RESEARCH CENTER

FIELD
Networks, Systems and Services, Distributed Computing

Activity Report 2012

Section Application Domains

Edition: 2013-04-24

Distributed Systems and Services

Distributed and High Performance Computing

Networks and Telecommunications

ACES Project-Team (section vide)

ADAM Project-Team

4. Application Domains

4.1. Introduction

The ADAM project-team targets the software engineering of adaptive service-oriented applications and middleware. The application domain covered by ADAM is broad and spans from distribution applications to middleware. In all these cases, adaptability is the property which is sought: applications and middleware must be adaptable to new execution contexts, they must react to changes in the environment and they must be able to discover and integrate new services.

The ADAM project-team produces software and middleware building blocks. This explains why the application domain is broad, yet targeting applications where adaptability is the key requirement. This includes electronic commerce, embedded systems, health care information systems, and terrestrial transport information systems. These domains are in direct relation with our currently funded activities. They act as testbeds for the solutions that we propose in terms of middleware services, middleware platforms, runtime kernels, component libraries, languages design or domain modeling.

4.2. Electronic Commerce

Applications in the domain of electronic commerce are by essence distributed. They involve many different participants with heterogeneous information systems which cannot be changed. The challenge is then to provide an adaptation layer to be able to compose and let these systems interoperate. In the context of the ANR TLog SCOrWare, the ICT SOA4All and the FUI CAPPUCINO projects, our activities in this domain will aim at supporting service-oriented architectures. We want to have adaptive architectures that can be composed and orchestrated seamlessly. In this domain, the business relationship with customers is vital and many different usage scenarios must be supported. Customers are roaming, and the services must be kept operational across different devices. This puts some constraints on the server tier where technical services must be adapted to manage, for instance, long lasting transactions. The application server infrastructure must then provide a support for adapting technical services.

4.3. Embedded Systems

Embedded systems form a domain where adaptation is a key requirement. The design and the implementation of modern embedded software uses advanced software engineering techniques such model-driven development or software component frameworks. In this domain, we are involved in several projects, such as the ANR TLog Flex-eWare, and the FUI MIND projects. Several challenges must be addressed here. For example, when a model-driven developed application is adapted, designers have to ensure that the models and the operational level are kept synchronized. The co-evolution of these two levels is one of the challenges that we are addressing. A second challenge is related to software components that need to be customized in order to fit the requirements imposed by constrained environments. It is, for example, a matter of providing component frameworks that can accommodate various granularities of services.

4.4. Health Care Information Systems

Health care information systems form a third application domain in which the ADAM project-team is involved, for instance through demonstrators which will be implemented in the context of the ANR TLog FAROS project. The challenge here is to provide a distributed infrastructure where information will be available to medical staff wherever they are. This imposes to be able to provide this information on many different devices (from high resolution screens to embedded devices on the scene of an accident), while ensuring the privacy of the medical data of a patient (several level of data access must be granted depending on the categories of medical staff). Given the vital role of such an information system, we want to provide guarantees that the services will be highly available and trustworthy. We envision to provide a service-oriented architecture which will be extended to support software contracts and multi-scale environments.

4.5. Information Systems for Terrestrial Transport

Information systems for terrestrial transport are also a domain that we are relying on, to apply our research activities in accordance with the ANR ARA REVE project and the INRETS collaboration. Applications are here characterized by frequent disconnections, poor quality network links, and high mobility. We want to provide an infrastructure where the technical services, and among others the communication services, can be adapted to support new requirements. One of the paths that we propose to investigate is to include such a scenario in the general context of the adaptiveness of component frameworks.

ARLES Project-Team

4. Application Domains

4.1. Application Domains

The ARLES project-team is interested in the application of pervasive computing, and as such considers various application domains. Indeed, our application domain is voluntarily broad since we aim at offering generic solutions. However, we examine exploitation of our results for specific applications, as part of the experiments that we undertake to validate our research results through prototype implementation. Applications that we consider in particular include demonstrators developed in the context of the European and National projects to which we contribute (§ 7.1 & 7.2).

ASAP Project-Team

4. Application Domains

4.1. Overview

The results of the research targeted in ASAP span a wide range of applications. Below are a few examples.

- Personalized Web Search.
- Recommendation.
- Social Networks.
- Notification Systems.
- Distributed Storage.
- Video Streaming.

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, in that 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 2012, we have developed techniques for the Service-Level Agreement (SLA) management for Cloud elasticity, see Sec. 6.3 , as well as models and type systems for service-oriented systems, see Sec. 6.1 . Furthermore, we take part in a starting new European project A4Cloud on accountability challenges, that is, the responsible stewardship of third-party data and computations, see Sec. 8.2 .

4.2. Cluster, Grid and Cloud Computing

Cluster, Grid and more recently Cloud computing platforms aim at delivering a larger capacity of computing power compared to a single computer configuration. This capacity can be used to improve performance (for scientific applications) or availability (e.g., for Internet services hosted by a data center). These distributed infrastructures consist of a group of coupled computers that work together. This group can 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 calls for automation of the adaptation process. 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 2012, we have developed the DVMS system, which contains one of the most highly scalable scheduling algorithm for virtual machines; we have also generated several results on the energy efficient management of Cloud applications and infrastructures, see Sec. 6.3 .

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.

In 2012, we have developed the language EScala, which integrates reactive programming through events with aspect-oriented and object-oriented mechanisms, see Sec. 6.1 .

ATLANMOD Team

4. Application Domains

4.1. Introduction

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 [50] and to system engineering [41]. There are a lot of problems in these application domains that may be addressed by means of modeling and model transformation techniques. For the core techniques, we are now focusing on solving the scalability problem to make sure they can be successfully adopted by our industrial partners in the context of large and complex software systems.

Nevertheless, the team has indeed selected a set of vertical and horizontal domains that AtlanMod finds specially interesting. In what follows we describe three of them.

4.2. Reverse Engineering

One important and original domain that is being investigated by the AtlanMod team is the reverse engineering of existing IT systems. We do believe that efficiently dealing with such legacy systems is one of the main challenges in Software Engineering and related industry today. Having a better understanding of these systems in order to document, maintain, improve or migrate them is thus a key requirement for both academic and industrial actors in this area. However, it is not an easy task and it still raises interesting challenging issues to be explored [49].

We have shown how reverse engineering practices may be advantageously revisited with the help of the MDE approach and techniques, applying (as base principle) the systematic representation as models of the required information discovered from the legacy software artifacts (e.g.; source code, configuration files, documentation, metadata, etc). The rise in abstraction allowed by MDE can bring new hopes that reverse engineering is now able to move beyond more traditional ad-hoc practices. For instance, an ongoing industrial PhD in partnership with IBM France aims to investigate the possibilities of conceptualizing a generic framework enabling the extraction of business rules from a legacy application, as much as possible, independently of the language used to code it. Moreover, different pragmatic solutions for improving the overall scalability when dealing with large-scale legacy systems (handling huge data volumes) are intensively studied by the team.

In this context, AtlandMod has set up and is developing accordingly the open source Eclipse-MDT MoDisco project (see 5.2). MoDisco is notably being referenced by the OMG ADM (Architecture Driven Modernization) normalization task force as the reference implementation for several of its standard metamodels.

We have also opened a novel research line focused on integration of APIs into MDE. In the application of reverse engineering processes while modernizing software system it is very common to face the need of integrating Application Programming Intefaces (APIs). Indeed, building any application usually involves managing a plethora of APIs to access different software assets such as: basic infrastructures (e.g., operating system, databases, or middleware), general-purpose or domain specific libraries, frameworks, software components, web services, and even other applications. Thus, to promote the interoperability between the API and model technical spaces, we have developed API2MoL, which is an approach aimed at automating the implementation of API-MDE bridges. This new language allows defining mappings between the artefacts of a given API (e.g., API classes in object-oriented APIs) and the elements of a metamodel that represents this API in the MDE technical space. A mapping definition provides the information which is necessary to build a bridge for a concrete API specification and metamodel. Thanks to the API-MDE bridges automatically created, developers are liberated from having to manually implement the tasks of obtaining models from API objects and generating such objects from models. Therefore, API2MoL may improve the productivity and quality of the part of the MDE application that deals with the APIs.

4.3. Security Engineering

Several components are required to build up a system security architecture, such as firewalls, database user access control, intrusion detection systems, and VPN (Virtual Private Network) routers. These components must be properly configured to provide an appropriate degree of security to the system. The configuration process is highly complex and error-prone. In most organizations, security components are either manually configured based on security administrators expertise and flair; or simply recycled from existing configurations already deployed in other systems (even if they may not be appropriated for the current one). These practices put at risk the security of the whole organization.

We have started a Phd thesis in this domain intended to investigate the construction of a model-driven automatic reverse engineering mechanism (implemented as an extension of the MoDisco project) capable of analyzing deployed security aspects of components (e.g., concrete firewall configurations) to derive the abstract model (e.g., network security global policy) that is actually enforced over the system. Once the model is obtained, it can be reconciled with the expected security directives, to check its compliance, can be queried to test consistency or used in a process of forward engineering to generate validated security configurations.

As a first step we intend to apply model-driven techniques for the extraction of high level model representations of security policies enforced by firewalls. Firewalls, core components in network security systems, are generally configured by using very low level vendor specific rule-based languages, difficult to understand and to maintain. As a consequence, as the configuration files grow, understanding which security policy is being actually enforced or checking if inconsistencies has been introduced becomes a very complex and time consuming task. We propose to raise the level of abstraction so that the user can deal directly with the high level policies. Once a model representation of the enforced policy is available, model-driven techniques will ease some of the tasks we need to perform, like consistency checking, validation, querying and visualization. Easy migration between different firewall vendors will be also enabled.

4.4. Software Quality

As with any type of production, an essential part of software production is determining the quality of the software. The level of quality associated to a software product is inevitably tied to properties such as how well it was developed and how useful it is to its users. The AtlanMod team is considering software quality aspects in two areas that are described in the following paragraphs: formal verification of models and corpus-based domain-specific language (DSL) analysis.

In the context of MDE, models are expressions of a conceptualization (metamodel) of their respective domain. Therefore precise definitions of our metamodels are important to get the subsequent MDE activities right, such as testing, concrete syntax definition, model interchange, code generation, and any more. In this sense, metamodels establish a single point of truth. The Object Constraint Language (OCL) of the OMG is widely accepted as a standard language to provide such precise definitions.

However, precise metamodels are complex artifacts. To get them right, systematic approaches to quality assurance are required. As a solution to this problem, we propose formal checking of model satisfiability (can we express the desired things), and formal checking of unsatisfiability (does the metamodel not have models with undesirable properties). As both activities constitute NP-hard problems, they have to be put carefully into algorithms. The team maintains the tool EMFtoCSP which translates the problem into the domain of constrain logic programming (CLP) for which sophisticated decision procedures exist. The tool integrates the described functionality in the Eclipse Modeling Framework (EMF) and the Eclipse Modeling Tools (MDT), making the functionality available for MDE in practice.

In terms of DSLs, an equally important aspect of software language engineering other than the initial development process is the identification of quality characteristics of the language. We are currently investigating the use of corpus-based analysis to identify language characteristics. A corpus in this case consists of artifacts or models of the DSL that have been generated by its users. We aim to extract information based on the corpus of a DSL to identify various characteristics that can potentially prove useful for the language engineer in his or her efforts to improve the language. Such information can complement other quality measurements including the formal verification method described in the previous paragraph and user feedback-based evaluations. In addition to an initial investigation of cloning in DSLs, our corpus-based analysis is also investigating metamodel element instance and relationship analyses

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, were security (and safety) is a major concern, may benefit from ideas such as dynamic security policy monitoring. On the other hand, systems used by general public (basically, the internet and services such as web services, social networks, etc.) can also benefit from results obtained by CIDRE, especially with regards to privacy. Systems are getting more and more complex, decentralized, distributed, or spontaneous. The emergence of cloud computing brings many challenges that could benefit from ideas, approaches and solutions studied by CIDRE in the context of distributed systems.

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.

4.3. Software Product Lines

A Software Product Line is a set of software systems that together address a particular market segment or fulfill a particular mission. Today, Software Product Lines are successfully applied in a range of industries, including telephony, medical imaging, financial services, car electronics, and utility control [54]. Customization and integration are keywords in Software Product Lines: a specific system in the family is constructed by selecting its properties (often technically called “features”), and, following such selection, by customizing and integrating the needed components and deploying them on the required platform.

INDES Project-Team

4. Application Domains

4.1. Web programming

Along with games, multimedia applications, electronic commerce, and email, the web has popularized computers in everybody's 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. House Automation

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 web-based home automation platform. For instance, we experiment with controlling heaters, air-conditioners, and electronic shutters with our mobile phones using web GUIs.

LOGNET Team

3. Application Domains

3.1. Introduction

We study overlay networks and peer-to-peer systems. Our skills are applied to studying protocols to interconnect heterogeneous networks, while guaranteeing backward compatibility. We experiment with those networks and systems in various fields, such as social networks and video streaming.

We design and implement a generic social platform, which is able to “program” and “run” (in a cloud based platform hosting a NoSQL data base) generic social networks. This is the first step towards a full decentralized P2P-based social network platform.

We also study Trust and Reputation Systems for P2P networks, and for Network Web Economy.

The final objective of those research veins is to move the computer and the computability at the edge of the network.

As another topic, we also study logics and type theory for improving proof assistants based on the Curry-Howard Isomorphism.

3.2. Applications

Because of its generality, our overlay network can target many applications. We would like to list a small number of useful programmable overlay-network-related case studies that can be considered as “LogNet Grand Challenges”, to help potential readers understand the interest of our research program.

- Interconnecting overlay networks transparently;
- building a programmable social network platform relying on a cloud + P2P architecture;
- experimenting with our interconnecting algorithm in the domain of video streaming;
- studying and integrating mobile devices and mobile networks 3G/4G as a real peer in actual P2P systems;
- studying trust and reputation systems applied to P2P and web economy;
- studying new distributed models of computation (long term objective);
- studying new type theories and lambda-calculi to be the basis of new proof assistants based on Curry-Howard isomorphism.

