



RESEARCH CENTER

FIELD

**Networks, Systems and Services,
Distributed Computing**

Activity Report 2015

Section Application Domains

Edition: 2016-03-21

DISTRIBUTED SYSTEMS AND MIDDLEWARE

1. ASAP Project-Team (section vide) 5
2. ATLANMODELS Team 6
3. CIDRE Project-Team 7
4. COAST Project-Team (section vide) 8
5. CTRL-A Team 9
6. MIMOVE Team 11
7. MYRIADS Project-Team 13
8. REGAL Project-Team (section vide) 14
9. SCALE Team 15
10. SPIRALS Project-Team 17
11. WHISPER Project-Team 18

DISTRIBUTED AND HIGH PERFORMANCE COMPUTING

12. ALPINES Project-Team 19
13. AVALON Project-Team 21
14. HIEPACS Project-Team 23
15. KERDATA Project-Team 26
16. MESCAL Project-Team 27
17. MOAIS Project-Team (section vide) 28
18. ROMA Project-Team 29
19. STORM Team 30
20. TADAAM Team 31

DISTRIBUTED PROGRAMMING AND SOFTWARE ENGINEERING

21. ASCOLA Project-Team 32
22. DIVERSE Project-Team 34
23. FOCUS Project-Team 35
24. INDES Project-Team 36
25. PHOENIX Project-Team 37
26. RMOD Project-Team 39
27. TACOMA Team 40

NETWORKS AND TELECOMMUNICATIONS

28. COATI Project-Team 42
29. DANTE Project-Team 43
30. DIANA Project-Team (section vide) 45
31. DIONYSOS Project-Team (section vide) 46
32. DYOGENE Project-Team 47
33. EVA Team 49
34. FUN Project-Team 51
35. GANG Project-Team 52
36. INFINE Team (section vide) 53
37. MADYNES Project-Team 54

| | |
|---|----|
| 38. MAESTRO Project-Team | 55 |
| 39. MUSE Team | 56 |
| 40. RAP Project-Team (section vide) | 59 |
| 41. SOCRATE Project-Team (section vide) | 60 |
| 42. URBANET Team | 61 |

ASAP Project-Team (section vide)

ATLANMODELS Team

4. Application Domains

4.1. Application Domain

By definition, MDE can be applied to any software domain. Core MDE techniques developed by the team have been successfully applied to a large variety of industrial domains from information systems to embedded systems. MDE is not even restricted to software engineering, but also applies to data engineering [45] and to system engineering [37]. There are a lot of problems in these application domains that may be addressed by means of modeling and model transformation techniques.

As a result, AtlanMod has collaborated with a great variety of different companies ranging from the Automotive to the Insurances domains and from SMEs to large enterprises through the projects described later on in this same report. AtlanMod hopes to continue this trend in the future.

CIDRE Project-Team

4. Application Domains

4.1. Application Domains

With the infiltration of computers and software in almost all aspects of our modern life, security can nowadays be seen as an absolutely general concern. As such, the results of the research targeted by CIDRE apply to a wide range of domains. It is clear that critical systems, in which security (and safety) is a major concern can benefit from ideas such as dynamic security policy monitoring. On the other hand, systems used by the general public (basically, the internet and services such as web or cloud services, social networks, location-based services, etc.) can also benefit from results obtained by CIDRE, in particular to solve some of the privacy issues raised by these systems that manipulate huge amount of personal data. In addition, systems are getting more and more complex, decentralized, distributed, or spontaneous. Cloud computing, in particular, brings many challenges that could benefit from ideas, approaches and solutions studied by CIDRE in the context of distributed systems.

Industrial Control Systems and in particular Supervisory Control and Data Acquisition are also new application domains for intrusion detection. The Stuxnet attack has emphasized the vulnerability of such critical systems which are not totally isolated anymore. Securing ICS is challenging since modifications of the systems, for example to patch them, are often not possible. High availability requirements also often conflict with preventive approaches. In this case, security monitoring is appealing to protect such systems against malicious activities. Intrusion detection in ICS is not fundamentally different from traditional approaches. However, new hypotheses and constraints need to be taken into account, which also bring interesting new research challenges.

COAST Project-Team (section vide)

CTRL-A Team

4. Application Domains

4.1. Distributed systems and High-Performance Computing

Distributed systems have grown to levels of scale and complexity where it is difficult to master their administration and resources management, in dynamic and open environments. One of the growing concerns is that the energy consumption has reached levels where it can not be considered negligible anymore, ecologically or economically. Data centers or high performance computing grids need to be controlled in order to combine minimized power needs with sustained performance and quality of service. As mentioned above, this motivates the automation of their management, and is the major topic of, amongst others, our ANR project Ctrl-Green.

Another challenge in distributed systems is in the fast growing amounts of data to process and store. Currently one of the most common ways of dealing with these challenges is the parallel programming paradigm MapReduce which is slowly becoming the de facto tool for Big Data analytics. While its use is already widespread in the industry, ensuring performance constraints while also minimizing costs provides considerable challenges. Current approaches to ensure performance in cloud systems can be separated into three categories: static, reactive, predictive and hybrid approaches. In the industry, static deployments are the standard and usually tuned based on the application peak demand and are generally over-provisioned. Reactive approaches are usually based on reacting to an input metric such as the current CPU utilisation, request rate, response time by adding and removing servers as necessary. Some public cloud providers offer reactive techniques such as the Amazon Auto Scaler. They provide the basic mechanisms for reactive controllers, but it is up to the user to define the static scaling thresholds which is difficult and not optimal. To deal with this issue, we propose a control theoretical approach, based on techniques that have already proved their usefulness for the control community.

In the domain of parallel systems and High Performance Computing, systems are traditionally less open and more controlled by administrators, but this trend is changing, as they are facing the same challenges in energy consumption, needs for adaptivity in reaction to changing workloads, and security issues in computation outsourcing. Topics of interest for us in this domain concern problem in dynamical management of memory and communications features, which we are exploring in the HPES project of the Labex Persybal-lab (see [9.1](#)).

4.2. Reconfigurable architectures in embedded systems

Dynamically reconfigurable hardware has been identified as a promising solution for the design of energy efficient embedded systems. A common argument in favor of this kind of architecture is the specialization of processing elements, that can be adapted to application functions in order to minimize the delay, the control cost and to improve data locality. Another key benefit is the hardware reuse to minimise the area, and therefore the static power and cost. Further advantages such as hardware updates in long-life products and self-healing capabilities are also often mentioned. In presence of context changes (e.g. environment or application functionality), self-adaptive technique can be applied as a solution to fully benefit from the runtime reconfigurability of a system.

Dynamic Partial Reconfiguration (DPR) of FPGA is another accessible solution to implement and experiment reconfigurable hardware. It has been widely explored and detailed in literature. However, it appears that such solutions are not extensively exploited in practice for two main reasons: i) the design effort is extremely high and strongly depends on the available chip and tool versions; and ii) the simulation process, which is already complex for non-reconfigurable systems, is prohibitively large for reconfigurable architectures. As a result, new adequate methods are required to fully exploit the potential of dynamically reconfigurable and self-adaptive architectures. We are working in this topic, especially on the reconfiguration control aspect, in cooperation with teams specialized in reconfigurable architectures such as the former DaRT team at Inria Lille, and LabSticc in Lorient, as in the recently ended ANR project Famous.

A new ANR project in this application domain, starting end of 2015, is called HPeC, in cooperation with amongst others LabSticc in Lorient and Clermont-Ferrand U., will consider embedded video processing on drones (see [9.2.1](#)).

4.3. Smart environments and Internet of Things

Another application domain for autonomic systems design and control is the Internet of Things, and especially the design of smart environments, at the level of homes, buildings, or cities. These domains are often considered at the level of sensors networks, with a strong emphasis on the acquisition of data in massive scales. The infrastructures are sometimes also equipped with actuators, with a wide range of applications, for example concerning lighting or heating, or access and security aspects. We are interested in closing the control loop in such environments, which is less often studied. In particular, rule-based languages are often used to define the automated systems, and we want to contribute to the safe design of such controllers with guarantees on their behaviors. We are working in this topic in cooperation with teams specialized in infrastructures for smart environments at CEA LETI/DACLE and Orange labs (see [8.1](#), [8.2](#)).

MIMOVE Team

4. Application Domains

4.1. Mobile urban systems for smarter cities

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

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

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

- **Democratization of urban data for healthy cities.** The objective here is to integrate the various urban data sources, especially by way of crowd-Xing, to better understand city nuisances from raw pollution sensing (e.g., sensing noise) to the sensing of its impact on citizens (e.g., how people react to urban noise and how this affects their health).
- **Socially-aware urban mobility.** Mobility within mega-cities is known as one of the major challenges to face urgently due to the fact that today's mobility patterns do not scale and to the negative effect on the environment and health. It is our belief that mobile social and physical sensing may significantly help in promoting the use of public transport, which we have started to investigate through empirical study based on the development and release of dedicated apps.
- **Social applications.** Mobile applications are being considered by sociologists as a major vehicle to actively involve citizens and thereby prompt them to become activists. This is especially studied with the Social Apps Lab at UC Berkeley. Our objective is to study such a vehicle from the ICT perspective and in particular elicit relevant middleware solutions to ease the development and development of such "*civic apps*".

Acknowledging the need for collaborative research in the application domain of smart cities, MiMove is heavily involved and actually leading CityLab@Inria ⁰. CityLab is focused on the study of ICT solutions promoting social sustainability in smart cities, and involves the following Inria project-teams in addition to MiMove: CLIME, DICE, FUN, MYRIADS, SMIS, URBANET and WILLOW. CityLab further involves strong collaboration with California universities affiliated with CITRIS (Center for Information Technology Research in the Interest of Society) and especially UC Berkeley, in relation with the *Inria@SiliconValley* program. We note that Valérie Issarny acts as scientific manager of Inria@SiliconValley and is currently visiting scholar at CITRIS at UC Berkeley. In this context, MiMove researchers are working closely with colleagues of UC Berkeley, including researchers from various disciplines interested in smart cities (most notably sociologists).

⁰<https://citylab.inria.fr>

MYRIADS Project-Team

4. Application Domains

4.1. Application Domains

The Myriads team investigates the design and implementation of system services. Thus its research activities address a broad range of application domains. We validate our research results with selected use cases in the following application domains:

- Web services, Service oriented applications,
- Business applications,
- Bio-informatics applications,
- Computational science applications,
- Numerical simulations,
- Energy and sustainable development,
- Smart cities.

REGAL Project-Team (section vide)

SCALE Team

4. Application Domains

4.1. Simulation

4.1.1. Discrete-event simulation

Simulation is an example of an application with ever increasing computation needs that would benefit from the SCALE research results. In emergency planning and response, for example, users need to access the power of large scale distributed computing facilities to run faster than real-time simulations of the situations they face on the field; Such a computation can mix heterogeneous distributing computing platforms (PDA and laptops on the field, Cloud and HPC in background) and use a number of external services (eg. weather forecast).

Simulations made of multi-party contributed software models also demonstrate the need for a unifying and user-friendly programming model. Indeed, since the early 70's, the simulation field have been the subject of many efforts in order to abstract the computation models from their actual application domain. DEVS (Discrete Event Systems specification), is an example of such a popular formalism in the simulation community that breaks-down the representation of a simulation model into hierarchical components.

Our objective is to focus on the operational support of execution for such simulation models. For example, considering that the model of a single node of a Peer-to-peer network requires several (and possibly many) DEVS components, it is easy to see that running simulations of a realistic large-scale peer-to-peer network rapidly ends-up involving millions of DEVS components. In addition to the problems posed by the execution of a distributed simulation application made of millions of components, such a use-case is also challenging in terms of analytics, because when millions of components are instrumented to collect observations, it becomes a typical instance of a big-data analytics problem.

