Activity Report 2017

Project-Team ASCOLA

Aspect and composition languages

IN COLLABORATION WITH: Laboratoire des Sciences du numerique de Nantes

Table of contents

Project-Team ASCOLA

Creation of the Project-Team: 2009 January 01, end of the Project-Team: 2017 December 31

Keywords:

Computer Science and Digital Science:

      A1.1.6. - Cloud
      A1.1.8. - Security of architectures
      A1.1.13. - Virtualization
      A1.3. - Distributed Systems
      A1.6. - Green Computing
      A2.1. - Programming Languages
      A2.1.1. - Semantics of programming languages
      A2.1.2. - Object-oriented programming
      A2.1.3. - Functional programming
      A2.1.4. - Aspect-oriented programming
      A2.1.6. - Concurrent programming
      A2.1.7. - Distributed programming
      A2.1.10. - Domain-specific languages
      A2.1.11. - Proof languages
      A2.2.1. - Static analysis
      A2.4.2. - Model-checking
      A2.4.3. - Proofs
      A2.5. - Software engineering
      A2.6.2. - Middleware
      A2.6.3. - Virtual machines
      A3.1.3. - Distributed data
      A3.1.5. - Control access, privacy
      A3.1.8. - Big data (production, storage, transfer)
      A4.5. - Formal methods for security
      A4.6. - Authentication
      A4.7. - Access control
      A4.8. - Privacy-enhancing technologies
      A7.1. - Algorithms
      A7.2. - Logic in Computer Science

Other Research Topics and Application Domains:

      B3.1. - Sustainable development
      B4.5. - Energy consumption
      B4.5.1. - Green computing
      B5.1. - Factory of the future
      B6.1. - Software industry
      B6.1.1. - Software engineering
      B6.1.2. - Software evolution, maintenance
      B6.5. - Information systems

1. Personnel

Research Scientists

Ismael Cuadrado Cordero [IMT Atlantique, Researcher, until Aug 2017]
Shadi Ibrahim [Inria, Researcher, from Apr 2017]
Guillaume Munch [Inria, Researcher, until Apr 2017]
Nicolas Tabareau [Inria, Researcher, until Apr 2017, HDR]

Faculty Members

Mario Südholt [Team leader, IMT Atlantique, Professor, HDR]
Frederico Alvares de Oliveira Junior [IMT Atlantique, Associate Professor, temporary position, until Aug 2017]
Julien Cohen [Univ de Nantes, Associate Professor, until Apr 2017]
Pierre Cointe [IMT Atlantique, Professor, HDR]
Hélène Coullon [IMT Atlantique, Associate Professor, HDR]
Remi Douence [IMT Atlantique, Associate Professor, until Apr 2017, HDR]
Herve Grall [IMT Atlantique, Associate Professor, until Apr 2017]
Adrien Lebre [Inria, Advanced Research Position until Sep 2017, IMT Atlantique Associate Professor, HDR]
Thomas Ledoux [IMT Atlantique, Associate Professor]
Jean-Marc Menaud [IMT Atlantique, Professor, HDR]
Jacques Noyé [IMT Atlantique, Associate Professor]
Jean-Claude Royer [IMT Atlantique, Professor, HDR]

Post-Doctoral Fellows

Benedikt Ahrens [Inria, until Apr 2017]
Zakarea Al Shara [IMT Atlantique, until Oct 2017]
Eric Finster [Inria, from Feb 2017 until Apr 2017]
Jonathan Pastor [Ecole des Mines de Nantes , from Aug 2017]
Dimitri Pertin [Inria]
Anthony Simonet [Inria, until Jun 2017]
Chi Zhou [Inria, from Apr 2017 until Aug 2017]
Pierre-Marie Pedrot [IMT Atlantique, until Apr 2017]

PhD Students

Mohamed Abderrahim [Orange Labs]
Mohammad-Mahdi Bazm [Orange Labs]
Paul Blouet [Ecole des Mines de Nantes , until Apr 2017]
Fatima-Zahra Boujdad [Univ de Nantes]
Simon Boulier [Ecole normale supérieure de Rennes, until Apr 2017]
Emile Cadorel [Armines, from Oct 2017]
Maverick Chardet [Inria, from Oct 2017]
Bastien Confais [CNRS]
David Espinel [Orange Labs, from Oct 2017]
Gaetan Gilbert [IMT Atlantique, until Apr 2017]
Md Sabbir Hasan [INSA, co-supervision with Myriads team, Inria, until Mar 2017]
Yacine Hebbal [Orange Labs, until Oct 2017]
Ambroise Lafont [IMT Atlantique, until Apr 2017]
Thuy Linh Nguyen [Inria]
Kevin Quirin [Ecole normale supérieure de Rennes, until Aug 2017]
Igor Zhirkov [Armines]

Technical staff

Ronan-Alexandre Cherrueau [Inria]
Remy Pottier [IMT Atlantique, until Aug 2017]

Alexandre Van Kempen [Inria, from Oct 2017]

Interns

Santiago Bautista [Inria, from May 2017 until Jul 2017]

Adel Benabadji [Ecole des Mines de Nantes , from Apr 2017 until Jul 2017]

Antoine Carat [Inria, from Jun 2017 until Aug 2017]

Guillaume Combette [Ecole Normale Supérieure Lyon, from Mar 2017 until Apr 2017]

Tudal Le Bot [Inria, from Jun 2017 until Aug 2017]

Xavier Montillet [Ecole normale supérieure de Rennes, Apr 2017]

Administrative Assistant

Anne Claire Binetruy [Inria]

Visiting Scientist

Amin Timany [IMT Atlantique, from Mar 2017 until Apr 2017]

2. Overall Objectives

2.1. Presentation

The research team addresses the general problem of evolving software by developing concepts, languages, implementations and tools for building software architectures based on components and aspects. Its long term goal is the development of new abstractions for the programming of software architectures, their representation in terms of expressive programming languages and their correct and efficient implementation.

We pursue the following objectives:

- New concepts and techniques for the compositional definition and implementation of complex software systems, notably involving crosscutting concerns that cannot be handled modularly using traditional software development approaches.

- New programming techniques and algorithms for resource management in mutualized environments. We provide language abstractions and implementation techniques for large-scale applications in cloud- and grid-based systems, both on the level of (service-based) applications and (virtualized) infrastructures. We develop solutions, in particular, for the optimization of the energy consumption in such environments (data centers ...)

- We develop new formal theories for and apply formal methods to the correctness of software systems. We aim at developing more powerful techniques for theorem proving and enable complex, often dynamic, software systems to be proven correct using program transformations and analysis techniques. We develop solutions, in particular, for the constructive enforcement of security properties on the level of software systems.

Finally, we apply and validate our results based on real-world applications from numerous domains, notably enterprise information systems, the Cloud, and pervasive systems.

3. Research Program

3.1. Overview

Since we mainly work on new concepts for the language-based definition and implementation of complex software systems, we first briefly introduce some basic notions and problems of software components (understood in a broad sense, that is, including modules, objects, architecture description languages and services), aspects, and domain-specific languages. We conclude by presenting the main issues related to distribution and concurrency, in particular related to capacity planning issues that are relevant to our work.

3.2. Software Composition

Modules and services. The idea that building *software components*, i.e., composable prefabricated and parameterized software parts, was key to create an effective software industry was realized very early [77]. At that time, the scope of a component was limited to a single procedure. In the seventies, the growing complexity of software made it necessary to consider a new level of structuring and programming and led to the notions of information hiding, *modules*, and module interconnection languages [84], [60]. Information hiding promotes a black-box model of program development whereby a module implementation, basically a collection of procedures, is strongly encapsulated behind an interface. This makes it possible to guarantee logical invariant *properties* of the data managed by the procedures and, more generally, makes *modular reasoning* possible.

In the context of today's Internet-based information society, components and modules have given rise to *software services* whose compositions are governed by explicit *orchestration or choreography* specifications that support notions of global properties of a service-oriented architecture. These horizontal compositions have, however, to be frequently adapted dynamically. Dynamic adaptations, in particular in the context of software evolution processes, often conflict with a black-box composition model either because of the need for invasive modifications, for instance, in order to optimize resource utilization or modifications to the vertical compositions implementing the high-level services.

Object-Oriented Programming. Classes and objects provide another kind of software component, which makes it necessary to distinguish between *component types* (classes) and *component instances* (objects). Indeed, unlike modules, objects can be created dynamically. Although it is also possible to talk about classes in terms of interfaces and implementations, the encapsulation provided by classes is not as strong as the one provided by modules. This is because, through the use of inheritance, object-oriented languages put the emphasis on *incremental programming* to the detriment of modular programming. This introduces a white-box model of software development and more flexibility is traded for safety as demonstrated by the *fragile base class* issue [80].

Architecture Description Languages. The advent of distributed applications made it necessary to consider more sophisticated connections between the various building blocks of a system. The *software architecture* [89] of a software system describes the system as a composition of *components* and *connectors*, where the connectors capture the *interaction protocols* between the components [48]. It also describes the rationale behind such a given architecture, linking the properties required from the system to its implementation. *Architecture Description Languages* (ADLs) are languages that support architecture-based development [78]. A number of these languages make it possible to generate executable systems from architectural descriptions, provided implementations for the primitive components are available. However, guaranteeing that the implementation conforms to the architecture is an issue.

Protocols. Today, protocols constitute a frequently used means to precisely define, implement, and analyze contracts, notably concerning communication and security properties, between two or more hardware or software entities. They have been used to define interactions between communication layers, security properties of distributed communications, interactions between objects and components, and business processes.

Object interactions [82], component interactions [95], [86] and service orchestrations [61] are most frequently expressed in terms of *regular interaction protocols* that enable basic properties, such as compatibility, substitutability, and deadlocks between components to be defined in terms of basic operations and closure properties of finite-state automata. Furthermore, such properties may be analyzed automatically using, e.g., model checking techniques [58], [67].

However, the limited expressive power of regular languages has led to a number of approaches using more expressive *non-regular* interaction protocols that often provide distribution-specific abstractions, e.g., session types [71], or context-free or turing-complete expressiveness [87], [55]. While these protocol types allow conformance between components to be defined (e.g., using unbounded counters), property verification can only be performed manually or semi-automatically.

3.3. Programming languages for advanced modularization

The main driving force for the structuring means, such as components and modules, is the quest for clean *separation of concerns* [62] on the architectural and programming levels. It has, however, early been noted that concern separation in the presence of crosscutting functionalities requires specific language and implementation level support. Techniques of so-called *computational reflection*, for instance, Smith's 3-Lisp or Kiczales's CLOS meta-object protocol [90], [74] as well as metaprogramming techniques have been developed to cope with this problem but proven unwieldy to use and not amenable to formalization and property analysis due to their generality. Methods and techniques from two fields have been particularly useful in addressing such advanced modularization problems: Aspect-Oriented Software Development as the field concerned with the systematic handling of modularization issues and domain-specific languages that provide declarative and efficient means for the definition of crosscutting functionalities.

Aspect-Oriented Software Development [73], [46] has emerged over the previous decade as the domain of systematic exploration of crosscutting concerns and corresponding support throughout the software development process. The corresponding research efforts have resulted, in particular, in the recognition of *crosscutting* as a fundamental problem of virtually any large-scale application, and the definition and implementation of a large number of aspect-oriented models and languages.

However, most current aspect-oriented models, notably AspectJ [72], rely on pointcuts and advice defined in terms of individual execution events. These models are subject to serious limitations concerning the modularization of crosscutting functionalities in distributed applications, the integration of aspects with other modularization mechanisms such as components, and the provision of correctness guarantees of the resulting AO applications. They do, in particular, only permit the manipulation of distributed applications on a per-host basis, that is, without direct expression of coordination properties relating different distributed entities [91]. Similarly, current approaches for the integration of aspects and (distributed) components do not directly express interaction properties between sets of components but rather seemingly unrelated modifications to individual components [59]. Finally, current formalizations of such aspect models are formulated in terms of low-level semantic abstractions (see, e.g., Wand's et al semantics for AspectJ [94]) and provide only limited support for the analysis of fundamental aspect properties.