MYRIADS Project-Team

4. Application Domains

4.1. Application Domains

The MYRIADS research activities address a broad range of applications domains. We validate our research results with selected use cases from the following application domains:

- Web services, Service oriented Applications,
- Business applications,
- Bio-informatics applications,
- Computational science applications,
- Numerical simulation.

OASIS Project-Team

4. Application Domains

4.1. Grid and Cloud Computing

Grid, peer-to-peer, group communication, mobile object systems, Cloud, fault tolerance, distribution, security, synchronisation

As distributed systems are becoming ubiquitous, Grid computing, and the more recent concept of Cloud computing are facing a major challenge for computer science: seamless access and use of large-scale computing resources, world-wide. It is believed that by providing pervasive, dependable, consistent and inexpensive access to advanced computational capabilities, large-scale computing infrastructures allow new classes of applications to emerge.

There is a need for models and infrastructures for grid and peer-to-peer computing, and we promote programming models based on communicating mobile and distributed objects and components that can allow to harness these infrastructures. Another challenge is to use, for a given computation, unused CPU cycles of desktop computers in a Local Area Network, or even from wide area interconnected nodes. This is local or wide area Computational Peer-To-Peer, a concept that can contribute to a global energy footprint reduction. This is a challenge that also appears to be pregnant in more stable and homogeneous environments compared to P2P systems, such as datacenters under the problematics known as Virtual Machines consolidation.

4.2. Service Oriented Architectures (SOA)

Distributed Service Composition, Distributed Service Infrastructure, Peer-to-Peer data storage and lookup, Autonomic Management, Large-Scale deployment and monitoring

Service Oriented Architectures aim at the integration of distributed services and more generally at the integration of distributed and heterogeneous data, at the level of the Enterprise or of the whole Internet.

The team seeks solutions to the problems encountered here, with the underlying motivation to demonstrate the usefulness of a large-scale distributed programming approach and runtime support as featured by ProActive and GCM:

- Interaction between services: the uniform usage of web services based *client-server* invocations, through the possible support of an Enterprise Service Bus, can provide a simple interoperability between them. GCM components can be exposed as web services [61], and we have conducted research and development to permit a GCM component to invoke an external web service through a client interface. We also have provided a Service Component Architecture (see SCA specifications at http://www.oasis-opencsa.org/sca) personality for GCM components (GCM/SCA) so they can be integrated in SCA-based applications relying on SCA bindings configured as web services. For more loosely coupled interactions between services (e.g. compliant to the Web Services Notification standard), we pursue efforts to support *publish-subscribe* interaction models. Scalability in terms of number of notified events per time unit, and full interoperability through the use of semantic web notations applied to these events are some of the key challenges the community is addressing and we too. Events also correspond to data that may be worth to store, for future analytics, besides being propagated to interested parties (in the form of the event content). Our research can thus also contribute to the Big Data domain, a currently hot topic in ICT.
- Services compositions on a possibly large set of machines: if service compositions can even be turned as autonomic activities, these capabilities will really make SOA ready for the Open Internet scale (because at such a scale, a global management of all services is not possible). For service compositions represented as GCM-based component assemblies, we are indeed exploring the use of control components put in the components membranes, acting as sensors or actuators, that can

drive the self-deployment and self-management of composite services, according to negotiated Service Level Agreements. For service orchestrations usually expressed as BPEL like processes, and expressing the *composition in time* aspect of the composition of services, supports for deployment, management, and execution capable to support dynamic adaptations are also needed. Here again we believe a middleware based upon distributed and autonomous components as GCM can really help.

4.3. Simulation tools and methodology

simulation, component-based modeling, parallel and distributed simulation, reproducibility, architecture description language

Components are being used in simulation since many years. However, given its many application fields and its high computation needs, simulation is still a challenging application for component-based programming techniques and tools.

We have been exploring the application of Oasis programming methods to simulation problems in various areas of engineering problems, but also of financial applications.

More recently, with the arrival of O. Dalle in the team, and following a work previously started in the Mascotte project-team in 2006 [59], we are pursuing research on applying distributed component-based programming techniques to simulation. More precisely, new results are sought in three directions:

- improvement of simulation tools and methodology with techniques and tools borrowed from latest research in component-based software engineering;
- improvement of simulation scalability using high performance and distributed computing techniques;
- application of the results in the previous directions, in particular to the simulation of very large-scale distributed systems, such as peer-to-peer networks.

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, avionics and assisted living.

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. Assisted Living

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

The emerging industry of digital 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 solution that is open-ended in terms of applications and entities. (1) An on-line 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 can be added to a platform description to enhance its functionalities and extend the scope of future applications.

4.4. Avionics

In avionics, an aircraft can be seen as an environment full of sensors (*e.g.,* accelerometers, gyroscopes, and GPS sensors) and actuators (*e.g.,* ailerons and elevator trim). For example, a flight guidance system controls the aircraft using data produced by sensors. In a critical platform such as an aircraft, software systems have to be certified. Moreover the safety-critical nature of the avionics domain takes the form of stringent non-functional requirements, resulting in a number of challenges in software development:

- *Traceability.* Traceability is the ability to trace all the requirements throughout the development process. In the avionics certification processes, traceability is mandatory for both functional and non-functional requirements.
- *Coherence.* Functional and non-functional aspects of an application are inherently coupled. For example, dependability mechanisms can potentially deteriorate the overall performance of the application. The coherence of the requirements is particularly critical when the software evolves: even minor modifications to one aspect may tremendously impact the others, leading to unpredicted failures.
- *Separation of concerns.* Avionics platforms involve the collaboration of several experts (from low-level system to software, safety, QoS), making requirements traceability significantly more challenging. Providing development methodologies that allow a clear separation of concerns can tremendously improve traceability.

Our approach consists of enriching a design language with non-functional decalarations. Such declarations allow the safety expert to specify at design time how errors are handled, guiding and facilitating the implementation of error handling code. The design is also enriched with Quality of Service (QoS) declarations such as time constraints. For each of these non-functional declarations, specific development support can be generated. We have validated this approach by developing flight guidance applications for avionics and drone systems.

REGAL Project-Team

4. Application Domains

4.1. Research domain

To address the evolution of distributed platforms in recent years, we focus on the following areas:

- *Distributed algorithms for dynamic and large networks.* Network topology is no more static; distributed systems are increasingly dynamic, i.e., nodes can join, fail, recover, disconnect and reconnect, and change location. Examples include IaaS cloud computing infrastructures, where virtual machines can be moved according to load peaks, opportunistic networks such as DTNs (Delay-Tolerant Networks), and networks of robots.
- *Management of distributed data.* In emerging architectures such as distributed hash tables (DHTs) and cloud computing, our research topics include replica placement, responsiveness, load balancing, consistency maintenance, consensus algorithms, and synchronisation. This research direction is funded by several new collaborative projects (ConcoRDanT, MyCloud, Nu@age, Odisea, Prose, Shaman, Spades, Streams, R-Discover) and by industrial funding (Google).
- *Performance and robustness of Systems Software in multicore architectures.* Our research focuses on the efficient management of system resources at the user level. Issues considered include efficient synchronization and memory management in large-scale multicore architectures. At the same time, we focus on the robustness of systems software, based on the Coccinelle technology. This work is funded by ANR ABL and InfraJVM.

RMOD Project-Team (section vide)

SARDES Project-Team (section vide)

SCORE Team

4. Application Domains

4.1. Introduction

Our research is inspired by potential problems and use cases from the real world. The rapid pace of technology evolution oblige us to adapt our perspective on some of the problems we consider. Recently, the raise of the "As a Service" movement has both validated our past work but also obliged us to rethink some of our directions. Then, we have chosen to apply our work to domains because we think that the problems that they exhibit resonate with our work. This is the case with Crisis Management and with Software Engineering that are in some ways extreme case for collaboration and service consumption. We are also exploring more general business oriented applications.

4.2. E-government

E-government is a well established domain that provides its own requirements in the field of service and information management. From our perspective, e-government applications have very strong requirements regarding security, privacy and interoperability between different organizations, belonging potentially to different countries. One of the prominent contributions we have made in this domain is related to our collaboration with SAP on the relationship between processes, security policies and the problem of delegation that we considered as a important for organizational flexibility. This work resonate also with its current continuation in crisis management.

Crisis management is a special case of e-government application as it involves mostly governmental agencies in coordination with other organizations like the Red Cross or other NGO. Moreover, it brings with it a lot of requirements that are very interesting for us in the domain of coordination: a crisis process shall be very flexible, adaptable and distributed. It is mostly human driven and can be critical. In this domain, we are collaborating with SAP to define a new model of coordination that should support people involved in crisis resolution.

4.3. Groupware Systems and Software Forge

Software engineering environment can be seen as distributed collaborative systems. Software Forges are social software (e.g. Github). They transform strangers into collaborators. Forges are online services that allow instantiating, composing and managing collaborative services. Traditionally, collaborative services are version control systems, issue trackers, forums, mailing lists or wikis. We are applying our research results on coordination and data sharing into this context foreseeing a forge, not as a monolithic framework but as a composition of services that can be deployed in a Cloud infrastructure.

TRISKELL Project-Team

4. Application Domains

4.1. Application Domains

SOA, telecommunication, distributed systems, Embedded Systems, software engineering, test, UML

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.

ALGORILLE Project-Team

4. Application Domains

4.1. Scientific computation

In the context of studying GPU cluster programming as well as asynchronous algorithms, we have developed a parallel code to solve a PDE problem that is representative of scientific computations. This is a 3D version of the *advection-diffusion-reaction* problem modeling the interactions of two chemical species in shallow waters. Different versions of the application have been implemented. A first version has been the subject of extensive studies about computing and energy performance. Other versions, especially the last one using MPI+OpenMP+CUDA, are still the subject of studies.

Also in a close collaboration with Fatmir Asllanaj from the LEMTA laboratory, Nancy, France, we launched the development of a code for Radiative Transport Equations, ETR. As a first step to higher efficiency the code (originally in FORTRAN) was completely rewritten in C. Already this sequential rewrite has largely improved the time and memory efficiency of the code. The next step will be a parallelization for multi-core machines and clusters that will make this a unique tool to tackle ETRs in 2 and 3 dimensions.

4.2. Financial computation

AmerMal application is a parallel *American Option Pricer* developed in collaboration with Lokman Abbas-Turki during his PhD thesis at University Paris-Est. This pricing mathematic algorithm has been designed in order to be parallelized on clusters of GPUs. It is based on a Maillavin calculus and on Monte Carlo stochastic computations. However, the Monte Carlo trajectories are coupled, and our final parallel algorithm includes many communications between the computing nodes. Nevertheless, we achieved good scalability on our 16-nodes GPU clusters.

The initial version of AmerMal has been implemented using MPI+OpenMP+CUDA in 2010 and 2011. In 2012, we have ported this application on ORWL+CUDA, both in order to test and improve ORWL semantic and ORWL performances.

4.3. Signal processing

ParSONS is a parallel *sound sources classifier* developed at SUPELEC by Stephane Rossignol. It allows to identify different kinds of sound sources in an audio signal file (like a radio or TV archive file), and to tag different part of the file. Then it is possible to know which part of the archive file include human voice, or music, or more accurately jazz or rock music. From an algorithmic point of view, this application performs a lot of Fourrier Transform computations and a lot of IO operations.

We have developed an MPI+OpenMP version, optimizing input data file reading and overlapping these IO operations with computations (signal processing). We run some experiments on a 256-nodes dual-core PC cluster and on a 16-nodes 4-hyperthreaded-cores experimentation PC cluster, both located at SUPELEC. We identified the most efficient configurations, considering the number of input file reading processes, the number of threads per process, the number of processes per node... Finally, we have designed and implemented a deployment tool, in order to hide these configuration issues to signal processing researchers and users.

In a near future, we aim to port this application to ORWL, to experiment, improve and validate ORWL on distributed applications achieving a high ratio of IO compared to computations and classic communications.

AVALON Team

4. Application Domains

4.1. Overview

The Avalon team targets applications with large computing and/or data storage needs, which are yet difficult to program and maintain while achieving performance. Those applications can be parallel and/or distributed applications. It is typically the case of large simulations and/or code coupling applications from "classical" computational problems such as climate forecasting or cosmology. The team also deals with bioinformatics as they raise some interesting research issues linked to data management. For example, some of such applications rely on MapReduce algorithmic skeletons.

4.2. Bioinformatics

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

The MapReduce programming model publicized by Google is a widely used model for deploying application services on platforms for data processing on a large scale (Grids and Clouds). MapReduce implementations extremely are scalable and efficient when located on the datacenters of big companies. However, technological (use of a parallel file system) and algorithmic (hard-coded replication) choices of an implementation such as Hadoop compromise performance on platforms like grids of heterogeneous clusters. Applying it to large scale and volatile environments such as desktop grids still remain a challenge.

4.3. Climate Forecasting Simulations

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.

As for most applications the team is targeting, our goal is to thoroughly analyze climate forecasting applications to model its needs in terms of programing model, execution model, data access pattern, and computing needs. Once a proper model of the application has been derived, appropriate scheduling heuristics could be designed, tested, and compared.

4.4. Cosmological Simulations

Ramses[1] is a typical computational intensive application used by astrophysicists to study the formation of galaxies. *Ramses* is used, among other things, to simulate the evolution of a collisionless, self-gravitating fluid called "dark matter" through cosmic time. Individual trajectories of macro-particles are integrated using a state-of-the-art "N body solver", coupled to a finite volume Euler solver, based on the Adaptive Mesh Refinement technique. The computational space is decomposed among the available processors using a *mesh partitioning* strategy based on the Peano-Hilbert cell ordering.

Cosmological simulations are usually divided into two main categories. Large scale periodic boxes requiring massively parallel computers are performed on a very long elapsed time (usually several months). The second category stands for much faster small scale "zoom simulations".

