Activity Report 2015

Team INFINE

INFormation NEtworks

Inria teams are typically groups of researchers working on the definition of a common project, and objectives, with the goal to arrive at the creation of a project-team. Such project-teams may include other partners (universities or research institutions).

RESEARCH CENTER
Saclay - Île-de-France

THEME
Networks and Telecommunications

Table of contents

Team INFINE

Creation of the Team: 2014 April 01

Keywords:

Computer Science and Digital Science:

1.1.7. - Peer to peer
1.2.5. - Internet of things
1.3. - Distributed Systems
3.4. - Machine learning and statistics
3.5. - Social networks
7.10. - Network science

Other Research Topics and Application Domains:

6.4. - Internet of things
8.2. - Connected city
9.4.1. - Computer science
9.4.2. - Mathematics

1. Members

Research Scientists

Laurent Massoulie [Team leader, Inria, HdR]
Cedric Adjih [Inria, Researcher]
Emmanuel Baccelli [Inria, Researcher, HdR]
Aline Carneiro Viana [Inria, Researcher, HdR]

PhD Students

Eduardo Mucelli Rezende Oliveira [Inria]
Oliver Hahm [Inria]
Guangshuo Chen [Inria]
Lennart Gulikers [Microsoft Research-Inria Joint Centre]
Rémi Varloot [Microsoft Research-Inria Joint Centre]
Roni Shigueta [PUC-PR]

Post-Doctoral Fellows

Naveen Kolar Purushothama [Inria]
Kuang Xu [Microsoft Research-Inria Joint Centre]

Administrative Assistant

Helene Bessin Rousseau [Inria]

Others

Michael Herrmann [Inria, from May 2015 until Oct 2015]
Ichrak Amdouni [Inria until March 2015, UPEC from Apr to Dec 2015]

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 â 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. Highlights of the Year

4.1. Highlights of the Year

1. In collaboration with Charles Bordenave (CNRS, Toulouse) and Marc Lelarge (Inria) we proved the so-called « spectral redemption conjecture » formulated by physicists in 2013, suggesting that a novel spectral method for community detection would perform non-trivial detection under optimal conditions. This has been presented in the IEEE FOCS conference, one of the top two theoretical computer science conferences.
2. In collaboration with Freie Universitaet Berlin we have further developed RIOT, which now aggregates open source contributions from 120+ people (and counting) from all over the world, coming both from academia and from industry.

4.1.1. Awards

Aline Viana was awarded the PEDR in 2015, the Inria award for research excellence.

5. New Software and Platforms

5.1. RIOT

KEYWORDS: Internet of things - Operating system - Internet protocols - Sensors - IoT - Wireless Sensor Networks

SCIENTIFIC DESCRIPTION

While requiring as low as 1,5kB of RAM and 5kB or ROM, the RIOT operating system offers both advanced capabilities such as real time and energy efficiency capabilities, as well as standard capabilities, such as a single API (partially POSIX compliant) across heterogeneous IoT hardware (8-bit, 16-bit and 32-bit low-power hardware). This API is developer-friendly in that it enables Linux-like developer experience which was not possible so far for embedded programming: multi-threading, standard C and C++ application programming and the use of standard debugging tools (gbd, valgrind etc.). On top of this, as expected from a modern OS, RIOT also provides standard connectivity capabilities, by means of several networks stacks that are readily available, such as a standard IPv6/6LoWPAN stack and a information-centric network stack (based on CCN).

FUNCTIONAL DESCRIPTION

RIOT is an open source operating system that provides an API and SDK for energy-, memory- and CPU-constrained IoT devices and similar communicating embedded systems. This API builds upon standard systems concepts (e.g. POSIX) and standard communication protocols (e.g. IPv6). RIOT thus allows the development of applications that collect sensor data and transmit it to the cloud over the Internet, using end-to-end communication means e.g. IPv6 communication from sensor to cloud servers, that can be protected by transport layer security (e.g. DTLS). This data can then be used for smart energy management for example.

An important part of the design of RIOT is that it can easily be ported to different hardware devices (tens of types of heterogeneous IoT devices are already supported), and it can easily be extended to support the latest evolution of communication standards (a wide range of protocols are already supported).

