbeds, that we have participated in designing and deploying recently. We plan to further develop IoT-LAB with more heterogeneous, up-to-date IoT hardware and radios that will provide a usable and realistic experimentation environment. The goal is to provide a tool that enables testing an validation of upcoming software platforms and network stacks targeting concrete IoT deployments.

In parallel, on the software side, IoT hardware available so far made it uneasy for developers to build apps that run across heterogeneous hardware platforms. For instance Linux does not scale down to small, energy-constrained devices, while microcontroller-based OS alternatives were so far rudimentary and yield a steep learning curve and lengthy development life-cycles because they do not support standard programming and debugging tools. As a result, another necessity for research in this domain is to allow the emergence of it more powerful, unifying IOT software platforms, to bridge this gap. For that, we plan to build upon RIOT, a new open source software platform which provides a portable, Linux-like API for heterogeneous IoT hardware. We plan to continue to develop the systems and network stacks aspects of RIOT, within the open source developer community currently emerging around RIOT, which we co-founded together with Freie Universitaet Berlin. The key challenge is to improve usability and add functionalities, while maintaining architectural consistency and a small enough memory footprint. The goal is to provide an IoT software platform that can be used like Linux is used for less constrained machines, both (i) in the context of research and/or teaching, as well as (ii) in industrial contexts. Of course, we plan to use it ourselves for our own experimental research activities in the domain of IoT e.g., as an API to implement novel network protocols running on IoT hardware, to be tested and validated on IoT-LAB testbeds.

3.3.2. Design & Standardization of Architectures and Efficient Protocols for Internet of Things

As described before, and by definition, the Internet of Things will integrate not only a massive number of homogeneous devices (e.g., networks of wireless sensors), but also heterogeneous devices using various communication technologies. Most devices will be very constrained resources (memory resources, computational resources, energy). Communicating with (and amongst) such devices is a key challenge that we will focus on. The ability to communicate efficiently, to communicate reliably, or even just to be able to communicate at all, is non-trivial in many IoT scenarios: in this respect, we intend to develop innovative protocols, while following and contributing to standardization in this area. We will focus and base most of our work on standards developed in the context of the IETF, in working groups such as 6lo, CORE, LWIG etc., as well as IRTF research groups such as NWCRG on network coding and ICNRG on Information Centric Networking. We note however that this task goes far beyond protocol design: recently, radical rearchitecturing of the networks with new paradigms such as Information Centric Networking, ICN, (or even in wired networks, software-defined networks), have opened exciting new avenues. One of our direction of research will be to explore these content-centric approaches, and other novel architectures, in the context of IoT.

4. Application Domains

4.1. Panorama

The research in INFINE spans a wide range of application areas ranging from Internet-based, wireless sensor-based, mobile wireless-based, and OSN-based applications. These applications are related to the three main research axes described in the previous sections.

4.2. Mobile wireless network

Smart portable devices such as smartphones, PDAs or tablet PCs are being considered as pervasive mobile sensing platforms due to their increasing proliferation and their wide range of embedded heterogeneous capabilities (in terms of type of communication and data gathering possibilities - e.g., 3G, WiFi, GPS, video, camera, etc). Such devices are changing the way people are communicating, generating, and exchanging data: They allow the free sensing/gathering of data of the surrounding environment anytime and anywhere. On the other hand, the projected increase of mobile data traffic demand pushes towards additional complementary offloading methods. Novel mechanisms are thus needed, which must fit both the new context that Internet users experience now, and their forecasted demands.

In these contexts, the application domains that we are targeting are related to traffic offloading in large-scale mobile wireless networks. Among the numerous offloading solutions fitting in this application domain, we are specially interested in the ones related to: infrastructure deployment, traffic modeling, opportunistic communication, or still task delegation. A core principle of such solutions is the understanding and modeling of users behavior in terms of their context (i.e., imposed by mobility) and their content demands.

4.3. Online Social Networks

Our high-level goal here is to help increase the relevance of content accessed by users, through the elaboration of contact and content recommendation mechanisms, as well as incentive mechanisms. The scientific context in which we phrase this goal is that of:

- modeling information propagation in OSN;
- statistical inference problems raised by the search for improved information propagation. In particular these include community detection for contact and content recommendation, and bandit-like algorithms for active learning of given content type at limited "spamming" cost;
- the mechanism design branch of economic theory, which can be leveraged to conceive reward mechanisms meant to incentivize efficient collaborative content filtering by OSN users.

4.4. Spontaneous Wireless Networks applications

The advances in hardware development have made possible the miniaturization of micro-electro-mechanical systems and consequently, the development of sensor networks. The combination of inexpensive, autonomous, low-power sensing, and compact devices has established the viability of deploying large and dense wireless sensor networks (WSNs) able to sense the physical world. By essence, such networks require fully decentralized solutions in which the load is evenly balanced in the system, merely because participating entities have limited in power, storage and communication capabilities. Thus one of the applications of Spontaneous Wireless Networks has been traditionally such wireless sensor networks, where some typical applications are to continuously monitor data (real-time data collection to a sink), and to be able to do manage network after deployment (for instance reflashing nodes with firmware other the air). The challenge is to operate this with standards (such as IP), constrained devices (battery, memory, power, ...), which requires sophisticated protocols, with reliable and tested implementations.

The applications of the more recent "Internet of Object" are much broader, since they literally consists of any application running on any object (in the industrial factories, in living spaces, ...). While some of the constraints in wireless sensor networks are a still present in IoT in general, what characterizes IoT is the heterogeneity of the platforms.

5. New Software and Platforms

5.1. Software

5.1.1. MACACOapp

Participant: Aline Carneiro Viana.

MACACOapp (https://macaco.inria.fr/macacoapp/) is developped in the context of the EU CHIST-ERA MACACO project (https://macaco.inria.fr/). It consists in a mobile phone application that periodically samples phone's information on the mobility (through, e.g., GPS sensor, accelerometer and WiFi/Bluetooth/Cellular environment, connectivity type) and on the data traffic it generates (through, e.g., Internet browser history and applications data consumption). The information collected will be time-stamped and will be periodically sent to the central servers for analysis and visualization. We expect that (1) the collected information will allow us studying the correlation between mobility and content demand patterns and that (2) the results of this analysis will allow us inferring the best times and places to transfer content from/to users' phones location and/or from/to the wireless infrastructure closest to the users' phones location. Users will be also invited to fill a non-mandatory questionnaire relevant to this study. Our questionnaire collects information about the personality traits and application preferences of people. We expect that the information collected from questionnaire will allow us to analyse the correlation between users personality traits and their application preferences and interests. User's application preferences and interests will be inferred from the Internet browsing history and running app information obtained from the MACACO App.