Different approaches have been put forward to tackle these problems, in particular, in the context of so-called *stateful* or *history-based aspect languages* [63], [64], which provide pointcut and advice languages that directly express rich relationships between execution events. Such languages have been proposed to directly express coordination and synchronization issues of distributed and concurrent applications [83], [53], [66], provide more concise formal semantics for aspects and enable analysis of their properties [49], [65], [63], [47]. Furthermore, first approaches for the definition of *aspects over protocols* have been proposed, as well as over regular structures [63] and non-regular ones [93], [81], which are helpful for the modular definition and verification of protocols over crosscutting functionalities.

They represent, however, only first results and many important questions concerning these fundamental issues remain open, in particular, concerning the semantics foundations of AOP and the analysis and enforcement of correctness properties governing its, potentially highly invasive, modifications.

Domain-specific languages (DSLs) represent domain knowledge in terms of suitable basic language constructs and their compositions at the language level. By trading generality for abstraction, they enable complex relationships among domain concepts to be expressed concisely and their properties to be expressed and formally analyzed. DSLs have been applied to a large number of domains; they have been particularly popular in the domain of software generation and maintenance [79], [97].

Many modularization techniques and tasks can be naturally expressed by DSLs that are either specialized with respect to the type of modularization constructs, such as a specific brand of software component, or to the compositions that are admissible in the context of an application domain that is targeted by a modular implementation. Moreover, software development and evolution processes can frequently be expressed by transformations between applications implemented using different DSLs that represent an implementation at different abstraction levels or different parts of one application.

Functionalities that crosscut a component-based application, however, complicate such a DSL-based transformational software development process. Since such functionalities belong to another domain than that captured by the components, different DSLs should be composed. Such compositions (including their syntactic expression, semantics and property analysis) have only very partially been explored until now. Furthermore, restricted composition languages and many aspect languages that only match execution events of a specific domain (e.g., specific file accesses in the case of security functionality) and trigger only domain-specific actions clearly are quite similar to DSLs but remain to be explored.

3.4. Distribution and Concurrency

While ASCOLA does not investigate distribution and concurrency as research domains per se (but rather from a software engineering and modularization viewpoint), there are several specific problems and corresponding approaches in these domains that are directly related to its core interests that include the structuring and modularization of large-scale distributed infrastructures and applications. These problems include crosscutting functionalities of distributed and concurrent systems, support for the evolution of distributed software systems, and correctness guarantees for the resulting software systems.

Underlying our interest in these domains is the well-known observation that large-scale distributed applications are subject to *numerous crosscutting functionalities* (such as the transactional behavior in enterprise information systems, the implementation of security policies, and fault recovery strategies). These functionalities are typically partially encapsulated in distributed infrastructures and partially handled in an ad hoc manner by using infrastructure services at the application level. Support for a more principled approach to the development and evolution of distributed software systems in the presence of crosscutting functionalities has been investigated in the field of *open adaptable middleware* [54], [76]. Open middleware design exploits the concept of reflection to provide the desired level of configurability and openness. However, these approaches are subject to several fundamental problems. One important problem is their insufficient, framework-based support that only allows partial modularization of crosscutting functionalities.

There has been some *criticism* on the use of *AspectJ-like aspect models* (which middleware aspect models like that of JBoss AOP are an instance of) for the modularization of distribution and concurrency related concerns, in particular, for transaction concerns [75] and the modularization of the distribution concern itself [91]. Both criticisms are essentially grounded in AspectJ's inability to explicitly represent sophisticated relationships between execution events in a distributed system: such aspects therefore cannot capture the semantic relationships that are essential for the corresponding concerns. History-based aspects, as those proposed by the ASCOLA project-team provide a starting point that is not subject to this problem.

From a point of view of language design and implementation, aspect languages, as well as domain specific languages for distributed and concurrent environments share many characteristics with existing distributed languages: for instance, event monitoring is fundamental for pointcut matching, different synchronization strategies and strategies for code mobility [69] may be used in actions triggered by pointcuts. However, these relationships have only been explored to a small degree. Similarly, the formal semantics and formal properties of aspect languages have not been studied yet for the distributed case and only rudimentarily for the concurrent one [49], [66].

3.5. Security

Security properties and policies over complex service-oriented and standalone applications become ever more important in the context of asynchronous and decentralized communicating systems. Furthermore, they constitute prime examples of crosscutting functionalities that can only be modularized in highly insufficient ways with existing programming language and service models. Security properties and related properties, such as accountability properties, are therefore very frequently awkward to express and difficult to analyze and enforce (provided they can be made explicit in the first place).

Two main issues in this space are particularly problematic from a compositional point of view. First, information flow properties of programming languages, such as flow properties of Javascript [51], and service-based systems [57] are typically specially-tailored to specific properties, as well as difficult to express and analyze. Second, the enforcement of security properties and security policies, especially accountability-related properties [85], [92], is only supported using ad hoc means with rudimentary support for property verification.

The ASCOLA team has recently started to work on providing formal methods, language support and implementation techniques for the modular definition and implementation of information flow properties as well as policy enforcement in service-oriented systems as well as, mostly object-oriented, programming languages.

3.6. Green IT

With the emergence of the Future Internet and the dawn of new IT architecture and computation models such as cloud computing, the usage of data centers (DC) as well as their power consumption increase dramatically [56]. Besides the ecological impact [70], energy consumption is a predominant criterion for DC providers since it determines the daily cost of their infrastructure. As a consequence, power management becomes one of the main challenges for DC infrastructures and more generally for large-scale distributed systems.

To address this problem, we study two approaches: a workload-driven [52] and power-driven one [88]. As part of the workload-driven solution, we adapt the power consumption of the DC depending on the application workload, and evaluate whether this workload to be more reactive. We develop a distributed system from the system to the service-oriented level mainly based on hardware and virtualization capabilities that is managed in a user-transparent fashion. As part of the power-driven approach, we address energy consumption issues through a strong synergy inside the infrastructure software stack and more precisely between applications and resource management systems. This approach is characterized by adapting QoS properties aiming at the best trade-off between cost of energy (typically from the regular electric grid), its availability (for instance, from renewable energy), and service degradation caused, for instance, by application reconfigurations to jobs suspensions.

3.7. Capacity Planning for Large Scale Distributed System

Since the last decade, cloud computing has emerged as both a new economic model for software (provision) and as flexible tools for the management of computing capacity [50]. Nowadays, the major cloud features have become part of the mainstream (virtualization, storage and software image management) and the big market players offer effective cloud-based solutions for resource pooling. It is now possible to deploy virtual infrastructures that involve virtual machines (VMs), middleware, applications, and networks in such a simple manner that a new problem has emerged since 2010: VM sprawl (virtual machine proliferation) that consumes valuable computing, memory, storage and energy resources, thus menacing serious resource shortages. Scientific approaches that address VM sprawl are both based on classical administration techniques like the lifecycle management of a large number of VMs as well as the arbitration and the careful management of all resources consumed and provided by the hosting infrastructure (energy, power, computing, memory, network etc.) [68], [96].

The ASCOLA team investigates fundamental techniques for cloud computing and capacity planning, from infrastructures to the application level. Capacity planning is the process of planning for, analyzing, sizing, managing and optimizing capacity to satisfy demand in a timely manner and at a reasonable cost. Applied to distributed systems like clouds, a capacity planning solution must mainly provide the minimal set of resources necessary for the proper execution of the applications (i.e., to ensure SLA). The main challenges in this context are: scalability, fault tolerance and reactivity of the solution in a large-scale distributed system, the analysis and optimization of resources to minimize the cost (mainly costs related to the energy consumption of datacenters), as well as the profiling and adaptation of applications to ensure useful levels of quality of service (throughput, response time, availability etc.).

Our solutions are mainly based on virtualized infrastructures that we apply from the IaaS to the SaaS levels. We are mainly concerned by the management and the execution of the applications by harnessing virtualization capabilities, the investigation of alternative solutions that aim at optimizing the trade-off between performance and energy costs of both applications and cloud resources, as well as arbitration policies in the cloud in the presence of energy-constrained resources.

4. Application Domains

4.1. Enterprise Information Systems and Services

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

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

- Loose-coupling: services are autonomous: they do not require other services to be executed;

- Ease of integration: Services communicate over standard protocols.

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

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

- Service level: as services can be composed to build processes, aspect weaving will deal with the orchestration and the choreography of services.

- Implementation level: as services are abstractly specified, aspect weaving will require to extend service interfaces in order to describe the effects of the executed services on the sensitive resources they control.

4.2. Capacity Planning in Cloud, Fog and Edge Computing

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

In 2017, we have been consolidating our expertise around the OpenStack ecosystem. We proposed in particular EnOS, a dedicated framework to conduct performance analyses of OpenStack at large-scale ina reproducible manner. The framework enables researchers to conduct experiments in an automize manner on top of different testbeds such as Grid'5000 and Chameleon. see Sec. 7.1.

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

4.3. Pervasive Systems

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

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

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

5. Highlights of the Year

5.1. Highlights of the Year

5.1.1. *Remarkable results: research and third-party funding*

Regarding scientific results, the team has produced a number of outstanding results on Fog/Edge architectures, notably on how to leverage renewable energy in this context [29], [9], [8], [33]. In the software engineering domain, particularly notable contributions have been made on software adaptability [4], [11].

Concerning third-party funding, 2017 has seen the acceptance of the large industrial/academic Hydda project as well as the start of two individual projects, the Kerdata and ConnectTalent projects, both of which issue of highly-competitive calls.

5.1.2. *The future: the Gallinette and Stack teams*

After a 10-year adventure, the research path of the Ascola team finishes at the end of 2017 after having given rise to two new teams in 2017: the Gallinette team in April and the Stack team in November. These new teams pursue and diversify Ascola's main research domains, respectively formal methods for programming languages and distributed software systems. Note that because of the rather early split of the Gallinette team, we have not included the corresponding results in this year's Ascola report.

5.1.3. *Awards*

In 2017 members of the team have been awarded three research-related awards: two personal awards and a best paper award:

- **Programme Jeunes Talents France Chine 2017:**

  Shadi Ibrahim was one of the 12 researchers selected for the "Programme Jeunes Talents France Chine" award (12 out of 54 applicants).

- **ICA3PP-2017 Outstanding Leadership Award:**

  Shadi Ibrahim received an Outstanding Leadership Award as program chair of the ICA3PP-2017.

BEST PAPER AWARD:

[27]

J. LEJEUNE, F. ALVARES, T. LEDOUX. *Towards a generic autonomic model to manage Cloud Services*, in "The 7th International Conference on Cloud Computing and Services Science (CLOSER 2017)", Porto, Portugal, April 2017, https://hal.archives-ouvertes.fr/hal-01511360

6. New Software and Platforms

6.1. btrCloud

KEYWORDS: Cloud computing - Data center - Cluster - Placement - Autonomic system - Orchestration - Energy - Grid - Virtualization - Scheduler

FUNCTIONAL DESCRIPTION: Orchestration, virtualization, energy, autonomic system, placement, cloud computing, cluster, data center, scheduler, grid

btrCloud is a virtual machine manager for clusters and provides a complete solution for the management and optimization of virtualized data centers. btrCloud (acronym of better cloud) is composed of three parts.

The analysis function enables operatives and people in charge to monitor and analyze how a data-center works - be it on a daily basis, on the long run, or in order to predict future trends. This feature includes boards for performance evaluation and analysis as well as trends estimation.

btrCloud, by the integration of btrScript, provides (semi-)automated VM lifecycle management, including provisioning, resource pool management, VM tracking, cost accounting, and scheduled deprovisioning. Key features include a thin client interface, template-based provisioning, approval workflows, and policy-based VM placement.