For example, RIOT already enables tests and experiements of arbitrary IoT applications and protocols on FIT IoT-Lab, which provides a large-scale infrastructure facility with 2700 nodes for testing remotely small wireless devices. Once validated, the code for such applications and protocols can then run on any other IoT device hardware supported by RIOT.

- Participants: Emmanuel Baccelli and Oliver Hahm
- Partner: Freie Universitaet Berlin
- Contact: Emmanuel Baccelli
- URL: http://www.riot-os.org

5.2. GardiNet (previously known as DragonNet)

FUNCTIONAL DESCRIPTION

GardiNet (previously known as 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, GardiNet 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.

- Participants: Cédric Adjih, Ichrak Amdouni, Hana Baccouch and Antonia Masucci
- Contact: Cédric Adjih
- URL: http://getrf.gforge.inria.fr/ https://github.com/GardiNet/gardinet

5.3. MACACO

Mobile context-Adaptive CAching for COntent-centric networking
FUNCTIONAL DESCRIPTION

MACACOapp is developped in the context of the EU CHIST-ERA MACACO project. It consists in a mobile phone application that periodically samples phoneâs 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 usersâ phones location and/or from/to the wireless infrastructure closest to the usersâ 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. Userâs application preferences and interests will be inferred from the Internet browsing history and running app information obtained from the MACACO App.

- Participants: Aline Carneiro Viana, Katia Jaffres and Marco Fiore
- Contact: Aline Carneiro Viana
- URL: https://macaco.inria.fr/macacoapp/

6. New Results

6.1. Online Social Networks (OSN)

Community detection; bandit algorithms; privacy preservation; reward mechanisms

6.1.1. Community detection

Participants: Laurent Massoulié, Marc Lelarge, Charles Bordenave.

We have progressed in the design of spectral methods for community detection and in the corresponding analysis, in particular by proving the so-called spectral redemption conjecture. This has been published in IEEE FOCS'15. The abstract of the paper is as follows. A non-backtracking walk on a graph is a directed path such that no edge is the inverse of its preceding edge. The nonbacktracking matrix of a graph is indexed by its directed edges and can be used to count non-backtracking walks of a given length. It has been used recently in the context of community detection and has appeared previously in connection with the Ihara zeta function and in some generalizations of Ramanujan graphs. In this work, we study the largest eigenvalues of the non-backtracking matrix of the Erdős-Rényi random graph and of the Stochastic Block Model in the regime where the number ´ of edges is proportional to the number of vertices. Our results confirm the "spectral redemption conjecture" that community detection can be made on the basis of the leading eigenvectors above the feasibility threshold.

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

Participants: Laurent Massoulié, Mesrob Ohanessian, Alexandre Proutière.

Progress on "bandit algorithms" for targeted news dissemination. 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. This work was published at ACM Sigmetrics'15.

6.1.3. Clustering and Inference From Pairwise Comparisons

Participants: Rui Wu, Jiaming Xu, Srikant Rayadurgam, Marc Lelarge, Laurent Massoulié, Bruce Hajek.

In a short publication at ACM Sigmetrics'15, we do the following. Given a set of pairwise comparisons, the classical ranking problem computes a single ranking that best represents the preferences of all users. In this paper, we study the problem of inferring individual preferences, arising in the context of making personalized recommendations. In particular, we assume users form clusters; users of the same cluster provide similar pairwise comparisons for the items according to the Bradley-Terry model. We propose an efficient algorithm to estimate the preference for each user: first, compute the net-win vector for each user using the comparisons; second, cluster the users based on the net-win vectors; third, estimate a single preference for each cluster separately. We show that the net-win vectors are much less noisy than the high dimensional vectors of pairwise comparisons, therefore our algorithm can cluster the users reliably. Moreover, we show that, when a cluster is only approximately correct, the maximum likelihood estimation for the Bradley-Terry model is still close to the true preference.

6.2. Spontaneous Wireless Networks and Internet of Things

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

6.2.1. Platform Design for the Internet of Things

Participants: Emmanuel Baccelli, Cedric Adjih, Oliver Hahm, Matthias Waehlisch, Thomas Schmidt, Hauke Petersen.