The data collection and on-the-phone storage of MACACOapp is designed in accordance with the state-of-the-art best practices in application development. The data collected on the phone is encrypted and inaccessible by any other application installed on the same phone or to any other third party, even in case your phone gets lost or stolen. Moreover, any user's identity information available in the collected data or in the questionnaire will be completely and irreversibly anonymised before its transfer to the central servers. The on-the-phone collected data and questionnaire data will be transferred via a secure transmission protocol to the central servers. This application is in phase of test.

5.1.2. RIOT

Participants: Emmanuel Baccelli, Oliver Hahm.

RIOT (http://www.riot-os.org) is a nano operating system for the Internet of Things. While requiring as low as 1,5kB of RAM and 5kB or ROM, RIOT offers real time and energy efficiency capabilities, as well as a single API (partially POSIX compliant) across heterogeneous 8-bit, 16-bit and 32-bit low-hardware. This API is developer-friendly in that it enables multi-threading, standard C and C++ application programming and the use of standard debugging tools (which was not possible so far for embedded programming). On top of this, RIOT includes several network stacks, such as a standard IPv6/6LoWPAN stack and a information-centric network stack (based on CCN).

RIOT is developed by an international community of open-source developers that was co-founded by Inria and Freie Universitaet Berlin. The goal of RIOT is to provide a powerful, free, open-source IoT software platform that can be used like Linux is for less constrained machines, both (i) in the context of research and/or teaching, as well as (ii) in industrial contexts.

5.1.3. DragonNet

Participants: Cédric Adjih, Ichrak Amdouni, Hana Baccouch, Antonia Masucci.

DragonNet is a generic framework for network coding in wireless networks. It is a initially result of the GETRF project of the Hipercom2 team.

It is based on intra-flow coding where the source divides the flow in a sequence of payloads of equal size (padding may be used). The design keys of DragonNet are simplicity and universality; DragonNet does not use explicit or implicit knowledge about the topology (such as the direction or distance to the source, the loss rate of the links, ...). Hence, it is perfectly suited to the most dynamic wireless networks. The protocol is distributed and requires minimal coordination. DragonNet architecture is modular, it is based on 5 building blocks (LIB, SIG, Protocol, SEW and DRAGON). Each block is almost independent. This makes DragonNet generic and hence adaptable to many application scenarios. DragonNet derives from a prior protocol called DRAGONCAST. Indeed, DragonNet shares the same principles and theoretical overview of DRAGONCAST. It enriches DRAGONCAST by the information base and signaling required to perform broadcast in wireless networks and in wireless sensor networks in particular.

5.2. Platforms

5.2.1. FIT IoT-LAB

Participants: Cedric Adjih, Emmanuel Baccelli, Ichrak Amdouni.

FIT IoT-LAB is a platform built to help foster the development, tuning and experimentation of protocols and applications for the Internet of Things and wireless sensor networks. IoT-LAB provides both dedicated IoT hardware deployments, a front-end webportal and backend management software. Using these elements, IoT-LAB enables users to share access to this IoT hardware, set-up and manage experiments. Remote use, and large scale experiments on concrete IoT deployments are thus made possible.

The Infine team is now managing the IoT-LAB site currently located in Rocquencourt, and which was publically opened in November 2014. It consists of the following:

- A set of GPS repeaters are relaying the GPS signal indoor (used for time synchronization)
- 200 A8 nodes, all equipped with GPS (10 deployed outside – identifiers between 166 and 175)
- 24 M3 nodes
- 120 WSN430 nodes

This platform was developped as part of the Equipex FIT (see section 8.1.1).

6. New Results

6.1. Highlights of the Year

- We proved a conjecture made in 2011 about the feasibility of non-trivial community detection just above a threshold below which it was known that only trivial detection could be done, see [13]. This was published in ACM STOC'14 and well-received, as the proof required the invention of new techniques to control the spectral properties of random matrices.
- The official opening of IoT-LAB of all sites through the "Workshop Internet Of Things/Equipex FIT IoT-LAB" held in Grenoble (on 6 and 6 november 2014), has been a major event for our team: it concludes several years of preparation of the IoT-LAB site located in Rocquencourt, currently managed by C. Adjih, E. Baccelli and I. Amdouni, which was itself opened the same month https://www.iot-lab.info/opening-of-the-paris-rocquencourt-site/.

6.2. Panorama

All the INFINE research activities encompass both theoretical or protocol designing research (to seek for conceptual advances or optimizations) and applied research (to validate and/or experiment the proposed concepts against real networking scenarios). The target applications range from Internet-based applications to mobile wireless networks. Taking an information- and user-centric perspective, we envision networks as means to convey relevant information to users, while adapting to customary practices (in terms of context, interests, or content demands) of such users. INFINE is thus organized along three main axes, namely Online Social Networks, Resource and Traffic Management, and Spontaneous Wireless Networks.

6.3. Online Social Networks (OSN)

Community detection; bandit algorithms; privacy preservation; reward mechanisms

6.3.1. Community detection

Participants: Laurent Massoulié, Marc Lelarge, Jiaming Xu.

We have progressed in the design of spectral methods for community detection and in the corresponding analysis (see above and references [3], [13], [22]).

6.3.2. Bandit algorithms for active learning of content type at low spam cost

Participants: Laurent Massoulié, Mesrob Ohanessian, Alexandre Proutière.

We developed a framework in which to cast the problem, and the so-called "greedy Bayes" algorithm to determine which user to expose to a given content. We proved corresponding optimality properties, and observed that "greedy Bayes" beats the so-called Thompson sampling approach, that is the state-of the-art method in bandit problems. Work currently under submission.

6.4. Resource and Traffic Management

Traffic offloading; infrastructure deployment; opportunistic routing; traffic modeling; intermittently connected networks.

6.4.1. From Routing to Network Deployment for Data Offloading in Metropolitan Areas

Participants: Eduardo Mucceli, Aline Carneiro Viana.

Smartphone sales are booming, nearly half billion were sold in 2011; more smartphones, more mobile data traffic, and Currently, 3G cellular networks in metropolitan areas are struggling to attend the recent boost up of mobile data consumption. Carefully deploying WiFi hotspots allow to maximize WiFi offloading and can both be cheaper than upgrade the cellular network structure and concede substantial improvement in the network capacity. In this context, in this work, we first propose a new way to map into a graph the *people behavior* (i.e., mobility context) in an urban scenario. Our proposed behavior-to-graph solution is simple, take into consideration the restrictions imposed by transportation modes to traffic demand, the space-time interaction between people and urban locations, and finally, is powerful to be used as input to any popular area identification problem (key points for an efficient network planning). Secondly, we propose a metric to identify locations more capable of providing coverage for people and consequently, more suitable for receiving hotspots. Deploying a small percentage of hotspots ranked by the herein proposed metric provides high percentages of coverage time for people moving around in the city. Using a real-life metropolitan trace, we show our routine-based strategy guarantees higher offload ratio than current approaches in the literature while using a realistic traffic model. Different parts of this work has been published in the international conferences IEEE SECON 2014 [14], IEEE WCNC 2014 [18] and IEEE WMNC 2014 [17], and in the international Student workshop IEEE Infocom [15]. An extended version of this paper is under submission in a trasaction. This version includes new characterization results of the used trace and new analysis of space-traffic correlation.