4.1.2. Stochastic simulation platform

Understanding how complex objects, as found in finance/insurance (option contracts), biology (proteins structure), etc. evolve is often investigated by stochastic simulations (e.g. Monte-Carlo based). These can be very computational intensive and the associated communities are always seeking adequate parallel computing infrastructures and simulation software. Being able to harness all the available computing power, while ensuring the simulation is at first performant but also robust, capable to self-adapt, e.g. to failures, is a real opportunity for research and validation of our approach. Many other simulation applications could also benefit from our models and techniques, and we may in the future set up specific collaborations, e.g. in biocomputing, data-center activity management, or other engineering domains. We have recently solved pricing of high-performance demanding financial products on heterogeneous GPUs and multicore CPUs clusters, mixing use of active objects and OpenCL codes. This kind of application could continue to serve as a benchmark for our multi-level programming model.

4.2. Big data

4.2.1. Big data analytics

The amount of data digitally produced is increasing at an exponential rate. Having a dedicated programming model and runtime, such as Hadoop-MapReduce, has proved very useful to build efficient big data mining and analysis applications albeit for very static environments. However, if we consider that not only the environment is dynamic (node sharing, failures...) but so are the data (variation in popularity, arrival rate...), it becomes a much more complex problem. This domain is thus a very good candidate as an application field for our work.

More precisely, we plan to contribute at the deployment level, runtime level, and at the analytics programming model for the end-user level. We already worked on close topics with the distributed P2P storage and publish/subscribe system for Semantic Web data (named *EventCloud*). However, expressing a particular interest about data through simple or even more complex subscriptions (CEP) is only a first step in data analytics. Going further requires the full expressivity of a programming language to express how to mine into the real-time data streams, aggregate intermediate analytics results, combine with past data when relevant, etc. We intend to enlarge this effort about extracting meaningful information by also creating tighter collaborations with groups specialized in data mining algorithms (e.g. the Mind team at I3S).

We think that the approach advocated in SCALE is particularly adapted to the programming and support of analytics. Indeed, the mix of computational aspects and of large amount of data make the computation of analytics the perfect target for our programming paradigms. We aim at illustrating the effectiveness of our approach by experimenting on different computations of analytics, but we will put a particular focus on the case of data streams, where the analysis is made of chains (even cyclic graphs) of parallel and distributed operators. These operators can naturally be expressed as coarse grained composition of fine grained parallel entities, both granularity levels featuring autonomic adaptation. Also, the underlying execution platform that supports this execution also has to feature autonomic adaptation in order to deal with an unstable and heterogeneous execution environment. Here autonomic adaptation is also crucial because the programmer of analytics is not expected to be an expert in distributed systems.

Overall, this second application domain target should illustrate the effectiveness of our runtime platform and of our methodology for dynamic and autonomic adaptation.

SPIRALS Project-Team

4. Application Domains

4.1. Introduction

Although our research is general enough to be applied to many application domains, we currently focus on applications and distributed services for the retail industry and for the digital home. These two application domains are supported by a strong expertise in mobile computing and in cloud computing that are the two main target environments on which our research prototypes are build, for which we are recognized, and for which we have already established strong collaborations with the industrial ecosystem.

4.2. Distributed software services for the retail industry

This application domain is developed in relation with the **PICOM** (*Pôle de compétitivité Industries du Commerce*) cluster. We have established strong collaborations with local companies in the context of former funded projects, such as Cappucino and **Macchiato**, which focused on the development of a new generation of mobile computing platforms for e-commerce. We are also involved in the Datalyse and OCCIware funded projects that define cloud computing environments with applications for the retail industry. Finally, our activities in terms of crowd-sensing and data gathering on mobile devices with the APISENSE[®] platform share also applications for the retail industry.

4.3. Distributed software services for the digital home

We are developing new middleware solutions for the digital home, in particular through our long standing collaboration with Orange Labs. We are especially interested in developing energy management and saving solutions with the POWERAPI software library for distributed environments such the ones that equip digital homes. We are also working to bridge the gap between distributed services hosted on home gateways and distributed services hosted on the cloud to be able to smoothly transition between both environments. This work is especially conducted with the SALOON platform.

WHISPER Project-Team

4. Application Domains

4.1. Linux

Linux is an open-source operating system that is used in settings ranging from embedded systems to supercomputers. The most recent release of the Linux kernel, v3.17, comprises over 12 million lines of code, and supports 29 different families of CPU architectures, 73 file systems, and thousands of device drivers. Linux is also in a rapid stage of development, with new versions being released roughly every 2.5 months. Recent versions have each incorporated around 13,500 commits, from around 1500 developers. These developers have a wide range of expertise, with some providing hundreds of patches per release, while others have contributed only one. Overall, the Linux kernel is critical software, but software in which the quality of the developed source code is highly variable. These features, combined with the fact that the Linux community is open to contributions and to the use of tools, make the Linux kernel an attractive target for software researchers. Tools that result from research can be directly integrated into the development of real software, where it can have a high, visible impact.

Starting from the work of Engler et al. [34], numerous research tools have been applied to the Linux kernel, typically for finding bugs [33], [51], [60], [68] or for computing software metrics [39], [75]. In our work, we have studied generic C bugs in Linux code [9], bugs in function protocol usage [45], [46], issues related to the processing of bug reports [63] and crash dumps [38], and the problem of backporting [15], illustrating the variety of issues that can be explored on this code base. Unique among research groups working in this area, we have furthermore developed numerous contacts in the Linux developer community. These contacts provide insights into the problems actually faced by developers and serve as a means of validating the practical relevance of our work. Section 6.1.2 presents our dissemination efforts to the Linux community.

4.2. Device Drivers

Device drivers are essential to modern computing, to provide applications with access, via the operating system, to physical devices such as keyboards, disks, networks, and cameras. Development of new computing paradigms, such as the internet of things, is hampered because device driver development is challenging and error-prone, requiring a high level of expertise in both the targeted OS and the specific device. Furthermore, implementing just one driver is often not sufficient; today's computing landscape is characterized by a number of OSes, *e.g.*, Linux, Windows, MacOS, BSD and many real time OSes, and each is found in a wide range of variants and versions. All of these factors make the development, porting, backporting, and maintenance of device drivers a critical problem for device manufacturers, industry that requires specific devices, and even for ordinary users.

The last fifteen years have seen a number of approaches directed towards easing device driver development. Réveillère, who was supervised by G. Muller, proposes Devil [7], a domain-specific language for describing the low-level interface of a device. Chipounov *et al.* propose RevNic, [28] a template-based approach for porting device drivers from one OS to another. Ryzhyk *et al.* propose Termite, [61], [62] an approach for synthesizing device driver code from a specification of an OS and a device. Currently, these approaches have been successfully applied to only a small number of toy drivers. Indeed, Kadav and Swift [40] observe that these approaches make assumptions that are not satisfied by many drivers; for example, the assumption that a driver involves little computation other than the direct interaction between the OS and the device. At the same time, a number of tools have been developed for finding bugs in driver code. These tools include SDV [19], Coverity [34], CP-Miner, [50] PR-Miner [51], and Coccinelle [8]. These approaches, however, focus on analyzing existing code, and do not provide guidelines on structuring drivers.

In summary, there is still a need for a methodology that first helps the developer understand the software architecture of drivers for commonly used operating systems, and then provides guidelines and tools for the maintenance and the development of new drivers. Section 3.2 describes this research direction.

ALPINES Project-Team

4. Application Domains

4.1. Compositional multiphase Darcy flow in heterogeneous porous media

We study the simulation of compositional multiphase flow in porous media with different types of applications, and we focus in particular on reservoir/bassin modeling, and geological CO₂ underground storage. All these simulations are linearized using Newton approach, and at each time step and each Newton step, a linear system needs to be solved, which is the most expensive part of the simulation. This application leads to some of the difficult problems to be solved by iterative methods. This is because the linear systems arising in multiphase porous media flow simulations cumulate many difficulties. These systems are non-symmetric, involve several unknowns of different nature per grid cell, display strong or very strong heterogeneities and anisotropies, and change during the simulation. Many researchers focus on these simulations, and many innovative techniques for solving linear systems have been introduced while studying these simulations, as for example the nested factorization [Appleyard and Cheshire, 1983, SPE Symposium on Reservoir Simulation].

4.2. Inverse problems

The research of F. Nataf on inverse problems is rather new since this activity was started from scratch in 2007. Since then, several papers were published in international journals and conference proceedings. All our numerical simulations were performed in FreeFem++.

We focus on methods related to time reversal techniques. Since the seminal paper by [M. Fink et al., Imaging through inhomogeneous media using time reversal mirrors. *Ultrasonic Imaging*, 13(2):199, 1991.], time reversal is a subject of very active research. The main idea is to take advantage of the reversibility of wave propagation phenomena such as it occurs in acoustics, elasticity or electromagnetism in a non-dissipative unknown medium to back-propagate signals to the sources that emitted them. Number of industrial applications have already been developed: touchscreen, medical imaging, non-destructive testing and underwater communications. The principle is to back-propagate signals to the sources that emitted them. The initial experiment, was to refocus, very precisely, a recorded signal after passing through a barrier consisting of randomly distributed metal rods. In [de Rosny and Fink. Overcoming the diffraction limit in wave physics using a time-reversal mirror and a novel acoustic sink. *Phys. Rev. Lett.*, 89 (12), 2002], the source that created the signal is time reversed in order to have a perfect time reversal experiment. Since then, numerous applications of this physical principle have been designed, see [Fink, Renversement du temps, ondes et innovation. Ed. Fayard, 2009] or for numerical experiments [Larmat et al., Time-reversal imaging of seismic sources and application to the great sumatra earthquake. *Geophys. Res. Lett.*, 33, 2006] and references therein.

4.3. Numerical methods for wave propagation in multi-scale media

We are interested in the development of fast numerical methods for the simulation of electromagnetic waves in multi-scale situations where the geometry of the medium of propagation may be described through characteristic lengths that are, in some places, much smaller than the average wavelength. In this context, we propose to develop numerical algorithms that rely on simplified models obtained by means of asymptotic analysis applied to the problem under consideration.

Here we focus on situations involving boundary layers and *localized* singular perturbation problems where wave propagation takes place in media whose geometry or material characteristics are submitted to a small scale perturbation localized around a point, or a surface, or a line, but not distributed over a volumic sub-region of the propagation medium. Although a huge literature is already available for the study of localized singular perturbations and boundary layer phenomena, very few works have proposed efficient numerical methods that rely on asymptotic modeling. This is due to their natural functional framework that naturally involves singular functions, which are difficult to handle numerically. The aim of this part of our research is to develop and analyze numerical methods for singular perturbation methods that are prone to high order numerical approximation, and robust with respect to the small parameter characterizing the singular perturbation.

4.4. Data analysis in astrophysics

We focus on computationally intensive numerical algorithms arising in the data analysis of current and forthcoming Cosmic Microwave Background (CMB) experiments in astrophysics. This application is studied in collaboration with researchers from University Paris Diderot, and the objective is to make available the algorithms to the astrophysics community, so that they can be used in large experiments.

In CMB data analysis, astrophysicists produce and analyze multi-frequency 2D images of the universe when it was 5% of its current age. The new generation of the CMB experiments observes the sky with thousands of detectors over many years, producing overwhelmingly large and complex data sets, which nearly double every year therefore following Moore's Law. Planck (<http://planck.esa.int/>) is a keystone satellite mission which has been developed under auspices of the European Space Agency (ESA). Planck has been surveying the sky since 2010, produces terabytes of data and requires 100 Petaflops per image analysis of the universe. It is predicted that future experiments will collect half petabyte of data, and will require 100 Exaflops per analysis as early as in 2020. This shows that data analysis in this area, as many other applications, will keep pushing the limit of available supercomputing power for the years to come.