Finally, several kinds of optimizations are currently available, such as energy and load balancing. The former can help save up to around 20

- Participants: Frédéric Dumont, Guillaume Le Louët and Jean-Marc Menaud
- Contact: Guillaume Le Louët
- URL: http://www.btrcloud.org/btrCloud/index_EN.html

6.2. SimGrid

KEYWORDS: Large-scale Emulators - Grid Computing - Distributed Applications

SCIENTIFIC DESCRIPTION: SimGrid is a toolkit that provides core functionalities for the simulation of distributed applications in heterogeneous distributed environments. The simulation engine uses algorithmic and implementation techniques toward the fast simulation of large systems on a single machine. The models are theoretically grounded and experimentally validated. The results are reproducible, enabling better scientific practices.

Its models of networks, cpus and disks are adapted to (Data)Grids, P2P, Clouds, Clusters and HPC, allowing multi-domain studies. It can be used either to simulate algorithms and prototypes of applications, or to emulate real MPI applications through the virtualization of their communication, or to formally assess algorithms and applications that can run in the framework.

The formal verification module explores all possible message interleavings in the application, searching for states violating the provided properties. We recently added the ability to assess liveness properties over arbitrary and legacy codes, thanks to a system-level introspection tool that provides a finely detailed view of the running application to the model checker. This can for example be leveraged to verify both safety or liveness properties, on arbitrary MPI code written in C/C++/Fortran.

RELEASE FUNCTIONAL DESCRIPTION:

- Four releases in 2017. Major changes:
    - S4U: many progress, toward SimGrid v4.0. About 80% of the features offered by SimDag and MSG are now integrated, along with examples. Users can now write plugins to extend SimGrid.
    - SMPI: Support MPI 2.2, RMA support, Convert internals to C++.
    - Java: Massive memleaks and performance issues fixed.
    - New models: Multi-core VMs, Energy consumption due to the network
    - All internals are now converted to C++, and most of our internally developped data containers were replaced with std::* constructs.
    - (+ bug fixes, cleanups and documentation improvements)
- Participants: Adrien Lèbre, Arnaud Legrand, Augustin Degomme, Florence Perronnin, Frédéric Suter, Jean-Marc Vincent, Jonathan Pastor, Jonathan Rouzaud-Cornabas, Luka Stanisic, Mario Südholt and Martin Quinson
- Partners: CNRS - ENS Rennes
- Contact: Martin Quinson
- URL: http://simgrid.gforge.inria.fr/

6.3. VMPlaces

FUNCTIONAL DESCRIPTION: VMPlaces is a dedicated framework to evaluate and compare VM placement algorithms. This framework is composed of two major components: the injector and the VM placement algorithm. The injector is the generic part of the framework (i.e. the one you can directly use) while the VM placement algorithm is the part you want to study (or compare with available algorithms). Currently, the VMPlaceS is released with three algorithms:

Entropy, a centralized approach using a constraint programming approach to solve the placement/reconfiguration VM problem

Snooze, a hierarchical approach where each manager of a group invokes Entropy to solve the placement/reconfiguration VM problem. Note that in the original implementation of Snooze, it is using a specific heuristic to solve the placement/reconfiguration VM problem. As the sake of simplicity, we have simply reused the entropy scheduling code.

DVMS, a distributed approach that dynamically partitions the system and invokes Entropy on each partition.

- Participants: Adrien Lèbre, Flavien Quesnel, Jonathan Pastor, Mario Südholt and Takahiro Hirofuchi
- Contact: Adrien Lèbre
- URL: http://beyondtheclouds.github.io/VMPlaceS/

6.4. ENOS

Experimental eNvironment for OpenStack

KEYWORDS: OpenStack - Experimentation - Reproducibility
FUNCTIONAL DESCRIPTION: Enos workflow :

A typical experiment using Enos is the sequence of several phases:

- enos up : Enos will read the configuration file, get machines from the resource provider and will prepare the next phase - enos os : Enos will deploy OpenStack on the machines. This phase rely highly on Kolla deployment. - enos init-os : Enos will bootstrap the OpenStack installation (default quotas, security rules, ...) - enos bench : Enos will run a list of benchmarks. Enos support Rally and Shaker benchmarks. - enos backup : Enos will backup metrics gathered, logs and configuration files from the experiment.

- Partner: Orange Labs
- Contact: Adrien Lèbre

7. New Results

7.1. Cloud programming and management

7.1.1. *Cloud infrastructures*

Our contributions regarding cloud infrastructures can be divided into three main topics described below: contributions related to (i) geo-distributed clouds (e.g., Fog and Edge computing), (ii) the convergence of Cloud and HPC infrastructures and (iii) the simulation of virtualized infrastructures.

7.1.1.1. *Geo-distributed Clouds*

Many academic and industry experts are now advocating a shift from large-centralized Cloud Computing infrastructures to massively small-geo-distributed data centers at the edge of the network. This new paradigm of utility computing is often called Fog and Edge Computing. Advantages of this paradigm are, among others, data-locality that enhances security aspects and response times for latency-critical applications, new energetic options because of reduced size of data centers (e.g., renewable energies), single point of failure avoidance etc. Among the obstacles to the adoption of this model though is the development of a convenient and powerful IaaS system capable of managing a significant number of remote data-centers in a unified way, including monitoring and data management issues in a decentralized environment.

In 2017, we achieved three main contributions toward this challenge.

In [12], we investigate how a holistic monitoring service for a Fog/Edge infrastructure, hosting next generation digital services, should be designed. Although several solutions have been proposed in the past for the monitoring of clusters, grids and cloud systems, none of those is well appropriate to the specific Fog and Edge Computing context. The contributions of this study are: (i) the problem statement, (ii) a classification and a qualitative analysis of major existing solutions, and (iii) a preliminary discussion of the impact of deployment strategies on the monitoring service.

In [6], [39], [17], we present successive studies related to the design and development of a first-class object store service for Fog/Edge facilities. After a deep analysis of major existing solutions (Ceph, Cassandra ...), we designed a proposal that combines Scale-out Network Attached Storage systems (NAS) and IPFS, a BitTorrent-based object store spread throughout the Fog/Edge infrastructure. Without impacting the IPFS advantages particularly in terms of data mobility, the use of a Scale-out NAS on each site reduces the inter-site exchanges that are costly but mandatory for the metadata management in the original IPFS implementation. Several experiments conducted on Grid'5000 testbed are analysed and corroborate, first, the benefit of using an object store service spread at the Edge, and second, the importance of mitigating inter-site accesses. Ongoing activities are related to the management of meta data information in order to benefit from data movements.

Finally, in [26], we introduce the premises of a fog/edge resource management system by leveraging the OpenStack software, a leading IaaS manager in the industry. The novelty of the presented prototype is to operate such an Internet-scale IaaS platform in a fully decentralized manner, using P2P mechanisms to achieve high flexibility and avoid single points of failure. More precisely, we revised the OpenStack Nova service (i.e., virtual machine management and allocation) by leveraging a distributed key/value store instead of the centralized SQL backend. We present experiments that validate the correct behavior and gives performance trends of our prototype through an emulation of several data-centers using Grid'5000 testbed.

7.1.1.2. Cloud and HPC convergence

Geo-distribution of Cloud Infrastructures is not the only current trend of utility computing. Another important challenge is to reach the convergence of Cloud and HPC infrastructures, in other words on-demand HPC. Among challenges of this convergence is, for example, the enhancement of the use of light virtualization techniques on HPC systems, as well as the enhancement of mechanisms to be able to consolidate those VMs without deteriorating the performance of HPC applications, thus minimizing interferences between applications.

In [36], we present Eley, a burst buffer solution that helps to accelerate the performance of Big Data applications while guaranteeing the QoS of HPC applications. To achieve this goal, Eley embraces interference-aware prefetching technique that makes reading data input faster while introducing low interference for HPC applications. Specifically, we equip the prefetcher with five optimization actions including No Action, Full Delay, Partial Delay, Scale Up and Scale Down. It iteratively chooses the best action to optimize the prefetching while guaranteeing the pre-defined QoS requirement of HPC applications (i.e., the deadline constraint for the completion of each I/O phase). Evaluations using a wide range of Big Data and HPC applications show the effectiveness of Eley in reducing the execution time of Big Data applications (shorter map phase) while maintaining the QoS of HPC applications.

7.1.1.3. Virtualization simulation

Finally, it is important to be able to simulate the behavior of proposals for the future architectures. However, current models for virtualized resources are not accurate.

In [32], we present our latest results regarding virtualization abstractions and models for cloud simulation toolkits. Cloud simulators still do not provide accurate models for most Virtual Machine (VM) operations. This leads to incorrect results in evaluating real cloud systems. Following previous works on live-migration, we discuss an experimental study we conducted in order to propose a first-class VM boot time model. Most cloud simulators often ignore the VM boot time or give a naive model to represent it. After studying the relationship between the VM boot time and different system parameters such as CPU utilization, memory usage, I/O and network bandwidth, we introduce a first boot time model that could be integrated into current cloud simulators. Through experiments, we also show that our model correctly reproduced the boot time of a VM under different resources contention.

7.1.2. Deployment and reconfiguration in the Cloud

Being able to manage the new generation of utility computing infrastructures is an important step to build useful system building blocks. The next step is to be able to perform initial deployment of any kind of distributed software (i.e., systems, frameworks or applications) on those infrastructures, thus dealing with a complex process that includes interactions between building blocks such as virtual machine management, optimized deployment plans, monitoring of deployment etc. Such deployment processes cannot be handled manually anymore, for this reason automatic deployments tools have to be designed according to the challenges of new infrastructures (e.g., geo-distribution, hybrid infrastructures etc.). Moreover, as distributed software are more and more dynamic (i.e., reconfiguring themselves at runtime), reconfiguration and self-management capabilities should be handled in an efficient and scalable manner.

7.1.2.1. Initial deployment and placement strategies

When focusing on the initial deployment, many challenges should already need to be addressed such as placement of distributed software onto virtual machines, themselves being placed onto physical resources. This

kind of placement problem can be modeled in many different ways, such as linear or constraint programming or graph partitioning. Most of the time a multi-objective NP-hard problem is formulated, and specific heuristics have to be built to reach scalable solutions.

In [18], we present new specific placement constraints and objectives adapted to hybrid clouds infrastructures, and we address this problem through constraint programming. Furthermore we evaluate the expressivity and performance of the solution on a real case study. In the Cloud, if public providers enable simple access to resources for companies and users who have sporadic computation or storage needs, private clouds could sometimes be preferred for security or privacy reasons, or for cost reasons due to a high frequency usage of services. However, in many cases a choice between public or private clouds does not fulfill all requirements of companies and hybrid cloud infrastructures should be preferred. Solutions have already been proposed to address hybrid cloud infrastructures, however most of the time the placement of a distributed software on such infrastructure has to be indicated manually.

In [37], we present a geo-aware graph partitioning method named G-Cut, which aims at minimizing the inter-DC data transfer time of graph processing jobs in geo-distributed DCs while satisfying the WAN usage budget. G-Cut adopts two novel optimization phases which address the two challenges in WAN usage and network heterogeneities separately. G-Cut can be also applied to partition dynamic graphs thanks to its light-weight runtime overhead. We evaluate the effectiveness and efficiency of G-Cut using real-world graphs with both real geo-distributed DCs and simulations. Evaluation results demonstrate that effectiveness of G-Cut in reducing the inter-DC data transfer time and the WAN usage with a low runtime overhead.

Many other challenges than placement rise from the initial deployment. In [20], we present a survey of existing deployment tools that have been used in production to deploy OpenStack, which is a complex distributed system composed of more than a hundred different services. To fully understand how IaaSes are deployed today, we propose in this paper an overall model of the application deployment process that describes each step with their interactions. This model then serves as the basis to analyse five different deployment tools used to deploy OpenStack in production: Kolla, Enos, Juju, Kubernetes, and TripleO. Finally, a comparison is provided and the results are discussed to extend this analysis.