Within this activity, we have further developed the platforms we champion for the Internet of Things: the open source operating system RIOT and open-access IoT-lab testbeds. RIOT now aggregates open source contributions from 120+ people (and counting) from all over the world, coming both from academia and from industry, and received financial backing from top companies including Cisco and Google in 2015. Revisiting concepts from the early Internet, we have designed and introduced a new software architecture that fits the (memory, CPU, energy) constraints of low-end IoT devices, while being full-featured and easily extensible, thus more future-proof that state of the art. This work was published in ACM MobiSys'15 (IoT-Sys workshop), and released as open source code, integrated in the latest version of RIOT 2015-12. We have also designed a distributed test framework which supports advanced continuous integration techniques, allows for the integration of project contributors to volunteer hardware and software resources to the test system, and

can function as a permanent distributed plugtest for network interoperability testing. This work was published in ACM MobiSys'15 (IoT-Sys workshop). Concerning IoT-lab, we have contributed to the completion of the design and the roll-out of IoT-lab testbeds in multiple sites in France and started deploying an additional one in Berlin. Description of completed work and design was published in IEEE IoT-WF'15.

6.2.2. Standards for Spontaneous Wireless Networks

Participant: Emmanuel Baccelli.

Within this activity, we have contributed to new network protocol standards for spontaneous wireless networking, applied to ad hoc networks and the Internet of Things. In particular, collaborating with Fraunhofer, we have published Directional Airtime Metric (DAT), a new wireless metric standard targeting wireless mesh networks. The standard is in the RFC editor's queue (which means the corresponding IETF standard, an RFC, will be published within weeks). Furthermore, collaborating with ARM and Sigma Designs, we published RFC 7733, which provides guidance in the configuration and use of protocols from the RPL protocol suite to implement the features required for control in building and home environments. In collaboration with various industrial partners, with have also published a number of other Internet drafts, including an analysis of the characteristics of multi-hop ad hoc wireless communication between interfaces in the context of IP networks, and an analysis of the challenges of information-centric networking in the Internet of Things.

6.3. Resource and Traffic Management

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

6.3.1. 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 was published in the ACM SIGCOMM'15 workshop on All Things Cellular.

6.3.2. From Routine 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. This work, including new characterization results of the used trace and new analysis of space-traffic correlation, is under submission in a trasaction.

6.3.3. 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 and to improve the quality of communication service provided. Hence, the definition of a *usage pattern* can allow telecommunication operators to better foreseen 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, including a new characterization results of the used trace, including analysis correlating age and gender to traffic demands, as well as new profiling results, is under submission in a transaction.

6.3.4. 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.

6.3.5. Designing Adaptive Replication Schemes in Distributed Content Delivery Networks

Participants: Mathieu Leconte, Marc Lelarge, Laurent Massoulié.

In a paper published at the ITC'15 conference, we address the problem of content replication in large distributed content delivery networks, composed of a data center assisted by many small servers with limited capabilities and located at the edge of the network. We aim at optimizing the placement of contents on the servers to offload the data center as much as possible. We model the sub-system constituted by the small servers as a loss network, each loss corresponding to a request to the data center. Based on large system / storage behavior, we obtain an asymptotic formula for the optimal replication of contents and propose adaptive schemes to attain it by reacting to losses, as well as faster algorithms which can react before losses occur. We show through simulations that our adaptive schemes outperform significantly standard replication strategies both in terms of loss rates and adaptation speed.

6.3.6. 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.

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.

6.3.7. 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.

6.3.8. 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. This work was published in 2015 at the Performance Evaluation Elsevier Journal [9].

7. Bilateral Contracts and Grants with Industry

7.1. Bilateral Contracts with Industry

1. Participation to Microsoft Research – Inria Joint Centre, which funds two PhD students (Lennart Gulikers and Remi Varloot) and funded postdoc Kuang Xu over 2015.
2. During 2015, Cisco Systems (through direct contract) and Google (through GSoC) have funded further development of RIOT.

7.2. 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 mobility analysis.

8. Partnerships and Cooperations

8.1. Regional Initiatives

8.1.1. LiveGrid

Participants: Cedric Adjih, Emmanuel Baccelli.