6.4.2. Mobile Data Traffic Modeling: Revealing Hidden Facets

Participants: Eduardo Mucceli, Aline Carneiro Viana, Kolar Purushothama Naveen, Carlos Sarraute.

Smartphone devices provide today the best means of gathering users information about content consumption behavior on a large scale. In this context, the literature is rich in work studying and modeling users mobility, but little is publicly known about users content consumption patterns. The *understanding of users' mobile data traffic demands* is of fundamental importance when looking for solutions to manage the recent boost up of mobile data usage [14] and to improve the quality of communication service provided. Hence, the definition of a *usage pattern* can allow telecommunication operators to better forseen future demanded traffic and consequently, to better (1) deploy data offloading hotspots or (2) timely plan network resources allocation and then, set subscription plans.

Using a large-scale dataset collected from a major 3G network in a big metropolitan area, in this work, we present the first detailed measurement-driven modeling of mobile data traffic usage of smartphone subscribers. Our main outcome is a synthetic measurement-based mobile data traffic generator, capable of imitating traffic-related activity patterns of different categories of subscribers and time periods of a routinary normal day in their lives. For this, we first characterize individual subscribers routinary behaviour, followed by the detailed investigation of subscribers' usage pattern (i.e., "when" and "how much" traffic is generated). Broadly, our observations bring important insights into network resource usage. We then classify the subscribers into six distinct profiles according to their usage pattern and model these profiles according to two different journey

periods: peak and non-peak hours. We show that the synthetic trace generated by our data traffic model consistently imitates different subscriber profiles in two journey periods, when compared to the original dataset. We discuss relevant issues in traffic demands and describe implications in network planning and privacy. This work has been published in the international conference IEEE PERCOM 2014 [16]. An extended version of this paper is under submission in a trasaction and a technical report is available in [26]. This version includes new characterization results of the used trace, including analysis correlating age and gender to traffic demands, as well as new profiling results.

6.4.3. On the Interaction between Content Caching and Routing

Participants: Kolar Purushothama Naveen, Laurent Massoulié, Emmanuel Baccelli, Aline Carneiro Viana, Don Towsley.

Nowadays Internet users are mobile over 60% of their time online, and mobile data traffic is expected to increase by more than 60% annually to reach 15.9 exabytes per month by 2018. This evolution will likely incur durably congested wireless access at the edge despite progress in radio technologies. To alleviate congestion at the Internet edge, one promising approach is to target denser deployments of wireless access points. As a result, mobile users are potentially within radio reach of several access points (AP) from which content may be directly downloaded. In this context, distinct AP's can have very different bandwidth and memory capacities. Such differences raise the following question: When requests can be sent to several such access points, how to optimize performance through both load balancing and content replication?

In this work, we introduce formal optimization models to address this question, where bandwidth availability is represented via a cost function, and content availability is represented either by a cost function or a sharp constraint. For both formulations we propose dynamic caching and request assignment algorithms. Crucially our request assignment scheme is based on a server price signal jointly reflecting content and bandwidth availability. Using mean field approximation and Lyapunov functions techniques, we prove that our algorithms are optimal and stable in a limiting fluid regime with large arrival rates and content chunking. Through simulations we exhibit the efficacy of our request assignment strategy in comparison to the common practices of assigning requests purely based on either bandwidth or content availability. Finally, using the popular LRU (Least Recently Used) strategy instead for cache replacements, we again demonstrate the superior performance of our request assignment strategies. This work in under submission in an international conference.

6.4.4. Data Delivery in Opportunistic and Intermittently Connected Networks

Participants: Ana Cristina Vendramin, Anelise Munaretto, Myriam Delgado, Aline Carneiro Viana, Mauro Fonseca.

The pervasiveness of computing devices and the emergence of new applications and cloud services are factors emphasizing the increasing need for adaptive networking solutions. In most cases, this adaptation requires the design of interdisciplinary approaches as those inspired by nature, social structures, games, and control systems. The approach presented in this work brings together solutions from different, yet complementary domains, i.e., networking, artificial intelligence, and complex networks, and is aimed at addressing the problem of efficient data delivery in intermittently connected networks.

As mobile devices become increasingly powerful in terms of communication capabilities, the appearance of opportunistic and intermittently connected networks referred to as Delay Tolerant Networks (DTNs) is becoming a reality. In such networks, contacts occur opportunistically in corporate environments such as conferences sites, urban areas, or university campuses. Understanding node mobility is of fundamental importance in DTNs when designing new communication protocols that consider opportunistic encounters among nodes. This work proposes the Cultural Greedy Ant (CGrAnt) protocol to solve the problem of data delivery in opportunistic and intermittently connected networks. CGrAnt is a hybrid Swarm Intelligence-based forwarding protocol designed to address the dynamic and complex environment of DTNs. CGrAnt is based on: (1) Cultural Algorithms (CA) and Ant Colony Optimization (ACO) and (2) operational metrics that characterize the opportunistic social connectivity between wireless users. The most promising message forwarders are selected via a greedy transition rule based on local and global information captured from the

DTN environment. Using simulations, we first analyze the influence of the ACO operators and CA knowledge on the CGrAnt performance. We then compare the performance of CGrAnt with the PROPHET and Epidemic protocols (two well known related protocols in the literature) under varying networking parameters. The results show that CGrAnt achieves the highest delivery ratio (gains of 99.12% compared with PROPHET and 40.21% compared with Epidemic) and the lowest message replication (63.60% lower than PROPHET and 60.84% lower than Epidemic). This work is under submission to an international journal. Some parts of this work were previously published in the international conference ACM GECCO 2012 and in the Elsevier Computer Networks journal.

6.4.5. Vehicular Network under a Social Perception

Participants: Felipe D. Cunha, Aline Carneiro Viana, Raquel A. F. Mini, Antonio A.f. Loureiro.