AVALON Project-Team

4. Application Domains

4.1. Overview

The Avalon team targets applications with large computing and/or data storage needs, which are still difficult to program, maintain, and deploy. Those applications can be parallel and/or distributed applications, such as large scale simulation applications or code coupling applications. Applications can also be workflow-based as commonly found in distributed systems such as grids or clouds.

The team aims at not being restricted to a particular application field, thus avoiding any spotlight. The team targets different HPC and distributed application fields, which bring use cases with different issues. This will be eased by our various collaborations: the team participates to the INRIA-Illinois Joint Laboratory for Petascale Computing, the Physics, Radiobiology, Medical Imaging, and Simulation French laboratory of excellence, the E-Biothon project, the INRIA large scale initiative Computer and Computational Sciences at Exascale (C2S@Exa), and to BioSyL, a federative research structure about Systems Biology of the University of Lyon. Moreover, the team members have a long tradition of cooperation with application developers such as CERFACS and EDF R&D. Last but not least, the team has a privileged connection with CC IN2P3 that opens up collaborations, in particular in the astrophysics field.

In the following, some examples of representative applications we are targeting are presented. In addition to highlighting some application needs, they also constitute some of the use cases we will use to validate our theoretical results.

4.2. Climatology

The world's climate is currently changing due to the increase of the greenhouse gases in the atmosphere. Climate fluctuations are forecasted for the years to come. For a proper study of the incoming changes, numerical simulations are needed, using general circulation models of a climate system. Simulations can be of different types: HPC applications (*e.g.*, the NEMO framework [69] for ocean modelization), code-coupling applications (*e.g.*, the OASIS coupler [75] for global climate modeling), or workflows (long term global climate modeling).

As for most applications the team is targeting, the challenge is to thoroughly analyze climate-forecasting applications to model their needs in terms of programming model, execution model, energy consumption, data access pattern, and computing needs. Once a proper model of an application has been set up, appropriate scheduling heuristics could be designed, tested, and compared. The team has a long tradition of working with CERFACS on this topic, for example in the LEGO (2006-09) and SPADES (2009-12) French ANR projects.

4.3. Astrophysics

Astrophysics is a major field to produce large volume of data. For instance, the Large Synoptic Survey Telescope (<http://www.lsst.org/lsst/>) will produce 15 TB of data every night, with the goals of discovering thousands of exoplanets and of uncovering the nature of dark matter and dark energy in the universe. The Square Kilometer Array (<http://www.skatelescope.org/>) produces 9 Tbits/s of raw data. One of the scientific projects related to this instrument called Evolutionary Map of the Universe is working on more than 100 TB of images. The Euclid Imaging Consortium will generate 1 PB data per year.

Avalon collaborates with the *Institut de Physique Nucléaire de Lyon* (IPNL) laboratory on large scale numerical simulations in astronomy and astrophysics. Contributions of the Avalon members have been related to algorithmic skeletons to demonstrate large scale connectivity, the development of procedures for the generation of realistic mock catalogs, and the development of a web interface to launch large cosmological simulations on GRID'5000.

This collaboration, that continues around the topics addressed by the CLUES project (<http://www.clues-project.org>), has been extended thanks to the tight links with the CC-IN2P3. Major astrophysics projects execute part of their computing, and store part of their data on the resources provided by the CC-IN2P3. Among them, we can mention SNFactory, Euclid, or LSST. These applications constitute typical use cases for the research developed in the Avalon team: they are generally structured as workflows and a huge amount of data (from TB to PB) is involved.

4.4. Bioinformatics

Large-scale data management is certainly one of the most important applications of distributed systems in the future. Bioinformatics is a field producing such kinds of applications. For example, DNA sequencing applications make use of MapReduce skeletons.

The Avalon team is a member of BioSyL (<http://www.biosyl.org>), a Federative Research Structure attached to University of Lyon. It gathers about 50 local research teams working on systems biology. Moreover, the team cooperates with the French Institute of Biology and Chemistry of Proteins (IBCP <http://www.ibcp.fr>) in particular through the ANR MapReduce project where the team focuses on a bio-chemistry application dealing with protein structure analysis. These collaborations bring scientific applications that are both dynamic and data-intensive.

HIEPACS Project-Team

4. Application Domains

4.1. Material physics

Participants: Pierre Blanchard, Olivier Coulaud, Arnaud Etcheverry.

Due to the increase of available computer power, new applications in nano science and physics appear such as study of properties of new materials (photovoltaic materials, bio- and environmental sensors, ...), failure in materials, nano-indentation. Chemists, physicists now commonly perform simulations in these fields. These computations simulate systems up to billion of atoms in materials, for large time scales up to several nanoseconds. The larger the simulation, the smaller the computational cost of the potential driving the phenomena, resulting in low precision results. So, if we need to increase the precision, there are two ways to decrease the computational cost. In the first approach, we improve algorithms and their parallelization and in the second way, we will consider a multiscale approach.

A domain of interest is the material aging for the nuclear industry. The materials are exposed to complex conditions due to the combination of thermo-mechanical loading, the effects of irradiation and the harsh operating environment. This operating regime makes experimentation extremely difficult and we must rely on multi-physics and multi-scale modeling for our understanding of how these materials behave in service. This fundamental understanding helps not only to ensure the longevity of existing nuclear reactors, but also to guide the development of new materials for 4th generation reactor programs and dedicated fusion reactors. For the study of crystalline materials, an important tool is dislocation dynamics (DD) modeling. This multiscale simulation method predicts the plastic response of a material from the underlying physics of dislocation motion. DD serves as a crucial link between the scale of molecular dynamics and macroscopic methods based on finite elements; it can be used to accurately describe the interactions of a small handful of dislocations, or equally well to investigate the global behavior of a massive collection of interacting defects.

To explore i.e. to simulate these new areas, we need to develop and/or to improve significantly models, schemes and solvers used in the classical codes. In the project, we want to accelerate algorithms arising in those fields. We will focus on the following topics (in particular in the currently under definition **OPTIDIS** project in collaboration with CEA Saclay, CEA Ile-de-france and SIMaP Laboratory in Grenoble) in connection with research described at Sections 3.4 and 3.5 .

- The interaction between dislocations is long ranged ($O(1/r)$) and anisotropic, leading to severe computational challenges for large-scale simulations. In dislocation codes, the computation of interaction forces between dislocations is still the most CPU time consuming and has to be improved to obtain faster and more accurate simulations.
- In such simulations, the number of dislocations grows while the phenomenon occurs and these dislocations are not uniformly distributed in the domain. This means that strategies to dynamically construct a good load balancing are crucial to achieve high performance.
- From a physical and a simulation point of view, it will be interesting to couple a molecular dynamics model (atomistic model) with a dislocation one (mesoscale model). In such three-dimensional coupling, the main difficulties are firstly to find and characterize a dislocation in the atomistic region, secondly to understand how we can transmit with consistency the information between the two micro and meso scales.

4.2. Co-design for scalable numerical algorithms in scientific applications

Participants: Pierre Brenner, Jean-Marie Couteyen, Mathieu Faverge, Xavier Lacoste, Guillaume Latu, Salli Moustafa, Pierre Ramet, Fabien Rozar, Jean Roman.

The research activities concerning the ITER challenge are involved in the Inria Project Lab (IPL) **C2S@EXA**.

4.2.1. MHD instabilities edge localized modes

The numerical simulations tools designed for ITER challenges aim at making a significant progress in understanding active control methods of plasma edge MHD instabilities Edge Localized Modes (ELMs) which represent particular danger with respect to heat and particle loads for Plasma Facing Components (PFC) in ITER. Project is focused in particular on the numerical modeling study of such ELM control methods as Resonant Magnetic Perturbations (RMPs) and pellet ELM pacing both foreseen in ITER. The goals of the project are to improve understanding the related physics and propose possible new strategies to improve effectiveness of ELM control techniques. The tool for the nonlinear MHD modeling (code **JOREK**) will be largely developed within the present project to include corresponding new physical models in conjunction with new developments in mathematics and computer science strategy in order to progress in urgently needed solutions for ITER.

The fully implicit time evolution scheme in the **JOREK** code leads to large sparse linear systems that have to be solved at every time step. The MHD model leads to very badly conditioned matrices. In principle the **PaStiX** library can solve these large sparse problems using a direct method. However, for large 3D problems the CPU time for the direct solver becomes too large. Iterative solution methods require a preconditioner adapted to the problem. Many of the commonly used preconditioners have been tested but no satisfactory solution has been found. The research activities presented in Section 3.3 will contribute to design new solution techniques best suited for this context.

4.2.2. Turbulence of plasma particules inside a tokamak

In the context of the ITER challenge, the **GYSELA** project aims at simulating the turbulence of plasma particules inside a tokamak. Thanks to a better comprehension of this phenomenon, it would be possible to design a new kind of source of energy based of nuclear fusion. Currently, **GYSELA** is parallalized in a MPI/OpenMP way and can exploit the power of the current greatest supercomputers. To simulate faithfully the plasma physic, **GYSELA** handles a huge amount of data. In fact, the memory consumption is a bottleneck on very large simulations (449 K cores). In this context, mastering the memory consumption of the code becomes critical to consolidate its scalability and to enable the implementation of new numerical and physical features to fully benefit from the extreme scale architectures.

The scientific objectives of these research activities are first the design of advanced generic tools to manage and to better predict and limit the memory consumption peak in order to reduce the memory footprint of **GYSELA**, and second to design a set of tools that analyses the performance and the topology of the targeted architecture to optimize the deployment of Gysela runs. This will allow the design of new advanced numerical methods (for the gyroaverage operator, for the source and collision operators) and efficient scalable parallel algorithms in order to be able to deal with new physics in **GYSELA**. In particular the objective is to tackle kinetic electron configurations for more realistic simulations.

4.2.3. SN Cartesian solver for nuclear core simulation

As part of its activity, EDF R&D is developing a new nuclear core simulation code named COCAGNE that relies on a Simplified PN (SPN) method to compute the neutron flux inside the core for eigenvalue calculations. In order to assess the accuracy of SPN results, a 3D Cartesian model of PWR nuclear cores has been designed and a reference neutron flux inside this core has been computed with a Monte Carlo transport code from Oak Ridge National Lab. This kind of 3D whole core probabilistic evaluation of the flux is computationally very demanding. An efficient deterministic approach is therefore required to reduce the computation effort dedicated to reference simulations.

In this collaboration, we work on the parallelization (for shared and distributed memories) of the DOMINO code, a parallel 3D Cartesian SN solver specialized for PWR core reactivity computations which is fully integrated in the COCAGNE system.

4.2.4. 3D aerodynamics for unsteady problems with moving bodies

Aribus Defence and Space has developed for 20 years the FLUSEPA code which focuses on unsteady phenomenon with changing topology like stage separation or rocket launch. The code is based on a finite volume formulation with temporal adaptive time integration and supports bodies in relative motion. The temporal adaptive integration classifies cells in several temporal levels, zero being the level with the slowest cells and each level being twice as fast as the previous one. This repartition can evolve during the computation, leading to load-balancing issues in a parallel computation context. Bodies in relative motion are managed through a CHIMERA-like technique which allows building a composite mesh by merging multiple meshes. The meshes with the highest priorities recover the least ones, and at the boundaries of the covered mesh, an intersection is computed. Unlike classical CHIMERA technique, no interpolation is performed, allowing a conservative flow integration.

The main objective of this research is to design a new scalable version of FLUSEPA from a task-based parallelization over a runtime system in order to run efficiently on modern multicore parallel architectures very large 3D simulations (for example ARIANE 5 and 6 booster separation).

KERDATA Project-Team

4. Application Domains

4.1. Joint genetic and neuroimaging data analysis on Azure clouds

Joint acquisition of neuroimaging and genetic data on large cohorts of subjects is a new approach used to assess and understand the variability that exists between individuals. Both neuroimaging- and genetic-domain observations include a huge amount of variables (of the order of millions). Performing rigorous statistical analyses on such amounts of data is a major computational challenge that cannot be addressed with conventional computational techniques only. On the one hand, sophisticated regression techniques need to be used in order to perform significant analysis on these large datasets; on the other hand, the cost entailed by parameter optimization and statistical validation procedures (e.g. permutation tests) is very high.