[1] http://www.itp.uzh.ch/~teyssier/Site/RAMSES.html

4.5. Code Coupling Applications

Different kinds of code coupling applications are considered. The simplest form is within a parallel code. For example, OpenAtom, a molecular dynamics simulation application exhibits several coupling of different pieces of codes, whose arrangement and configuration depend on many parameters, and some depend on input simulation parameters. The challenge is to let application designers express the functionality of the application while the actual execution code can be automatically configured.

Another class of code coupling applications is the coupling at temporal level. Our collaboration with EDF R&D provides us with several use cases including scenarios with *Code_Aster* (thermo-mechanics simulations), *Syrthes* (transient thermal simulations in complex solid geometries), etc. These couplings face other issues such as combining several large scale resources, managing data transfers, etc.

CEPAGE Project-Team

4. Application Domains

4.1. Resource Allocation and Scheduling

4.1.1. Project-team positioning

CEPAGE has undertaken tasks related to the *high level modeling* of heterogeneous networks, both at logical level (overlay networks design) and performance level (latency, bandwidth prediction, connectivity artifacts) in order to optimize tasks such as *resource allocation* and *scheduling* of computations and communications. Objectives include replica placement, broadcasting (streaming) of large messages, independent tasks scheduling and optimization of OLAP databases. Such problems have received at lot of attention in research centers in the USA (Armherst, Colorado, ...), in Spain (Madrid), Poland (Wroclaw), Germany (Dortmund), and others. Papers on algorithmic aspects of platform modeling, scheduling and resource allocation appear at parallel processing conferences and journals in Parallel and Distributed Computing (IPDPS, EuroPar, HIPC, SPAA, IEEE TPDS, JPDC) and members of CEPAGE are strongly involved in many of these events (IPDPS, EuroPar, TPDS) as well as helping to animate well-established specialized workshops, such as HCW and HeteroPar.

Within Inria, studies on overlay networks are performed in the ASAP and GANG projects, and studies related to scheduling and resource allocation are done within the ROMA and the MOAIS projects (and to some extent within ALGORILLE).

4.1.2. Scientific achievements

The approach followed in the CEPAGE project, and our main originality, is to consider the whole chain, from gathering actual data on the networks to platform modeling and complexity analysis. Indeed, many complexity analysis studies are performed on models whose parameters cannot actually be evaluated (this applies, for instance, to all algorithms that assume that the topology of a platform running over the Internet is known in advance) and many platform models are intractable from an algorithmic perspective (this applies, for instance, to all models that represent latencies or bandwidths between all pairs of nodes as a general matrix). Our general goal is to provide models whose parameters can be evaluated at runtime using actual direct measurements, to propose algorithms whose worst-case (or average-case) behavior can be proved for this model, and finally to evaluate the whole chain (model + algorithm + implementation).

From an applicative perspective, in the framework of the PhD Thesis of Hejer Rejeb, we have considered several storage and resource allocation problems in collaboration with Cyril Banino-Rokkones at Yahoo! Trondheim (dealing with actual datasets enabled us to improve known approximation results in this specific context). We have in particular studied the modeling of TCP mechanism for handling contentions and its influence on the performance of several scheduling algorithms and advocated the use of QoS mechanisms for prescribed bandwidth sharing (IPDPS 2010 [67], ICPADS 2008 [52], AlgoTel 2009 [64], ICPADS 2009 [63], PDP 2010 [65]). In the PhD thesis of Hubert Larchevêque, we have considered the problem of aggregating resources (or placing replicas) in a distributed network (Sirocco 2008 [54], Opodis 2008 [55], ICPP 2011 [60], AlgoTel 2011 [56]) so that each group satifies some properties (in terms of aggregated memory, CPU and maximal distance in terms of latency within a group). We proved several multi-criteria approximation results for this problem, and we compared several embedding tools (Vivaldi, Sequoia) in the context of resource aggregation. For these applications, we have also provided when possible distributed algorithms based on sophisticated overlay networks, in particular in order to deal with heterogeneity (IPDPS 2008 [61]). In the PhD Thesis of Przemyslaw Uznanski, we focus on the design of efficient streaming and broadcasting strategies, in particular in presence of connectivity artifacts like firewalls (IPDPS 2010 [62], ICPADS 2011 [59]). We have also worked on establishing under the bounded multiport model several new complexity results for classical distributed computing models such as divisible load theory (HCW 2008 [57], IPDPS 2008 [107], IPDPS 2012 [58]) that have been later extended to Continuous Integration (HCW 2012 [53]).

In the context of database query optimization, materializing some queries results for optimization is a standard solution when execution time performance is crucial. In the datacube context, the problem has been studied for a long time under the storage space limit constraint. Here also, we were able to reformulate this problem by considering instead the execution time as the hard constraint while the objective is to reduce the storage space. Even if the problem turns to be NP-hard, this reformulation allowed us to provide effective approximate solutions with both space and performance bounded guarantees (EDBT 2009 [96]). Moreover, reducing the storage space tends to reduce the maintenance time since the latter is linearly proportional to the former. Finally, we characterized the minimal number of updates to be performed before performance becomes no more guaranteed and a new solution must be recomputed (ADBIS 2008 [97]). One of the key concepts we used for solving this problem was that of a *border*. It turns out that this notion is equivalent to e.g., maximal frequent itemsets or minimal functional dependencies extensively studied by data mining community. In contrast to all previous proposals, we proposed the only parallel algorithm computing these borders with a speed-up guarantee regarding the number of processing units (CIKM 2011 [95]). Besides the analytical study, its implementation in maximal frequent itemset mining outperforms state of the art implementations (see Section 5.1).

To achieve these results, our efforts have also focused on analyzing and building realistic datasets (AlgoTel 2012 [86]) and proposing data analysis results for specific distributions (ISAAC 2011 [48]). On the modeling side, in general, for bandwidth and contention modeling, we have proved that the bounded multi-port model (where each node is associated to an incoming bandwidth, an outgoing bandwidth and a maximal number of simultaneous TCP connexions) is both implementable, realistic and tractable (EuroPar 2011 [66]). In particular, we have proved in strongly different contexts (allocation of virtual machines to physical machines, overlay design for broadcasting, server allocation for volunteer computing) that the use of resource augmentation enables to obtain quasi-optimal results. All our modeling efforts and algorithms have been included into the SimGRID Software (http://simgrid.gforge.inria.fr), which enables us both to compare several algorithms under the same exact conditions and to compare the results obtained with several communication models (see Section 5.1)..

Perspectives: We believe that our approach based on sound models, approximation algorithms for these models, followed by experimental validation is a strong one and we intend to continue in this direction in the following years. Our goal of designing realistic solutions pushes towards considering average case analysis of our algorithms, as well as robust optimization techniques. Furthermore, the recent strong interest in Cloud systems from the community entices us to use our expertise in resource allocation for the optimization of Cloud systems, both from the provider and from the user points of view. We already have some interesting contacts with local companies to share start collaborating on these topics. In this context, reliability issues are very important, and we believe that robust optimization is a very relevant approach for these problems.

4.2. Compact Routing

4.2.1. Project-team positioning

In this axis, CEPAGE mainly works on the design on distributed and light data structures. One of the techniques consists in summarizing the topology and metric of the networks allowing to route or to approximate the original distances within the network. Such structures, often called *spanners*, does not require the storage of all the original network links. Then we get economic distributed data structures that can be updated without a high communication cost. Our main collaborations are done with the best specialists world-wide, in particular: Israel (Weizmann), USA (MIT, Microsoft, Chicago), Belgium (Alcatel Lucent-Bell), France (Paris, Nice).

Algorithms and Routing are also intensively studied in research labs in the USA (CAIDA). Our contributions appear regularly at all of the major conferences in Distributed Computing (PODC, DISC, SPAA), as well as at top venues with a more general algorithmic audience (STOC, SODA, ICALP, ESA). Members of CEPAGE actively participate in these events (ICALP 2010 and DISC 2009 were organized by members of CEPAGE).

Within Inria, studies of mobile agents are also performed in the GANG project and to some extent also within MASCOTTE within the european project EULER.

4.2.2. Scientific achievements

There are several techniques to manage sub-linear size routing tables (in the number of nodes of the platform) while guaranteeing almost shortest paths. Some techniques provide routes of length at most $1+\epsilon$ times the length of the shortest one while maintaining a poly-logarithmic number of entries per routing table. However, these techniques are not universal in the sense that they apply only on some class of underlying topologies. Universal schemes exist. Typically they achieve $O(\sqrt{n})$-entry local routing tables for a stretch factor of 3 in the worst case. Some experiments have shown that such methods, although universal, work very well in practice, in average, on realistic scale-free or existing topologies.

The space lower bound of $O(\sqrt{n})$-entry for routing with *multiplicative* stretch 3 is due to the existence of dense graphs with large girth. Dense graphs can be sparsified to subgraphs (spanners), with various stretch guarantees. There are spanners with *additive* stretch guarantees (some even have constant additive stretch) but only very few additive routing schemes are known.

In (SPAA 2012 [90]), we give reasons why routing in unweighted graphs with *additive* stretch is difficult in the form of space lower bounds for general graphs and for planar graphs. On the positive side, we give an almost tight upper bound: we present the first non-trivial compact routing scheme with $o(\lg^2 n)$-bit addresses, *additive* stretch $O(n^{1/3})$, and table size $O(n^{1/3})$ bits for planar graphs.

We have recently considered the *forbidden-set* extension of distance oracles and routing schemes. Given an arbitrary set of edge/node failure F, a source s and a target t such that $s,t \notin F$, the goal is to route (or evaluate the distance) between s and t in the graph $G \setminus F$, so avoiding F. The classical problem is for $F = \varnothing$. This extension is considered as a first step toward fully dynamic data-structures, a challenging goal. For graphs of low doubling dimension we have shown in (PODC 2012 [47]) that it is possible to route from s to t in $G \setminus F$ with stretch $1+\epsilon$, for all s,t,F, given poly-logarithmic size labels of all the nodes invoked in the query (s,t,F). This has been generalized to all planar graphs achieving similar stretch and label size performences. As a byproduct we have designed a fully dynamic algorithm for maintaining $1+\epsilon$ approximate distances in planar graphs supporting edge/node addition/removal within update and query time $\sqrt{n}$ in the worst-case (STOC 2012 [46]).

Θ_k-graphs are geometric graphs that appear in the context of graph navigation. The shortest-path metric of these graphs is known to approximate the Euclidean complete graph up to a factor depending on the cone number k and the dimension of the space. We have introduced in (WG 2010 [68]) a specific subgraph of the Θ_6-graph defined in the 2D Euclidean space, namely the half-Θ_6-graph, composed of the even-cone edges of the Θ_6-graph. Our main contribution is to show that these graphs are exactly the TD-Delaunay graphs, and are strongly connected to the geodesic embeddings of orthogonal surfaces of coplanar points in the 3D Euclidean space. We also studied the asymptotic behavior of these spanners (*Adv. in Appl. Proba.* [105]) and in collaboration with Ljubomir Perković, we worked on the question of bounded degree planar spanner. We proposed an algorithm that computes a plane 6-spanner of degree at most 6 in (ICALP 2010 [69]). The previous best bound on the maximum degree for constant stretch plane spanners was only 14.

In order to cope with network dynamism and failures, and motivated by multipath routing, we introduce a multi-connected variant of spanners. For that purpose we introduce in (OPODIS 2011 [91]) the p-multipath cost between two nodes u and v as the minimum weight of a collection of p internally vertex-disjoint paths between u and v. Given a weighted graph G, a subgraph H is a p-multipath s-spanner if for all u,v, the p-multipath cost between u and v in H is at most s times the p-multipath cost in G. The s factor is called the stretch. Building upon recent results on fault-tolerant spanners, we show how to build p-multipath spanners of constant stretch and of $O(n^{1+1/k})$ edges, for fixed parameters p and k, n being the number of nodes of the graph. Such spanners can be constructed by a distributed algorithm running in $O(k)$ rounds. Additionally, we give an improved construction for the case $p=k=2$. Our spanner H has $O(n^{3/2})$ edges and the p-multipath cost in H between any two node is at most twice the corresponding one in G plus $O(W)$, W being the maximum edge weight.

We also worked on compact coding in data warehouses: in order to get quick answer in large data, we have to estimate, select and materialize (store) partial data structures. We got several solutions with a prescribed guarantee in different models for the following problems: view size estimation with small samples, view selection, parallel computation of frequent itemsets. In (*Theor. Comp. Sci.* [94]) a new algorithm that allow the administrator or user of a DBMS to choose which part of the data cube to optimize (known as the *the views selection problem*), that takes as input a fact table and computes a set of views to store in order to speed up queries.

Perspectives: The compact coding activity in data-warehouse is promising since the amount of data collected keeps on increasing and being able to answer in real-time complex requests (data mining) is still challenging.

Some robust data structures already exist which, given a small number of k changes of topology or k faults, tolerate these faults, i.e., alternative routes with bounded stretch can be provided without any updates. This is a first step toward dynamic networks but the updates of these data structures are currently still quite complicated with a high communication cost.

4.3. Mobile Agents

4.3.1. Project-team positioning

CEPAGE has undertaken tasks related to the design of algorithms which control the behavior of so called *mobile agents*, moving around a network or a geometric environment, with the goal of achieving a specified objective. Objectives of central importance to the study include: exploration of unknown environments, terrain patrolling, network maintenance, and coordination of activities with other agents. Such problems have in recent years been the object of interest of numerous research teams working on Distributed Computing worldwide, in particular, at research centers in Canada (Quebec), Israel (Tel Aviv, Haifa), France (Paris, Marseille), the UK (London, Liverpool), and Switzerland (Zurich). Algorithms for mobile agents in social networking applications are also intensively studied in research labs in the USA (Stanford, Facebook). Papers on mobile agents appear regularly at all of the major conferences in Distributed Computing (PODC, DISC, SPAA), as well as at top venues with a more general algorithmic audience (SODA, ICALP, ESA). Members of CEPAGE actively participate in these events, and are also a recognizable part of the European community focused around mobile agents, helping to animate well-established specialized conferences, such as SIROCCO and OPODIS.