Infine is one of the teams from Inria participating to LiveGrid: LiveGrid is a consortium of the main actors of industry, research organisations, local authorities and competitive cluster from the Paris-Saclay campus. The goal of LiveGrid is make the Paris-Saclay campus one of the leader regions of smart grids. Infine expertise is in infrastructure: testbeds, communication protocols, embedded open source OS.

8.2. National Initiatives

8.2.1. Equipex FIT

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

Partners: Inria (Lille, Sophia-Antipolis, Grenoble), INSA, UPMC, Institut Telecom Paris, Institut TÃ©lÃ©com 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 Researchâs "Ãquipements dâExcellence" (Equipex) research grant program, in 2011.

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

The Infine team is more specifically managing the FIT IoT-LAB site formerly at Rocquencourt, which recently moved to Saclay (on-going re-deployment), and is participating in the deployment of an additional IoT-lab testbed in Berlin (at Freie Universitaet Berlin).

The Infine team is actively collaborating with UPEC on wireless sensor network testbeds (and protocols): in 2015, the testbed from UPEC with 45 Arduino nodes has been integrated with the FIT IoT-LAB testbed.

8.3. European Initiatives

8.3.1. FP7 & H2020 Projects

8.3.1.1. AGILE (H2020 project)

Participants: Emmanuel Baccelli, Cedric Adjih.

Program: H2020 ICT-30-2015 Topic: Internet of Things and Platforms for Connected Smart Objects

Project acronym: AGILE

Project title: Adoptive Gateways for dIverse muLtiple Environments

Duration: 2015-2017

Coordinator: Emmanuel Baccelli

Other partners: Canonical (UK), Eclipse IoT Foundation (IE), Mobistar (BE), Libelium (ES), Startupbootcamp IoT (SP), CREATE-NET (IT), iMinds (BE), Atos (SP), Rulemotion (UK), Jolocom (DE), Passau University (DE), Sky-Watch (DN), BioAssist (GR), Graz Technical University (AT), Eurotech (IT), IoTango (US).

Abstract:

The AGILE project is a 3-year H2020 project started in January 2016, which will deliver an integrated framework of open source tools and platforms that interoperate for enabling the delivery of adaptive, self-configurable and secure IoT elements (both software and hardware) that can be utilized in a variety of scenarios. Such tools target actors with heterogeneous skills, including entrepreneurs, researchers, and individuals, aiming to enable the realization of IoT applications respecting user privacy and data ownership.

8.3.2. Collaborations in European Programs, except FP7 & H2020

8.3.2.1. 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. âwhatâ, âwhenâ and âwhereâ 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.

8.3.2.2. ANR/BMBF SAFEST

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. Project completed in 2015. One of its main result was the emergence of RIOT.

8.4. International Initiatives

8.4.1. Inria International Partners

8.4.1.1. Declared Inria International Partners

1. On-going formal collaboration with Freie Universitaet Berlin around the long-term stay of Emmanuel Baccelli in Berlin, around the topics of the Internet of Things and Information-Centric Networking.
2. The Inria teams Infine and Eva are part of the "D2D Communication for LTE Advanced Cellular Network" , a project funded by the Indo-French Centre for the Promotion of Advanced Research (CEFIPRA). With industrial partners, and also with Indian partners, this project is focusing on the evolution of cellular networks towards 5G: this includes exploration of device-to-device (D2D) communication, and more generally IoT communication in a cellular context. Research directions include efficient access for IoT devices (massive numbers of devices with low volume communication); combination of random access protocols/error coding/physical layer ; efficient neighbor discovery, ldots.

8.4.1.2. Informal International Partners

1. On-going collaboration with Hamburg University of Applied Science around RIOT.
2. Informal collaborations with UIUC and UMass.

8.4.2. Participation In other International Programs

8.4.2.1. 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 IngenierÃ-a, 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 â describing changes in the user positioning and the current environment he/she is in â and the traffic demand he/she generates.

8.4.2.2. PHC PESSOA 2015

Participant: Aline Carneiro Viana.

PHC PESSOA 2015 with University of Coimbra (2015-2016).