7.1.2.2. Capacity planning and scheduling

While a placement problem is a discrete problem at a given instant, some other challenges of deployment and reconfiguration may include the time dimension leading to scheduling optimization.

in [30] we have proposed two original workload prediction models for Cloud infrastructures. These two models, respectively based on constraint programming and neural networks, focus on predicting the CPU usage of physical servers in a Cloud data center. The predictions could then be exploited for designing energy-efficient resource allocation mechanisms like scheduling heuristics or over-commitment policies. We also provide an efficient trace generator based on constraint satisfaction problem and using a small amount of real traces. Such a generator can overcome availability issues of extensive real workload traces employed for optimization heuristics validation. While neural networks exhibit higher prediction capabilities, constraint programming techniques are more suitable for trace generation, thus making both techniques complementary.

7.1.2.3. Reconfiguration and self-management

Being able to handle the dynamicity of hardware, system building blocks, middleware and applications is a great challenge of today's and future utility computing systems. On large infrastructures such as Cloud, Fog or Edge Computing, manual administration of such dynamicity is not feasible. The automatic management of reconfiguration, or self-management of software is of great importance to guarantee reliability, fault tolerance, security, and cost and energy optimization.

In [4], in order to improve the self-adaptive behaviors in the context of Component-based Architecture, we design self-adaptive software components based on logical discrete control approaches, in which the self-adaptive behavioural models enrich component controllers with a knowledge not only on events, configurations and past history, but also with possible future configurations. This article provides the description, implementation and discussion of Ctrl-F, a Domain-specific Language whose objective is to provide high-level support for describing these control policies. In [13], we extended Ctrl-F with modularity capabilities. Apart

from the benefits of reuse and substitutability of Ctrl-F programs, modularity allows to break down the combinatorial explosion intrinsic to the generation of correct-by-construction controllers in the compilation process of Ctrl-F. A further advantage of modularity is that the executable code, that is, the controllers resulting from that compilation, are loss-coupled and can therefore be deployed and executed in a distributed fashion.

However, higher abstraction-level tools also have to be proposed for reconfiguration. In [21], we introduce ElaScript, a Domain Specific Language (DSL) which offers Cloud administrators a simple and concise way to define complex elasticity-based reconfiguration plans. ElaScript is capable of dealing with both infrastructure and software elasticities, independently or together, in a coordinated way. We validate our approach by first showing the interest to have a DSL offering multiple levels of control for Cloud elasticity, and then by showing its integration with a realistic well-known application benchmark deployed in OpenStack and the Grid'5000 infrastructure testbed.

Finally, self-management can be applied at many different levels of the Cloud paradigm, from infrastructure reconfigurations to application topology reconfigurations. In practice these reconfiguration mechanisms are tightly coupled. For example, a change in the infrastructure could lead to the re-deployment of virtual machines upon it that could lead itself to application reconfigurations. In [27], we advocate that Cloud services, regardless of the layer, may share the same consumer/provider-based abstract model. From that model, we can derive a unique and generic Autonomic Manager (AM) that can be used to manage any XaaS (Everything-as-a-Service) layer defined with that model. The paper proposes such an abstract (although extensible) model along with a generic constraint-based AM that reasons on abstract concepts, service dependencies as well as SLA (Service Level Agreements) constraints in order to find the optimal configuration for the modeled XaaS. The genericity of our approach are shown and discussed through two motivating examples and a qualitative experiment has been carried out in order to show the applicability of our approach as well as to discuss its limitations.

7.2. Energy-aware computing

7.2.1. *Renewable energy*

In his PhD thesis [1], Md Sabbir Hasan proposes – across three different contributions – how to smartly use green energy at the infrastructure and application levels for further reduction of the corresponding carbon footprints. First, he investigates the options and challenges to integrate different renewable energy sources in a realistic way and proposes a *Cloud energy broker*, which can adjust the availability and price combination to buy Green energy dynamically from the energy market in advance to make a data center partially green. Then, he introduces the concept of *virtualization of green energy*, which can be seen as an alternative to energy storage used in data centers to eliminate the intermittency problem to some extent. With the adoption of this virtualization concept, we can maximize the usage of green energy contrary to energy storage which induces energy losses, while introducing a notion of Green Service Level Agreement based on green energy for service provider and end-users. Finally, he proposes an energy adaptive autoscaling solution to exploit application internals to create green energy awareness in the interactive SaaS applications, while respecting traditional QoS properties.

In [9], we present a scheme for green energy management in the presence of explicit and implicit integration of renewable energy in data center. More specifically we propose three contributions: i) we introduce the concept of *virtualization of green energy* to address the uncertainty of green energy availability, ii) we extend the Cloud Service Level Agreement (CSLA) language [1] to support Green SLA by introducing two new threshold parameters and iii) we introduce green SLA algorithm which leverages the concept of virtualization of green energy to provide per interval specific Green SLA. Experiments were conducted with real workload profile from PlanetLab and server power model from SPECpower to demonstrate that Green SLA can be successfully established and satisfied without incurring higher cost.

---

[1] http://web.imt-atlantique.fr/x-info/csla

In [8], we investigate a thorough analysis of energy consumption and performance trade-off by allowing smart usage of green energy for interactive cloud application. Moreover, we propose an auto-scaler, named as SaaScaler, that implements several control loop based application controllers to satisfy different performance (i.e., response time, availability and user experience) and resource aware metrics (i.e., quality of energy). Based on extensive experiments with RUBiS benchmark and real workload traces using single compute node in Openstack/Grid'5000, results suggest that 13% brown energy consumption can be reduced without deprovisioning any physical or virtual resources at IaaS layer while 29% more users can access the application by dynamically adjusting capacity requirements. In [23], we add to the previous paper the capability of the infrastructure layer to be elastic. We propose a PaaS solution which efficiently utilize the elasticity nature at both infrastructure and application levels, by leveraging adaptation in facing to changing condition i.e., workload burst, performance degradation, quality of energy, etc. While applications are adapted by dynamically re-configuring their service level based on performance and/or green energy availability, the infrastructure takes care of addition/removal of resources based on application's resource demand. Both adaptive behaviors are implemented in separated modules and are coordinated in a sequential manner. We validate our approach by extensive experiments and results obtained over Grid'5000 testbed. Results show that, application can reduce significant amount of brown energy consumption by 35% and daily instance hour cost by 37% compared to a baseline approach.

in [28] we address the problem of improving the utilization of renewable energy for a single data center by using two approaches: opportunistic scheduling and energy storage. Our first result deals with analyzing the workload to find ideal solar panel dimension and battery size, this is used to power the entire workload without any brown energy consumption. However, in reality, either the solar panel dimension or the battery size are limited, and we still have to address the problem of matching the workload consumption and renewable energy production. The second result shows that opportunistic scheduling can reduce the demand for battery size while the renewable energy is sufficient. The last results demonstrate that for different battery sizes and solar panel dimensions, we can find an optimal solution combining both approaches that balances the energy losses due to different causes such as battery efficiency and VM migrations due to consolidation algorithms.

In [5] we presented the EPOC project, focus on energy-aware task execution from the hardware to application's components in the context of a mono-site data center (all resources are in the same physical location) which is connected to the regular electric Grid and to renewable energy sources (such as windmills or solar cells). we have presented the EpoCloud principles, architecture and middleware components. EpoCloud is our prototype, which tackles three major challenges: 1) To optimize the energy consumption of distributed infrastructures and service compositions in the presence of ever more dynamic service applications and ever more stringent availability requirements for services; 2) To design a clever cloud's resource management, which takes advantage of renewable energy availability to perform opportunistic tasks, then exploring the trade-off between energy saving and performance aspects in large-scale distributed system; 3) To investigate energy-aware optical ultra high-speed interconnection networks to exchange large volumes of data (VM memory and storage) over very short periods of time.

in [31] we extend our previous work on PIKA (focus 2 in the EPOC project) and introduced the green energy aware scheduling problem (GEASP) to optimize the energy consumption of a small/medium size data center. Using our model to solve the GEASP, we could optimize the energy consumption of a small/medium size data center in three ways. First, we slightly decrease its overall energy consumption, second we considerably decrease its brown energy consumption and finally we significantly increase its green energy consumption.

7.2.2. *Energy-aware consolidation and reconfiguration*

In [41] we compared the performance of VMs and containers when consolidating multiple services, in terms of QoS and EE. Our experiments compared two broadly recognized virtualization technologies: KVM for the VM approach, and Docker for the containers. We conclude that Docker outperforms KVM both in QoS and EE. According to our measurements, Docker allows running up to a 21% more services than KVM, when setting a maximum latency of 3,000 ms. In this configuration, Docker offers this service while using a 11.33% less energy than KVM. At a datacenter level, the same computation could run using less servers and less energy per server, accounting for a total of a 28% energy savings inside the datacenter.

The emergence of Internet of Things (IoT) is participating to the increase of data- and energy-hungry applications. As connected devices do not yet offer enough capabilities for sustaining these applications, users perform computation offloading to the cloud. To avoid network bottlenecks and reduce the costs associated to data movement, edge cloud solutions have started being deployed, thus improving the Quality of Service. In [29] , we advocated for leveraging on-site renewable energy production in the different edge cloud nodes to green IoT systems while offering improved QoS compared to core cloud solutions. We proposed an analytic model to decide whether to offload computation from the objects to the edge or to the core Cloud, depending on the renewable energy availability and the desired application QoS. This model is validated on our application use-case that deals with video stream analysis from vehicle cameras.

In [33], we address the problem of stragglers (i.e., slow tasks) in Big Data applications. In particular, we introduce a novel straggler detection mechanism to improve the energy efficiency of speculative execution in Hadoop, namely a hierarchical detection mechanism. The goal of this detection mechanism is to identify critical stragglers which strongly affect the job execution times and reduce the number of killed speculative copies which lead to energy waste. We also present an energy-aware copy allocation method to reduce the energy consumption of speculative execution. The core of this allocation method is a performance model and an energy model which expose the trade-off between performance and energy consumption when scheduling a copy. We evaluate our hierarchical detection mechanism and energy-aware copy allocation method on the Grid'5000 testbed using three representative MapReduce applications. Experimental results show a good reduction in the resource wasted on killed speculative copies and an improvement in the energy efficiency compared to state-of-the-art mechanisms.

The increasing size of main memories has lead to the advent of new types of storage systems. These systems propose to keep all data in distributed main memories. In [35], we present a study to characterize the performance and energy consumption of a representative in-memory storage system, namely RAMCloud, to reveal the main factors contributing to performance degradation and energy-inefficiency. Firstly, we reveal that although RAMCloud scales linearly in throughput for read-only applications, it has a non-proportional power consumption. Mainly because it exhibits the same CPU usage under different levels of access. Secondly, we show that prevalent Web workloads i.e., read-heavy and update-heavy workloads, can impact significantly the performance and the energy consumption. We relate it to the impact of concurrency, i.e., RAMCloud poorly handles its threads under highly-concurrent accesses. Thirdly, we show that replication can be a major bottleneck for performance and energy. Finally, we quantify the overhead of the crash-recovery mechanism in RAMCloud on both energy-consumption and performance.

7.3. Software engineering

7.3.1. Security and privacy

This year, we have developed new results on the security and privacy of cloud systems on all layers of abstraction: a first notion of distributed side-channel attacks on the system-level, privacy-aware middleware storage systems and accountability specifications and implementations on the application level.

7.3.1.1. System-level security for virtualized environments

Isolation on the system-level is a core security challenge for Cloud infrastructures. Similarly, fog and edge infrastructures are based on virtualization to share physical resources among several self-contained execution environments like virtual machines and containers. Yet, isolation may be threatened due to side-channels, created by the virtualization layer or due to the sharing of physical resources like the processor. Side-channel attacks (SCAs) exploit and use such leaky channels to obtain sensitive data. Previous SCAs are local and exploit isolation challenges of virtualized environments to retrieve sensitive information. We have introduced, as a first, the concept of *distributed side-channel attack (DSCA)* that is based on coordinating local attack techniques. We have explored how such attacks can threaten isolation of any virtualized environments such as fog and edge computing. Finally, we have proposed a first set of applicable countermeasures for attack mitigation of DSCAs. [14], [44]

In [24] we presented how the increasing adoption of cloud environments operated with virtualization technology opened the way to a promising hypervisor-based security monitoring approach named Virtual Machine Introspection (VMI). We investigated in Kbin-ID the application of binary code introspection at hypervisor level and analysis mechanisms on all VM kernel binary code, namely all kernel functions, to widely narrow the semantic gap in an automatic and largely OS independent way. Kbin-ID [40] is a novel hypervisor-based main kernel binary code disassembler which enables the hypervisor to locate all VM main kernel binary code and divide it into code blocks given only the address of one arbitrary kernel instruction. In [24] we presented a security use case, we are able to detect running processes that are hidden from Linux task list and ps command output, and more generally that our solution can be used for designing easily automatic and largely kernel portable VMI applications that detect and safely react against malicious activities thanks to the instrumentation of kernel functions.