Within Inria, studies of mobile agents are also performed in the GANG project, and to some extent also within MASCOTTE. CEPAGE has active research links with both of these teams.

4.3.2. Scientific achievements

The work of CEPAGE has focused on contributing new decentralized algorithms for controlling mobile entities known as *agents*, deployed in unknown environments. We mainly considered the network setting, in which agents moving around the nodes of the network graph may be used to analyze the structure of the network and to perform maintenance tasks, such as detecting dynamic faults, improving/monitoring dissemination of information, etc. Our theoretical studies focused on designing new strategies for controlling the behavior of agents and answering crucial questions concerning the feasibility of solving fundamental problems, subject to different model assumptions and constraints on the knowledge and computational power of agents.

One major line of our research focused on the so called *anonymous graph model* in which an agent is unable to determine the identifier of the node of its current location, but can only see a local ordering of the links around it. Such a study is motivated e.g. by scenarios in which the identifiers of nodes may be too large for the agent to process using its bounded resources, or may change in time. In this model, we studied two of the most fundamental problems: that of traversing all of the nodes of the network (exploration) and of meeting another agent in the network (rendezvous), so as to coordinate with it. Our contributions include a precise characterization of the space requirements for agents solving both of these problems deterministically: exploration in (*Trans. Alg.* 2008 [73]) and rendezvous in (*Dist. Comp.* 2012 [81]), in a paper presented at the Best Paper Session of PODC 2010. We have also studied fast solutions for specific scenarios of the rendezvous problem (DISC 2010 [49], DISC 2011 [74], SPAA 2012 [82]) and the problem of approximate map construction within an anonymous graph (OPODIS 2010 [71]). A separate problem, intensively studied in recent years by several research teams, concerns the exploration of a network with pre-configured ports so as to assist the agent. In our work on the topic, our team has proposed several new techniques for graph decomposition, leading in particular to the shortest currently known strategies of periodic exploration for both the case of memoryless (*Algorithmica* 2012 [101]) and small-memory agents (SIROCCO 2009 [77]).

A closely related line of research was devoted to the design of network exploration strategies which guarantee a fast and fair traversal of all the nodes, making use of agents with extremely restricted capabilities. Such strategies were inspired by the random walk, but had the additional advantage of deterministic and desirable behavior in worst-case scenarios. We presented a series of results in the area at notable conferences, involving both the design of new exploration strategies (ICALP 2009 [75]) and completely new insights into previously known approaches such as the so called "rotor-router model" (DISC 2009 [50], OPODIS 2009 [51]). All of the proposed algorithms were shown to be viable alternatives to the random walk, competing in terms of such parameters as cover time, steady-state exploration frequency, and stabilization in the event of faults.

Our efforts have also focused on the theory of coordinating activities of large groups of agents. We have conducted pioneering work in the so called look-compute-move model in networks, in which extremely restricted (asynchronous and oblivious) agents, relying on snapshot views of the system, are nevertheless able to perform useful computational tasks. Our solutions to the problems of collective exploration in trees (*Theor. Comp. Sci.* 2010 [88]) and gathering agents on a ring (*Theor. Comp. Sci.* 2008 [99] and 2010 [98]) have sparked a long line of follow-up research, accumulating more than 120 citations in total (according to Google Scholar). In a slightly different scenario, we have considered computations with teams of agents whose task is to collaboratively detect and mark potentially dangerous (faulty) links of the network, called "black holes", which are capable of destroying agents which enter them. We have provided important contributions to the theory of black hole search in both undirected (SIROCCO 2008 [76], DISC 2008 [89]) and directed (*Theor. Comp. Sci.* [102]) graphs.

It is expected that the mobile agent theme of CEPAGE will give rise to 2 PhD theses. In 2013, Ahmed Wade will defend his thesis on mobile agent protocols for dynamic networks, whereas in 2014 Dominik Pajak will defend his thesis on multi-agent protocols for efficient graph exploration. Our scientific interests also include mobile agent protocols for geometric applications, more remote from the central themes of CEPAGE, but having extensive applications in robotics (providing protocols, e.g., for efficient patrolling and guarding of terrains, traversing terrains using groups of robots, etc.). We have already published several papers in this area (SIROCCO 2010 [79], SWAT 2010 [80], ESA 2011 [78]), building up the theoretical fundamentals of a new field, and already attracting the attention of a wider community of researchers working in robotics and AI.

Perspectives: Our goal is to explore applications of mobile agent techniques in domains of growing importance, namely, social networks and robotics. We are currently discussing applications of our techniques in problems of brand recognition on the web with a local industrial partner (Systonic KeepAlert), and other companies (through our research collaborators in Liverpool). We intend to undertake collaboration with European/American research labs and industrial partners.

GRAND-LARGE Project-Team (section vide)

HIEPACS Project-Team

4. Application Domains

4.1. Introduction

Currently, we have one major application which is material physics, and for which we contribute to all steps that go from modelling aspects to the design and the implementation of very efficient algorithms and codes for very large multi-scale simulations. Moreover, we apply our algorithmic research about linear algebra (see Section 3) in the context of several collaborations with industrial and academic partners. Our high performance libraries are or will be integrated in several complex codes and will be used and validated for very large simulations.

4.2. Material physics

Participants: Bérenger Bramas, Olivier Coulaud, Aurélien Esnard, Pierre Fortin, Luc Giraud, Jean Roman.

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 is two ways to decrease the computational cost. In the first approach, we improve classical methods and algorithms and in the second way, we will consider a multiscale approach.

Many applications in material physics need to couple several models like quantum mechanic and molecular mechanic models, or molecular and mesoscopic or continuum models. These couplings allow scientists to treat larger solids or molecules in their environment. Many of macroscopic phenomena in science depend on phenomena at smaller scales. Full simulations at the finest level are not computationally feasible in the whole material. Most of the time, the finest level is only necessary where the phenomenon of interest occurs; for example in a crack propagation simulation, far from the tip, we have a macroscopic behavior of the material and then we can use a coarser model. The idea is to limit the more expensive level simulation to a subset of the domain and to combine it with a macroscopic level. This implies that atomistic simulations must be speeded up by several orders of magnitude.

We will focus on two applications; the first one concerns the computation of optical spectra of molecules or solids in their environment. In the second application, we will develop faster algorithms to obtain a better understanding of the metal plasticity, phenomenon governing by dislocation behavior. Moreover, we will focus on the improvement of the algorithms and the methods to build faster and more accurate simulations on modern massively parallel architectures.

4.2.1. Hybrid materials

There is current interest in hybrid pigments for cosmetics, phototherapy and paints. Hybrid materials, combining the properties of an inorganic host and the tailorable properties of organic guests, particularly dyes, are also of wide interest for environmental detection (oxygen sensors) and remediation (trapping and elimination of dyes in effluents, photosensitised production of reactive oxygen species for reduction of air and water borne contaminants). A thorough understanding of the factors determining the photo and chemical stability of hybrid pigments is thus mandated by health, environmental concerns and economic viability.

Many applications of hybrid materials in the field of optics exploit combinations of properties such as transparency, adhesion, barrier effect, corrosion, protection, easy tuning of the colour and refractive index, adjustable mechanical properties and decorative properties. It is remarkable that ancient pigments, such as Maya Blue and lacquers, fulfill a number of these properties. This is a key to the attractiveness of such materials. These materials are not simply physical mixtures, but should be thought of as either miscible organic and inorganic components, or as a heterogeneous system where at least one of the component exhibits a hierarchical order at the nanometer scale. The properties of such materials no longer derive from the sum of the individual contributions of both phases, since the organic/inorganic interface plays a major role. Either organic and inorganic components are embedded and only weak bonds (hydrogen, van der Waals, ionic bonds) give the structure its cohesion (class I) or covalent and iono-covalent bonds govern the stability of the whole (class II).

These simulations are complex and costly and may involve several length scales, quantum effects, components of different kinds (mineral-organic, hydro-philic and -phobic parts). Computer simulation already contributes widely to the design of these materials, but current simulation packages do not provide several crucial functions, which would greatly enhance the scope and power of computer simulation in this field.

The computation of optical spectra of molecules and solids is the greatest use of the Time Dependent Density Functional Theory (TDDFT). We compute the ground state of the given system as the solution of the Kohn-Sham equations (DFT). Then, we compute the excited states of the quantum system under an external perturbation - electrical field of the environment - or thanks to the linear theory, we compute only the response function of the system. In fact, physicists are not only interesting by the spectra for one conformation of the molecule, but by an average on its available configurations. To do that, they sample the trajectory of the system and then compute several hundred of optical spectra in one simulation. But, due to the size of interesting systems (several thousands of atoms) and even if we consider linear methods to solve the Kohn-Sham equations arising from the Density Functional Theory, we cannot compute all the system at this scale. In fact, such simulations are performed by coupling Quantum mechanics (QM) and Molecular mechanic (MM). A lot of works are done on the way to couple these two scales, but a lot of work remains in order to build efficient methods and efficient parallel couplings.

The most consuming time in such coupling is to compute optical spectra is the TDDFT. Unfortunately, examining optical excitations based on contemporary quantum mechanical methods can be especially challenging because accurate methods for structural energies, such as DFT, are often not well suited for excited state properties. This requires new methods designed for predicting excited states and new algorithms for implementing them. Several tracks will be investigated in the project:

- Typically physicists or chemists consider spectral functions to build a basis (orbital functions) and all the computations are performed in a spectral way. Due to our background, we want to develop new methods to solve the system in the real space by finite differences or by wavelets methods. The main expectation is to construct error estimates based on for instance the grid-size h parameter.
- For a given frequency in the optical spectra, we have to solve a symmetric non Hermitian system. With our knowledge on linear solvers, we think that we can improve the methods commonly used (Lanczos like) to solve the system (see Section 3.3).
- Improving the parallel coupling is crucial for large systems because the computational cost of the atomic and quantum models are really different. In parallel we have the following order of magnitude: one second or less per time step for the molecular dynamics, several minutes or more for the DFT and the TDDFT. The challenge to find the best distribution in order to have the same CPU time per time step is really important to reach high performance. Another aspect in the coupling is the coupling with the visualization to obtain online visualization or steerable simulations. Such steerable simulations help the physicists to construct the system during the simulation process by moving one or a set of molecules. This kind of interaction is very challenging in terms of algorithmic and this is a good field for our software platform `EPSN` .

4.2.2. Material failures

Another 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 modelling 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) modelling. 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 starting OPTIDIS ANR-COSINUS 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 acheive 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.

4.3. Application framework customers of high performance linear algebra solvers

Participants: Emmanuel Agullo, Luc Giraud, Abdou Guermouche, Stojce Nakov, Jean Roman, Xavier Vasseur.

We are currenlty collaborating with various research groups involved in geophysics, electromagnetics and structural mechanics. For all these application areas, the current bottleneck is the solution of huge sparse linear systems often involving multiple right-hand sides either available simultaneously or given in sequence. The robustness, efficiency and scalability of the numerical tools designed in Section 3.3 will be preliminary investigated in the parallel simulation codes of these partners.

For the solution of large systems arsing from PDE discretization, the geometric full multigrid technique based on a few levels in the grid hierarchy and an efficient parallel sparse direct solver on the coarsest level can be considered. Originally developped for 3D Maxwell solution in collaboration with CEA-CESTA, the approach can be extended to other application fields.

Many simulation codes need the solution with simultaneous right-hand sides but also with right-hand sides given in sequence. The first situation arises in RCS calculations, but is generic in many parametric studies, while the second one comes from the nature of the solver such as implicit time tepping schemes or inverse iterations. Many of the numerical approaches and possible outcoming software are well suited to tackle these challenging problems.

On more academic sides, some ongoing collaborations with other Inria EPIs will be continued and others will be started. In collaboration with the NACHOS Inria project team, we will continue to investigate the use of efficient linear solvers for the solution of the Maxwell equations in the time and frequency domains where discontinuous Galerkin discretizations are considered. Additional funding will be sought out in order to foster this research activity in connection with actions described in Section 3.3 .

The efficient solution of linear systems strongly relies on the activities described in Section 3.2 (e.g. complex load balancing problem) and in Section 3.3 (for the various parallel linear algebra kernels).

4.4. Scalable numerical schemes for scientific applications

Participants: Olivier Coulaud, Yohann Dudouit, Luc Giraud, Guillaume Latu, Alexis Praga, Jean Roman, Pablo Salas Medina, Xavier Vasseur.

We are also collaborating with application research group to design or improve numerical schemes in the view of large scale parallel simulations.

Seismic wave propagation in heterogeneous media requires to properly capture the local heterogeneity and consequently requires locally refined meshes. In close collaboration with TOTAL we study new parallelizable schemes for the solution of the elastodynamic system with local spatial refinments based on discontinuous Galerkin techniques. The objective is to design novel parallel scalable implementations for large 3D simulations. A second work is currently carried on with TOTAL for Seismic modeling and Reverse Time Migration (RTM) based on the full wave equation discretization. These tools are of major importance since they give an accurate representation of complex wave propagation areas. Unfortunately, they are highly compute intensive. To address this challenge we have designed a fast parallel simulator that solves the acoustic wave equation on a GPU cluster.

Thermoacoustic instabilities are an important concern in the design of gas turbine combustion chambers. Most modern combustion chambers have annular shapes and this leads to the appearance of azimuthal acoustic modes. These modes are often powerful and can lead to structural vibrations being sometimes damaging. Therefore, they must be identified at the design stage in order to be able to eliminate them. However, due to the complexity of industrial combustion chambers with a large number of burners, numerical studies of real configurations are a challenging task. Such a challenging calculations performed in close collaboration with the Computational Fluid Dynamic project at CERFACS.