Program: -FCT - Programa PESSOA

Project title: Routine-based Enhanced Connectivity under User Mobility

Duration: 2015-2016

Coordinator: Aline Carneiro Viana and João Paulo da Silva Machado Garcia Vilela (University of Coimbra)

Abstract: The main goal of this project is to improve WiFi connectivity of users under mobility. The steady growth of smart-phones usage has put cellular networks under great strain, justifying the need for WiFi offloading as a solution that transfers part of the demand on cellular networks to WiFi hotspots that are in many cases already available. However, this must be performed in a way that provides benefits to the cellular operator while ensuring users a similar level of connectivity that they would achieve with cellular networks, even under user mobility (e.g. walking, taking a bus/train, etc). In this work we aim at (1) developing prediction mechanisms for selection of best hotspots by users under mobility, and (2) develop lightweight security schemes to reduce the burden of the association/authentication process of WiFi networks, therefore making WiFi offloading an effective and secure alternative to the growing demand on cellular networks.

8.5. International Research Visitors

8.5.1. Visits of International Scientists

1. Artur Ziviani, National Laboratory for Scientific Computing (LNCC), Brazil, October 2015.
2. José Ignacio Alvarez-Hamelin, Facultad de Ingeniería, Universidad de Buenos Aires, Argentina, November 2015.
3. Joao P. Vilela, University of Coimbra, July 2015

8.5.2. Visits to International Teams

8.5.2.1. Research stays abroad

Emmanuel Baccelli has been visiting Freie Universitaet (FU) Berlin in 2015, 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.

9. Dissemination

9.1. Promoting Scientific Activities

9.1.1. Scientific events organisation

9.1.1.1. General chair, scientific chair

- Emmanuel Baccelli co-chair of ACM MobileHealth 2015

9.1.1.2. Member of the organizing committees

- Aline Viana was Publication Chair of ACM MobiCom 2015, IEEE SECON 2015 and part of Poster and Demo committee of ACM SIGCOMM 2015

9.1.2. Scientific events selection

9.1.2.1. Chair of conference program committees

- In 2015 Aline Viana was TPC co-chair of the Symposium on Selected Topics in Communications.

9.1.2.2. Member of the conference program committees

- In 2015 Aline Viana was on the conference programm committee for IFIP NTMS, IEEE SECON

9.1.2.3. Reviewer

- In 2015 Cedric Adjih was reviewer for InterIoT 2015, ICC 2016, PEMWN 2015, Ad Hoc & Sensor Wireless Networks.

9.1.3. Journal

9.1.3.1. Member of the editorial boards

- In 2015 Aline Viana was Associate Editor of ACM Computer Communication Review (ACM CCR)

9.1.3.2. Reviewer - Reviewing activities

- In 2015 Aline Viana was reviewer for IEEE Transaction on Mobile Computing (TMC), IEEE Transactions on Parallel and Distributed Systems (TPDS), Elsevier Pervasive and Mobile Computing, Elsevier AdHoc Networks.

9.2. Teaching - Supervision - Juries

9.2.1. Teaching

- **Masters**
- Engineering School: Cédric Adjih, “Microcontrollers: from embedded systems to the Internet of Things”, 3 hours TD, Telecom SudParis.
- Engineering school, third year students at Ecole Polytechnique: Laurent Massoulié taught an 18 hour course on “Networks: distributed control and emerging phenomena”, and gave 18 hours of exercise classes for the Operations Research course taught by Stéphane Gaubert.
- PhD students course: Laurent Massoulié taught a 10 hour course to PhD students (mostly from EPFL and ETH but not only) on community detection, as part of a summer school in Les Diablerets.
- M2 course: Laurent Massoulié taught 10 hours in an M2 course of the Master of Probability and Statistics at Jussieu, on random graphs and epidemic models.

9.2.2. Supervision

1. Cedric Adjih currently advises 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).
2. Laurent Massoulie currently advises 2 PhD students: Lennart Gulikers and Remi Varloot.