To address the above challenges, the A-Brain (AzureBrain) Project was carried out within the Microsoft Research-Inria Joint Research Center. It was co-led by the KerData (Rennes) and Parietal (Saclay) Inria teams. They jointly address this computational problem using cloud related techniques on the Microsoft Azure cloud infrastructure. The two teams brought together their complementary expertise: KerData in the area of scalable cloud data management, and Parietal in the field of neuroimaging and genetics data analysis.

This application scenario is a typical multi-disciplinary Data Science project which serves as background for several on-going research activities, beyond the end of the A-Brain project.

4.2. Structural protein analysis on Nimbus and IBM clouds

In the framework of the MapReduce ANR project led by KerData (2010-2014), we have focused on the FastA bioinformatics application used for massive protein sequence similarity searching. This is a typical data-intensive application that can leverage the Map-Reduce model for a scalable execution on large-scale distributed platforms. FastA remains an interesting use case that we are considering beyond the end of the MapReduce project, for benchmarking our research results in the the area of optimized MapReduce processing.

4.3. I/O intensive climate simulations for the Blue Waters post-petascale machine

A major research topic in the context of HPC simulations running on post-petascale supercomputers is to explore how to record and visualize data during the simulation efficiently without impacting the performance of the computation generating that data. Conventional practice consists in storing data on disk, moving them off-site, reading them into a workflow, and analyzing them. This approach becomes increasingly harder to use because of the large data volumes generated at fast rates, in contrast to limited back-end performance. Scalable approaches to deal with these I/O limitations are thus of utmost importance. This is one of the main challenges explicitly stated in the roadmap of the [Blue Waters Project](#), which aims to build one of the most powerful supercomputers in the world.

In this context, the KerData project-team is exploring innovative ways to remove the limitations mentioned above through collaborative work in the framework of the Joint Inria-Illinois-ANL-BSC-JSC-RIKEN/AICS Laboratory for Extreme-Scale Computing (JLESC, formerly called JLPC), whose research activity focuses on the Blue Waters project. An example is the atmospheric simulation code CM1 (Cloud Model 1), one of the target applications of the Blue Waters machine. State-of-the-art I/O approaches, which typically consist in periodically writing a very large number of small files are inefficient: they cause bursts of I/O in the parallel file system, leading to poor performance and extreme variability (*jitter*). The challenge here is to investigate how to make an efficient use of the underlying file system, by avoiding synchronization and contention as much as possible. In collaboration with the JLESC, we are addressing these challenges through the Damaris approach.

MESCAL Project-Team

4. Application Domains

4.1. Cloud, Grid, Multi-core and Desktop Computing

Participants: Arnaud Legrand, Olivier Richard, Jean-Marc Vincent.

Software tools were developed to carry experiments on clouds and grids (Kameleon and Expo). Other tools (Pajé, Viva, Framesoc and Ocelotl) have been designed to monitor, trace and analyse applications running on multi-core and grid computers. Such traces have also been used in SIMGRID to simulate volunteer computing systems at unprecedented scale.

4.2. Wireless Networks

Participants: Bruno Gaujal, Panayotis Mertikopoulos.

MESCAL is involved in the common laboratory between Inria and Alcatel-Lucent. Bruno Gaujal is leading the Selfnets research action. This action was started in 2008 and was renewed for four more years (from 2012 to 2016). In our collaboration with Alcatel we use game theory techniques as well as evolutionary algorithms to compute optimal configurations in wireless networks (typically 3G or LTE networks) in a distributed manner. We have also been working on optimal spectrum management of MIMO systems, routing in ad-hoc works and power allocation in future 5G networks.

4.3. On-demand Geographical Maps

Participant: Jean-Marc Vincent.

This joint work involves the UMR 8504 Géographie-Cité, LIG, UMS RIATE and the Maisons de l'Homme et de la Société.

Improvements in the Web developments have opened new perspectives in interactive cartography. Nevertheless existing architectures have some problems to perform spatial analysis methods that require complex calculus over large data sets. Such a situation involves some limitations in the query capabilities and analysis methods proposed to users. The HyperCarte consortium with LIG, Géographie-cité and UMR RIATE proposes innovative solutions to these problems. Our approach deals with various areas such as spatio-temporal modeling, parallel computing and cartographic visualization that are related to spatial organizations of social phenomena.

4.4. Energy and Transportation

Participant: Nicolas Gast.

This work is mainly done within the Quanticol European project.

Smart urban transport systems and smart grids are two examples of collective adaptive systems. They consist of a large number of heterogeneous entities with decentralised control and varying degrees of complex autonomous behaviour. Within the QUANTICOL project, we develop an analysis tools to help to reason about such systems. Our work relies on tools from fluid and mean-field approximation to build decentralized algorithms that solve complex optimization problems. We focus on two problems: decentralized control of electric grids and capacity planning in vehicle-sharing systems to improve load balancing.

MOAIS Project-Team (section vide)

ROMA Project-Team

4. Application Domains

4.1. Applications of sparse direct solvers

Sparse direct (multifrontal) solvers have a wide range of applications as they are used at the heart of many numerical methods in computational science: whether a model uses finite elements or finite differences, or requires the optimization of a complex linear or nonlinear function, one often ends up solving a linear system of equations involving sparse matrices. There are therefore a number of application fields, among which some of the ones cited by the users of our sparse direct solver MUMPS (see Section 6.1) are: structural mechanics, biomechanics, medical image processing, tomography, geophysics, electromagnetism, fluid dynamics, econometric models, oil reservoir simulation, magneto-hydro-dynamics, chemistry, acoustics, glaciology, astrophysics, circuit simulation, and work on hybrid direct-iterative methods.

STORM Team

4. Application Domains

4.1. Supporting Numeric Libraries and Scientific Simulation Applications

The application of our work covers linear algebra, solvers, fast-multipole methods, in collaboration with other Inria teams and with industry. This allows the scientific development of new techniques adapted to these applications, and opens its use in a large variety of physic simulations in high performance computing.

In terms of direct application, the software developped in the team have allowed applications in various physics fields, ranging from sismic, mechanic of fluids, molecular dynamics, high energy physics or material simulations. Similarly, the domains of image processing and signal processing can take advantage of the expertise and software of the team.

TADAAM Team

4. Application Domains

4.1. Mesh-based applications

TADAAM targets scientific simulation applications on large-scale systems, as these applications present huge challenges in terms of performance, locality, scalability, parallelism and data management. Many of these HPC applications use meshes as the basic model for their computation. For instance, PDE-based simulations using finite differences, finite volumes, or finite elements methods operate on meshes that describe the geometry and the physical properties of the simulated objects. This is the case for at least two thirds of the applications selected in the 9th PRACE. call ⁰, which concern quantum mechanics, fluid mechanics, climate, material physic, electromagnetism, etc.

Mesh-based applications not only represent the majority of HPC applications running on existing supercomputing systems, yet also feature properties that should be taken into account to achieve scalability and performance on future large-scale systems. These properties are the following:

Size Datasets are large: some meshes comprise hundreds of millions of elements, or even billions.

Dynamicity In many simulations, meshes are refined or coarsened at each time step, so as to account for the evolution of the physical simulation (moving parts, shockwaves, structural changes in the model resulting from collisions between mesh parts, etc.).

Structure Many meshes are unstructured, and require advanced data structures so as to manage irregularity in data storage.

Topology Due to their rooting in the physical world, meshes exhibit interesting topological properties (low dimensionality embedding, small maximum degree, large diameter, etc.). It is very important to take advantage of these properties when laying out mesh data on systems where communication locality matters.

All these features make mesh-based applications a very interesting and challenging use-case for the research we want to carry out in this project. Moreover, we believe that our proposed approach and solutions will contribute to enhance these applications and allow them to achieve the best possible usage of the available resources of future high-end systems.

⁰<http://www.prace-ri.eu/prace-9th-regular-call/>

ASCOLA Project-Team

4. Application Domains

4.1. Enterprise Information Systems and Services

Large IT infrastructures typically evolve by adding new third-party or internally-developed components, but also frequently by integrating already existing information systems. Integration frequently requires the addition of glue code that mediates between different software components and infrastructures but may also consist in more invasive modifications to implementations, in particular to implement crosscutting functionalities. In more abstract terms, enterprise information systems are subject to structuring problems involving horizontal composition (composition of top-level functionalities) as well as vertical composition (reuse and sharing of implementations among several top-level functionalities). Moreover, information systems have to be more and more dynamic.

Service-Oriented Computing (SOC) that is frequently used for solving some of the integration problems discussed above. Indeed, service-oriented computing has two main advantages:

- Loose-coupling: services are autonomous: they do not require other services to be executed;
- Ease of integration: Services communicate over standard protocols.

Our current work is based on the following observation: similar to other compositional structuring mechanisms, SOAs are subject to the problem of crosscutting functionalities, that is, functionalities that are scattered and tangled over large parts of the architecture and the underlying implementation. Security functionalities, such as access control and monitoring for intrusion detection, are a prime example of such a functionality in that it is not possible to modularize security issues in a well-separated module. Aspect-Oriented Software Development is precisely an application-structuring method that addresses in a systemic way the problem of the lack of modularization facilities for crosscutting functionalities.

We are considering solutions to secure SOAs by providing an aspect-oriented structuring and programming model that allows security functionalities to be modularized. Two levels of research have been identified:

- Service level: as services can be composed to build processes, aspect weaving will deal with the orchestration and the choreography of services.
- Implementation level: as services are abstractly specified, aspect weaving will require to extend service interfaces in order to describe the effects of the executed services on the sensitive resources they control.

In 2015, we have published results on constructive mechanisms for security and accountability properties in service-based systems as well as results on service provisioning problems, in particular, service interoperability and mediation, see Sec. 6.3 . Furthermore, we take part in the European project A4Cloud on accountability challenges, that is, the responsible stewardship of third-party data and computations, see Sec. 8.3 .

4.2. Capacity Planning in Cluster, Grid and Cloud Computing

Cluster, Grid and more recently Cloud computing platforms aim at delivering large capacities of computing power. These capacities can be used to improve performance (for scientific applications) or availability (e.g., for Internet services hosted by datacenters). These distributed infrastructures consist of a group of coupled computers that work together and may be spread across a LAN (cluster), across a WAN (Grid), and across the Internet (Clouds). Due to their large scale, these architectures require permanent adaptation, from the application to the system level and call for automation of the corresponding adaptation processes. We focus on self-configuration and self-optimization functionalities across the whole software stack: from the lower levels (systems mechanisms such as distributed file systems for instance) to the higher ones (i.e. the applications themselves such as J2EE clustered servers or scientific grid applications).

In 2015, we have proposed VMPlaces, a dedicated framework to evaluate and compare VM placement algorithms. Globally the framework is composed of two major components: the injector and the VM placement algorithm. The injector constitutes the generic part of the framework (i.e. the one you can directly use) while the VM placement algorithm is the component a user wants to study (or compare with other existing algorithms), see Sec. 6.4 .

In the energy field, we have designed a set of techniques, named OptiPlace, for cloud management with flexible power models through constraint programming. OptiPlace supports external models, named views. Specifically, we have developed a power view, based on generic server models, to define and reduce the power consumption of a datacenter's physical servers. We have shown that OptiPlace behaves at least as good as our previous system, Entropy, requiring as low as half the time to find a solution for the constrained-based placement of tasks for large datacenters.

4.3. Pervasive Systems

Pervasive systems are another class of systems raising interesting challenges in terms of software structuring. Such systems are highly concurrent and distributed. Moreover, they assume a high-level of mobility and context-aware interactions between numerous and heterogeneous devices (laptops, PDAs, smartphones, cameras, electronic appliances...). Programming such systems requires proper support for handling various interfering concerns like software customization and evolution, security, privacy, context-awareness... Additionally, service composition occurs spontaneously at runtime.