7.3.1.2. *Privacy-Aware Data Storage.*

In [34] we propose a cloud storage service that protects the privacy of users by breaking user documents into blocks in order to spread them on several cloud providers. As cloud providers only own a part of the blocks and they do not know the block organization, they can not read user documents. Moreover, the storage service connects directly users and cloud providers without using a third-party as is generally the practice in cloud storage services. Consequently, users do not give critical information (security keys, passwords, etc.) to a third-party.

7.3.1.3. *Accountability for Cloud applications.*

Nowadays we are witnessing the democratization of cloud services, as a result, more and more end-users (individuals and businesses) are using these services in their daily life. In such scenarios, personal data is generally flowed between several entities. end-users need to be aware of the management, processing, storage and retention of personal data, and to have necessary means to hold service providers accountable for the use of their data. In Walid Benghabrit's thesis we present an accountability framework called Accountability Laboratory (AccLab) that allows to consider accountability from design time to implementation. We developed a language called Abstract Accountability Language (AAL) that allows to write obligations and accountability policies. This language is based on a formal logic called First Order Linear Temporal Logic (FOTL) which allows to check the consistency of the accountability policies and the compliance between two policies. These policies are translated into a temporal logic called FO-DTL 3, which is associated to a monitoring technique based on formula rewriting. Finally we developed a monitoring tool called Accountability Monitoring (AccMon) which provides means to monitor accountability policies in the context of a real system. These policies are based on FO-DTL 3 logic and the framework can act in both centralized and distributed modes and can run in on-line and off-line modes.

Accountability means to obey a contract and to ensure responsibilities in case of violations. In previous work we defined the Abstract Accountability Language and its AccLab tool support. In order to evaluate the suitability of our language and tool we experiment with the laptop user agreement, one of the policies of the Hope University in Liverpool. While this experiment is still incomplete we are able to draw some preliminary conclusions. The use of FOTL is rather tricky and the only existing prover is not maintained we think to target a first-order logic approach in the future. Natural specifications have traditional issues, for instance missing information, noises, ambiguities etc. But in case of these policies we can say much more. The information system is missing but also most of the details about the auditing process and the rectification aspects (sanction, compensation, explanation, etc). There is also a mixture of proper user behavior with the usage policy which confuses the specifier. A mean to structure the specification is important, we suggest to use templates, and it is also convenient to capture usage and accountability practices.

7.3.2. Software development and programming languages

7.3.2.1. *Industrial Internet*

In [19], we present a first "vision" paper toward Cloud Manufacturing. More precisely we try to reconsider relationships between Cloud Computing and Cloud Manufacturing based on basic definitions and historical evolution of both worlds. History shows many relations between computer science and manufacturing

processes, starting with the initial idea of "digital manufacturing" in the '70s. Since then, advances in computer science have given birth to the *Cloud Computing* (CC) paradigm, where computing resources are seen as a *service* offered to various end-users. Of course, CC has been used as such to improve the IT infrastructure associated to a manufacturing infrastructure, but its principles have also inspired a new manufacturing paradigm *Cloud Manufacturing* (CMfg) with the perspective of many benefits for both the manufacturers and their customers. However, despite the usefulness of CC for CMfg, we advocate that considering CC as a core enabling technology for CMfg, as is often put forth in the literature, is limited and should be reconsidered. This paper presents a new core-enabling vision toward CMfg, called *Cloud Anything* (CA). CA is based on the idea of abstracting low-level resources, beyond computing resources, into a set of core control building blocks providing the grounds on top of which any domain could be "cloudified".

7.3.2.2. Cloud and HPC programming

In [43], we deal with testing reproducibility in the context of Cloud elasticity, which requires control of the elasticity behavior, the possibility to select specific resources to be allocated/unallocated, and the coordination of events parallel to the elasticity process. We propose an approach fulfilling those requirements in order to make elasticity testing reproducible. To validate our approach, we perform three experiments on representative bugs on MongoDB and Zookeeper Cloud applications, where our approach succeeds in reproducing all the bugs.

In [7], the Multi-Stencil Framework (MSF) is presented. Even though this framework is applied on HPC numerical simulations, this work can be transposed to many different domains, for instance smart-* applications of Fog and Edge computing infrastructures, where heterogeneity of computations and programming models have to be handled. As the computation power of modern high performance architectures increases, their heterogeneity and complexity also become more important. One of the big challenges of exascale is to reach programming models that give access to high performance computing (HPC) to many scientists and not only to a few HPC specialists. One relevant solution to ease parallel programming for scientists is Domain Specific Language (DSL). However, one problem to avoid with DSLs is to mutualized existing codes and libraries instead of implementing each solution from scratch. For example, this phenomenon occurs for stencil-based numerical simulations, for which a large number of languages has been proposed without code reuse between them. The Multi-Stencil Framework (MSF) presented in this paper combines a new DSL to component-based programming models to enhance code reuse and separation of concerns in the specific case of stencils. MSF can easily choose one parallelization technique or another, one optimization or another, as well as one back-end implementation or another. It is shown that MSF can reach same performances than a non component-based MPI implementation over 16.384 cores. Finally, the performance model of the framework for hybrid parallelization is validated by evaluations.

8. Bilateral Contracts and Grants with Industry

8.1. Bilateral Contracts with Industry

Participants: Adrien Lebre [Contact point], Ronan-Alexandre Cherrueau, Alexandre Van Kempen.

During 2017, we agreed with Orange Labs (Lannion) to conduct a dedicated study on the evaluation of AMQP message bus alternatives within the OpenStack ecosystem. This bilateral contract ("Contrat de Recherche Externalisé") officially started in Sept 2017 for one year. With the allocated budget (100K), we hired a new research engineer, Alexandre Van Kempen. Alexandre Van Kempen works with Ronan-Alexandre Cherrueau (Temporary Resarch Engineer, hired in the context of the MERCURY InriHub) and Matthieu Simonin (Permanent Research Engineer from the Rennes Bretagne Atlantique Center) on conducting this analysis. In addition to extending the EnOS framework previously presented, they are performing several experiments with the support of the OpenStack open-source community (in particulat RedHat). The goal of the study is to identify major drawbacks of the default RabbitMQ solution with respect to the Fog/Edge requirements and evaluate whether some alternatives are available in the open-source ecosytem.

9. Partnerships and Cooperations

9.1. Regional Initiatives

9.1.1. RFI Atlanstic 2020

9.1.1.1. CoMe4ACloud
Participants: Thomas Ledoux [coordinator], Frederico Alvares de Oliveira Junior, Zakarea Al Shara.

The high-level objective of the 1-year CoMe4ACloud (Constraints and Model Engineering for Autonomic Clouds) project is to provide an end-to-end solution for autonomic Cloud services. To that end, we rely on techniques of Constraint Programming so as a decision-making tool and Model-driven Engineering to ease the automatic generation of the so-called autonomic managers as well as their synchronization with the managed system (i.e., the Cloud layers).

This year, we got the best paper award of CLOSER 2017 (the 7th International Conference on Cloud Computing and Services Science) [27]. We have also submitted two publications and provided two video-demonstrations of the early results.

CoMe4ACloud is an Atlanstic2020 funded project and supports a post-doc position. The project is led by Ascola research team and involves also AtlanModels and TASC, all of them from the LS2N and situated at IMT Atlantique. See https://come4acloud.github.io for more information.

9.1.1.2. SyMeTRIC
Participant: Jean-Marc Menaud [coordinator].

SyMeTRIC is a regional federated project in Systems Medicine funded by the Pays de la Loire french region. Systems Medicine approaches can be compared to Systems Biology. They aim at integrating several information sources to design and validate bio-models and biomarkers to anticipate and enhance patients follow-up (diagnosis, treatment response prediction, prognosis).

9.2. National Initiatives

9.2.1. CominLabs laboratory of excellence

9.2.1.1. EPOC
Participants: Jean-Marc Menaud [coordinator], Thomas Ledoux, Md Sabbir Hasan, Yunbo Li.

The project EPOC (Energy Proportional and Opportunistic Computing system) is a project running for 4 years. Four other partners collaborate within the project that is coordinated by ASCOLA: Myriads team, and the three institutions ENIB, ENSTB and University of Nantes. In this project, the partners focus on energy-aware task execution from the hardware to application components in the context of a *mono-site* data center (all resources are in the same physical location) which is connected to the *regular electric Grid and to renewable energy sources* (such as windmills or solar cells). Three major challenges are addressed in this context: optimize the energy consumption of distributed infrastructures and service compositions in the presence of ever more dynamic service applications and ever more stringent availability requirements for services; design a clever cloud's resource management which takes advantage of renewable energy availability to perform opportunistic tasks, then exploring the trade-off between energy saving and performance aspects in large-scale distributed system; investigate energy-aware optical ultra high-speed interconnection networks to exchange large volumes of data (VM memory and storage) over very short periods of time.

One of the strengths of the project is to provide a systematic approach, and use a single model for the system (from hard to soft) by mixing constraint programming and behavioral models to manage energy consumption in data centers.

9.2.1.2. PrivGen
Participants: Fatima-Zahra Boujdad, Mario Südholt [coordinator].

PrivGen ("Privacy-preserving sharing and processing of genetic data") is a three-year project that has been started in Oct. 2016 and is conducted by three partners: a team of computer scientists from the LATIM Inserm institute in Brest mainly working on data watermarking techniques, a team of geneticians from an Inserm institute in Rennes working on the gathering and interpretation of genetic data, and the Ascola team. The project provides funding of 330 KEUR altogether with an Ascola share of 120 KEUR.

The project considers challenges related to the outsourcing of genetic data that is in the Cloud by different stakeholders (researchers, organizations, providers, etc.). It tackles several limitations of current security solutions in the cloud, notably the lack of support for different security and privacy properties at once and computations executed at different sites that are executed on behalf of multiple stakeholders.

The partners are working on three main challenges:

- Mechanisms for a continuous digital content protection
- Composition of security and privacy-protection mechanisms
- Distributed processing and sharing of genetic data

The Ascola team is mainly involved in providing solutions for the second and third challenges.

9.2.2. ANR

9.2.2.1. GRECO (ANR)
Participant: Adrien Lebre [Contact point].

The GRECO project (Resource manager for cloud of Things) is an ANR project (ANR-16-CE25-0016) running for 42 months (starting in January 2017 with an allocated budget of 522KEuros, 90KEuro for ASCOLA).

The consortium is composed of 4 partners: Qarnot Computing (coordinator) and 3 academic research group (DATAMOVE and AMA from the LIG in Grenoble and ASCOLA from Inria Rennes Bretagne Atlantique).

The goal of the GRECO project (https://anr-greco.net) is to design a manager for cloud of things. The manager should act at the IaaS, PaaS and SaaS layer of the cloud. One of the principal challenges will consist in handling the execution context of the environment in which the cloud of things operates. Indeed, unlike classical resource managers, connected devices imply to consider new types of networks, execution supports, sensors and new constraints like human interactions. The great mobility and variability of these contexts complexify the modelling of the quality of service. To face this challenge, we intend to innovate in designing scheduling and data management systems that will use machine learning techniques to automatically adapt their behaviour to the execution context. Adaptation here requires a modelling of the recurrent cloud of things usages, the modelling of the dynamics of physical cloud architecture.