The chemistry and transport models (CTM) play a central role in global geophyscal models. The solution of the CTM represents up-to 50 % on the computing ressources involved in global geophyscal simulations. Therefore, the availability of efficient scalable parallel numerical schemes on emerging and future supercomputers is crucial. The purpose of this research activity is to study, design and implement novel numerical schemes following the work initiated by D. Cariolle in the framework of the ANR Solstice project. Alexi Praga, PhD hired by CERFACS, is conducing this research action under the joint supervision of HiePACS and the Aviation and Environment project at CERFACS in close collaboration with CNRM/Meteo-France.

KERDATA Project-Team

4. Application Domains

4.1. Application Domains

Below are three examples which illustrate the needs of large-scale data-intensive applications with respect to storage, I/O and data analysis. They illustrate the classes of applications that can benefit from our research activities.

4.1.1. 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, and that has remained poorly understood so far. As both neuroimaging- and genetic-domain observations represent a huge amount of variables (of the order of millions), performing statistically rigorous 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.

The A-Brain (AzureBrain) Project started in October 2010 within the Microsoft Research-Inria Joint Research Center. It is co-led by the KerData (Rennes) and Parietal (Saclay) Inria teams. They jointly address this computational problem using cloud related techniques on Microsoft Azure cloud infrastructure. The two teams bring together their complementary expertise: KerData in the area of scalable cloud data management, and Parietal in the field of neuroimaging and genetics data analysis.

In particular, KerData brings its expertise in designing solutions for optimized data storage and management for the Map-Reduce programming model. This model has recently arisen as a very effective approach to develop high-performance applications over very large distributed systems such as grids and now clouds. The computations involved in the statistical analysis designed by the Parietal team fit particularly well with this model.

4.1.2. Structural protein analysis on Nimbus clouds

Proteins are major components of the life. They are involved in lots of biochemical reactions and vital mechanisms for the living organisms. The three-dimensional (3D) structure of a protein is essential for its function and for its participation to the whole metabolism of a living organism. However, due to experimental limitations, only few protein structures (roughly, 60,000) have been experimentally determined, compared to the millions of proteins sequences which are known. In the case of structural genomics, the knowledge of the 3D structure may be not sufficient to infer the function. Thus, an usual way to make a structural analysis of a protein or to infer its function is to compare its known, or potential, structure to the whole set of structures referenced in the *Protein Data Bank* (PDB).

In the framework of the MapReduce ANR project led by KerData, we focus on the SuMo application (*Surf the Molecules*) proposed by Institute for Biology and Chemistry of the Proteins from Lyon (IBCP, a partner in the MapReduce project). This application performs structural protein analysis by comparing a set of protein structures against a very large set of structures stored in a huge database. This is a typical data-intensive application that can leverage the Map-Reduce model for a scalable execution on large-scale distributed platforms. Our goal is to explore storage-level concurrency-oriented optimizations to make the SuMo application scalable for large-scale experiments of protein structures comparison on cloud infrastructures managed using the Nimbus IaaS toolkit developed at Argonne National Lab (USA).

If the results are convincing, then they can immediately be applied to the derived version of this application for drug design in an industrial context, called MED-SuMo, a software managed by the MEDIT SME (also a partner in this project). For pharmaceutical and biotech industries, such an implementation run over a cloud computing facility opens several new applications for drug design. Rather than searching for 3D similarity into biostructural data, it will become possible to classify the entire biostructural space and to periodically update all derivative predictive models with new experimental data. The applications in that complete chemo-proteomic vision concern the identification of new druggable protein targets and thereby the generation of new drug candidates.

4.1.3. 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 efficiently record and visualize data during the simulation without impacting the performance of the computation generating that data. Conventional practice consists in storing data on disk, moving it off-site, reading it into a workflow, and analyzing it. It becomes increasingly harder to use because of the large data volumes generated at fast rates, in contrast to limited back-end speeds. 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 (http://www.ncsa.illinois.edu/BlueWaters/), which aims to build one of the most powerful supercomputers in the world when it comes online in 2012.

In this context, the KerData project-team started to explore ways to remove the limitations mentioned above through a collaborative work in the framework of the Joint Inria-UIUC Lab for Petascale Computing (JLPC, Urbana-Champaign, Illinois, USA), whose research activity focuses on the Blue Waters project. As a starting point, we are focusing on a particular tornado simulation code called CM1 (Cloud Model 1), which is intended to be run on the Blue Waters machine. Preliminary investigation demonstrated the inefficiency of the current I/O approaches, which typically consists in periodically writing a very large number of small files. This causes burst of I/O in the parallel file system, leading to poor performance and extreme variability (jitter) compared to what could be expected from the underlying hardware. 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 JLPC, we started to address those challenges through an approach based on dedicated I/O cores.

MESCAL Project-Team

4. Application Domains

4.1. 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.

Nowadays, analysis are done on huge heterogeneous data set. For example, demographic data sets at nuts 5 level, represent more than 100.000 territorial units with 40 social attributes. Many algorithms of spatial analysis, in particular potential analysis are quadratic in the size of the data set. Then adapted methods are needed to provide "user real time" analysis tools.

4.2. Wireless Networks

Participants: Bruno Gaujal, Corinne Touati, 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.

4.3. Cloud and Desktop Computing

Participants: Derrick Kondo, Arnaud Legrand, Olivier Richard.

The research of MESCAL on desktop grids has been very active and fruitful during the evaluation period. The main achievements concern the collection and statistical exploitation of traces in volunteer computing systems. Such models have enabled to optimize the behavior of volunteer computing systems or to extend the scope of their applicability. Such traces have also been used in SIMGRID to simulate volunteer computing systems at unprecedented scale.

MOAIS Project-Team

4. Application Domains

4.1. Outline

The scientific methodology of MOAIS consists in:

- designing algorithms with provable performance on generic theoretical models. In particular we develop randomized algorithms for distributed scheduling and approximate multi-objective optimization theory..
- implementing and evaluating those algorithms with our main softwares:
  - Kaapi for fine grain scheduling of compute-intensive applications;
  - FlowVR for coarse-grain scheduling of interactive applications;
  - TakTuk, a tool for large scale remote executions deployment.
  - Triva, for the visualization of heterogeneous parallel executions.
  - KRASH, to generate reproducible CPU load on many-cores machines.
- customizing our softwares for their use in real applications studied and developed by other partners. Applications are essential to the validation and further development of MOAIS results. Application fields are: virtual reality and scientific computing (simulation, visualization, combinatorial optimization, biology, computer algebra). Depending on the application the target architecture ranges from MPSoCs (multi-processor system on chips), multicore and GPU units to clusters and heterogeneous grids. In all cases, the performance is related to the efficient use of the available, often heterogeneous, parallel resources.

MOAIS research is not only oriented towards theory but also focuses on applicative software and hardware platforms developed with external partners. Significant efforts are made to build, manage and maintain these platforms. We are involved with other teams in four main platforms:

- SOFA, a real-time physics simulation engine (http://www.sofa-framework.org/;
- Grimage, a 3D modeling and high performance 3D rendering platform (http://www.inrialpes.fr/grimage) and its evolution with the new Kinovis platform.
- Digitalis, a 780 core cluster based on Intel Nehalem processors and Infiniband network. Digitalis is used both for batch computations and interactive applications;
- Grid'5000, the exprimental national grid (http://www.grid5000.fr/).

4.2. Virtual Reality

Participants: Thierry Gautier, Bruno Raffin, Jean-Louis Roch.

We are pursuing and extending existing collaborations to develop virtual reality applications on PC clusters and grid environments:

- Real time 3D modeling. An on-going collaboration with the MORPHEO project focuses on developing solutions to enable real time 3D modeling from multiple cameras using a PC cluster. This work is tightly coupled to the FlowVR software. Our recent developments take two main directions:
  - to provide the user a high level of interaction and immersion in the mixed reality environment. This work is focused on the Grimage platform and its successor, the new equipex Kinovis managed by Morpheo team. The camera position and orientation need to be precisely known at anytime, requiring to develop on-line calibration approaches. The background subtraction cannot anymore be based on a static background learning for the mobile camera, required here too new algorithms.

- Distributed collaboration across distant sites. In the context of the ANR DALIAn we have developped a collaborative application where multiple users, distributed in several sites each using a real time 3D modeling platform, can meet in a virtual world with a user in Grenoble also using a similar platform. The main issues are related to data transfers that need to be carefully managed to ensure a good latency while keeping a good quality, and the development of new interaction paradigms. Focusing on distributed scientific simulation, we extend those technologies in the context of the FVNANO and PetaFlow contracts.

- Real time physical simulation. We are collaborating with the Imagine project on the SOFA simulation framework. Marie Durand a Ph.D. co-advised by François Faure (IMAGINE) and Bruno Raffin, works on parallelizing SOFA using the KAAPI programming environment. The challenge is to provide SOFA with a parallelization that is efficient (real-time) while not being invasive for SOFA programmers (usually not parallel programmer). We developed a first version using the Kaapi environment for SMP machines that relies on a mix of work-stealing and dependency graph analysis and partitioning. A second version targets machines with multiples CPUs and multiple GPUs. We extended the initial framework to support a work stealing based load balancing between CPUs and GPUs. It required to extend Kaapi to support heterogeneous tasks (GPU and CPU ones) and to adapt the work stealing strategy to limit data transfers between CPUs and GPUs (the main bottleneck for GPU computing).
- Distant collaborative work. We conduct experiments using FlowVR for running applications on Grid environments. Two kinds of experiments will be considered: collaborative work by coupling two or more distant VR environments ; large scale interactive simulation using computing resources from the grid. For these experiments, we are collaborating with the LIFO and the LABRI.
- Parallel cache-oblivious algorithms for scientific visualization. In collaboration with the CEA DAM, we have developed a cache-oblivious algorithm with provable performance for irregulars meshes. Based on this work, we are studying parallel algorithms that take advantage of the shared cache usually encountered on multi-core architectures (L3 shared cache) and of harware accelerators. In collaboration with EDF, we develop new parallel algorithms for scientific vizualization (eg VTK) on multicore (phD thesis of Mathias Ettinger). We are also considering adaptive algorithms to take advantage of the new trend of computers to integrate several computing units that may have different computing abilities (multicore arithmetic and graphical processing units, eventually integrated on one chip). We study balancing workload on multi GPU and CPU architectures for scientific visualization problems.

4.3. Code Coupling and Grid Programming

Participants: François Broquedis, Thierry Gautier, Jean-Louis Roch, Vincent Danjean, Frédéric Wagner.

Code coupling aim is to assemble component to build distributed applications by reusing legacy code. The objective here is to build high performance applications for cluster and grid infrastructures.

- **Grid programming model and runtime support.** Programming the grid is a challenging problem. The MOAIS Team has a strong knowledge in parallel algorithms and develop a runtime support for scheduling grid program written in a very high level interface. The parallelism from recursive divide and conquer applications and those from iterative simulation are studied. Scheduling heuristics are based on online work stealing for the former class of applications, and on hierarchical partitioning for the latter. The runtime support provides capabilities to hide latency by computation thanks to a non-blocking one-side communication protocol and by re-ordering computational tasks.
- **Grid application deployment.** To test grid applications, we need to deploy and start programs on all used computers. This can become difficult if the real topology involves several clusters with firewall, different runtime environments, etc. The MOAIS Team designed and implemented a new tool called `karun` that allows a user to easily deploy a parallel application wrote with the KAAPI software. This KAAPI tool relies on the `TakTuk` software to quickly launch programs on all nodes. The user only

needs to describe the hierarchical networks/clusters involved in the experiment with their firewall if any.

- **Visualization of grid applications execution.** The analysis of applications execution on the grid is challenging both because of the large scale of the platform and because of the heterogeneous topology of the interconnections. To help users to understand their application behavior and to detect potential bottleneck or load unbalance, the MOAIS team designed and implemented a tool named **Triva**. This tool proposes a new three dimensional visualization model that combines topological information to space time data collected during the execution. It also proposes an aggregation mechanism that eases the detection of application load unbalance.

4.4. Safe Distributed Computations

Participants: Vincent Danjean, Thierry Gautier, Clément Pernet, Jean-Louis Roch.

Large scale distributed platforms, such as the GRID and Peer-to-Peer computing systems, gather thousands of nodes for computing parallel applications. At this scale, component failures, disconnections (fail-stop faults) or results modifications (malicious faults) are part of operation, and applications have to deal directly with repeated failures during program runs. Indeed, since failure rate in such platform is proportional to the number of involved resources, the mean time between failure is dramatically decreased on very large size architectures. Moreover, even if a middleware is used to secure the communications and to manage the resources, the computational nodes operate in an unbounded environment and are subject to a wide range of attacks able to break confidentiality or to alter the resources or the computed results. Beyond fault-tolerancy, yet the possibility of massive attacks resulting in an error rate larger than tolerable by the application has to be considered. Such massive attacks are especially of concern due to Distributed Denial of Service, virus or Trojan attacks, and more generally orchestrated attacks against widespread vulnerabilities of a specific operating system that may result in the corruption of a large number of resources. The challenge is then to provide confidence to the parties about the use of such an unbound infrastructure. The MOAIS team addresses two issues:

- fault tolerance (node failures and disconnections): based on a global distributed consistent state , for the sake of scalability;
- security aspects: confidentiality, authentication and integrity of the computations.

Our approach to solve those problems is based on the efficient checkpointing of the dataflow that described the computation at coarse-grain. This distributed checkpoint, based on the local stack of each work-stealer process, provides a causally linked representation of the state. It is used for a scalable checkpoint/restart protocol and for probabilistic detection of massive attacks.

Moreover, we study the scalability of security protocols on large scale infrastructures. Within the SHIVA contract (global competitiveness cluster Minalogic in Grenoble) and in collaboration with C-S company, the Ph.D. of Ludovic Jacquin (coadvised with the PLANETE EPI) we developed a high-rate systematic ciphering platform based on the coupling of a multicore architecture with security components (FPGA and smart card) developed by industrial partners.

4.5. Embedded Systems

Participants: Jean-Louis Roch, Guillaume Huard, Denis Trystram, Vincent Danjean.