Vehicular Mobility is strongly influenced by the speed limits, destinations, traffic conditions, period of the day, and direction of the public roads. At the same time, the driver's behavior produces great influences in vehicular mobility. People tend to go to the same places, at the same day period, through the same trajectories, which le ad them to the appearance of driver's daily routines. These routines lead us to the study of mobility in VANETs under a social perspective and to investigate how effective is to explore social interactions in this kind of network. In this work, we thus characterize and evaluate social properties of a realistic vehicular trace found in literature. Our aim is to study the vehicles' mobility in accordance to social behaviors. Social metrics are computed and the obtained results are compared to random graphs. With our analysis, we could verify the existence of regularity and common interests among the drivers in vehicular networks. This work was published in the international conference IEEE ISCC 2014 [10], in the international Student workshop of IEEE Infocom 2014 [9], and at the internation workshop Internet of Things Communications and Technologies (IoT 2013) held in conjunction with IEEE WiMob 2013.

After having identified routine in vehicles mobility patterns and their correlation with the period of the day, we then leverage the identified social aspects to design a *Socially Inspired Broadcast Data Dissemination* for VANETs. We claim that protocols and applications designed for Vehicular Ad Hoc Networks need to adapt to vehicles routines in order to provide better services. With this issue in mind, we designed a data dissemination solution for these networks that considers the daily road traffic variation of large cities and the relationship among vehicles. The focus of our approach is to select the best vehicles to rebroadcast data messages according to social metrics, in particular, the clustering coefficient and the node degree. Moreover, our solution is designed in such a way that it is completely independent of the perceived road traffic density. Simulation results show that, when compared to related protocols, our proposal provides better delivery guarantees, reduces the network overhead and possesses an acceptable delay. This work was published as a short paper at the international conference ACM MSWiM 2014 [8]

6.4.6. Design and Analysis of an Efficient Friend-to-Friend Content Dissemination System

Participants: Kanchana Thilakarathna, Aline Carneiro Viana, Aruna Seneviratne, Henrik Petander.

In this work, we focus on dissemination of content for delay tolerant applications/services, (i.e. content sharing, advertisement propagation, etc.) where users are geographically clustered into communities. Due to emerging security and privacy concerns, majority of users are becoming more reluctant to interact with strangers and are only willing to share information/content with the users who are previously identified as friends. As a result, despite its promise, opportunistic communications systems have not been widely adopted. In addition, in this environment, opportunistic communication will not be effective due to the lack of known friends within the communication range. We thus propose a novel architecture which combines the advantages of distributed decentralized storage and opportunistic communications. The proposed system addresses the trust and privacy concerns of opportunistic communications systems, and enables the provision of efficient distributed mobile social networking services. We exploit the fact that users will trust their friends, and the friends will help in disseminating content by temporarily storing and forwarding content. This can be done by replicating content on friends' devices who are likely to consume that content and provide the content to other friends when the device has access to low cost networks. The fundamental challenge then is to minimize the number of replicas, to ensure high and timely availability. We provide a formal definition of this content replication problem, and

show that it is NP hard. Then, we propose a community based greedy heuristic algorithm with novel dynamic centrality metrics that replicates the content on a minimum number of friends' devices, and maximizes the availability of content. Using both real world and synthetic traces, we validate effectiveness of the proposed scheme. In addition, we demonstrate the practicality of the the proposed system, through an implementation on Android smartphones. This work is under submission in an international transaction. An initial version of this work was published at the international conference ACM MobiHoc 2013, and an extended version is under submission in an international trasaction.

6.4.7. Telling Apart Social and Random Relationships in Dynamic Networks

Participants: Pedro Olmo Vaz de Melo, Aline Carneiro Viana, Marco Fiore, Katia Jaffrès-Runser, Frédéric Le Mouël, Antonio A. F. Loureiro, Lavanya Addepalli, Guangshuo Chen.

Recent studies have analyzed data generated from mobile individuals in urban regions, such as cab drivers or students in large campuses. Particular attention has been paid to the dynamics of user movement, whose real-world complexity cannot be fully captured through synthetic models. Indeed, understanding user mobility is of fundamental importance when designing new communication protocols that exploit opportunistic encounters among users. In this case, the problem mainly lies in correctly forecasting future contacts. To that end, the regularity of daily activities comes in handy, as it enforces periodic (and thus predictable) space-time patterns in human mobility. Although human behavior is characterized by an elevated rate of regularity, random events are always possible in the routines of individuals. Those are hardly predictable situations that deviate from the regular pattern and are unlikely to repeat in the future.

We argue that the ability to accurately spot random and social relationships in dynamic networks is essential to network applications that rely on a precise description of human routines, such as recommendation systems, forwarding strategies and opportunistic dissemination protocols. We thus propose a strategy to analyze users' interactions in mobile networks where users act according to their interests and activity dynamics. Our strategy, named *Random rElationship ClASsifier sTrategy (RECAST)*, allows classifying users' wireless interactions, separating random interactions from different kinds of social ties. To that end, RECAST observes how the real system differs from an equivalent one where entities' decisions are completely random. We evaluate the effectiveness of the RECAST classification on five real-world user contact datasets collected in diverse networking contexts. Our analysis unveils significant differences among the dynamics of users' wireless interactions in the datasets, which we leverage to unveil the impact of social ties on opportunistic routing. We show that, for such specific purpose, the relationships inferred by classifier are more relevant than, e.g., self-declared friendships on Facebook. An initial version of this work was published in the international conference ACM MSWiM 2013 (selected as one of the five better papers of this venue) and an extented version bringing new analysis (e.g., the contact duration-related analysis performed by the internship Lavanya Addepalli and the PhD student Guangshuo Chen) was accepted to be published in 2015 at the Performance Evaluation Elsevier Journal [2].

6.5. Spontaneous Wireless Networks and Internet of Things

internet of things; wireless sensor networks; dissemination; resource management

6.5.1. Network Coding in Large Scale IoT Networks

Participants: Cedric Adjih, Ichrak Amdouni, Hana Baccouch, Antonia Masucci.

We had designed a generic broadcast protocol, called DragonNet, based on network coding and designed for constrained networks such as wireless sensor networks and internet of things. It minimizes the assumptions made of the networks. A variant of this protocol was implemented and run on IoT-LAB: some results were initially presented at IRTF, and a live demo was presented in MASS in october 2014.

6.5.2. Information-Centric Networking in the Internet of Things

Participants: Emmanuel Baccelli, Oliver Hahm, Matthias Waehlisch, Thomas Schmidt, Christian Mehlis.

Within this activity, we explored the feasibility, advantages, and challenges of an ICN-based approach in the Internet of Things. We report on the first NDN experiments in a life-size IoT deployment, spread over tens of rooms on several floors of a building. Based on the insights gained with these experiments, we have analysed the shortcomings of CCN applied to IoT. Several interoperable CCN enhancements are then proposed and evaluated. We significantly decreased control traffic (i.e., interest messages) and leverage data path and caching to match IoT requirements in terms of energy and bandwidth constraints. Our optimizations increase content availability in case of IoT nodes with intermittent activity. Within this activity, we also provided the first experimental comparison of CCN with the common IoT standards 6LoWPAN/RPL/UDP.