3. Emmanuel Baccelli currently advises 2 PhD students: Oliver Hahm, on ”Dependable IoT protocols and platforms”, (Nov. 2012-) and Hauke Petersen, on ”Holistic and generic IoT network stacks”, (Nov. 2013-).
4. Aline Viana currently advises 3 PhD students: Felipe Domingos, on ”Identifying Social Attributes in VANETs” (January 2013-). Roni Shigueta, on ”Resource allocation in highly mobile wireless networks” (January 2012-). Guangshuo Chen, on ”Understanding and predicting human demanded content and mobility” (September 2014-).
5. PhD completed: Eduardo Mucelli, ”Revealing new facets of humain behavior”, November 2011-May 2015. Advisor: Aline Carneiro Viana.

9.2.3. Juries

1. Aline Viana was on the PhD jury of Filippo Rebecchi, ”Device-to-Device Data Offloading: From Model to Implementation”, Univ. Paris 6 - Sorbone Universite, September 2015;
2. Aline Viana was on the PhD jury of N. R. Zema, ”Spontaneous mobility and autonomours (Re)configuration techiniques to support next generation networks”, Univ. of Reggio Calabria, April 2015;

9.3. Popularization

1. An interview from Laurent Massoulie was published in «La Recherche », concerning our work on community detection.
2. Cedric Adjih participated to the "Inria-Industry Meeting" dedicated to energy transition, with a demo of RIOT, an open-source OS for the Internet of Things. Among different uses of the RIOT, RIOT enables IoT devices and embedded systems to easily participate in an interoperable, IP-based, communication infrastructure. Such a standard-based infrastructure is essential for data collection (e.g. sensors) providing the information feeding smart energy systems (big data analytics).
3. Emmanuel Baccelli was the editor (with Dave Raggett) of ERCIM News Special Issue on Internet of Things and Web of Things (ERCIM News No 101).

10. Bibliography

Publications of the year

Doctoral Dissertations and Habilitation Theses

[1] E. MUCELLI REZENDE OLIVEIRA. *From Human Routine to More Efficient Mobile Networks*, École Polytechnique, May 2015, https://pastel.archives-ouvertes.fr/tel-01160280

Articles in International Peer-Reviewed Journals

[2] C. ADJIH, E. BACCELLI, E. FLEURY, G. HARTER, N. MITTON, T. NOEL, R. PISSARD-GIBOLLET, F. SAINT-MARCEL, G. SCHREINER, J. VANDAELE, T. WATTEYNE. *FIT IoT-LAB: A Large Scale Open Experimental IoT Testbed – A valuable tool for IoT deployment in Smart Factories*, in "IEEE ComSoc Multimedia Technical Committee E-Letter", December 2015, https://hal.inria.fr/hal-01208387

[3] E. BACCELLI, O. HAHM, H. PETERSEN, K. SCHLEISER. *RIOT and the Evolution of IoT Operating Systems and Applications*, in "ERCIM News", April 2015, n^{o} 101, https://hal.inria.fr/hal-01244722

[4] S. BANERJEE, N. HEGDE, L. MASSOULIÉ. *The Price of Privacy in Untrusted Recommender Systems*, in "IEEE Journal of Selected Topics in Signal Processing", 2015, vol. 9, n^{o} 7, pp. 1319 - 1331 [*DOI :* 10.1109/JSTSP.2015.2423254], https://hal.archives-ouvertes.fr/hal-01226756

[5] O. HAHM, E. BACCELLI, H. PETERSEN, N. TSIFTES. *Operating Systems for Low-End Devices in the Internet of Things: a Survey*, in "IEEE Internet of Things Journal", December 2015 [*DOI :* 10.1109/JIOT.2015.2505901], https://hal.inria.fr/hal-01245551

[6] A. KARBASI, S. IOANNIDIS, L. MASSOULIÉ. *From Small-World Networks to Comparison-Based Search*, in "IEEE Transactions on Information Theory", 2015, vol. 61, n° 6, 19 p. [*DOI :* 10.1109/TIT.2015.2418284], https://hal.archives-ouvertes.fr/hal-01226763

[7] M. S. KURAN, A. CARNEIRO VIANA, L. IANNONE, D. KOFMAN, G. MERMOUD, J.-P. VASSEUR. *A Smart Parking Lot Management System for Scheduling the Recharging of Electric Vehicles*, in "IEEE Transactions on Smart Grid", November 2015, pp. 1-11, https://hal.inria.fr/hal-01112692