To improve the performance of current embedded systems, Multiprocessor System-on-Chip (MPSoC) offers many advantages, especially in terms of flexibility and low cost. Multimedia applications, such as video encoding, require more and more intensive computations. The system should be able to exploit the resources as much as possible to save power and time. This challenge may be addressed by parallel computing coupled with performant scheduling. On-going work focuses on developing the scheduling and monitoring technologies developed in MOAIS for embedded systems.

In the framework of our cooperation with STM (Miguel Santana) and within the SocTrace project, we are developping tools to manage distributed large scale traces. We especially focus on visualization, developping visual aggregation techniques (Phd Damien Dosimont, started in 2/2012 advised by Guillaume Huard in collaboration with Jean-Marc Vincent).

ROMA Team

3. Application Domains

3.1. Application of sparse direct solvers

Sparse direct (multifrontal) solvers in distributed-memory environments have a wide range of applications as they are used at the heart of many numerical methods in simulation: 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 4.1) are: structural mechanics, biomechanics, medical image processing, tomography, geophysics, ad-hoc networking modeling (e.g., Markovian processes), electromagnetics, fluid dynamics, econometric models, oil reservoir simulation, magneto-hydro-dynamics, chemistry, acoustics, glaciology, astrophysics, circuit simulation.

RUNTIME Project-Team

4. Application Domains

4.1. Application Domains

HPC, simulation

The RUNTIME group is working on the design of efficient runtime systems for parallel architectures. We are currently focusing our efforts on High Performance Computing applications that merely implement numerical simulations in the field of Seismology, Weather Forecasting, Energy, Mechanics or Molecular Dynamics. These time-consuming applications need so much computing power that they need to run over parallel machines composed of several thousands of processors.

Because the lifetime of HPC applications often spreads over several years and because they are developed by many people, they have strong portability constraints. Thus, these applications are mostly developed on top of standard APIs (e.g. MPI for communications over distributed machines, OpenMP for shared-memory programming). That explains why we have long standing collaborations with research groups developing parallel language compilers, parallel programming environments, numerical libraries or communication software. Actually, all these “clients” are our primary target.

Although we are currently mainly working on HPC applications, many other fields may benefit from the techniques developed by our group. Since a large part of our efforts is devoted to exploiting multicore machines and GPU accelerators, many desktop applications could be parallelized using our runtime systems (e.g. 3D rendering, etc.).

DANTE 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 makes 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 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-Christopge Lucet (Professeur des université Paris VII ? Praticien hospitalier APHP).

4.2. Network Science / Complex networks

In the last ten years, the study of complex networks has received an important boost with large interdisciplinary efforts aimed at their analysis and characterization. Two main points explain this large activity: on the one hand, many systems coming from very different disciplines (from biology to computer science) have a convenient representation in terms of graphs; on the other hand, the ever-increasing availability of large data sets and computer power have allowed their storage and manipulation. Many maps have emerged, describing many networks of practical interest in social science, critical infrastructures, networking, and biology. The DANTE project targets the study of dynamically evolving networks, from the point both of their structure and of the dynamics of processes taking place on them.

DIONYSOS Project-Team (section vide)

DISTRIBCOM Project-Team

4. Application Domains

4.1. Telecommunication network management

The management of telecommunication networks is traditionally a human performed activity that covers the five FCAPS functions: Fault management, network Configuration, Accounting, Performances and Security. This simple classification has exploded in the last decade, under the pressure of several phenomena. The first one concerns the growth in size and complexity of networks, with the emergence of new (possibly virtual) operators, the multiplication of vendors, new core and (wireless) access technologies, the variety of terminal devices, the convergence of phone/computer/radio/TV networks, the multiplication of services over the top, the necessity to provide QoS for a wide variety of traffic demands, etc. As a consequence, the management task is reaching the limits of human operators and demands automation. It is estimated that telecommunication companies spend over 50% of their manpower on management tasks. They naturally want to reduce it and dedicate their effort to the design and offer of innovative services, where the added value is more important (as witnessed by the success of some over-the-top companies). The result of these trends is that network management now covers a much wider variety of problems, for which automatic solutions are requested. This takes the name of self-management, or autonomic management: one wishes to manage networks by high-level objectives, and networks should be able to adapt themselves automatically to fulfill these objectives. DistribCom is contributing to this field with its background on the modeling of distributed/concurrent systems, and its expertise in distributed algorithms. Networks are perfect examples of large distributed and concurrent systems, with specific features like the dynamicity (their structure evolves) and a hierarchical structure (multiple layers, multiple description granularities). We have proposed model-based distributed algorithms to solve problems like failure diagnosis, negotiation of QoS (quality of service) parameters, parameter optimization, graceful shutdown of OSPF routers for maintenance operations... The present activities in this domain are related to the joint diagnosis for access network + core network + services, within the European IP UniverSelf. The challenges cover self-modelling methods (how to obtain the network model that is used by the management algorithms), active diagnosis methods that both adapt the scope of their network model and perform tests to explain a fault situation, and self-healing methods.

4.2. Web services and active structured documents

Keywords: Active documents, Web services, choreographies, orchestrations, QoS.

Web services architectures are usually composed of distant services, assembled in a composite framework. This raises several practical issues: one of them is how to choose services, assemble them, and coordinate their executions in a composite framework. Another issue is to guarantee good properties of a composite framework (safety but also QoS properties). All this has to be done in a context where a distant service provided by a subcontractor is only perceived as an interface, specifying legal inputs and outputs, and possibly a quality contract. The standard in industry for Web-services is now BPEL [43] but most of the problems listed above are untractable for this language. Composition of services can also be performed using choreography languages such as ORC [55]. The implementation of orchestration and choreography description languages raises a number of difficulties related to efficiency, clean semantics, and reproducibility of executions, issues of composite QoS associated with orchestrations. We develop studies in these areas, with the aim of proposing service composition frameworks equipped with tools to specify, but also to monitor and analyze the specified architectures. Another issue is the convergence between data and workflows. Web Services architectures are frequently considered exclusively as workflows, or as information systems. Many approaches to Web Service orchestration and choreography abstract data away. Symmetrically, modern approaches to Web data management typically based on XML and Xqueries rely on too simplistic forms of control. We develop a line of research on Active documents. Active documents are structured data embedding references to services, which allow for the definitions of complex workflows involving data aspects. The original model was proposed by S. Abiteboul (see for instance [42]), but the concept of active document goes beyond AXML, and offers a document oriented alternative to Web services orchestrations and choreographies. This approach is in particular well adapted to the modeling of E-business processes, or information processing in organizations, etc. Our aim is to extend and promote the concept of active document. This means developing verification and composition tools for document-based architectures, considered not only as theoretical models but also as effectively running systems. To this extend, we develop an active document platform.

FUN Team (section vide)

GANG Project-Team

3. Application Domains

3.1. 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 o f such studies.

HIPERCOM Project-Team

4. Application Domains

4.1. Introduction

The HIPERCOM project-team is mainly concerned by six domains:

- wireless mobile ad hoc networks,
- services over mobile networks,
- community networks,
- vehicular networks,
- large ad hoc networks with sensor nodes,
- energy-efficient wireless sensor networks.

4.2. Wireless mobile ad hoc networks

Wireless mobile ad hoc networks, Services over mobile networks, Community Networks.

Abstract. Mobile wireless networks have numerous applications in rescue and emergency operation, military tactical networking and in wireless high speed access to the internet.

A mobile ad hoc network is a network made of a collection of mobile nodes that gather spontaneously and communicate without requiring a pre-existing infrastructure. Of course a mobile ad hoc network use a wireless communication medium. They can be applied in various contexts:

- military;
- rescue and emergency;
- high speed access to internet.

The military context is the most obvious application of mobile ad hoc networks.

Soldiers invading a country won't susbscribe in advance to the local operator. On the reverse side, home units won't use their local operators firstly because they will likely be disrupted in the first hours of the conflict, and secondly because a wireless communication via an operator is not stealth enough to protect the data and the units. In Checheny, a general has been killed by a missile tracking the uplink signal of his portable phone.

The rescue context is halfway between military and civilian applications. In the september 11 disaster, most of the phone base station of the area have knocked out in less than twenty minutes. The remaining base stations were unable to operate because they could not work in ad hoc mode. The Wireless Emergency Rescue Team recommanded afterward that telecom operators should provide ad hoc mode for their infrastructure in order to operate in emergency situation in plain cooperation with police, firemen and hospital networks.

Mobile ad hoc network provide an enhanced coverage for high speed wireless access to the internet. The now very popular WLAN standard, WiFi, provides much larger capacity than mobile operator networks. Using a mobile ad hoc network around hot spots will offer high speed access to much larger community, including cars, busses, trains and pedestrians.

4.3. Services over mobile networks

Abstract. New wireless network calls for new services that fullfil the requirement in terms of mobility and capacity.

The generalization of a new generation of mobile networks calls for a new set of services and applications. For example:

- Indoor and outdoor positioning
- Service discovery and localisation
- Multicast and quality of services

Quality of service has become the central requirement that users expect from a network. High throughput, service continuity are critical issue for multimedia application over the wireless internet where the bandwidth is more scarce than in the wired world. A significant issue in the ad-hoc domain is that of the integrity of the network itself. Routing protocols allow, according to their specifications, any node to participate in the network - the assumption being that all nodes are behaving well and welcome. If that assumption fails - then the network may be subject to malicious nodes, and the integrity of the network fails. An important security service over mobile networks is to ensure that the integrity of the network is preserved even when attacks are launched against the integrity of the network.

4.4. Community Networks

Abstract. There is an increasing demand to deploy network within a community, rural or urban, with cabled or wireless access.

Community networks or citizen network are now frequent in big cities. In America most of the main cities have a community network. A community network is using the communication resource of each member (ADSL, Cable and wireless) to provide a general coverage of a city. Pedestrian in the street or in city mails can communicate via a high speed mobile mesh network. This new trend now appears in Europe with many experiments of the OLSR routing protocol in Paris, Lille, Toulouse, Berlin, Bruxelles, Seattle. The management of such networks is completely distributed and makes them very robust to faults. There is room for smart operators in this business.

4.5. Vehicular Networks

Abstract. Intelligent transport systems require efficient wireless telecommunications.

Vehicular ad hoc networks (VANET) are based on short- to medium-range transmission systems that support both vehicle-to-vehicle and vehicle-to-roadside communications. Vehicular networks will enable vehicular safety applications (safety warnings) as well as non-safety applications (real-time traffic information, routing support, mobile entertainment, and many others). We are interested in developing an efficient routing protocol that takes advantage of the fixed network infrastructure deployed along the roads. We are also studying MAC layer issues in order to provide more priority for security messages which have stringent delivery constraints.

4.6. Large ad hoc networks with sensor nodes

Abstract. Large autonomous wireless sensors in the internet of the things need very well tuned algorithms.

Self-organization is considered as a key element in tomorrow's Internet architecture. A major challenge concerning the integration of self-organized networks in the Internet is the accomplishment of light weight network protocols in large ad hoc environments.

In this domain, Hipercom's activity with wireless sensor nodes in collaboration with the Freie Universitaet in Berlin explores various solutions, including extensions of OLSR (for example DHT-OLSR) using programmable sensor nodes co-designed by the Freie Universitaet, and provides one of the largest testbeds of this kind, to date.

4.7. Energy efficient wireless sensor networks

Abstract. Energy efficiency is a key property in wireless sensor networks.

Various techniques are used to contribute to energy efficiency. In the OCARI network, an industrial wireless sensor network, we have designed and implemented an energy efficient routing protocol and a node activity scheduling algorithm allowing router nodes to sleep. We have applied a cross-layering approach allowing the optimization of MAC and network protocols taking into account the application requirements and the environment in which the network operates. This activity has been done in collaboration with our partners EDF, LIMOS and TELIT.

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 service 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,
- ambiant environments.

MAESTRO Project-Team

4. Application Domains

4.1. Application Domains

MAESTRO's main application area is networking and in particular, modeling, performance evaluation, optimization and control of protocols and network architectures. It includes:

- 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.

MASCOTTE Project-Team

4. Application Domains

4.1. Application Domains

In the last year the main application domain of the project remained Telecommunications. 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 SMEs like 3-Roam and Avisto .

MASCOTTE is mainly interested in the design and management of heterogeneous networks. The project has kept working on the design of backbone networks (optical networks, radio networks, IP networks).

The project has also been working on routing algorithms such as dynamic and compact routing schemes in the context of the FP7 EULER leaded by Alcatel-Lucent Bell-Labs (Belgium). It also studied the evolution of the routing in case of any kind of topological modifications (maintenance operations, failures, capacity variations, etc.). Finally, an emphasis is done on green networks with low power consumption. This work is in collaboration with Orange Labs and the SME 3-Roam and partly supported by the ANR DIMAGREEN.

PLANETE Project-Team

4. Application Domains

4.1. Next Generation Networks

The next-generation network must overcome the limitations of existing networks and allow adding new capabilities and services. Future networks should be available anytime and anywhere, be accessible from any communication device, require little or no management overhead, be resilient to failures, malicious attacks and natural disasters, and be trustworthy for all types of communication traffic. Studies should therefore address a balance of theoretical and experimental researchs that expand the understanding of large, complex, heterogeneous networks, design of access and core networks based on emerging wireless and optical technologies, and continue the evolution of Internet. On the other hand, it is also highly important to design a next-generation Internet which we will call the "Future Internet" from core functionalities in order to ensure security and robustness, manageability, utility and social need, new computing paradigms, integration of new network technologies and higher-level service architectures.

To meet emerging requirements for the Internet's technical architecture, the protocols and structures that guide its operation require coordinated, coherent redesign. A new approach will require rethinking of the network functions and addressing a range of challenges. These challenges include, but are not limited to, the following examples:

- New models for efficient data dissemination;
- Coping with intermittent connectivity;
- The design of secured, privacy protecting, and robust networked systems;
- Understanding the Internet behavior;
- Building network evaluation platforms.

The following research directions are essential building blocks we are contributing to the future Internet architecture.