Like Pervasive systems, Internet of thing is a major theme of these last ten years. Many research works has been led on the whole chain, from communicating sensors to big data management, through communication middlewares. Few of these works have addressed the problem of gathered data access.

The more a sensor networks senses various data, the more the users panel is heterogeneous. Such an heterogeneity leads to a major problem about data modeling: for each user, to aim at precisely addressing his needs and his needs only; ie to avoid a data representation which would overwhelm the user with all the data sensed from the network, regardless if he needs it or not. To leverage this issue, we propose in [24], [35] a multitree modeling for sensor networks which addresses each of these specific usages. With this modeling comes a domain specific language (DSL) which allows users to manipulate, parse and aggregate information from the sensors.

In 2014, we have extended the language EScala, which integrates reactive programming through events with aspect-oriented and object-oriented mechanisms, see Sec. 6.3 .

DIVERSE Project-Team

4. Application Domains

4.1. From Embedded Systems to Service Oriented Architectures

From small embedded systems such as home automation products or automotive systems to medium sized systems such as medical equipment, office equipment, household appliances, smart phones; up to large Service Oriented Architectures (SOA), building a new application from scratch is no longer possible. Such applications reside in (group of) machines that are expected to run continuously for years without unrecoverable errors. Special care has then to be taken to design and validate embedded software, making the appropriate trade-off between various extra-functional properties such as reliability, timeliness, safety and security but also development and production cost, including resource usage of processor, memory, bandwidth, power, etc.

Leveraging ongoing advances in hardware, embedded software is playing an evermore crucial role in our society, bound to increase even more when embedded systems get interconnected to deliver ubiquitous SOA. For this reason, embedded software has been growing in size and complexity at an exponential rate for the past 20 years, pleading for a component based approach to embedded software development. There is a real need for flexible solutions allowing to deal at the same time with a wide range of needs (product lines modeling and methodologies for managing them), while preserving quality and reducing the time to market (such as derivation and validation tools).

We believe that building flexible, reliable and efficient embedded software will be achieved by reducing the gap between executable programs, their models, and the platform on which they execute, and by developing new composition mechanisms as well as transformation techniques with a sound formal basis for mapping between the different levels.

Reliability is an essential requirement in a context where a huge number of softwares (and sometimes several versions of the same program) may coexist in a large system. On one hand, software should be able to evolve very fast, as new features or services are frequently added to existing ones, but on the other hand, the occurrence of a fault in a system can be very costly, and time consuming. While we think that formal methods may help solving this kind of problems, we develop approaches where they are kept “behind the scene” in a global process taking into account constraints and objectives coming from user requirements.

Software testing is another aspect of reliable development. Testing activities mostly consist in trying to exhibit cases where a system implementation does not conform to its specifications. Whatever the efforts spent for development, this phase is of real importance to raise the confidence level in the fact that a system behaves properly in a complex environment. We also put a particular emphasis on on-line approaches, in which test and observation are dynamically computed during execution.

FOCUS Project-Team

4. Application Domains

4.1. Ubiquitous Systems

The main application domain for Focus are ubiquitous systems, broadly systems whose distinctive features are: mobility, high dynamicity, heterogeneity, variable availability (the availability of services offered by the constituent parts of a system may fluctuate, and similarly the guarantees offered by single components may not be the same all the time), open-endedness, complexity (the systems are made by a large number of components, with sophisticated architectural structures). In Focus we are particularly interested in the following aspects.

- *Linguistic primitives* for programming dialogues among components.
- *Contracts* expressing the functionalities offered by components.
- *Adaptability and evolvability* of the behaviour of components.
- *Verification* of properties of component systems.
- Bounds on component *resource consumption* (e.g., time and space consumed).

4.2. Service Oriented Computing and Cloud Computing

Today the component-based methodology often refers to Service Oriented Computing. This is a specialized form of component-based approach. According to W3C, a service-oriented architecture is “a set of components which can be invoked, and whose interface descriptions can be published and discovered”. In the early days of Service Oriented Computing, the term services was strictly related to that of Web Services. Nowadays, it has a much broader meaning as exemplified by the XaaS (everything as a service) paradigm: based on modern virtualization technologies, Cloud computing offers the possibility to build sophisticated service systems on virtualized infrastructures accessible from everywhere and from any kind of computing device. Such infrastructures are usually examples of sophisticated service oriented architectures that, differently from traditional service systems, should also be capable to elastically adapt on demand to the user requests.

INDES Project-Team

4. Application Domains

4.1. Web programming

Along with games, multimedia applications, electronic commerce, and email, the web has popularized computers for daily life. The revolution is engaged and we may be at the dawn of a new era of computing where the web is a central element. The web constitutes an infrastructure more versatile, polymorphic, and open, in other words, more powerful, than any dedicated network previously invented. For this very reason, it is likely that most of the computer programs we will write in the future, for professional purposes as well as for our own needs, will extensively rely on the web. In addition to allowing reactive and graphically pleasing interfaces, web applications are de facto distributed. Implementing an application with a web interface makes it instantly open to the world and accessible from much more than one computer. The web also partially solves the problem of platform compatibility because it physically separates the rendering engine from the computation engine. Therefore, the client does not have to make assumptions on the server hardware configuration, and vice versa. Lastly, HTML is highly durable. While traditional graphical toolkits evolve continuously, making existing interfaces obsolete and breaking backward compatibility, modern web browsers that render on the edge web pages are still able to correctly display the web pages of the early 1990's. For these reasons, the web is arguably ready to escape the beaten track of n-tier applications, CGI scripting and interaction based on HTML forms. However, we think that it still lacks programming abstractions that minimize the overwhelming amount of technologies that need to be mastered when web programming is involved. Our experience on reactive and functional programming is used for bridging this gap.

4.2. Multimedia

Electronic equipments are less and less expensive and more and more widely spread out. Nowadays, in industrial countries, computers are almost as popular as TV sets. Today, almost everybody owns a mobile phone. Many are equipped with a GPS or a PDA. Modem, routers, NASes and other network appliances are also commonly used, although they are sometimes sealed under proprietary packaging such as the Livebox or the Freebox. Most of us evolve in an electronic environment which is rich but which is also populated with mostly isolated devices. The first multimedia applications on the web have appeared with the Web 2.0. The most famous ones are Flickr, YouTube, or Deezer. All these applications rely on the same principle: they allow roaming users to access the various multimedia resources available all over the Internet via their web browser. The convergence between our new electronic environment and the multimedia facilities offered by the web will allow engineers to create new applications. However, since these applications are complex to implement this will not happen until appropriate languages and tools are available. In the Indes team, we develop compilers, systems, and libraries that address this problem.

4.3. Robotics

The web is the de facto standard of communication for heterogeneous devices. The number of devices able to access the web is permanently increasing. Nowadays, even our mobile phones can access the web. Tomorrow it could even be the turn of our wristwatches! The web hence constitutes a compelling architecture for developing applications relying on the ambient computing facilities. However, since current programming languages do not allow us to develop easily these applications, ambient computing is currently based on ad-hoc solutions. Programming ambient computing via the web is still to be explored. The tools developed in the Indes team allow us to build prototypes of a robot as a web entity, and the use of remote web services to manage, monitor or extend the features of the robot. Among the direct benefits of relying on a web framework for robotics are the ability to use any web enabled device such as a smartphone or tablet to drive the robot.

PHOENIX Project-Team

4. Application Domains

4.1. Introduction

Building on our previous work, we are studying software development in the context of communication services, in their most general forms. That is, going beyond human-to-human interactions, and covering human-to-machine and machine-to-machine interactions. Software systems revolving around such forms of communications can be found in a number of areas, including telephony, pervasive computing, and assisted living; we view these software systems as coordinating the communication between networked entities, regardless of their nature: human, hardware or software. In this context, our three main application domains are pervasive computing, internet of things and assistive computing.

4.2. Pervasive Computing

Pervasive computing systems are being deployed in a rapidly increasing number of areas, including building automation and supply chain management. Regardless of their target area, pervasive computing systems have a typical architectural pattern. They aggregate data from a variety of distributed sources, whether sensing devices or software components, analyze a context to make decisions, and carry out decisions by invoking a range of actuators. Because pervasive computing systems are standing at the crossroads of several domains (e.g., distributed systems, multimedia, and embedded systems), they raise a number of challenges in software development:

Heterogeneity. Pervasive computing systems are made of off-the-shelf entities, that is, hardware and software building blocks. These entities run on specific platforms, feature various interaction models, and provide non-standard interfaces. This heterogeneity tends to percolate in the application code, preventing its portability and reusability, and cluttering it with low-level details.

Lack of structuring. Pervasive computing systems coordinate numerous, interrelated components. A lack of global structuring makes the development and evolution of such systems error-prone: component interactions may be invalid or missing.

Combination of technologies. Pervasive computing systems involve a variety of technological issues, including device intricacies, complex APIs of distributed systems technologies and middleware-specific features. Coping with this range of issues results in code bloated with special cases to glue technologies together.

Dynamicity. In a pervasive computing system, devices may either become available as they get deployed, or unavailable due to malfunction or network failure. Dealing with these issues explicitly in the implementation can quickly make the code cumbersome.

Testing. Pervasive computing systems are complicated to test. Doing so requires equipments to be acquired, tested, configured and deployed. Furthermore, some scenarios cannot be tested because of the nature of the situations involved (e.g., fire and smoke). As a result, the programmer must resort to writing specific code to achieve ad hoc testing.

4.3. Internet of Things

The Internet of Things (IoT) has become a reality with the emergence of Smart Cities, populated with large amounts of smart objects which are used to deliver a range of citizen services (e.g., security, well being, etc.) The IoT paradigm relies on the pervasive presence of smart objects or “things”, which raises a number of new challenges in the software engineering domain.

We introduce a *design-driven development approach* that is dedicated to the domain of orchestration of masses of sensors. The developer declares what an application does using a domain-specific language (DSL), named DiaSwarm. Our compiler processes domain-specific declarations to generate a customized programming framework that guides and supports the programming phase.

DiaSwarm addresses the main phases of an application orchestrating masses of sensors.

Service discovery. Standard service discovery at the individual object level does not address the needs of applications orchestrating large numbers of smart objects. Instead, a high-level approach which provides constructs to specifying subsets of interest is needed. Our approach allows developers to introduce application-specific concepts (e.g., regrouping parking spaces into lots or districts) at the design time and then these can be used to express discovery operations. Following our design-driven development approach, these concepts are used to generate code to support and guide the programming phase.

Data gathering. Applications need to acquire data from a large number of objects through a variety of delivery models. For instance, air pollution sensors across a city may only push data to the relevant applications when pollution levels exceed tolerated levels. Tracking sensors, however, might determine the location of vehicles and send the acquired measurements to applications periodically (e.g., 10 min. intervals). Data delivery models need to be introduced at design time since they have a direct impact on the application's program structure. In doing so, the delivery models used by an application can be checked against sensor features early in the development process.

Data processing. Data that is generated from hundreds of thousands of objects and accumulated over a period of time calls for efficient processing strategies to ensure the required performance is attained. Our approach allows for an efficient implementation of the data processing stage by providing the developer with a framework based on the MapReduce [34] programming model which is intended for the processing of large data sets.

4.4. Assistive Computing

Cognitive impairments (memory, attention, time and space orientation, *etc.*) affect a large part of the population, including older adults, patients with brain injuries (traumatic brain injury, stroke, etc), and people exhibiting cognitive incapacities, such as Down syndrome.