7. Bilateral Contracts and Grants with Industry

7.1. GranData

Participants: Aline Carneiro Viana, Eduardo Mucelli.

Since June 2014, we have a collaboration with GranData (http://grandata.com/), Buenos Aires, Argentina on traffic vs mobility modeling of smartphone users. GranData is a small company that integrates first-party and telco partner data to understand key market trends, to predict customer behavior, and to deliver business results. Its products integrates and analyzes diverse data traces (e.g., telco, social media, or mobile data) to generate behavioral insights and deliver targeted mobile marketing. Part of the thesis of Eduardo Mucelli analysis data traffic using telco traces provided by GranDatas. While this collaboration allow us collaborating with machine learning experts, GranData has the opportunity to get our expertise in mobiliy analysis.

8. Partnerships and Cooperations

8.1. National Initiatives

8.1.1. Equipex FIT

Participants: Cedric Adjih, Emmanuel Baccelli, Ichrak Amdouni, Alaeddine Weslati, Vincent Ladeveze.

Partners: Inria (Lille, Sophia-Antipolis, Grenoble), INSA, UPMC, Institut Télécom Paris, Institut Télécom Evry, LSIIT Strasbourg.

FIT (Future Internet of Things) aims to develop an experimental facility, a federated and competitive infrastructure with international visibility and a broad panel of customers. It provides this facility with a set of complementary components that enable experimentation on innovative services for academic and industrial users. The project gives french internet stakeholders a means to experiment on mobile wireless communications at the network and application layers thereby accelerating the design of advanced networking technologies for the future internet. FIT was one of 52 winning projects from the first wave of the French Ministry of Higher Education and Research's "Équipements d'Excellence" (Equipex) research grant program, in 2011.

One component of the FIT platform is the sets of IoT-LAB testbeds (voir le site IoT-LAB). These were motivated by the observation that the world is moving towards an "Internet of Things", in which most communication over networks will be between objects rather than people.

The Infine team is now managing the FIT IoT-LAB site currently at Rocquencourt whose development has been started in 2010 (in the Hipercom team).

8.2. European Initiatives

8.2.1. EU CHIST-ERA MACACO

Participants: Aline Carneiro Viana, Emmanuel Baccelli, Eduardo Mucelli.

Program: EU CHIST-ERA, topic Context- and Content-Adaptive Communication Networks
Project acronym: MACACO
Project title: Mobile context-Adaptive CAching for COntent-centric networking
Duration: 2013-2016
Coordinator: Aline Carneiro Viana
Other partners: INPT-ENSEEIHT at University of Toulouse, University of Birmingham (UK), SUPSI (Switzerland), CNR (Italy) and Federal University of Minas Gerais (Brazil)
Abstract:

MACACO (Mobile context-Adaptive CAching for COntent-centric networking) is a 3-year CHIST-ERA European Project addressing the topic Context- and Content-Adaptive Communication Networks. It is funded by ANR in France, SNSF in Switzerland, and ESPRC in UK. It focus on data offloading mechanisms that take advantage of context and content information. Our intuition is that if it is possible to extract and forecast the behaviour of mobile network users in the threedimensional space of time, location and interest (i.e. 'what', 'when' and 'where' users are pulling data from the network), it is possible to derive efficient data offloading protocols. Such protocols would pre-fetch the identified data and cache it at the network edge at an earlier time, preferably when the mobile network is less charged, or offers better quality of service. This project has officially started in November 2013. The first annual report will be delivered before January 25, 2015.

8.2.2. Collaborations in European Programs, except FP7 & H2020

Participants: Emmanuel Baccelli, Cedric Adjih, Oliver Hahm.

Program: ANR/BMBF French-German partnership within CSOSG Framework
Project acronym: SAFEST
Project title: Social Area Framework for Early Security Triggers
Duration: 2012-2015
Coordinators: Emmanuel Baccelli (France), Jochen Schiller (Germany)
Other partners: Freie Universitat Berlin, Fraunhofer, Hamburg University, Sagem, Daviko, FOS
Abstract: Public spaces, such as airports, railway stations, or stadiums bring together large numbers of people on limited space to use security-sensitive infrastructure. These spaces pose two distinct challenges to public security: (a) detecting unauthorized intrusions and (b) monitoring large crowds in order to provide guidance in case of unexpected events (e.g., mass panic). To ensure the safety of the general public as well as individuals, we thus require a flexible and intelligent method for area surveillance. One example in which current monitoring systems proved to be dangerously inefficient is the Love Parade music festival in Duisburg, Germany, July 2010. Crowd control failed to provide guidance to a large crowd, resulting in a mass panic with 21 deaths and several hundred injured. In this particular case, overloaded communication infrastructure led to a lack of information about the density and the movement of the crowd, which in turn resulted in misjudgments on appropriate strategies to resolve the situation. This incident highlights the need for more sophisticated and reliable methods for area surveillance. The SAFEST project aims to analyse the social context of area surveillance and to develop a system that can fulfill this task, both in terms of technology as well as acceptance by the general public. The system will operate in distributed way, collect anonymised data, securely transfer this data to a central location for evaluation, and if necessary notify the operator and/or issue alerts directly to the general public. SAFEST addresses the following topics: (i) it proposes a solution for crisis management, addressing social, technical, and economic issues, (ii) it enhances the protection of the population against risks and dangers, including the evaluation of acceptance of said solution, and (iii) it addresses the protection of critical infrastructures by the means of a comprehensive technical solution.

8.3. International Initiatives

8.3.1. Participation In other International Programs

8.3.1.1. STIC AmSud UCOOL

Participants: Aline Carneiro Viana, Eduardo Mucelli.

Program: STIC AmSud

Project acronym: UCOOL

Project title: Understanding and predicting human demanded COntent and mObiLity

Duration: 2013-2015

Coordinator: Aline Carneiro Viana

Other partners: National Laboratory for Scientific Computing (Brazil), Facultad de Ingeniería, Universidad de Buenos Aires (Argentina), Universidad Tecnica Federico Santa Maria (Chile), Telecom Sud Paris, and Inria (with INFINE at Saclay and DANTE at Rhone-Alpes)

Abstract: The UCOOL (Understanding and predicting human demanded COntent and mObiLity, https://macaco.inria.fr/) project is granted by STIC-AmSUD, it is a 2-year project, and has officially started in January 2014. The main goal of this project is to define solutions for the identification and modelling of correlations between the user mobility – describing changes in the user positioning and the current environment he/she is in – and the traffic demand he/she generates.