9.2.2.2. KerStream (ANR)
Participant: Shadi Ibrahim [Coordinator].

The KerStream project (Big Data Processing: Beyond Hadoop!) is an ANR JCJC (Young Researcher) project (ANR-16-CE25-0014-1) running for 48 months (starting in January 2017 with an allocated budget of 238KEuros).

The goal of the KerStream project is to address the limitations of Hadoop when running Big Data stream applications on large-scale clouds and do a step beyond Hadoop by proposing a new approach, called KerStream, for scalable and resilient Big Data stream processing on clouds. The KerStream project can be seen as the first step towards developing the first French middleware that handles Stream Data processing at Scale.

9.2.3. FSN

9.2.3.1. Hosanna (FSN)
Participants: Jean-Marc Menaud [coordinator], Remy Pottier.

The Hosanna project aims to scientifically and technically addresses the problem of deploying applications on a distributed multi-cloud virtual infrastructure (private cloud, Amazon, OVH, CloudWatt, Numergy etc.) This recent need is an important topic issue highlighted by recent major Outages in 2013 by the biggest players in the cloud such as Amazon or Netflix. This project aims to provide services that allow users to deploy their cloud multi-tier applications on hybrid Clouds infrastructures without any separation between IaaS. The Ascola team is extending its optimization solution to address the task placement problem in a multi-cloud environment and will develop a case study on a secure distributed file system. The project started in 2015 for a duration of 2 years.

9.2.3.2. *Hydda (FSN)*
Participants: Jean-Marc Menaud [coordinator], Hélène Coullon.

The HYDDA project aims to develop a software solution allowing the deployment of Big Data applications (with hybrid design (HPC/CLoud)) on heterogeneous platforms (cluster, Grid, private Cloud) and orchestrators (Task scheduler like Slurm, Virtual orchestrator (like Nova for OpenStack or Swarm for Docker). The main challenges addressed by the project are: how to propose an easy-to-use service to host (from deployment to elimination) application components that are both typed Cloud and HPC? How propose a service that unifies the HPCaaS (HPC as a service) and the Infrastructure as a Service (IaaS) in order to offer resources on demand and to take into account the specificities of scientific applications? How optimize resources usage of these platforms (CPU, RAM, Disk, Energy, etc.) in order to propose solutions at the least cost?

9.2.4. CPER

9.2.4.1. *SeDuCe*
Participants: Jean-Marc Menaud [coordinator], Adrien Lebre.

The SeDuCe project (Sustainable Data Centers: Bring Sun, Wind and Cloud Back Together), aims to design an experimental infrastructure dedicated to the study of data centers with low energy footprint. This innovative data center will be the first experimental data center in the world for studying the energy impact of cloud computing and the contribution of renewable energy (solar panels, wind turbines) from the scientific, technological and economic viewpoints. This project is integrated in the national context of grid computing (Grid'5000), and the Constellation project, which will be an inter-node (Pays de la Loire, Brittany).

9.2.5. Inria Project Labs

9.2.5.1. *DISCOVERY*
Participants: Hélène Coullon, Shadi Ibrahim, Adrien Lebre [coordinator], Dimitri Pertin, Ronan-Alexandre Cherrueau, Alexandre Van Kempen, Mario Südholt.

To accommodate the ever-increasing demand for Utility Computing (UC) resources, while taking into account both energy and economical issues, the current trend consists in building larger and larger Data Centers in a few strategic locations. Although such an approach enables UC providers to cope with the actual demand while continuing to operate UC resources through centralized software system, it is far from delivering sustainable and efficient UC infrastructures for future needs.

The DISCOVERY initiative [26] aims at exploring a new way of operating Utility Computing (UC) resources by leveraging any facilities available through the Internet in order to deliver widely distributed platforms that can better match the geographical dispersal of users as well as the ever increasing demand. Critical to the emergence of such locality-based UC (also referred as Fog/Edge Computing) platforms is the availability of appropriate operating mechanisms. The main objective of DISCOVERY is to design, implement, demonstrate and promote a new kind of Cloud Operting System (OS) that will enable the management of such a large-scale and widely distributed infrastructure in an unified and friendly manner.

The consortium is composed of experts in the following research areas: large-scale infrastructure management systems, networking and P2P algorithms. Moreover, two key network operators, namely Orange and RENATER, are involved in the project.

By deploying and using a Fog/Edge OS on backbones, our ultimate vision is to enable large parts of the Internet to be hosted and operated by its internal structure itself: a scalable set of resources delivered by any computing facilities forming the Internet, starting from the larger hubs operated by ISPs, governments and academic institutions, to any idle resources that may be provided by end users.

ASCOLA leads the DISCOVERY IPL and contributes mainly around two axes: VM life cycle management and security concerns.

9.2.6. InriaHub

9.2.6.1. MERCURY

Participants: Ronan-Alexandre Cherrueau, Adrien Lebre [coordinator].

ASCOLA, in particular within the framework of the DISCOVERY initiative has been working on the massively distributed use case since 2013. With the development of several proof-of-concepts around OpenStack, the team has had the opportunity to start an InriaHub action. Named MERCURY, the goal of this action is twofold: (i) support the research development made within the context of DISCOVERY and (ii) favor the transfer toward the OpenStack community.
Further information available at: http://beyondtheClouds.github.io.

9.2.7. Fond d'amorçage IMT Industrie du Futur 2017

9.2.7.1. aLIFE

Participants: Hélène Coullon [coordinator], Jacques Noyé.

The French engineering school IMT Atlantique is organizing the aLIFE workshop between industry and academia, in Nantes during two days on January, 30-31 2018. The objective of this workshop is to share various experiences and success stories, as well as open challenges related to the contribution of software-related research to Factories of the Future, in French apport de l'industrie du Logiciel à l'Industrie du Futur Européenne (aLIFE). To this end, big multinational companies, as well as SMEs and academics will exchange through plenary sessions and discussion panels.

9.2.8. Connect Talent

9.2.8.1. Apollo (Connect Talent)

Participant: Shadi Ibrahim [Coordinator].

The Apollo project (Fast, efficient and privacy-aware Workflow executions in massively distributed Data-centers) is an individual research project "'Connect Talent" running for 36 months (starting in November 2017 with an allocated budget of 201KEuros).

The goal of the Apollo project is to investigate novel scheduling policies and mechanisms for fast, efficient and privacy-aware data-intensive workflow executions in massively distributed data-centers.

9.3. European Initiatives

9.3.1. FP7 & H2020 Projects

9.3.1.1. CoqHoTT

Title: Coq for Homotopy Type Theory

Programm: H2020

Type: ERC

Duration: June 2015 - May 2020

Coordinator: Inria

Inria contact: Nicolas TABAREAU

Every year, software bugs cost hundreds of millions of euros to companies and administrations. Hence, software quality is a prevalent notion and interactive theorem provers based on type theory have shown their efficiency to prove correctness of important pieces of software like the C compiler of the CompCert project. One main interest of such theorem provers is the ability to extract directly the code from the proof. Unfortunately, their democratization suffers from a major drawback, the mismatch between equality in mathematics and in type theory. Thus, significant Coq developments have only been done by virtuosos playing with advanced concepts of computer science and mathematics. Recently, an extension of type theory with homotopical concepts such as univalence is gaining traction because it allows for the first time to marry together expected principles of equality. But the univalence principle has been treated so far as a new axiom which breaks one fundamental property of mechanized proofs: the ability to compute with programs that make use of this axiom. The main goal of the CoqHoTT project is to provide a new generation of proof assistants with a computational version of univalence and use them as a base to implement effective logical model transformation so that the power of the internal logic of the proof assistant needed to prove the correctness of a program can be decided and changed at compile time—according to a trade-off between efficiency and logical expressivity. Our approach is based on a radically new compilation phase technique into a core type theory to modularize the difficulty of finding a decidable type checking algorithm for homotopy type theory. The impact of the CoqHoTT project will be very strong. Even if Coq is already a success, this project will promote it as a major proof assistant, for both computer scientists and mathematicians. CoqHoTT will become an essential tool for program certification and formalization of mathematics.

9.3.1.2. BigStorage

Title: BigStorage: Storage-based Convergence between HPC and Cloud to handle Big Data

Programm: H2020

Duration: January 2015 - December 2018

Coordinator: Universidad politecnica de Madrid

Partners:

> Barcelona Supercomputing Center - Centro Nacional de Supercomputacion (Spain)
>
> Ca Technologies Development Spain (Spain)
>
> Commissariat A L Energie Atomique et Aux Energies Alternatives (France)
>
> Deutsches Klimarechenzentrum (Germany)
>
> Foundation for Research and Technology Hellas (Greece)
>
> Fujitsu Technology Solutions (Germany)
>
> Johannes Gutenberg Universitaet Mainz (Germany)
>
> Universidad Politecnica de Madrid (Spain)
>
> Seagate Systems Uk (United Kingdom)

Inria contact: G. Antoniu & A. Lebre

The consortium of this European Training Network (ETN) 'BigStorage: Storage-based Convergence between HPC and Cloud to handle Big Data' will train future data scientists in order to enable them and us to apply holistic and interdisciplinary approaches for taking advantage of a data-overwhelmed world, which requires HPC and Cloud infrastructures with a redefinition of storage architectures underpinning them - focusing on meeting highly ambitious performance and energy usage objectives. There has been an explosion of digital data, which is changing our knowledge about the world. This huge data collection, which cannot be managed by current data management systems, is known as Big Data. Techniques to address it are gradually combining with what has been traditionally known as High Performance Computing. Therefore, this ETN will focus on the convergence of Big Data, HPC, and Cloud data storage, ist management and analysis. To gain value from Big Data it must be addressed from many different angles: (i) applications, which

can exploit this data, (ii) middleware, operating in the cloud and HPC environments, and (iii) infrastructure, which provides the Storage, and Computing capable of handling it. Big Data can only be effectively exploited if techniques and algorithms are available, which help to understand its content, so that it can be processed by decision-making models. This is the main goal of Data Science. We claim that this ETN project will be the ideal means to educate new researchers on the different facets of Data Science (across storage hardware and software architectures, large-scale distributed systems, data management services, data analysis, machine learning, decision making). Such a multifaceted expertise is mandatory to enable researchers to propose appropriate answers to applications requirements, while leveraging advanced data storage solutions unifying cloud and HPC storage facilities.'

9.4. International Initiatives

9.4.1. Inria International Partners

9.4.1.1. Informal International Partners

National University of Singapore (NUS): We collaborate on resource management for workflows in the cloud and optimizing graph processing in geo-distributed data-centers.

9.5. International Research Visitors

9.5.1. Visits to International Teams

9.5.1.1. Research Stays Abroad

HUST and ShenZhen University, China: From October 28 to November 11, S. Ibrahim visited the Services Computing Technology and System Lab at Huazhong university of Science and Technology and the National High Performance Computing Center at Shenzhen University.

10. Dissemination

10.1. Promoting Scientific Activities

10.1.1. Scientific Events Organisation

10.1.1.1. Member of the Organizing Committees

- J. Noyé has co-organized with M. Aksit the symposium Modularity 2017, colocated with <Programming> 2017 (Brussels).
- A. Lebre has been Publicity Chair of the Big Graph Processing workshop (co-located with ICDCS'17).
- A. Lebre has co-organized the CNRS Rescom Summer School 2017 edition (70 persons).

10.1.2. Scientific Events Selection

10.1.2.1. Chair of Conference Program Committees

- A. Lebre and A. Simonet were co-program chair of ICFEC'17.
- S. Ibrahim was program co-chair of the 17th International Conference on Algorithms and Architectures for Parallel Processing (ICA3PP-2017), Helsinki, Finland, August, 2017.
- S. Ibrahim was program co-chair of the1st Workshop on the Integration of Extreme Scale Computing and Big Data Management and Analytics (EBDMA 2017), co-located with CCGrid'17, Madrid, Spain, May 2017.