The emerging industry of assistive technologies provide hardware devices dedicated to specific tasks, such as a telephone set with a keyboard picturing relatives (<http://www.doro.fr>), or a device for audio and video communication over the web (<http://www.technosens.fr>). These assistive technologies apply a traditional approach to personal assistance by providing an equipment dedicated to a single task (or a limited set of tasks), without leveraging surrounding devices. This traditional approach has fundamental limitations that must be overcome to significantly improve assistive technologies:

- They are not adaptable to one's needs. They are generally dedicated to a task and have very limited functionalities: no networking, limited computing capabilities, a limited screen and rudimentary interaction modalities. This lack of functionality may cause a proliferation of devices, complicating the end-user life. Moreover, they are rarely designed to adapt to the cognitive changes of the user. When the requirements evolve, the person must acquire a new device.
- They are often proprietary, limiting innovation. As a result, they cannot cope with the evolution of users' needs.
- They have limited or no interoperability. As a result, they cannot rely on other devices and software services to offer richer applications.

To break this model, we propose to offer an assistive platform that is open-ended in terms of applications and entities. (1) An online catalog of available applications enables every user and caregiver to define personalized assistance in the form of an evolving and adapted set of applications; this catalog provides a community of developers with a mechanism to publish applications for specific daily-activity needs. (2) New types of entities (whether hardware or software) can be added to a platform description to enhance its functionalities and extend the scope of future applications.

RMOD Project-Team

4. Application Domains

4.1. Programming Languages and Tools

Many of the results of RMoD are improving programming languages or development tools for such languages. As such the application domain of these results is as varied as the use of programming languages in general. Pharo, the language that RMoD develops, is used for a very broad range of applications. From pure research experiments to real world industrial use (the Pharo Consortium has around 20 company members) <http://consortium.pharo.org> Examples are web applications, server backends for mobile applications or even graphical tools and embedded applications.

4.2. Software Reengineering

Moose is a language-independent environment for reverse- and re-engineering complex software systems. Moose provides a set of services including a common meta-model, metrics evaluation and visualization. As such Moose is used for analysing software systems to support understanding and continuous development as well as software quality analysis.

TACOMA Team

4. Application Domains

4.1. Pervasive applications in Smart Home

A smart home is a residence equipped with information-and-communication-technology (ICT) devices conceived to collaborate in order to anticipate and respond to the needs of the occupants, working to promote their comfort, convenience, security and entertainment while preserving their natural interaction with the environment.

The idea of using the Ubiquitous Computing paradigm in the smart home domain is not new. However, the state-of-the-art solutions only partially adhere to its principles. Often the adopted approach consists in a heavy deployment of sensor nodes, which continuously send a lot of data to a central elaboration unit, in charge of the difficult task of extrapolating meaningful information using complex techniques. This is a *logical approach*. TACOMA proposed instead the adoption of a *physical approach*, in which the information is spread in the environment, carried by the entities themselves, and the elaboration is directly executed by these entities "inside" the physical space. This allows performing meaningful exchanges of data that will thereafter need a less complicate processing compared to the current solutions. The result is a smart home that can, in an easier and better way, integrate the context in its functioning and thus seamlessly deliver more useful and effective user services. Our contribution aims at implementing the physical approach in a domestic environment, showing a solution for improving both comfort and energy savings.

4.2. Metamorphic House

The motivation for metamorphic houses is that many countries, including France, are going through socio-demographic evolutions, like growth of life expectancy and consequent increase in the number of elderly people, urbanization and resource scarcity. Households experience financial restrictions, while housing costs increase with the raise of real estate and energy prices [4].

Important questions arise concerning the future of housing policies and ways of living. We observe novel initiatives like participative housing and developing behaviors, including house-sharing, teleworking and longer stay of children in parents' homes.

To tackle the challenges raised by these emerging phenomena, future homes will have to be modular, upgradeable, comfortable, sparing of resources. They should be integrated in the urban context and exchange information with other homes, contribute to reducing the distances to be covered daily and respect the characteristics of the territory where they are located.

To reach these goals, metamorphic domestic environments will modify their shape and behavior to support activities and changes in life cycle of occupants, increase comfort and optimize the use of resources. Thanks to Information and Communication Technologies (ICT) and adaptive building elements, the same physical spaces will be transformed for different uses, giving inhabitants the illusion of living in bigger, more adapted and more comfortable places.

4.3. Pervasive applications in uncontrolled environnements

Some limitations of existing RFID technology become challenging: unlike standard RFID application scenarios, pervasive computing often involves uncontrolled environment for RFID, where tags and reader have to operate in much more difficult situations that those usually encountered or expected for classical RFID systems.

RFID technology is to avoid missing tags when reading multiple objects, as reading reliability is affected by various effects such shadowing or wave power absorption by some materials. The usual applications of RFID operate in a controlled environment in order to reduce the risk of missing tags while scanning objects.

In pervasive computing applications, a controlled reading environment is extremely difficult to achieve, as one of the principle is to enhance existing processes "in situ", unlike the controlled conditions that can be found in industrial processes. Consider for example a logistic application, where RFID tags could be used on items inside a package in order to check for its integrity along the shipping process. Tags would likely be placed randomly on items inside the package, and reading conditions would be variable depending on where the package is checked.

RFID operation in uncontrolled environments is challenging because RFID performance is affected by multiple parameters, in particular:

- Objects materials (on which tags are attached to),
- Materials in the surrounding environment,
- RFID frequency spectrum,
- Antenna nature and placement with respect to the tags.

In controlled environment, the difficulty to read tags can be limited by using the appropriate parameters to maximize the RFID performance for the application. But in many cases, it is needed to read large number of objects of various nature, arranged randomly in a given area or container. **Most pervasive computing applications fall in this context.**

COATI Project-Team

4. Application Domains

4.1. Telecommunication Networks

COATI is mostly interested in telecommunications networks. Within this domain, we consider applications that follow the needs and interests of our industrial partners, in particular Orange Labs or Alcatel-Lucent Bell-Labs, but also SME like 3-Roam.

We focus on the design and management of heterogeneous networks. The project has kept working on the design of backbone networks (optical networks, radio networks, IP networks). We also study routing algorithms such as dynamic and compact routing schemes, as we did in the context of the FP7 EULER led by Alcatel-Lucent Bell-Labs (Belgium), and the evolution of the routing in case of any kind of topological modifications (maintenance operations, failures, capacity variations, etc.).

4.2. Other Domains

Our combinatorial tools may be well applied to solve many other problems in various areas (transport, biology, resource allocation, chemistry, smart-grids, speleology, etc.) and we intend to collaborate with experts of these other domains.

For instance, we have recently started a collaboration in Structural Biology with EPI ABS (Algorithms Biology Structure) from Sophia Antipolis (described in Section 7.2). Furthermore, we are working on robot moving problems coming from Artificial Intelligence/Robotic in collaboration with Japan Advanced Institute of Science and Technology. In the area of transportation networks, we have started a collaboration with Amadeus on complex trip planning, and a collaboration with SME Instant-System on dynamic car-pooling combined with multi-modal transportation systems. Last, we have started a collaboration with GREDEG (Groupe de Recherche en Droit, Economie et Gestion, Univ. Nice Sophia Antipolis) on the analysis and the modeling of systemic risks in networks of financial institutions.

DANTE Project-Team

4. Application Domains

4.1. Life Science & Health

In parallel to the advances in modern medicine, health sciences and public health policy, epidemic models aided by computer simulations and information technologies offer an increasingly important tool for the understanding of transmission dynamics and of epidemic patterns. The increased computational power and use of Information and Communication Technologies make feasible sophisticated modeling approaches augmented by detailed in vivo data sets, and allow to study a variety of possible scenarios and control strategies, helping and supporting the decision process at the scientific, medical and public health level. The research conducted in the DANTE project finds direct applications in the domain of LSH since modeling approaches crucially depend on our ability to describe the interactions of individuals in the population. In the MOSAR/iBird project we are collaborating with the team of Pr. Didier Guillemot (Inserm/Institut. Pasteur/Université de Versailles). Within the TUBEXPO and ARIBO projects, we are collaborating with Pr. Jean-Christophe Lucet (Professeur des université Paris VII, Praticien hospitalier APHP).

4.2. Network Science / Complex networks

In the last ten years the science of complex networks has been assigned an increasingly relevant role in defining a conceptual framework for the analysis of complex systems. Network science is concerned with graphs that map entities and their interactions to nodes and links. For a long time, this mathematical abstraction has contributed to the understanding of real-world systems in physics, computer science, biology, chemistry, social sciences, and economics. Recently, however, enormous amounts of detailed data, electronically collected and meticulously catalogued, have finally become available for scientific analysis and study. This has led to the discovery that most networks describing real world systems show the presence of complex properties and heterogeneities, which cannot be neglected in their topological and dynamical description. This has called forth a major effort in developing the methodology to characterize the topology and temporal behavior of complex networks, to describe the observed structural and temporal heterogeneities, to detect and measure emerging community structure, to see how the functionality of networks determines their evolving structure, and to determine what kinds of correlations play a role in their dynamics. All these efforts have brought us to a point where the science of complex networks has become advanced enough to help us to disclose the deeper roles of complexity and gain understanding about the behavior of very complicated systems.

In this endeavor the DANTE project targets the study of dynamically evolving networks, concentrating on questions about the evolving structure and dynamical processes taking place on them. During the last year we developed several projects along these lines concerning three major datasets:

- Mobile telephony data: In projects with academic partners and Grandata we performed projects based on two large independent datasets collecting the telephone call and SMS event records for million of anonymized individuals. The datasets record the time and duration of mobile phone interactions and some coarse grained location and demographic data for some users. In addition one of the dataset is coupled with anonymised bank credit information allowing us to study directly the socioeconomic structure of a society and how it determines the communication dynamics and structure of individuals.
- Skype data: Together with Skype Labs/STACC and other academic groups we were leading projects in the subject of social spreading phenomena. These projects were based on observations taken from a temporally detailed description of the evolving social network of (anonymized) Skype users registered between 2003 and 2011. This data contains dates of registration and link creation together with gradual information about their location and service usage dynamics.

- Twitter data: In collaboration with ICAR-ENS Lyon we collected a large dataset about the microblogs and communications of millions of Twitter users in the French Twitter space. This data allows us to follow the spreading of fads/opinions/hashtags/ideas and more importantly linguistic features in online communities. The aim of this collaboration is to set the ground for a quantitative framework studying the evolution of linguistic features and dialects in an social-communication space mediated by online social interactions.

DIANA Project-Team (section vide)

DIONYSOS Project-Team (section vide)

DYOGENE Project-Team

4. Application Domains

4.1. Wireless Networks

Wireless networks can be efficiently modelled as dynamic stochastic geometric networks. Their analysis requires taking into account, in addition to their geometric structure, the specific nature of radio channels and their statistical properties which are often unknown a priori, as well as the interaction through interference of the various individual point-to-point links. Established results contribute in particular to the development of network dimensioning methods and some of them are currently used in Orange internal tools for network capacity calculations.

4.2. Embedded Networks

Critical real-time embedded systems (cars, aircrafts, spacecrafts) are nowadays made up of multiple computers communicating with each other. The real-time constraints typically associated with operating systems now extend to the networks of communication between sensors/actuators and computers, and between the computers themselves. Once a media is shared, the time between sending and receiving a message depends not only on technological constraints, but also, and mainly from the interactions between the different streams of data sharing the media. It is therefore necessary to have techniques to guarantee maximum network delays, in addition to local scheduling constraints, to ensure a correct global real-time behaviour to distributed applications/functions.

Moreover, pessimistic estimate may lead to an overdimensioning of the network, which involves extra weight and power consumption. In addition, these techniques must be scalable. In a modern aircraft, thousands of data streams share the network backbone. Therefore algorithm complexity should be at most polynomial.