[8] G. MAIA, L. VILLAS, A. CARNEIRO VIANA, A. L. L. AQUINO, A. BOUKERCHE, A. A. F. LOUREIRO. *A Rate Control Video Dissemination Solution for Extremely Dynamic Vehicular Ad hoc Networks*, in "Performance Evaluation", January 2015, 16 p. , https://hal.inria.fr/hal-01118850

[9] 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, vol. 87, n° 19-36 [*DOI :* 10.1016/J.PEVA.2015.01.005], https://hal.inria.fr/hal-01109969

Invited Conferences

[10] T. WATTEYNE, C. ADJIH, X. VILAJOSANA. *Lessons Learned from Large-scale Dense IEEE802.15.4 Connectivity Traces*, in "CASE 2015 - IEEE International Conference on Automation Science and Engineering", Gothenburg, Sweden, August 2015 [*DOI :* 10.1109/COASE.2015.7294053], https://hal.inria.fr/hal-01208425

International Conferences with Proceedings

[11] C. BORDENAVE, M. LELARGE, L. MASSOULIÉ. *Non-backtracking spectrum of random graphs: community detection and non-regular Ramanujan graphs*, in "2015 IEEE 56th Annual Symposium on Foundations of Computer Science", Berkeley, United States, October 2015 [*DOI :* 10.1109/FOCS.2015.86], https://hal.archives-ouvertes.fr/hal-01226796

[12] M. LECONTE, M. LELARGE, L. MASSOULIÉ. *Designing Adaptive Replication Schemes in Distributed Content Delivery Networks*, in "Teletraffic Congress (ITC 27), 2015 27th International", Ghent, Belgium, 2015 [*DOI :* 10.1109/ITC.2015.11], https://hal.archives-ouvertes.fr/hal-01226791

[13] L. MASSOULIÉ, A. PROUTIÈRE, M. OHANNESSIAN. *Greedy-Bayes for Targeted News Dissemination*, in "SIGMETRICS '15 Proceedings of the 2015 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems", Portland, United States, June 2015, vol. 43, n° 1, 12 p. [*DOI :* 10.1145/2745844.2745868], https://hal.archives-ouvertes.fr/hal-01226779

[14] 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

[15] W. RUI, J. XU, S. RAYADURGAM, M. LELARGE, L. MASSOULIÉ, B. HAJEK. *Clustering and Inference From Pairwise Comparisons*, in "SIGMETRICS '15 Proceedings of the 2015 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems", Portland, United States, 2015, vol. 43, n° 1, 2 p. [*DOI :* 10.1145/2796314.2745887], https://hal.archives-ouvertes.fr/hal-01226785

Conferences without Proceedings

[16] C. ADJIH, E. BACCELLI, E. FLEURY, G. HARTER, N. MITTON, T. NOEL, R. PISSARD-GIBOLLET, F. SAINT-MARCEL, G. SCHREINER, J. VANDAELE, T. WATTEYNE. *FIT IoT-LAB: A Large Scale Open Experimental IoT Testbed*, in "IEEE World Forum on Internet of Things (IEEE WF-IoT)", Milan, Italy, December 2015, https://hal.inria.fr/hal-01213938

[17] I. AMDOUNI, C. ADJIH, T. PLESSE. *Network coding in military wireless ad hoc and sensor networks: Experimentation with GardiNet*, in "International Conference on Military Communications and Information Systems ICMCIS (former MCC)", Cracow, Poland, May 2015 [*DOI :* 10.1109/ICMCIS.2015.7158701], https://hal.inria.fr/hal-01244826

[18] E. BACCELLI, A. DANILKINA, S. MÜLLER, A. VOISARD, M. WÄHLISCH. *Privacy-Preserving Crowd Incident Detection: A Holistic Experimental Approach*, in "ACM SIGSPATIAL Workshop on the Use of GIS in Emergency Management (EM-GIS-2015)", Seattle, United States, November 2015, https://hal.inria.fr/hal-01244673