10.1.2.2. Member of the Conference Program Committees

- H. Coullon was member of the program committees of the following conferences: CloudCOM'17, ICFEC'17, ICCS'17, SAC PAPP'17, Compas'17

- S. Ibrahim was member of the program committees of SC'17, Cluster'17, HiPC'17, CCGrid'17, ISPA'17, I-SPAN'17, CloudCom'17, FCST'17, PDSW-DISCS@SC'17, NetBOS@ICNP'17, HPBDC@IPDPS'17, SCRAMBL@CCGrid'17.

- A. Lebre was member of the program committees of HPDC'17, SC'17, CC-GRID'17,CloudCom'2017, and NoF'17.

- T. Ledoux was member of the program committees of the following workshops: ARM'17@Middleware, CrossCloud'17@EuroSys

- J.-M. Menaud was member of the program committees of SDS'17, Xgreen2017, AICT'17, CEIS'17, EEEP'17, Energy'17, SMARTGREENS'17

- J-C. Royer was member of the program committees of CAReMAS, ICIS'2017, SCAI'2017 and WETICE'2017.

- M. Südholt was a member of the program committees of CloudCom'17, ProWeb'17, and Programming'17

- C. Zhou was member of the program committees of FCST'17, ICA3PP'17, EBDMA'@CCGrid'17.

10.1.3. Journal

10.1.3.1. Member of the Editorial Boards

- A. Lebre is an Associate Editor of the IEEE Transactions on BigDat

- S. Ibrahim is a Guest Editor of IEEE Transactions on Big Data – Special Issue on the Integration of Extreme Scale Computing and Big Data Management and Analytics. a Journal.

- M. Südholt is an Associate Editor of the journals Programming and Modularity (Springer).

10.1.3.2. Reviewer - Reviewing Activities

- H. Coullon has been a reviewer for the following journals: Annals of Telecommunication, Future Generation Computer Systems, Transactions on BigData.

- H. Coullon has been a reviewer for the following conferences: Cluster'17, HPDC'17.

- A. Lebre has been reviewer for the following journals: IEEE Transactions on Network and Service Management, and Journal of Parallel and Distributed Computing.

- A. Lebre has been reviewer for Europar 2017.

- T. Ledoux has been a reviewer for the journal IEEE Transactions on Services Computing.

- J. Noyé has been a reviewer for the Journal of Object Technology.

- S. Ibrahim has been a reviewer for the following journals: IEEE Transactions on Parallel and Distributed Systems, and IEEE Transactions on Big Data.

10.1.4. Invited Talks

- S. Ibrahim has been invited to present a talk at the ResCom summer school (Le Croisic) : " Big Data Processing in the Cloud: Hadoop and Beyond".

- S. Ibrahim has been invited to present a talk at CGCL (Huazhong University of Science and Technology, China, 02/11/2017): "Scalable Big Data Management on clouds and HPC systems".

- S. Ibrahim has been invited to present a talk at ShenZhen University (China, 09/11/2017): "Scalable Big Data Management on clouds and HPC systems".

- A. Lebre has been invited to present a talk at the 11th edition of the CloudControl Workshop serie (Sweden) : "Enos: a Holistic Framework for Conducting Scientific Evaluations of OpenStack".

- A. Lebre has been invited to present a talk at the ResCom summer school (Le Croisic) : " Utility Computing: From Mainframes to Clouds and Beyond!".

- T. Ledoux has been invited for a talk about *frugal* Cloud by the GDS EcoInfo (CNRS) (Grenoble, France, 02/03/2017)
- T. Ledoux has been invited to present the CoMe4ACloud project at the 4th Grenoble Workshop on Autonomic Computing and Control (Grenoble, France, 10/23/2017)
- M. Südholt has been invited to present a talk at the IMT Cybersecurity day on "Privacy and sharing of genomic data."

10.1.5. *Scientific Expertise*

- S. Ibrahim is Leading the *Resource Management and Scheduling for Data-Intensive HPC Workflows* activity within the JLESC, Joint Inria-Illinois-ANL-BSC-JSC-RIKEN/AICS Laboratory for Extreme-Scale Computing.
- S. Ibrahim is member of Grid'5000 Sites Committee – Responsible for the Rennes site.
- A. Lebre is member of the executive committee of the GDR CNRS RSD "Réseau et Système distribué" and Co-leader of the transversal action Virtualization and Clouds of this GDR since 2015.
- A. Lebre is leading the OpenStack "Fog/Edge/Massively Distributed Clouds" Special Interest Group (further information at: https://wiki.openstack.org/wiki/Massively_Distributed_Clouds).
- A. Lebre is member of the executive and architect committees of the Grid'5000 GIS (Groupement d'intérêt scientifique).
- J.-M. Menaud is the organizer of "Pôle Science du Logiciel et des Systèmes Distribués" in Laboratoire des Sciences du Numérique à Nantes (LS2N) since June 2015.
- J.-M. Menaud member of the thesis committee Gilles Kahn Award, sponsored by the French Academy of Sciences awarded by the SiF

10.1.6. *Research Administration*

J. Noyé is deputy head of the Automation, Production and Computer Sciences department of IMT Atlantique.

10.2. Teaching - Supervision - Juries

10.2.1. *Supervision*

- PhD: Md Sabbir Hasan, "Smart management of renewable energy in Clouds: from infrastructure to application", INSA Rennes, 05 March 2017, advisor: T. Ledoux
- HdR : Adrien Lebre, "Contributions to Large Scale Distributed Systems: The infrastructure viewpoint", University of Nantes, 1 Sept 2017.
- PhD: Emile Cadorel, director: J-M. Menaud, advisor: H. Coullon
- PhD: Fatima-zahra Boujdad, advisor: Mario Südholt
- PhD: Maverick Chardet, advisors: H. Coullon and A. Lebre
- PhD: Yewan Wang, director: J-M. Menaud.
- PhD: Mohammad Mahdi Bazm: codirectors: Mario Sudholt, J-M. Menaud.
- Postdoc: Dimitri Pertin, advisor: H. Coullon
- Postdoc: Chi Zhou, advisor: S. Ibrahim

10.2.2. *Juries*

- S. Ibrahim was member of the PhD Committee of Thomas Lambert, "Etude de l'effet de la replication de fichiers d'entree sur l'efficacite et la robustesse d'un ensemble de calcul", University of Bordeaux, September 2017.
- A. Lebre was member of the PhD Committee of Ismael Cuadrado-Cordero, "Microclouds: An Approach for a Network-Aware Energy-Efficient Decentralized Cloud", University of Rennes 1, Feb 2017.

- A. Lebre was member of the PhD Committee of Luis Pineda, "Efficient Support fo Data-Intensive Scientific Workflows on Geo-Distributed Clouds", University of Rennes 1, May 2017.

- A. Lebre was member of the PhD Committee of Aymen Jlassi, "Optimisation de la gestion des ressources sur une plateforme inforamtique de type Big-data basée sur le logiciel Hadoop", University of Tours, Dec 2017.

- T. Ledoux was a member of the PhD committee of Xuan Sang Le, "Software/FPGA Co-design for Edge-computing: Promoting Object-oriented Design", Univ. Bretagne Occidentale, May 2017

- J.-M. Menaud was reviewer of the PhD of Ines de Courchelle, "Vers une meilleure utilisation des énergies renouvelables : application à des bâtiments scientifiques",(Nov. 20, 2017), Toulouse ; Boris Teabe (Oct 12, 2017) "Performance et qualité de service de l'ordonnanceur dans un environnement virtualisé", Toulouse ; Alexis Martin (Jan. 13, 2017) "Infrastructure pour la gestion générique et optimisée des traces d'exécution pour les systèmes embarqués", Grenoble.

- J-C. Royer was member of the PhD comittee of Walid Benghabrit, "A Formal Model for Accountability", IMT Atlantique, October 27 2017 and Gwendal Daniet, "Efficient Persistence, Query and Tranformation of Large Models", IMT Atlantique, November 14, 2017.

- M. Südholt was a member of the PhD committee of Pauline Bolignano, "Formal Models and Verification of Memory Management in a Hypervisor", Inria, May 2017.

11. Bibliography

Publications of the year

Doctoral Dissertations and Habilitation Theses

[1] M. S. HASAN. *Smart management of renewable energy in clouds : from infrastructure to application*, INSA de Rennes, May 2017, https://tel.archives-ouvertes.fr/tel-01617963

[2] A. LEBRE. *Contributions to Large-scale Distributed Systems The infrastructure Viewpoint - Toward Fog and Edge Computing as The Next Utility Computing Paradigm?*, Université de Nantes, September 2017, Habilitation à diriger des recherches, https://tel.archives-ouvertes.fr/tel-01664343

[3] Y. LI. *Resource allocation in a Cloud partially powered by renewable energy sources*, Ecole nationale supérieure Mines-Télécom Atlantique, June 2017, https://tel.archives-ouvertes.fr/tel-01595953

Articles in International Peer-Reviewed Journals

[4] F. ALVARES, E. RUTTEN, L. SEINTURIER. *A Domain-specific Language for The Control of Self-adaptive Component-based Architecture*, in "Journal of Systems and Software", January 2017, https://hal.archives-ouvertes.fr/hal-01450517

[5] N. BELDICEANU, B. DUMAS FERIS, P. GRAVEY, S. HASAN, C. JARD, T. LEDOUX, Y. LI, D. LIME, G. MADI-WAMBA, J.-M. MENAUD, P. MOREL, M. MORVAN, M.-L. MOULINARD, A.-C. ORGERIE, J.-L. PAZAT, O. H. ROUX, A. SHARAIHA. *Towards energy-proportional Clouds partially powered by renewable energy*, in "Computing", January 2017, vol. 99, n^o 1, 20 p. [*DOI : 10.1007/S00607-016-0503-Z*], https://hal.inria.fr/hal-01340318

[6] B. CONFAIS, A. LÈBRE, B. PARREIN. *Performance Analysis of Object Store Systems in a Fog and Edge Computing Infrastructure*, in "Transactions on Large-Scale Data- and Knowledge-Centered Systems", August 2017 [*DOI : 10.1007/978-3-662-55696-2_2*], https://hal.archives-ouvertes.fr/hal-01587459

[7] H. COULLON, J. BIGOT, C. PÉREZ. *Extensibility and Composability of a Multi-Stencil Domain Specific Framework*, in "International Journal of Parallel Programming", November 2017 [*DOI :* 10.1007/S10766-017-0539-5], https://hal.archives-ouvertes.fr/hal-01650998

[8] M. S. HASAN, F. ALVARES, T. LEDOUX, J.-L. PAZAT. *Investigating Energy Consumption and Performance Trade-off for Interactive Cloud Application*, in "IEEE Transactions on Sustainable Computing", June 2017 [*DOI :* 10.1109/TSUSC.2017.2714959], https://hal.archives-ouvertes.fr/hal-01540159

[9] M. S. HASAN, Y. KOUKI, T. LEDOUX, J.-L. PAZAT. *Exploiting Renewable sources: when Green SLA becomes a possible reality in Cloud computing*, in "IEEE Transactions on Cloud Computing", June 2017, vol. 5, n⁰ 2, pp. 249-262 [*DOI :* 10.1109/TCC.2015.2459710], https://hal.archives-ouvertes.fr/hal-01187907

[10] Y. LI, A.-C. ORGERIE, I. RODERO, B. LEMMA AMERSHO, M. PARASHAR, J.-M. MENAUD. *End-to-end Energy Models for Edge Cloud-based IoT Platforms: Application to Data Stream Analysis in IoT*, in "Future Generation Computer Systems", 2017, pp. 1-12, forthcoming [*DOI :* 10.1016/J.FUTURE.2017.12.048], https://hal.archives-ouvertes.fr/hal-01673501

[11] S. MARTINEZ, M. TISI, R. DOUENCE. *Reactive Model Transformation with ATL*, in "Science of Computer Programming", March 2017, vol. 136, pp. 1 - 16 [*DOI :* 10.1016/J.SCICO.2016.08.006], https://hal.inria.fr/hal-01627991