4.3. Distributed Content Delivery Networks

A content distribution network (CDN) is a globally distributed network of proxy servers deployed in multiple data centers. The goal of a CDN is to serve content to end-users with high availability and high performance. CDNs serve a large fraction of the Internet content today, including web objects (text, graphics and scripts), downloadable objects (media files, software, documents), applications (e-commerce, portals), live streaming media, on-demand streaming media, and social networks. In [32], we address the problem of content replication in large distributed content delivery networks.

4.4. Probabilistic Algorithms for Renewable Integration in Smart grid

This reserach is developed by the Associate Team PARIS; <http://www.di.ens.fr/~busic/PARIS/>.

Challenges to Renewable Integration. With greater penetration of renewables, there is a need for tremendous shock absorbers to smooth the volatility of renewable power. An example is the balancing reserves obtained today from fossil-fuel generators, that ramp up and down their power output in response to a command signal from a grid balancing authority - an example of an ancillary service. In the absence of large, expensive batteries, we may have to increase our inventory of responsive fossil-fuel generators, negating the environmental benefits of renewable energy.

The goal of our research is to demonstrate that we do not need to rely entirely on expensive batteries or fast-responding fossil fuel generators to track regulation signals or balancing reserves. There is enormous flexibility in the power consumption of the majority of electric loads. This flexibility can be exploited to create “virtual batteries”. The best example of this is the heating, ventilation, and air conditioning (HVAC) system of a building: There is no perceptible change to the indoor climate if the airflow rate is increased by 10% for 20 minutes, and decreased by 10% for the next 20 minutes. Power consumption deviations follow the airflow deviations closely, but indoor temperature will be essentially constant.

A starting point in our research is the fact that many of the ancillary services needed today are defined by a power deviation reference signal that has zero mean. Examples are PJM's RegD signal, or BPA's balancing reserves⁰. We have demonstrated that loads can be classified based on the frequency bandwidth of ancillary service that they can offer. If demand response from loads respects these frequency limitations, it is possible to obtain highly reliable ancillary service to the grid, while maintaining strict bounds on the quality of service (QoS) delivered by each load [22].

Control Design with Local Intelligence at the Loads. An emphasis of our research is the creation of Smart Communities to complement a Smart Grid: intelligence is created at each load in the community. For example, a water heater may be equipped with a simple device that measures the grid frequency – a measure of power mismatch that is regulated to stabilize the power grid. Larger loads may receive a signal from a balancing authority.

A challenge in residential communities is that many loads are either on or off. How can an on/off load track the continuously varying regulation signal broadcast by a grid operator? The answer proposed in our recent work is based on probabilistic algorithms: A single load cannot track a regulation signal such as the balancing reserves. A collection of loads can, provided they are equipped with local control. The value of probabilistic algorithms is that a) they can be designed with minimal communication, b) they avoid synchronization of load responses, and c) it is shown in our recent work that they can be designed to simplify control at the grid level (see the survey [22] and [19], [29]). Other researchers have introduced randomization (see in particular the thesis of J. Mathieu [58]), but without the use of “local intelligence” (distributed control).

4.5. Algorithms for finding communities

In the study of complex networks, a network is said to have community structure if the nodes of the network can be easily grouped into (potentially overlapping) sets of nodes such that each set of nodes is densely connected internally. Community structures are quite common in real networks. Social networks include community groups (the origin of the term, in fact) based on common location, interests, occupation, etc. Metabolic networks have communities based on functional groupings. Citation networks form communities by research topic. Being able to identify these sub-structures within a network can provide insight into how network function and topology affect each other. We propose several algorithms for this problem and extensions [45], [35], [26], [36]

⁰BPA balancing authority. Online, <http://tinyurl.com/BPAgenload><http://tinyurl.com/BPAbalancing>.

EVA Team

4. Application Domains

4.1. Generalities

Wireless networks have become ubiquitous and are an integral part of our daily lives. These networks are present in many application domains; the most important are detailed in this section.

4.2. Industrial process automation

Networks in industrial process automation typically perform **monitoring and control** tasks. Wired industrial communication networks, such as HART⁰, have been around for decades and, being wired, are highly reliable. Network administrators tempted to “go wireless” expect the same reliability. Reliable process automation networks – especially when used for control – often impose stringent latency requirements. Deterministic wireless networks can be used in critical systems such as control loops, however, the unreliable nature of the wireless medium, coupled with their large scale and “ad-hoc” nature raise some of the most important challenges for low-power wireless research over the next 5-10 years.

Through the involvement of team members in standardization activities, the protocols and techniques will be proposed for the standardization process with a view to becoming the *de-facto* standard for wireless industrial process automation. Besides producing top level research publications and standardization activities, EVA intends this activity to foster further collaborations with industrial partners.

4.3. Environmental Monitoring

Today, outdoor WSNs are used to monitor vast rural or semi-rural areas and may be used to detect fires. Another example is detecting fires in outdoor fuel depots, where the delivery of alarm messages to a monitoring station in an upper-bounded time is of prime importance. Other applications consist in monitoring the snow melting process in mountains, tracking the quality of water in cities, registering the height of water in pipes to foresee flooding, etc. These applications lead to a vast number of technical issues: deployment strategies to ensure suitable coverage and good network connectivity, energy efficiency, reliability and latency, etc.

We will work on such applications in an associate team “REALMS” comprising members from EVA, the university of Berkeley and the university of Michigan.

4.4. The Internet of Things

The general agreement is that the Internet of Things (IoT) is composed of small, often battery-powered objects which measure and interact with the physical world, and encompasses smart home applications, wearables, smart city and smart plant applications.

The Internet of Things (IoT) has received continuous attention since 2013, and has been a marketing tool for industry giants such as IBM and Cisco, and the focal point of major events such the Consumer Electronics Show and the IETF. The danger of such exposure is that any under-performance may ultimately disappoint early adopters.

It is absolutely essential to (1) clearly understand the limits and capabilities of the IoT, and (2) develop technologies which enable user expectation to be met.

With the general public becoming increasingly familiar with the term “Internet of Things”, its definition is broadening to include all devices which can be interacted with from a network, and which do not fall under the generic term of “computer”.

⁰Highway Addressable Remote Transducer, <http://en.hartcomm.org/>.

The EVA team is dedicated to understanding and contributing to the IoT. In particular, the team will maintain a good understanding of the different technologies at play (Bluetooth, IEEE 802.15.4, WiFi, cellular), and their trade-offs. Through scientific publications and other contributions, EVA will help establish which technology best fits which application.

4.5. Military, Energy and Aerospace

Through the HIPERCOM project, EVA has developed cutting-edge expertise in using wireless networks for military, energy and aerospace applications. Wireless networks are a key enabling technology in the application domains, as they allow physical processes to be instrumented (e.g. the structural health of an airplane) at a granularity not achievable by its wired counterpart. Using wireless technology in these domains does however raise many technical challenges, including end-to-end latency, energy-efficiency, reliability and Quality of Service (QoS). Mobility is often an additional constraint in energy and military applications. Achieving scalability is of paramount importance for tactical military networks, and, albeit to a lesser degree, for power plants. EVA will work in this domain.

4.6. Smart Cities

It has been estimated that by 2030, 60% of the world's population will live in cities. On the one hand, smart cities aim at making everyday life more attractive and pleasant for citizens; on the other hand, they facilitate how those citizens can participate in the life of the city.

Smart cities share the constraint of mobility (both pedestrian and vehicular) with tactical military networks. Vehicular Ad-hoc NETWORKS (VANETs) will play an important role in the development of smarter cities.

The coexistence of different networks operating in the same radio spectrum can cause interference that should be avoided. Cognitive radio provides secondary users with the frequency channels that are temporarily unused (or unassigned) by primary users. Such opportunistic behavior can also be applied to urban wireless sensor networks. Smart cities raise the problem of transmitting, gathering, processing and storing big data. Another issue is to provide the right information at the place where it is most needed.

4.7. Emergency Applications

In an "emergency" application, heterogeneous nodes of a wireless network cooperate to recover from a disruptive event in a timely fashion, thereby possibly saving human lives. These wireless networks can be rapidly deployed and are useful to assess damage and take initial decisions. Their primary goal is to maintain connectivity with the humans or mobile robots (possibly in a hostile environment) in charge of network deployment. The deployment should ensure the coverage of particular points or areas of interest. The wireless network has to cope with pedestrian mobility and robot/vehicle mobility. The environment, initially unknown, is progressively discovered and may contain numerous obstacles that should be avoided. The nodes of the wireless network are usually battery-powered. Since they are placed by a robot or a human, their weight is very limited. The protocols supported by these nodes should be energy-efficient to maximize network lifetime. In such a challenging environment, sensor nodes should be replaced before their batteries are depleted. It is therefore important to be able to accurately determine the battery lifetime of these nodes, enabling predictive maintenance.

FUN Project-Team

4. Application Domains

4.1. Application Domains

The set of applications enabled through FUN and IoT is very large and can apply in every application area. We can thus not be exhaustive but among the most spread applications, we can name every area, event or animal monitoring, understanding and protection. To illustrate this, we may refer to the use cases addressed by our PREDNET project which goals is to equip rhinoceros with smart communicating devices to fight against poaching.

Other field of application is exploration of hostile and/or unknown environment by a fleet of self-organizing robots that cooperate with RFID and sensors to ensure a continue monitoring afterwards.

Also, IoT and FUN ca play a key role in logistics and traceability by relying on the use of sensors or RFID technologies as implemented in our TRACAVERRRE project or our collaboration with the start up TRAXENS.

Finally, IoT and FUN leverage a lot of applications in Smart City concept , ranging from parking aid to a better energy consumption going through air quality monitoring, traffic fluidizing etc. (See our CityLab Inria and VITAL projects).

GANG Project-Team

4. Application Domains

4.1. Application Domains

Application domains include evaluating Internet performances, the design of new peer-to-peer applications, enabling large scale ad hoc networks and mapping the web.

- The application of measuring and modeling Internet metrics such as latencies and bandwidth is to provide tools for optimizing Internet applications. This concerns especially large scale applications such as web site mirroring and peer-to-peer applications.
- Peer-to-peer protocols are based on a all equal paradigm that allows to design highly reliable and scalable applications. Besides the file sharing application, peer-to-peer solutions could take over in web content dissemination resistant to high demand bursts or in mobility management. Envisioned peer-to-peer applications include video on demand, streaming, exchange of classified ads,...
- Wifi networks have entered our every day life. However, enabling them at large scale is still a challenge. Algorithmic breakthrough in large ad hoc networks would allow to use them in fast and economic deployment of new radio communication systems.
- The main application of the web graph structure consists in ranking pages. Enabling site level indexing and ranking is a possible application of such studies.

INFINE Team (section vide)

MADYNES Project-Team

4. Application Domains

4.1. Mobile, ad-hoc and constrained networks

The results coming out from MADYNES can be applied to any dynamic infrastructure that contributes to the delivery of value added services. While this is a potentially huge application domain, we focus on the following environments at the network level:

1. multicast services,
2. ad-hoc networks,
3. mobile devices and IPv6 networks,
4. voice over IP infrastructure.

All these selected application areas exhibit different dynamicity features. In the context of multicast services, we focus on distribution, monitoring and accounting of key distribution protocols. On *ad-hoc* and dynamic networks we are investigating the provisioning, monitoring, configuration and performance management issues.

Concerning mobile devices, we are interested in their configuration, provisioning and monitoring. IPv6 work goes on in Information Models and on self-configuration of the agents.

4.2. Dynamic services infrastructures

At the service level, dynamics is also increasing very fast. We apply the results of our work on autonomous management on infrastructures which support dynamic composition and for which self-instrumentation and management automation is required.

The target service environments are:

- sensor networks,
- peer-to-peer infrastructures,
- information centric networks,
- ambient environments.

MAESTRO Project-Team

4. Application Domains

4.1. Main Application Domains

MAESTRO's main application area is networking, to which we apply modeling, performance evaluation, optimization and control. Our primary focus is on protocols and network architectures, and recent evolutions include the study of the Web and social networks, as well as models for Green IT.