8.3.1.2. STIC Asie URSA

Participant: Aline Carneiro Viana.

Program: STIC Asie

Project acronym: URSA

Project title: Urban Sensing for Ads Networks

Duration: 2012-2014

Coordinator: Stephane MAag (Telecom SudParis)

Other partners: Telecom SudParis, Inria, Fu Jen Catholic University (Taiwan), Institute for Infocomm Research (Singapore), and IFI (Vietnam).

Abstract: The URSA (Urban Sensing for Ads Networks, http://www-public.tem-tsp.eu/ maag/URSA/) project was granted by STIC-Asie. It was a 2-year project which has started in January 2013 and has just finished. URSA aims at defining an urban sensing network based on the mobility of users and the diffusion of advertisings by fixed elements.

8.4. International Research Visitors

8.4.1. Visits of International Scientists

Artur Ziviani, National Laboratory for Scientific Computing (LNCC), Brazil, October 2014.

José Ignacio Alvarez-Hamelin, Facultad de Ingeniería, Universidad de Buenos Aires, Argentina, October 2014.

Jorge Brea, GranData and Universidad de Buenos Aires, Argentina, October 2014.

Anelise Munaretto, Federal Technological University of the Parana (UTFPR), Brazil, October 2014.

Carlos Sarraute, GranData, Argentina, October 2014.

Don Towsley, University of Massachusetts - Amherst, USA, June 2014.

8.4.1.1. Internships

Lavanya Addepalli, Universidad Politecnica de Valencia, Spain, from Jun 2014 to Nov 2014.

8.4.2. Visits to International Teams

8.4.2.1. Research stays abroad

Emmanuel Baccelli has been visiting Freie Universitaet (FU) Berlin in 2014, within the context of the SAFEST project. The closer collaboration enabled by this stay allowed the initial development of the RIOT community http://www.riot-os.org, and the development of new activities around Information-centric networking in the Internet of Things. During his stay, Emmanuel Baccelli also taught a course on Computer Networking for graduate level. More details in the corresponding sections of this document.

9. Dissemination

9.1. Promoting Scientific Activities

9.1.1. Scientific events organisation

9.1.1.1. Member of the organizing committee

Aline Carneiro Viana served as:

- Workshop co-Chair of the 11th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON 2014);
- Grant Chair of the N2Women 2014 Workshop (an ACM SIGMOBILE program that is supported by the IEEE Communications Society, Microsoft Research and HP Labs)
- Publicity Chair of 12th Annual IEEE International Conference on Sensing, Communication, and Networking (IEEE SECON 2015), 20th Annual ACM International Conference on Mobile computing and Networking (ACM MobiCom 2014) and IEEE IFIP Wireless Days 2014;
- Publication co-Chair of 21st Annual ACM International Conference on Mobile computing and Networking (ACM MobiCom 2015);

Emmanuel Baccelli served as:

- co-Chair of the 4th ACM MOBIHOC Workshop on Pervasive Wireless Healthcare (MobileHealth 2014);

9.1.2. Scientific events selection

9.1.2.1. Responsable of the conference program committee

- **Aline Carneiro Viana** is the TPC co-chair of the Symposium on Selected Topics in Communications together Jun Luo (from Nanyang Technological University, Singapore) to be organized jointly with IEEE/CIC International Conference on Communications in China (ICCC 2015)

9.1.2.2. Member of the conference program committee

Aline Carneiro Viana has served in the program committees of the following conferences/workshops:

- 7th IFIP International Conference on New Technologies, Mobility and Security (IFIP NTMS 2015)
- 12th Annual IEEE Conference on Sensing, Communication and Networking (IEEE SECON 2015)
- IEEE IFIP Wireless Days 2014
- IEEE 10th International Conference on Wireless and Mobile Computing, Networking and Communications (IEEE WiMob'2014)
- 32th Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos (SBRC 2014)

Cédric Adjih has been TPC member/reviewer for: IEEE Globecom 2014 (Global Communications Conference), and IEEE ICC 2015 (International Conference on Communications)

Emmanuel Baccelli has been TPC member for: IEEE SECON 2014 (IEEE Conference on Sensing, Communication and Networking), and SaSeIoT 2014 (International Conference on Safety and Security in Internet of Things).

9.1.3. Journal

9.1.3.1. Member of the editorial board

- **Aline Carneiro Viana** is Associate Editor of ACM Computer Communication Review (ACM CCR)

9.1.3.2. Reviewer

Aline Carneiro Viana has performed reviews for the following journals/transactions:

- IEEE Transaction on Mobile Computing (TMC)
- IEEE Transactions on Parallel and Distributed Systems (TPDS)
- IEEE Wireless Communications Magazine
- Elsevier: Computer Communications; Elsevier Computer Networks; Elsevier Ad Hoc Networks; Pervasive Mobile Computing (PMC); Sustainable Computing, Informatics and Systems (SUSCOM)
- IEEE Communications Letters

Cédric Adjih has performed reviews for the IEEE Journal on Selected Areas in Communications (JSAC), and the journal "Ad Hoc Networks".

Emmanuel Baccelli has performed reviews for the following journals/transactions:

- IEEE Transactions on Sensor Networks,
- IEEE Transactions on Emerging Telecommunications Technologies.

9.2. Teaching - Supervision - Juries

9.2.1. Teaching

Engineering School: Cédric Adjih, "Microcontrollers: from embedded systems to the Internet of Things", 3 hours TD, Telecom SudParis.

Masters

Emmanuel Baccelli, "Telematics", M1 Advanced Computer Networks, at Freie Universitaet Berlin; 15 semaines, 50 étudiants.

E-learning

Master: Laurent Massoulié, "epidemic algorithms", M2 ACN (Advanced Communication Networks) entre Ecole Polytechnique et Institut Telecom; 11 semaines, 4 étudiants. "Graphes aléatoires", M2 Probabilités et Statistiques, 4 semaines, UPMC, 10 étudiants.

9.2.2. Supervision

Aline Carneiro Viana:

- PhD: **Thiago Henrique Silva**, *"Large scale study of city dynamics and urban social behavior using participatory sensor networks"*, Federal University of Minas Gerais, May 2014.
- PhD in progress: **Eduardo Mucelli**, *"Revealing new facets of humain behavior"*, November 2011, Aline Carneiro Viana. **Felipe Domingos**, *"Identifying Social Attributes in VANETs"*, January 2013, Aline Carneiro Viana. **Roni Shigueta**, *"Resource allocation in highly mobile wireless networks"*, January 2012, Aline Carneiro Viana. **Guangshuo Chen**, *"Understanding and predicting human demanded content and mobility"*, September 2014, Aline Carneiro Viana.
- Post-doc in progress: K. P. Naveen on content caching and content-centric networking.