Towards Data-Centric Networking

From the Internet design, back to 1970, the resources to be addressed and localized are computers. Indeed, at that time there were few machines interconnected, and nobody believed this number would ever be larger than a few tens of thousand of machines. Moreover, those machines where static machines with well identified resources (e.g., a given hierarchy of files) that were explicitly requested by the users. Today, the legacy of this architecture is the notion of URLs that explicitly address specific resources on a specific machine. Even if modern architectures use caches to replicate contents with DNS redirection to make those caches transparent to the end-users, this solution is only an hack that do not solve today's real problem: Users are only interested in data and do not want anymore to explicitly address where those data are. Finding data should be a service offered by the network. In this context of data-centric network, which means that the network architecture is explicitly built to transparently support the notion of content, a data can be much more than a simple content. In such a network you can, of course, request a specific file without explicitly specifying its location, the network will transparently return the closest instance of the content. You can also request a specific service from a person without knowing its explicit network location. This is in particular the case of a VoIP or an instant messaging conversation. A data-centric architecture is much more than a simple modification of the naming scheme currently used in the Internet. It requires a major rethinking a many fundamental building blocks of the current Internet. Such networking architecture will however allow seamless handling of the tricky problematic of *episodic connectivity*. It also shifts the focus from transmitting data by geographic location, to *disseminating* it via named content. In the Planète project-team, we start to work on such data-centric architectures as a follow-up and federating axe for three of our current activities (adaptive multimedia transmission protocols for heterogeneous networks, data dissemination paradigms and peer-to-peer systems). It is important to study such data-centric architectures considering in particular the corresponding naming problem, routing and resource allocation, reliable transport, data security and authentication, content storage.

Today's Internet is characterized by high node and link heterogeneity. Nodes may vary substantially in terms of their processing, storage, communication, and energy capabilities. They may also exhibit very different mobility characteristics, from static nodes to nodes that are considerably mobile (e.g., vehicles). Links may be wired or wireless and thus operate at widely varying rates and exhibit quite different reliability characteristics. One of the challenges of data-centric architecture is to provide access to data anytime anywhere in the presence of high degree of heterogeneity. This means that the network will not be connected all the time, due to a number of factors such as node mobility, link instability, power-aware protocols that, for example, turn nodes off periodically, etc. Additionally, disconnections may last longer than what "traditional" routing protocols (e.g., MANET routing) can handle. These types of network, a.k.a, intermittently connected networks, or even episodically connected networks, have recently received considerable attention from the networking research community. Several new routing paradigms have been proposed to handle possibly frequent, long-lived disconnections. However, a number of challenges remain, including: (1) The support of scalable and transparent integration with "traditional" routing mechanisms including wired infrastructure, infrastructure-based wireless and MANET routing. (2) The study of heuristics for selecting forwarding nodes (e.g., based on node's characteristics such as node's speed, node's resources, sociability level, node's historic, etc. (3) The design of unicast and multicast transmission algorithms with congestion and error control algorithms tailored for episodically connected networks and taking into account the intrinsic characteristics of flows. (4) The design of incentive-based mechanisms to ensure that nodes forward packets while preventing or limiting the impact of possible misbehaving nodes. The solutions proposed, which are likely to extensively use cross-layer mechanisms, will be evaluated using the methodology and the tools elaborated in our new *Experimental Platform* research direction.

On the other hand, multicast/broadcast content delivery systems are playing an increasingly important role in data-centric networking. Indeed, this is an optimal dissemination technology, that enables the creation of new commercial services, like IPTV over the Internet, satellite-based digital radio and multimedia transmission to vehicles, electronic service guide (ESG) and multimedia content distribution on DVB-H/SH networks. This is also an efficient way to share information in WiFi, WiMax, sensor networks, or mobile ad hoc infrastructures. Our goal here is to take advantage of our strong background in the domain to design an *efficient, robust (in particular in case of tough environments) and secure (since we believe that security considerations will play an increasing importance) broadcasting system*. We address this problem by focusing on the following activities: (1) The protocols and applications that enable the high level control of broadcasting sessions (like the FLUTE/ALC sessions) are currently missing. The goal is to enable the content provider to securely control the underlying broadcasting sessions, to be able to launch new sessions if need be, or prematurely stop an existing session and to have feedback and statistics on the past/current deliveries. (2) The AL-FEC building block remains the cornerstone on which the whole broadcasting system relies. The goal is to design and evaluate new codes, capable of producing a large amount of redundancy (thereby approaching rateless codes), over very large objects, while requiring a small amount of memory/processing in order to be used on lightweight embedded systems and terminals. (3) The security building blocks and protocols that aim at providing content level security, protocol level security, and network level security must be natively and seamlessly integrated. This is also true of the associated protocols that enable the initialization of the elementary building blocks (e.g. in order to exchange security parameters and keys). Many components already exist. The goal here is to identify them, know how to optimally use them, and to design/adapt the missing components, if any. (4) It is important seamlessly integrated these broadcasting systems to the Internet, so that users can benefit from the service, no matter where and how he is attached to the network. More precisely we will study the potential impacts of a merge of the broadcasting networks and the Internet, and how to address them. For instance there is a major discrepancy when considering flow control aspects, since broadcasting networks are using a constant bit rate approach while the Internet is congestion controlled.

When a native broadcasting service is not enabled by the network, data should still be able to be disseminated to a large population in a scalable way. A peer-to-peer architecture supports such an efficient data dissemination. We have gained a fundamental understanding of the key algorithms of BitTorrent on the Internet. We plan to continue this work in two directions. First, we want to study how a peer-to-peer architecture can be natively supported by the network. Indeed, the client-server architecture is not robust to increase in load. The consequence is that when a site becomes suddenly popular, it usually becomes unreachable. The peer-to-peer architecture is robust to increase in load. However, a native support in the network of this architecture is a hard problem as it has implications on many components of the network (naming, addressing, transport, localization, etc.). Second, we want to evaluate the impact of wireless and mobile infrastructures on peer-to-peer protocols. This work has started with the European project Expeshare. The wireless medium and the mobility of nodes completely change the properties of peer-to-peer protocols. The dynamics becomes even more complex as it is a function of the environment and of the relative position of peers.

Network security and Privacy

The Internet was not designed to operate in a completely open and hostile environment. It was designed by researchers that trust each other and security at that time was not an issue. The situation is quite different today and the Internet community has drastically expanded. The Internet is now composed of more than 300 millions computers worldwide and the trust relationship has disappeared. One of the reason of the Internet success is that it provides ubiquitous inter-connectivity. This is also one of the its main weakness since it allows to launch attacks and to exploit vulnerabilities in a large-scale basis. The Internet is vulnerable to many different attacks, for example, Distributed Denial-of Service (DDoS) attacks, epidemic attacks (Virus/Worm), spam/phishing and intrusion attacks. The Internet is not only insecure but it also infringes users' privacy. Those breaches are due to the Internet protocols but also to new applications that are being deployed (VoIP, RFID,...). A lot of research is required to improve the Internet security and privacy. For example, more research work is required to understand, model, quantify and hopefully eliminate (or at least mitigate) existing attacks. Furthermore, more and more small devices (RFIDs or sensors) are being connected to the Internet. Current security/cryptographic solutions are too expensive and current trust models are not appropriate. New protocols and solutions are required : security and privacy must be considered in the Internet architecture as an essential component. The whole Internet architecture must be reconsidered with security and privacy in mind. Our current activities in this domain are on security in wireless, ad hoc and sensor networks, mainly the design of new key exchange protocols and of secured routing protocols. We also work on location privacy techniques, authentication cryptographic protocols and opportunistic encryption. We plan to continue our research on wireless security, and more specifically on WSN and RFID security focusing on the design of real and deployable systems. We started a new research topic on the security of the Next-Generation Internet. The important goal of this new task is to rethink the architecture of the Internet with security as a major design requirement, instead of an after-thought.

Wireless Sensor Networks: A lot of work has been done in the area of WSN security in the last years, but we believe that this is still the beginning and a lot of research challenges need to be solved. On the one hand it is widely believed that the sensor networks carry a great promise: Ubiquitous sensor networks will allow us to interface the physical environment with communication networks and the information infrastructure, and the potential benefits of such interfaces to society are enormous, possibly comparable in scale to the benefits created by the Internet. On the other hand, as with the advent of the Internet, there is an important associated risk and concern: How to make sensor network applications resilient and survivable under hostile attacks? We believe that the unique technical constraints and application scenarios of sensor networks call for new security techniques and protocols that operate above the link level and provide security for the sensor network application as a whole. Although this represents a huge challenge, addressing it successfully will result in a very high pay-off, since targeted security mechanisms can make sensor network operation far more reliable and thus more useful. This is the crux of our work. Our goal here is to design new security protocols and algorithms for constrained devices and to theoretically prove their soundness and security. Furthermore, to complement the fundamental exploration of cryptographic and security mechanisms, we will simulate and evaluate these mechanisms experimentally.

RFID: As already mentioned, the ubiquitous use of RFID tags and the development of what has become termed "the Internet of things" will lead to a variety of security threats, many of which are quite unique to RFID deployment. Already industry, government, and citizens are aware of some of the successes and some of the limitations or threats of RFID tags, and there is a great need for researchers and technology developers to take up some of daunting challenges that threaten to undermine the commercial viability of RFID tags on the one hand, or to the rights and expectations of users on the other. We will focus here on two important issues in the use of RFID tags: (1) *Device Authentication*: allows us to answer several questions such as: Is the tag legitimate? Is the reader a tag interacts with legitimate? (2) *Privacy*: is the feature through which information pertaining to a tag's identity and behavior is protected from disclosure by unauthorized parties or by unauthorized means by legitimate parties such as readers. In a public library, for example, the information openly communicated by a tagged book could include its title or author. This may be unacceptable to some readers. Alternatively, RFID- protected pharmaceutical products might reveal a person's pathology. Turning to authenticity, if the RFID tag on a batch of medicines is not legitimate, then the drugs could be counterfeit and dangerous. Authentication and privacy are concepts that are relevant to both suppliers and consumers. Indeed, it is arguable that an RFID deployment can only be successful if all parties are satisfied that the integrity between seller and buyer respects the twin demands of authentication and privacy. Our main goal here, therefore, is to propose and to prototype the design of cryptographic algorithms and secure protocols for RFID deployment. These algorithms and protocols may be used individually or in combination, and we anticipate that they will aid in providing authentication or privacy. One particular feature of the research in the RFID-AP project is that the work must be practical. Many academic proposals can be deeply flawed in practice since too little attention has been paid to the realities of implementation and deployment. This activity will therefore be notable for the way theoretical work will be closely intertwined with the task of development and deployment. The challenges to be addressed in the project are considerable. In particular there are demanding physical limits that apply to the algorithms and protocols that can be implemented on the cheapest RFID tags. While there often exist contemporary security solutions to issues such as authentication and privacy, in an RFID-based deployment they are not technically viable. And while one could consider increasing the technical capability of an RFID-tag to achieve a better range of solutions, the solution is not economically viable.

Next Generation Internet Security: The current Internet has reached its limits; a number of research groups around the world are already working on future Internet architectures. The new Internet should have built-in security measures and support for wireless communication devices, among other things. A new network design is needed to overcome unwanted traffic, malware, viruses, identity theft and other threats plaguing today's Internet infrastructure and end hosts. This new design should also enforce a good balance between privacy and accountability. Several proposals in the area have been made so far, and we expect many more to appear in the near future. Some mechanisms to mitigate the effects of security attacks exist today. However, they are far from perfect and it is a very open question how they will behave on the future Internet. Cyber criminals are very creative and new attacks (e.g. VoIP spam, SPIT) appear regularly. Furthermore, the expectation is that cyber criminals will move into new technologies as they appear, since they offer new attack opportunities, where existing countermeasures may be rendered useless. The ultimate goal of this research activity is to contribute to the work on new Internet architecture that is more resistant to today's and future security attacks. This goal is very challenging, since some of future attacks are unpredictable. We are analyzing some of the established and some of the new architectural proposals, attempting to identify architectural elements and patterns that repeat from one architectural approach to another, leading to understanding how they impact the unwanted traffic issue and other security issues. Some of the more prominent elements are rather easy to identify and understand, such as routing, forwarding, end-to-end security, etc. Others may well be much harder to identify, such as those related to data-oriented networking, e.g., caching. The motivation for this work is that the clean slate architectures provide a unique opportunity to provide built in security capabilities that would enable the prevention of phenomenon like unwanted traffic. New architectures will most likely introduce additional name-spaces for the different fundamental objects in the network and in particular for routing objects. These names will be the fundamental elements that will be used by the new routing architectures and security must be a key consideration when evaluating the features offered by these new name-spaces.

Network Monitoring

The Planète project-team contributes to the area of network monitoring. Our focus is on the monitoring of the Internet for the purpose of access quality assessment, problem detection and troubleshooting. Indeed, in the absence of an advanced management and control plan in the Internet, and given the simplicity of the service provided by the core of the network and the increase in its heterogeneity, it is nowadays common that users experience a service degradation and are unable to understand the reasons for the access quality they perceive. Problems at the access can be in the form of a pure disconnection, a decrease in the bandwidth or an increase in the delay or loss rate of packets. Service degradation can be caused by protocol anomalies, an attack, an increase in the load, or simply a problem at the source or destination machines. Actually, it is not easy to diagnose the reasons for service degradation. Basic tools exist as ping and trace-route, but they are unable to provide detailed answers on the source of the problem nor on its location. From operator point of view, the situation is not better since an operator has only access to its own network and can hardly translate local information into end-to-end measurements. The increase in the complexity of networks as is the case of wireless mesh networks will not ease the life of users and operators. The purpose of our work in this direction is to study to which extent one can troubleshoot the current Internet and estimate the quality at the access either with end-to-end solutions or core network solutions. Our aim is to propose an architecture that allows end-users by collaborating together to infer the reasons for service degradation and to estimate the quality of access they perceive. This architecture can be purely end-to-end or can rely on some information from the core of the network as BGP routing information. We will build on this study to understand the limitations in the current Internet architecture and propose modifications that will ease the troubleshooting and make it more efficient in future network architectures. The proposed architecture will be the subject of validation over large scale experimental platforms as PlanetLab and OneLab.