- Wireless (cellular, ad hoc, sensor) networks: WLAN, WiMAX, UMTS, LTE, HSPA, delay tolerant networks (DTN), power control, medium access control, transmission rate control, redundancy in source coding, mobility models, coverage, routing, green base stations,
- Internet applications: social networks, content distribution systems, peer-to-peer systems, overlay networks, multimedia traffic, video-on-demand, multicast;
- Information-Centric Networking (ICN) architectures: Content-Centric Network (CCN, also called Content-Oriented Networks);
- Internet infrastructure: TCP, high speed congestion control, voice over IP, service differentiation, quality of service, web caches, proxy caches.

MUSE Team

4. Application Domains

4.1. Home Network Diagnosis

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

Our long-term goal is to assist users with concrete indicators of the causes of potential problems and—ideally—ways to fix them. We intend to develop a set of easy-to-use home network diagnosis tools that can reliably identify performance and functionality shortcomings rooted in the home. The development of home network diagnosis tools brings a number of challenges. First, home networks are heterogenous. The set of devices, configurations, and applications in home networks vary significantly from one home to another. We must develop sophisticated techniques that can learn and adapt to any home network as well as to the level of expertise of the user. Second, there are numerous ways in which applications can fail or experience poor performance in home networks. Often there are a number of explanations for a given symptom. We must devise techniques that can identify the most likely cause(s) for a given problem from a set of possible causes. Third, even if we can identify the cause of the problem, we must then be able to identify a solution. It is important that the output of the diagnosis tools we build is “actionable”. Users should understand the output and know what to do.

We are conceiving methods for two application scenarios: (i) when the end user in the home deploys our diagnostic tools either on the home gateway (the gateway often combines a DSL/cable modem and an access point; it connects the home network to the ISP) or on devices connected to the home network and (ii) when ISPs collect measurements from homes of subscribers and then correlate these measurements to help identify problems.

Assisting end users. We are developing algorithms to determine whether network performance problems lie inside or outside the home network. Given that the home gateway connects the home with the rest of the Internet, we are designing an algorithm (called *HoA*) that analyzes traffic that traverses the gateway to distinguish access link and home network bottlenecks. A measurement vantage point on the gateway is key for determining if the performance bottleneck lies within the home network or the access ISP, but we also need to deploy diagnosis tools in end-devices. First, some users may not want (or not know how) to deploy a new home gateway in their homes. Second, some problems will be hard to diagnose with only the vantage point of the gateway (for example, when a device cannot send traffic or when the wireless is poor in certain locations of a home). We can obtain more complete visibility by leveraging *multiple* measurement nodes around the home, potentially including the home gateway, all participating jointly in the measurement task. We have an ongoing project to realize a home network analyzer as a web-based measurement application built on top of our team’s recently developed browser-based measurement platform, *Fathom*. To integrate the home gateway in the analyzer, we plan to engage the BISmark Project. BISmark already provides a web server as well as extensive configurability, allowing us to experiment freely with both passive as well as active measurements. We must develop a home network analyzer that can first discover the set of devices connected to the home network that can collaborate on the diagnosis task. We will then develop tomography algorithms to infer where performance problems lie given measurements taken from the set of available vantage points.

Assisting Internet Service Providers (ISPs). Our discussions with several large access ISPs reveal that service calls are costly, ranging from \$9–25 per call, and as many as 75% of service calls from customers are usually caused by problems that have nothing to do with the ISP. Therefore, ISPs are eager to deploy techniques to assist in home network diagnosis. In many countries ISPs control the home gateway and set-top-boxes in the home. We plan to develop more efficient mechanisms for home users to report trouble to their home ISP and consequently reduce the cost of service calls. This project is in collaboration with Technicolor and Portugal Telecom. Technicolor is a large manufacturer of home gateways and set-top-boxes. Portugal Telecom is the largest broadband access provider in Portugal. Technicolor already collects data from 200 homes in Portugal. We are working with the data collected in this deployment together with controlled experiments to develop methods to diagnose problems in the home wireless.

4.2. Quality of Experience

Understanding how users react to different levels of network performance presents two main challenges:

1. User perception is subjective and contextual. Different users may have different tolerance levels to network performance and the same user may have different expectations under different circumstances. Take for example the round-trip time (RTT), a typical network performance metric. If RTTs are larger than usual, a user who is doing remote login may feel that the connection is unusable, whereas another who is watching YouTube may notice no problem (because YouTube has a playout buffer to mask some network delay). Take another example of a user downloading her email. This user may tolerate some delay when she is leisurely checking her email at home, but she may become extremely frustrated with the same delay if she is in an airplane and needs to download her email just before takeoff.
2. It is challenging to “measure” users. We must develop methods to measure the user perception of network performance as users perform their routine online tasks. It is hence important that these methods are not too intrusive. Otherwise, users are unlikely to participate in the experiment. In addition, we must capture user perception at different levels of performance and in a variety of scenarios.

We will develop tools that run on end systems to collect network performance data annotated with the user perception. These tools will adopt a hybrid measurement methodology that combines network measurement techniques to infer application performance with techniques from HCI to measure user perception. We will later use the resulting datasets to build models of user perception of network performance based only on data that we can obtain automatically from the user device or from user’s traffic observed in the network. Models of user perception of network performance can be used to detect when performance is poor to trigger diagnosis or to adapt network/application performance to better serve users.

4.3. Crowd-sourced content recommendation

The Internet today serves as a large content distribution platform (online content varies from traditional news, TV series, and movies to specialized blogs and family pictures shared over social networks) as well as a platform for users to exchange opinions about practically everything (from movies to services and restaurants). The amount of information available online today overwhelms most users and selecting which content to watch or what do has become a challenge. We are applying passive measurement methods and content summarisation techniques to help users to identify relevant content in two scenarios. First, we are developing a system called WeBrowse that passively observes network traffic to extract user clicks (i.e., the URLs users visit). A user click is a good measure of interest, as users often have an idea of the type of content they are about to access (e.g., because they saw a preview or because a friend recommended it). Intuitively, the more users click on a URL, the higher the interest in the content on the corresponding page. WeBrowse then promotes “hottest” and most popular content to users of a network. We have a deployment of WeBrowse in a campus network. Second, we are working on techniques to summarise user feedback (for example, movie or restaurant reviews) with semi-structured feedback. Today reviews are either free-form text or star rating. Star rating is too coarse to capture the nuances of why a user likes or dislikes something, whereas free text is hard for users to parse and

extract a clear opinion. We are instead working with semi-structured reviewing where users enter *tags* (a short sequence of words describing the user experience). We are working with Technicolor on the summarisation of movie reviews and on building a mobile app (called TagIt) where users can review movies directly with tags.

RAP Project-Team (section vide)

SOCRATE Project-Team (section vide)

URBANET Team

4. Application Domains

4.1. Smart urban infrastructure

Unlike the communication infrastructure that went through a continuous development in the last decades, the distribution networks in our cities including water, gas and electricity are still based on 19th century infrastructure. With the introduction of new methods for producing renewable but unpredictable energy and with the increased attention towards environmental problems, modernizing distribution networks became one of the major concerns in the urban world. An essential component of these enhanced systems is their integration with information and communications technology, the result being a smart distribution infrastructure, with improved efficiency and reliability. This evolution is mainly based on the increased deployment of automatic equipment and the use of machine-to-machine and sensor-to-actuator communications that would allow taking into account the behavior and necessities of both consumers and suppliers

Another fundamental urban infrastructure is the transportation system. The progress made in the transportation industry over the last century has been an essential factor in the development of today's urban society, while also triggering the birth and growth of other economic branches. However, the current transportation system has serious difficulties coping with the continuous growth in the number of vehicles, especially in an urban environment. As a major increase in the capacity of a city road infrastructure, already in place for tens or even hundreds of years, would imply dissuasive costs, the more realistic approach is to optimize the use of the existing transportation system. As in the case of distribution networks, the intelligence of the system can be achieved through the integration of information and communication capabilities. However, for smart transportation the challenges are somehow different, because the intelligence is no longer limited to the infrastructure, but propagates to vehicles themselves. Moreover, the degree of automation is reduced in transportation systems, as most actions resulting in reduced road congestion, higher reliability or improved safety must come from the human driver (at least in the foreseeable future)

Finally, smart spaces are becoming an essential component of our cities. The classical architecture tools used to design and shape the urban environment are more and more challenged by the idea of automatically modifying private and public spaces in order to adapt to the requirements and preferences of their users. Among the objectives of this new urban planning current, we can find the transformation of the home in a proactive health care center, fast reconfigurable and customizable workplaces, or the addition of digital content in the public spaces in order to reshape the urban scene. Bringing these changing places in our daily lives is conditioned by a major shift in the construction industry, but it also involves important advancements in digital infrastructure, sensing, and communications

4.2. Urban participatory sensing

Urban sensing can be seen as the same evolution of the environment digitalization as social networking has been for information flows. Indeed, besides dedicated and deployed sensors and actuators, still required for specific sensing operations such as the real-time monitoring of pollution levels, there is a wide range of relevant urban data that can be collected without the need for new communication infrastructures, leveraging instead on the pervasiveness of smart mobile terminals. With more than 80% of the population owning a mobile phone, the mobile market has a deeper penetration than electricity or safe drinking water. Originally designed for voice transmitted over cellular networks, mobile phones are today complete computing, communication and sensing devices, offering in a handheld device multiple sensors and communication technologies.

Mobile devices such as smartphones or tablets are indeed able to gather a wealth of informations through embedded cameras, GPS receivers, accelerometers, and cellular, WiFi and bluetooth radio interfaces. When collected by a single device, such data may have small value per-se, however its fusion over large scales could prove critical for urban sensing to become an economically viable mainstream paradigm.

This is even more true when less traditional mobile terminals are taken into account: privately-owned cars, public transport means, commercial fleets, and even city bikes are starting to feature communication capabilities and the Floating Car Data (FCD) they generate can bring a dramatic contribution to the cause of urban sensing. Indeed, other than enlarging the sensing scope even further, e.g., through Electronic Control Units (ECUs), these mobile terminals are not burdened by strong energy constraints and can thus significantly increase the granularity of data collection. This data can be used by authorities to improve public services, or by citizens who can integrate it in their choices. However, in order to kindle this hidden information, important problems related to data gathering, aggregation, communication, data mining, or even energy efficiency need to be solved.

4.3. Human-centric networks

Combining location awareness and data recovered from multiple sources like social networks or sensing devices can surface previously unknown characteristics of the urban environment, and enable important new services. As a few examples, one could think of informing citizens about often disobeyed (and thus risky) traffic signs, polluted neighborhoods, or queue waiting times at current exhibitions in the urban area.

Beyond letting their own devices or vehicles autonomously harvest data from the environment through embedded or onboard sensors, mobile users can actively take part in the participatory sensing process because they can, in return, benefit from citizen-centric services which aim at improving their experience of the urban life. Crowdsourcing applications have the potential to turn citizens into both sources of information and interactive actors of the city. It is not a surprise that emerging services built on live mobile user feedback are rapidly meeting a large success. In particular, improving everyone's mobility is probably one of the main services that a smart city shall offer to its inhabitants and visitors. This implies providing, through network broadcast data or urban smart-furniture, an accurate and user-tailored information on where people should head in order to find what they are looking for (from a specific kind of shop to a free parking slot), on their current travel time estimates, on the availability of better alternate means of transport to destination. Depending on the context, such information may need to be provided under hard real-time constraints, e.g., in presence of road accidents, unauthorized public manifestations, or delayed public transport schedules.

In some cases, information can also be provided to mobile users so as to bias or even enforce their mobility: drivers can be alerted of the arrival of an emergency vehicle so that they leave the leftmost lane available, or participants leaving vast public events can be directed out of the event venue through diverse routes displayed on their smartphones so as to dynamically balance the pedestrian flows and reduce their waiting times.