Cédric Adjih:

- PhD in progress: **Fatma Somaa**, on handling mobility in routing in wireless sensor networks, thesis started in 2013 and co-advised with Inès El Korbi (Institut Supérieur d'Informatique et de Gestion de Kairouan, Tunisie).

Emmanuel Baccelli:

- PhD in progress: **Oliver Hahm** on the optimization of energy-efficient communication protocols for the Internet of Things, **Hauke Petersen** on the resilience of emergency spontaneous wireless networks.
- Post-doc in progress: **K. P. Naveen** on content caching and content-centric networking.
- Masters thesis in the domain of IoT, during 2014: 1 completed, 3 on-going.
- Bachelor thesis in the domain of IoT, during 2014: 2 completed.

9.2.3. Juries

Aline Carneiro Viana:

- **Examiner:***Oana Iova*, "Algorithms and Protocols for all-IP Wireless Sensor Networks in the Internet of Things", Université de Strasbourg, France. December 2014. *Tiphaine Phe-Neau*, "Properties and impact of vicinity in mobile opportunistic networks ", UPMC – Sorbonne Universités, France. January 2014.
- **Invited:***Thiago Henrique Silva* "Large scale study of city dynamics and urban social behavior using participatory sensor networks" at the Federal University of Minas Gerais, Brazil, May 2014.
- **PhD Qualification advisor:***Felipe Domingos*, *"Identifying Social Attributes in VANETs"*, at the Federal University of Minas Gerais, Brazil, August 29, 2014. *Roni Shigueta*, *"Reseource allocation in mobile networks"*, at the Pontifical University of Paraná (PUC-PR), Brazil, November 26, 2014.

9.3. Popularization

Aline Carneiro Viana:

- **Presentation at the "Futur en Seine" festival:** Title: *"Understanding human behavior for context- and content- networking"*, June 12, 2014.
- **Presentation at the ESSOURIAU High School at Ulis in the context of the Math week:** Title: *"Le smartphone, votre 6e sens ?"*, Mars 21, 2014.
- **Interview for the Interstices scientific magazine**.

Cédric Adjih, Emmanuel Baccelli, Oliver Hahm:

- **Demonstration of RIOT** at CeBIT 2014, the world's leading IT trade show.
- **Description of IoT-LAB, and test of IoT-LAB nodes** at the IETF 6TiSCH plugfest during IETF-89 in London.
- **Demonstration of RIOT on heterogeneous IoT hardware** At the IETF 6TiSCH plugfest during IETF-89 in London (https://bitbucket.org/6tisch/meetings/wiki/140306a_ietf89_london_plugfest).
- **Demonstration of IoT-LAB (Contiki RPL experiments)** at the Bits-n-Bytes event of IETF-90 in Toronto.
- **Demonstration of RIOT (CCN - 6LoWPAN interoperating)** at the Bits-n-Bytes event of IETF-90 in Toronto.
- **Demonstration of IoT-LAB** also at the LLN plugfest of the same IETF-90 in Toronto.
- **Demonstration of RIOT tsch functionalities with OpenWSN integration** at the LLN plugfest of the same IETF-90 in Toronto.

Oliver Hahm co-authored an article on RIOT and IoT device programming in the iX Developer magazine's 2014 special issue on Embedded Software. O. Hahm, C. Mehlis, M. Lenders, T. Eichinger, L. Ortmann, "Betriebssysteme für eingebettete Systeme im Internet der Dinge: Freie Fahrt für Experimentierfreudige," iX Developer Magazine, Special issue on Embedded Software, February 2014.

10. Bibliography

Publications of the year

Articles in International Peer-Reviewed Journals

[1] Z. ABBASSI, N. HEGDE, L. MASSOULIÉ. *Distributed Content Curation on the Web*, in "ACM Transactions on Internet Technology", October 2014, 9 p. [*DOI :* 10.1145/2663489], https://hal.archives-ouvertes.fr/hal-01107836

[2] P. OLMO VAZ DE MELO, A. CARNEIRO VIANA, M. FIORE, K. JAFFRÈS-RUNSER, F. LE MOUËL, A. A. F. LOUREIRO, L. ADDEPALLI, C. GUANGSHUO. *RECAST: Telling Apart Social and Random Relationships in Dynamic Networks*, in "Performance Evaluation", March 2015, 35 p. , https://hal.inria.fr/hal-01109969

[3] D.-C. TOMOZEI, L. MASSOULIÉ. *Distributed user profiling via spectral methods*, in "Stochastic Systems", 2014, vol. 4, pp. 1-43, http://www.i-journals.org/ssy/viewarticle.php?id=36 [*DOI :* 10.1214/11-SSY036], https://hal.archives-ouvertes.fr/hal-00976912

[4] M. VALERIO BARBERA, A. CARNEIRO VIANA, M. DIAS DE AMORIM, J. STEFA. *Data Offloading in social mobile networks through VIP Delegation*, in "Ad-Hoc Networks", August 2014, vol. 19, pp. 92-110, https://hal.inria.fr/hal-01083387

International Conferences with Proceedings

[5] E. BACCELLI, G. BARTL, A. DANILKINA, V. EBNER, F. GENDRY, C. GUETTIER, O. HAHM, U. KRIEGEL, G. HEGE, M. PALKOW, H. PERTERSEN, T. SCHMIDT, A. VOISARD, M. WÄHLISCH, H. ZIEGLER. *Area & Perimeter Surveillance in SAFEST using Sensors and the Internet of Things*, in "Workshop Interdisciplinaire sur la Sécurité Globale (WISG2014)", Troyes, France, January 2014, https://hal.inria.fr/hal-00944907

[6] E. BACCELLI, O. HAHM, M. WÄHLISCH. *Spontaneous Wireless Networking to Counter Pervasive Monitoring*, in "W3C / IAB workshop on Strengthening the Internet Against Pervasive Monitoring (STRINT)", London, United Kingdom, IAB and W3C, February 2014, https://hal.inria.fr/hal-00945081

[7] E. BACCELLI, C. MEHLIS, O. HAHM, T. SCHMIDT, M. WÄHLISCH. *Information Centric Networking in the IoT: Experiments with NDN in the Wild*, in "1st ACM Conference on Information-Centric Networking (ICN-2014)", Paris, France, ACM, September 2014, https://hal.inria.fr/hal-01058627

[8] F. DOMINGOS DA CUNHA, G. MAIA, A. C. VIANA, R. A. F. MINI, L. VILLAS, A. A. F. LOUREIRO. *Socially Inspired Data Dissemination for Vehicular Ad Hoc Networks*, in "MSWiM'14 Proceedings of the 17th ACM international conference on Modeling, analysis and simulation of wireless and mobile systems", Montreal, Canada, September 2014 [*DOI :* 10.1145/2641798.2641834], https://hal.inria.fr/hal-01082984