Experimental Environment for future Internet architecture

The Internet is relatively resistant to fundamental change (differentiated services, IP multicast, and secure routing protocols have not seen wide-scale deployment). A major impediment to deploy these services is the need for coordination: an Internet service provider (ISP) that deploys the service garners little benefit until other domains follow suit. Researchers are also under pressure to justify their work in the context of a federated network by explaining how new protocols could be deployed one network at a time, but emphasizing incremental deployability does not necessarily lead to the best architecture. In fact, focusing on incremental deployment may lead to solutions where each step along the path makes sense, but the end state is wrong. The substantive improvements to the Internet architecture may require fundamental change that is not incrementally deployable.

Network virtualisation has been proposed to support realistic large scale shared experimental facilities such as PlanetLab and GENI. We are working on this topic in the context of the European OneLab project.

Testing on PlanetLab has become a nearly obligatory step for an empirical research paper on a new network application or protocol to be accepted into a major networking conference or by the most prestigious networking journals. If one wishes to test a new video streaming application, or a new peer-to-peer routing overlay, or a new active measurement system for geo-location of internet hosts, hundreds of PlanetLab nodes are available for this purpose. PlanetLab gives the researcher login access to systems scattered throughout the world, with a Linux environment that is consistent across all of them.

However, network environments are becoming ever more heterogeneous. Third generation telephony is bringing large numbers of handheld wireless devices into the Internet. Wireless mesh and ad-hoc networks may soon make it common for data to cross multiple wireless hops while being routed in unconventional ways. For these new environments, new networking applications will arise. For their development and evaluation, researchers and developers will need the ability to launch applications on endhosts located in these different environments.

It is sometimes unrealistic to implement new network technology, for reasons that can be either technological - the technology is not yet available -, economical - the technology is too expensive -, or simply pragmatical - e.g. when actual mobility is key. For these kinds of situations, we believe it can be very convenient and powerful to resort to emulation techniques, in which real packets can be managed as if they had crossed, e.g., an ad hoc network.

In our project-team, we work to provide a realistic environment for the next generation of network experiments. Such a large scale, open, heterogeneous testbed should be beneficial to the whole networking academic and industrial community. It is important to have an experimental environment that increases the quality and quantity of experimental research outcomes in networking, and to accelerate the transition of these outcomes into products and services. These experimental platforms should be designed to support both research and deployment, effectively filling the gap between small-scale experiments in the lab, and mature technology that is ready for commercial deployment. As said above, in terms of experimental platforms, the well-known PlanetLab testbed is gaining ground as a secure, highly manageable, cost-effective world-wide platform, especially well fitted for experiments around New Generation Internet paradigms like overlay networks. The current trends in this field, as illustrated by the germinal successor known as GENI, are to address the following new challenges. Firstly, a more modular design will allow to achieve federation, i.e. a model where reasonably independent Management Authorities can handle their respective subpart of the platform, while preserving the integrity of the whole. Secondly, there is a consensus on the necessity to support various access and physical technologies, such as the whole range of wireless or optical links. It is also important to develop realistic simulators taking into account the tremendous growth in wireless networking, so to include the many variants of IEEE 802.11 networking, emerging IEEE standards such as WiMax (802.16), and cellular data services (GPRS, CDMA). While simulation is not the only tool used for data networking research, it is extremely useful because it often allows research questions and prototypes to be explored at many orders-of-magnitude less cost and time than that required to experiment with real implementations and networks.

Simulations allow a fast evaluation process, fully controlled scenarios, and reproducibility. However, they lack realism and the accuracy of the models implemented in the simulators is hard to assess. Emulation allows controlled environment and reproducibility, but it also suffers from a lack of realism. Experiments allow more realistic environment and implementations, but they lack reproducibility and ease of use. Therefore, each evaluation technique has strengths and weaknesses. However, there is currently no way to combine them in a scientific experimental workflow. Typical evaluation workflows are split into four steps: topology description and construction, traffic pattern description and injection, trace instrumentation description and configuration, and, analysis based on the result of the trace events and the status of the environment during the experimentation. To achieve the integration of experimental workflows among the various evaluation platforms, the two following requirements must be verified:

- Reproducibility: A common interface for each platform must be defined so that a same script can be run transparently on different platforms. This also implies a standard way to describe scenarios, which includes the research objective of the scenario, topology description and construction, the description of the traffic pattern and how it is injected into the scenario, the description and configuration of the instrumentation, and the evolution of the environment during the experimentation
- Comparability: As each platform has different limitations, a way to compare the conclusions extracted from experiments run on different platforms, or on the same platform but with different conditions (this is in particular the case for the wild experimental platforms) must be provided.

Benchmarking is the function that provides a method of comparing the performance of various subsystems across different environments. Both reproducibility and comparability are essential to benchmarking. In order to facilitate the design of a general benchmarking methodology, we plan to integrate and automate a networking experiments workflow within the OneLab platform. This requires that we:

- Automate the definition of proper scenario definition taking in consideration available infra-structure to the experiment.
- Automate the task of mapping the experimentation topology on top of the available OneLab topology. We propose to first focus on a simple one-to-one node and link mapping the beginning.
- Define and provide extensive instrumentation sources within the OneLab system to allow users to gather all interesting trace events for offline analysis
- Measure and provide access to "environment variables" which measure the state of the OneLab system during an experimentation
- Define an offline analysis library which can infer experimentation results and comparisons based on traces and "environment variables".

To make the use of these components transparent, we plan to implement them within a simulation-like system which should allow experiments to be conducted within a simulator and within the OneLab testbed through the same programming interface. The initial version will be based on the ns-3 programming interface.

RAP Project-Team (section vide)

SOCRATE Team

4. Application Domains

4.1. Example of SDR applications

SDR concept is not new and many research teams have been working on its implementation and use in various contexts, however two elements are in favor of Socrate's orientation towards this technology:

1. The mobile SDR technology is becoming mature. Up to now, Software-Defined Radio terminals were too expensive and power consuming for mobile terminal, this should change soon. For instance, CEA's Magali platform has demonstrated part of LTE-Advanced standard recently. It is important for applied researchers to be ready when a new technology rises up, opening to many new software issues.
2. Rhône-Alpes is a strategic place for this emerging technology with important actors such as ST-Microelectronics, CEA, Minalogic and many smaller actors in informatics for telecommunication and embedded systems.

SDR technologies enables the following scenarios:

- *Transparent radio adaptation:* Depending on the available wireless protocols in the air (e.g. Wifi versus UMTS), a terminal may choose to communicate on the cheapest, or the fastest channel.
- *Radio resource allocation:* In order to minimize expensive manual cell planning and achieve "tighter" frequency reuse patterns, resulting in improved system spectral efficiency, dynamic radio resource management is a promising application of SDR.
- *White space:* By sensing the air, a terminal is able to communicate using a particular frequency which is not used even if it is reserved for another kind of application.
- *Cooperation:* Using the neighboring terminals, a user can reduce power consumption by using relay communication with the base station.
- *Saturated bands:* A fixed wireless object, e.g. a gas meter sending regular data through the air, might check if the frequency it uses is saturated and choose, alone or in a distributed manner with other gas meters, to use another frequency (or even protocol) to communicate.
- *Radars:* With numerical communications, passive radar technology is changing, these radars will have to be updated regularly to be able to listen to new communication standards.
- *Internet of things:* With the predicted huge venue of wireless object, some reconfigurability will be needed even on the simplest smart object as mentionned above for facing the band saturation problem or simply communicating in a new environment.

4.2. Public wireless access networks

The commercial markets for wireless technologies are the largest markets for SDR and cognitive radio. these markets includes i) the cellular market (4G, LTE), ii) the Wireless Local Area Network market (WLAN, e.g. Wifi), and iii) the Broadband Wireless Access market (e.g. WiMax). The key objective here is to improve spectrum efficiency and availability, and to enable cognitive radio and SDR to support multimedia and multi-radio initiatives.

The future mobile radio access network referred to as 4G (4th generation) is expected to provide a wireless access of 100 Mbps in extended mobility and up to 1Gbps in reduced mobility as defined by the group IMT-Advanced of the ITU-R(adiocommunication) section. On the road towards the 4G, IMT-2000 standards evolutions are driven by the work of the WiMAX forum (IEEE 802.16e) on the one hand and by those of the LTE (Long Term Evolution) group of the 3GPP on the other hand. Both groups announced some targeted evolutions that could comply with the 4G requirements, namely the Gigabit Wimax (802.16m) and the LTE-Advanced proposal from the 3GPP.

In both technologies, the scarcity of the radio spectrum is taken care of by the use of MIMO and OFDMA technologies, combining the dynamic spatial and frequency multiple access. However, a better spectral efficiency will be achieved if the radio spectrum can be shared dynamically between primary and secondary networks, and if the terminals are reconfigurable in real-time. Socrate is active in this domain because of its past activity in Swing and its links to the telecommunication teaching department of Insa. The developpement of the FIT plateform [36] is a strong effort in this area.

4.3. Military SDR and Public Safety

Military applications have developed specific solutions for SDR. In France, Thales is a major actor (e.g. project Essor defining inter-operability between European military radio) and abroad the Join Tactical Radio System, and Darpa focus on Mobile Ad-hoc Networks (MANETs) have brought important deliverables, like the Software Communications Architecture (SCA) for instance [37].

Recent natural disasters have brought considerable attention to the need of enhanced public safety communication abroad [35]. Socrate in not currently implied in any military or public safety research programs but is aware of the potential importance this domain may take in Europe in a near future.

4.4. Ambient Intelligence: WSN and IoT

Sensor networks have been investigated and deployed for decades already; their wireless extension, however, has witnessed a tremendous growth in recent years. This is mainly attributed to the development of wireless sensor networks (WSNs): a large number of sensor nodes, reliably operating under energy constraints. It is anticipated that within a few years, sensors will be deployed in a variety of scenarios, ranging from environmental monitoring to health care, from the public to the private sector. Prior to large-scale deployment, however, many problems have to be solved, such as the extraction of application scenarios, design of suitable software and hardware architectures, development of communication and organization protocols, validation and first steps of prototyping, etc. The Citi laboratory has a long experience in WSN which leaded recently to the creation of a start-up company, leaded by two former Citi members: HIKOB(http://openlab.hikob.com).

The Internet of Things (IoT) paradigm is defined as a very large set of systems interconnected to provide a virtual twin world interacting with the real world. In our work we will mostly focus on wireless systems since the wireless link is the single media able to provide a full mobility and ubiquitous access. Wireless IoT is not a reality yet but will probably result from the convergence between mobile radio access networks and wireless sensor networks. If radio access networks are able to connect almost all humans, they would fail to connect a potential of several billions of objects. Nevertheless, the mutation of cellular systems toward more adaptive and autonomous systems is on going. This is why Socrate develops a strong activity in this applicative area, with its major industrial partners: Orange Labs and Alcatel-Lucent Bell labs.

For instance, the definition of a *smart node* intermediate between a WSN and a complex SDR terminal is one of the research direction followed in Socrate, explicitly stated in the ADT Snow project. Other important contributions are made in the collaboration with SigFox and Euromedia and in the EconHome project.

4.5. Body Area Networks

Body Area Network is a relatively new paradigm which aims at promoting the development or wireless systems in, on and around the human body. Wireless Body Area Networks (BAN) is now a well known acronym which encompasses scenarios in which several sensors and actuators are located on or inside the human body to sense different data, e.g. physiological information, and transfer them wirelessly towards a remote coordination unit which processes, forwards, takes decisions, alerts, records, etc. The use of BAN spans a wide area, from medical and health care to sport through leisure applications, which definitely makes the definition of a standard air interface and protocol highly challenging. Since it is expected that such devices and networks would have a growing place in the society and become more stringent in terms of quality of service, coexistence issues will be critical. Indeed, the radio resource is known to be scarce. The recent regulation difficulties of UWB systems as well as the growing interest for opportunistic radios show that any new system have to make an efficient use of the spectrum. This also applies to short range personal and body area network systems which are subject to huge market penetrations.

Socrate was involved in the Banet ANR project (2008-2010), in which we contributed to the development of a complete PHY/MAC standard in cooperation with Orange Labs and CEA Leti, who participated to the standardization group 802.15.6. Recently, Inria has been added as a partner the FET flagship untitled *Guardian Angels* (http://www.fet-f.eu/), an important european initiative to develop the BANs of the futur.

We consider that BANs will probably play an important role in the future of Internet as the multiple objects connected on body could also be connected to Internet by the mobile phone hosted by each human. Therefore the BAN success really depends on the convergence of WSN and radio access networks, which makes it a very interesting applicative framework for Socrate team.

TREC Project-Team

4. Application Domains

4.1. Application Domains

We have investigated various applications of our research results with the following industrial partners and user associations:

- **Wireless Networks**
  - Alcatel-Lucent Bell Laboratories (L. Thomas and L. Roullet) on self optimization in cellular networks.
  - Qualcomm (T. Richardson and his group) on improvements of CSMA CA.
  - Orange (M. Karray) on cellular networks.
- **Network Dynamics**
  - Thalès and Real-Time-at-Work on embedded networks.
  - Grenouille on probing in access networks.
- **Networks Economics**
  - Technicolor (J. Bolot) on economic incentives.

URBANET Team

4. Application Domains

4.1. Smart infrastructure

Unlike the communication infrastructure that went through a continuous development in the last decades, the distribution networks in our cities, whether we are talking about water, gas, or 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 becomes 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 needs of both consumers and suppliers.

Another fundamental urban infrastructure is the transportation system. The progress achieved by 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 will be obtained by 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 architectural 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 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. User-centric services

What is the most disobeyed traffic sign in your city? How does the level of pollution on your street compare with the one in other neighborhoods? How long is the queue at that exhibition you were planning to attend today? Combining location awareness and data recovered from multiple sources like social networks or sensing devices can provide answers to all these questions, making visible previously unknown characteristics of the urban environment.

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.