[19] A. DANILKINA, G. ALLARD, E. BACCELLI, G. BARTL, F. GENDRY, O. HAHM, G. HEGE, U. KRIEGEL, M. PALKOW, H. PETERSEN, T. C. SCHMIDT, A. VOISARD, M. WÄHLISCH, H. ZIEGLER. *Multi-Camera Crowd Monitoring: The SAFEST Approach*, in "Workshop Interdisciplinaire sur la Sécurité Globale", Troyes, France, February 2015, https://hal.inria.fr/hal-01244781

[20] C. GRECO, M. KIEFFER, C. ADJIH. *NeCoRPIA: Network Coding with Random Packet-Index Assignment for Mobile Crowdsensing*, in "ICC 2015", London, United Kingdom, June 2015 [*DOI :* 10.1109/ICC.2015.7249334], https://hal.inria.fr/hal-01244856

[21] N. KOLAR PURUSHOTHAMA, M. LAURENT, E. BACCELLI, A. CARNEIRO VIANA, D. TOWSLEY. *On the Interaction between Content Caching and Request Assignment in Cellular Cache Networks*, in "ACM SIGCOMM Workshop on All Things Cellular: Operations, Applications and Challenges", Londres, United Kingdom, August 2015 [*DOI :* 10.1145/2785971.2785975], https://hal.inria.fr/hal-01244774

[22] H. PETERSEN, C. ADJIH, O. HAHM, E. BACCELLI. *IoT Meets Robotics - First Steps, RIOT Car, and Perspectives*, in "ACM International Conference on Embedded Wireless Systems and Networks (EWSN)", Graz, Austria, February 2016, https://hal.inria.fr/hal-01262638

[23] H. PETERSEN, M. LENDERS, M. WÄHLISCH, O. HAHM, E. BACCELLI. *Old Wine in New Skins? Revisiting the Software Architecture for IP Network Stacks on Constrained IoT Devices*, in "ACM MobiSys Workshop on IoT challenges in Mobile and Industrial Systems (IoT-Sys-2015)", Florence, Italy, May 2015, pp. 31-35 [*DOI :* 10.1145/2753476.2753485], https://hal.inria.fr/hal-01244749

[24] P. ROSENKRANZ, M. WÄHLISCH, E. BACCELLI, L. ORTMANN. *A Distributed Test System Architecture for Open-source IoT Software*, in "ACM MobiSys Workshop on IoT challenges in Mobile and Industrial Systems (IoT-Sys-2015)", Florence, Italy, May 2015 [*DOI :* 10.1145/2753476.2753481], https://hal.inria.fr/hal-01244763

Scientific Books (or Scientific Book chapters)

[25] E. BACCELLI, D. RAGGETT. *The Internet of Things and The Web of Things*, ERCIM News Special Issue on The Internet of Things and The Web of Things, ERCIM, April 2015, n^{o} 101, https://hal.inria.fr/hal-01244735

Research Reports

[26] N. KOLAR PURUSHOTHAMA, L. MASSOULIE, E. BACCELLI, A. CARNEIRO VIANA, D. TOWSLEY. *On the Interaction between Content Caching and Request Assignment in Cellular Cache Networks*, Inria Saclay ; Inria, March 2015, n^{o} RR-8707, https://hal.inria.fr/hal-01138204

Scientific Popularization

[27] E. FLEURY, N. MITTON, T. NOEL, C. ADJIH. *FIT IoT-LAB: The Largest IoT Open Experimental Testbed*, in "ERCIM News", April 2015, n^{o} 101, 14 p. , https://hal.inria.fr/hal-01138038

Other Publications

[28] K. BHUMIP, S. SENTHIL, H. EVANGELOS, C. ADJIH. *Impact of Virtualization and SDN on Emerging Network Coding*, October 2015, Internet Draft, https://hal.inria.fr/hal-01254894

[29] C. GUANGSHUO, A. CARNEIRO VIANA. *Poster: On the Quest for Representative Mobile Datasets*, September 2015, S3'15 Proceedings of the 2015 Workshop on Wireless of the Students, by the Students & for the Students, Poster [*DOI :* 10.1145/2801694.2802147], https://hal.inria.fr/hal-01244990