[9] F. DOMINGOS DA CUNHA, R. A. F. MINI, A. C. VIANA, A. A. F. LOUREIRO. *Are Vehicular Networks Small World?*, in "33rd Annual IEEE International Conference on Computer Communications (INFOCOM'14) - Workshop for students", Toronto, Canada, April 2014, https://hal.inria.fr/hal-00952451

[10] F. DOMINGOS DA CUNHA, A. C. VIANA, R. A. F. MINI, A. A. F. LOUREIRO. *Is it Possible to Find Social Properties in Vehicular Networks?*, in "The Nineteenth IEEE Symposium on Computers and Communications", Ilha Madeira, Portugal, June 2014, https://hal.inria.fr/hal-00981127

[11] O. HAHM, E. BACCELLI, H. PETERSEN, M. WÄHLISCH, T. SCHMIDT. *Simply RIOT: Teaching and Experimental Research in the Internet of Things*, in "13th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN 2014)", Berlin, Germany, ACM, April 2014, https://hal.inria.fr/hal-01058634

[12] O. HAHM, S. PFEIFFER, J. SCHILLER. *On real-time requirements in constrained wireless networks for mobile health*, in "The 4th ACM MobiHoc Workshop on Pervasive Wireless Healthcare (ACM MobileHealth)", Philadelphia, United States, August 2014, pp. 1 - 6 [*DOI :* 10.1145/2633651.2637475], https://hal.inria.fr/hal-01077792

[13] L. MASSOULIÉ. *Community detection thresholds and the weak Ramanujan property*, in "STOC 2014: 46th Annual Symposium on the Theory of Computing", New York, United States, June 2014, pp. 1-10, https://hal.archives-ouvertes.fr/hal-00969235

[14] E. MUCELLI REZENDE OLIVEIRA, A. CARNEIRO VIANA. *From Routine to Network Deployment for Data Offloading in Metropolitan Areas*, in "SECON - 11th Internation Conference on Sensing, Communication and Networking", Singapore, Singapore, IEEE, June 2014, https://hal.inria.fr/hal-00931786

[15] E. MUCELLI REZENDE OLIVEIRA, A. CARNEIRO VIANA. *Routine-based Network Deployment*, in "INFOCOM Student Workshop", Toronto, Canada, April 2014, https://hal.inria.fr/hal-00953700

[16] E. MUCELLI REZENDE OLIVEIRA, A. CARNEIRO VIANA, K. P. NAVEEN, C. SARRAUTE. *Measurement-driven mobile data traffic modeling in a large metropolitan area*, in "PerCom 2015- 13th Conference on Pervasive Computing and Communications", St. Louis, Missouri, United States, IEEE, March 2015, https://hal.inria.fr/hal-01089434

[17] E. MUCELLI REZENDE OLIVEIRA, A. C. VIANA. *From Routine To Better Network Services*, in "WMNC - Wireless and Mobile Networking Conference - 2014", Algarve, Portugal, WMNC, IFIP/IEEE, April 2014, https://hal.inria.fr/hal-00931801

[18] E. MUCELLI REZENDE OLIVEIRA, A. C. VIANA. *Routine-based network deployment for data offloading in metropolitan area*, in "WCNC - Wireless Communications and Networking Conference - 2014", Istanbul, Turkey, WCNC, IEEE, April 2014, https://hal.inria.fr/hal-00862916

[19] H. PETERSEN, E. BACCELLI, M. WÄHLISCH. *Interoperable Services on Constrained Devices in the Internet of Things*, in "W3C Workshop on the Web of Things", Berlin, Germany, W3C (editor), June 2014, https://hal.inria.fr/hal-01058636

[20] H. PETERSEN, E. BACCELLI, M. WÄHLISCH, T. SCHMIDT, J. SCHILLER. *The Role of the Internet of Things in Network Resilience*, in "First International Conference on Safety and Security in Internet of Things (SaSeIoT)", Rome, Italy, October 2014, https://hal.inria.fr/hal-01078059

[21] T. H. SILVA, P. O. VAZ DE MELO, A. C. VIANA, J. SALLES, J. M. ALMEIDA, A. A. F. LOUREIRO. *Participatory Sensor Networks as Sensing Layers*, in "7th IEEE Conference on Social Computing and Networking (SocialCom'14)", Sydney, Australia, December 2014, https://hal.inria.fr/hal-01083660

[22] J. XU, L. MASSOULIÉ, M. LELARGE. *Edge Label Inference in Generalized Stochastic Block Models: from Spectral Theory to Impossibility Results*, in "Conference on Learning Theory", Barcelona, Spain, June 2014, pp. 903-920, https://hal.archives-ouvertes.fr/hal-01066047

National Conferences with Proceedings

[23] F. DOMINGOS DA CUNHA, A. C. VIANA, T. RODRIGUES, R. A. F. MINI, A. A. F. LOUREIRO. *Extração de Propriedades Sociais em Redes Veiculares*, in "XXXII Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos", Florianópolis, Brazil, SBC, May 2014, https://hal.inria.fr/hal-00981189

Research Reports

[24] E. BACCELLI, C. MEHLIS, O. HAHM, T. SCHMIDT, M. WÄHLISCH. *Information Centric Networking in the IoT: Experiments with NDN in the Wild*, Inria, June 2014, n^{o} RR-8551, https://hal.inria.fr/hal-01010879

[25] F. DOMINGOS DA CUNHA, A. BOUKERCHE, L. VILLAS, A. C. VIANA, A. A. F. LOUREIRO. *Data Communication in VANETs: A Survey, Challenges and Applications*, Inria Saclay, March 2014, n^{o} RR-8498, https://hal.inria.fr/hal-00981126

[26] E. MUCELLI REZENDE OLIVEIRA, A. CARNEIRO VIANA, K. P. NAVEEN, C. SARRAUTE. *Measurement-driven mobile data traffic modeling in a large metropolitan area*, Inria, October 2014, n^{o} RR-8613, 20 p. , https://hal.inria.fr/hal-01073129

[27] H. PETERSEN, E. BACCELLI, M. WÄHLISCH, T. SCHMIDT, J. SCHILLER. *The Role of IoT in Network Resilience*, Inria, June 2014, n^{o} RR-8552, https://hal.inria.fr/hal-01010891

Scientific Popularization

[28] O. HAHM, C. MEHLIS, M. LENDERS, T. EICHINGER, L. ORTMANN. *Betriebssysteme für Eingebettete Systeme im Internet der Dinge: Freie Fahrt für Experimentierfreudige*, in "iX Developer Magazine", February 2014, 5 p. , https://hal.inria.fr/hal-01109198