International Conferences with Proceedings

[12] M. ABDERRAHIM, M. OUZZIF, K. GUILLOUARD, J. FRANCOIS, A. LÈBRE. *A Holistic Monitoring Service for Fog/Edge Infrastructures: a Foresight Study*, in "The IEEE 5th International Conference on Future Internet of Things and Cloud (FiCloud 2017)", Prague, Czech Republic, August 2017, pp. 337-344 [*DOI :* 10.1109/FICLOUD.2017.30], https://hal.archives-ouvertes.fr/hal-01591161

[13] F. ALVARES, G. DELAVAL, E. RUTTEN, L. SEINTURIER. *Language Support for Modular Autonomic Managers in Reconfigurable Software Components*, in "SeAC 2017 - 2nd Workshop on Self-Aware Computing @ ICAC'17", Colombus, United States, July 2017, pp. 271 - 278 [*DOI :* 10.1109/ICAC.2017.48], https://hal.inria.fr/hal-01575353

[14] M.-M. BAZM, M. LACOSTE, M. SÜDHOLT, J.-M. MENAUD. *Side-Channels Beyond the Cloud Edge : New Isolation Threats and Solutions*, in "IEEE International Conference on Cyber Security in Networking (CSNet) 2017", Rio de Janeiro, Brazil, 1st Cyber Security in Networking Conference (CSNet'17), October 2017, https://hal.inria.fr/hal-01593144

[15] N. BERTHIER, F. ALVARES, H. MARCHAND, G. DELAVAL, E. RUTTEN. *Logico-numerical Control for Software Components Reconfiguration*, in "CCTA 2017 - IEEE Conference on Control Technology and Applications", Mauna Lani, HI, United States, IEEE, August 2017, pp. 1599 - 1606 [*DOI :* 10.1109/CCTA.2017.8062685], https://hal.inria.fr/hal-01644754

[16] R.-A. CHERRUEAU, D. PERTIN, A. SIMONET, A. LEBRE, M. SIMONIN. *Toward a Holistic Framework for Conducting Scientific Evaluations of OpenStack*, in "The 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID)", Madrid, Spain, May 2017, pp. 544-548 [*DOI :* 10.1109/CCGRID.2017.87], https://hal.archives-ouvertes.fr/hal-01664515

[17] B. CONFAIS, A. LÈBRE, B. PARREIN. *An Object Store Service for a Fog/Edge Computing Infrastructure based on IPFS and Scale-out NAS*, in "1st IEEE International Conference on Fog and Edge Computing - ICFEC'2017", Madrid, Spain, 2017 IEEE 1st International Conference on Fog and Edge Computing (ICFEC 2017), May 2017, pp. 41-50 [*DOI :* 10.1109/ICFEC.2017.13], https://hal.archives-ouvertes.fr/hal-01483702

[18] H. COULLON, G. LE LOUËT, J.-M. MENAUD. *Virtual Machine Placement for Hybrid Cloud using Constraint Programming*, in "ICPADS 2017", Shenzhen, China, December 2017, https://hal.archives-ouvertes.fr/hal-01581969

[19] H. COULLON, J. NOYÉ. *Reconsidering the Relationship between Cloud Computing and Cloud Manufacturing*, in "SOHOMA 2017", Nantes, France, October 2017, https://hal.inria.fr/hal-01591113

[20] H. COULLON, D. PERTIN, C. PÉREZ. *Production Deployment Tools for IaaSes: an Overall Model and Survey*, in "The IEEE 5th International Conference on Future Internet of Things and Cloud (FiCloud)", Prague, Czech Republic, August 2017, pp. 183-190 [*DOI :* 10.1109/FICLOUD.2017.51], https://hal.inria.fr/hal-01532489

[21] S. DUPONT, S. BOURI, F. ALVARES DE OLIVEIRA, T. LEDOUX. *ElaScript: a DSL for Coding Elasticity in Cloud Computing*, in "32nd ACM Symposium on Applied Computing - Track on Cloud Computing", Marrakesh, Morocco, Proceedings of the 32nd ACM Symposium on Applied Computing - Track on Cloud Computing, April 2017, https://hal.archives-ouvertes.fr/hal-01400236

[22] G. GILBERT. *Formalising Real Numbers in Homotopy Type Theory*, in "6th ACM SIGPLAN Conference on Certified Programs and Proofs", Paris, France, January 2017, pp. 112 - 124, https://arxiv.org/abs/1610.05072 [*DOI :* 10.1145/3018610.3018614], https://hal.inria.fr/hal-01449326

[23] M. S. HASAN, F. ALVARES DE OLIVEIRA JR., T. LEDOUX. *GPaaScaler: Green Energy aware Platform Scaler for Interactive Cloud Application*, in "IEEE/ACM International Conference on Utility and Cloud Computing", Austin, United States, December 2017, https://hal.archives-ouvertes.fr/hal-01650136

[24] Y. HEBBAL, L. SYLVIE, J.-M. MENAUD. *Hidden Process Detection using Kernel Functions Instrumentation*, in "IEEE Conference on Dependable and Secure Computing (DSC)", Taiwan, Taiwan, August 2017, https://hal.inria.fr/hal-01520864

[25] Y. HEBBAL, L. SYLVIE, J.-M. MENAUD. *K-binID: Kernel Binary Code Identification for Virtual Machine Introspection*, in "IEEE Conference on Dependable and Secure Computing (DSC)", Taiwan, Taiwan, August 2017, https://hal.inria.fr/hal-01520867

[26] A. LEBRE, J. PASTOR, A. SIMONET, F. DESPREZ. *Revising OpenStack to Operate Fog/Edge Computing infrastructures*, in "IEEE International Conference on Cloud Engineering", Vancouver, France, April 2017, https://hal.inria.fr/hal-01273427

[27] *Best Paper*
J. LEJEUNE, F. ALVARES, T. LEDOUX. *Towards a generic autonomic model to manage Cloud Services*, in "The 7th International Conference on Cloud Computing and Services Science (CLOSER 2017)", Porto, Portugal, April 2017, https://hal.archives-ouvertes.fr/hal-01511360.

[28] Y. LI, A.-C. ORGERIE, J.-M. MENAUD. *Balancing the use of batteries and opportunistic scheduling policies for maximizing renewable energy consumption in a Cloud data center*, in "PDP 2017 - 25th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing", St Petersburg, Russia, March 2017, https://hal.inria.fr/hal-01432752

[29] Y. LI, A.-C. ORGERIE, I. RODERO, M. PARASHAR, J.-M. MENAUD. *Leveraging Renewable Energy in Edge Clouds for Data Stream Analysis in IoT*, in "CCGrid 2017 - IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing", Madrid, Spain, May 2017, https://hal.inria.fr/hal-01472358

[30] G. MADI WAMBA, Y. LI, A.-C. ORGERIE, N. BELDICEANU, J.-M. MENAUD. *Cloud workload prediction and generation models*, in "SBAC-PAD: International Symposium on Computer Architecture and High Performance Computing", Campinas, Brazil, October 2017, https://hal.archives-ouvertes.fr/hal-01578354

[31] G. MADI WAMBA, Y. LI, A.-C. ORGERIE, N. BELDICEANU, J.-M. MENAUD. *Green energy aware scheduling problem in virtualized datacenters*, in "ICPADS 2017", Shenzen, China, December 2017, https://hal.archives-ouvertes.fr/hal-01582936

[32] T. L. NGUYEN, A. LEBRE. *Virtual Machine Boot Time Model*, in "PDP 2017 - 25th Euromicro International Conference on Parallel, Distributed and Network-based Processing", St Peterbourg, Russia, March 2017, pp. 430 - 437 [*DOI :* 10.1109/PDP.2017.58], https://hal.inria.fr/hal-01586932

[33] T.-D. PHAN, S. IBRAHIM, A. C. ZHOU, G. AUPY, G. ANTONIU. *Energy-Driven Straggler Mitigation in MapReduce*, in "Euro-Par'17 - 23rd International European Conference on Parallel and Distributed Computing", Santiago de Compostela, Spain, August 2017, https://hal.inria.fr/hal-01560044

[34] R. POTTIER, J.-M. MENAUD. *Privacy-Aware Data Storage in Cloud Computing*, in "International Conference on Cloud Computing and Services Science (CLOSER)", Porto, Portugal, April 2017, 8 p. , https://hal.inria.fr/hal-01472133

[35] Y. TALEB, S. IBRAHIM, G. ANTONIU, T. CORTES. *Characterizing Performance and Energy-Efficiency of The RAMCloud Storage System*, in "The 37th IEEE International Conference on Distributed Computing Systems (ICDCS 2017)", Atlanta, United States, June 2017, https://hal.inria.fr/hal-01496959

[36] O. YILDIZ, A. C. ZHOU, S. IBRAHIM. *Eley: On the Effectiveness of Burst Buffers for Big Data Processing in HPC systems*, in "Cluster'17-2017 IEEE International Conference on Cluster Computing", Hawaii, United States, September 2017, https://hal.inria.fr/hal-01570737

[37] A. C. ZHOU, S. IBRAHIM, B. HE. *On Achieving Efficient Data Transfer for Graph Processing in Geo-Distributed Datacenters*, in "ICDCS'17- The 37th IEEE International Conference on Distributed Computing Systems (ICDCS 2017)", Atlanta, United States, June 2017, https://hal.inria.fr/hal-01560187

National Conferences with Proceedings

[38] S. BOULIER, P.-M. PÉDROT, N. TABAREAU. *Modèles de la théorie des types donnés par traduction de programme*, in "28ièmes Journées Francophones des Langages Applicatifs", Gourette, France, January 2017, https://hal.archives-ouvertes.fr/hal-01503089

[39] B. CONFAIS, A. LÈBRE, B. PARREIN. *An object store for Fog infrastructures based on IPFS and a Scale-Out NAS*, in "RESCOM 2017", Le Croisic, France, RESCOM 2017 Summer school, June 2017, 2 p. , https://hal.archives-ouvertes.fr/hal-01559065

[40] Y. HEBBAL, L. SYLVIE, J.-M. MENAUD. *K-binID : Identification de code binaire noyau à des fins d'introspection de machine virtuelle*, in "Conférence d'informatique en Parallélisme, Architecture et Système (ComPAS)", Sophia Antipolis, France, June 2017, https://hal.inria.fr/hal-01520863

Conferences without Proceedings

[41] I. CUADRADO-CORDERO, A.-C. ORGERIE, J.-M. MENAUD. *Comparative Experimental Analysis of the Quality-of-Service and Energy-Efficiency of VMs and Containers' Consolidation for Cloud Applications*, in "SoftCOM: International Conference on Software, Telecommunications and Computer Networks", Split, Croatia, September 2017, https://hal.archives-ouvertes.fr/hal-01578325

Scientific Books (or Scientific Book chapters)

[42] D. SAUCEZ, A. LÈBRE, S. SECCI (editors). *RESCOM 2017 Summer school*, CNRS, June 2017, https://hal.inria.fr/hal-01558074

Scientific Popularization

[43] M. ALBONICO, J.-M. MOTTU, G. SUNYÉ, F. ALVARES. *Making Cloud-based Systems Elasticity Testing Reproducible*, in "7th International Conference on Cloud Computing and Services Science", Porto, Portugal, April 2017, https://hal.inria.fr/hal-01471916

Other Publications

[44] M.-M. BAZM, M. LACOSTE, M. SÜDHOLT, J.-M. MENAUD. *Side Channels in the Cloud: Isolation Challenges, Attacks, and Countermeasures*, March 2017, working paper or preprint, https://hal.inria.fr/hal-01591808

[45] S. BOULIER, N. TABAREAU. *Model structure on the universe in a two level type theory*, 2017, working paper or preprint, https://hal.archives-ouvertes.fr/hal-01579822

References in notes

[46] M. AKŞIT, S. CLARKE, T. ELRAD, R. E. FILMAN (editors). *Aspect-Oriented Software Development*, Addison-Wesley Professional, September 2004

[47] C. ALLAN, P. AVGUSTINOV, A. S. CHRISTENSEN, L. HENDREN, S. KUZINS, O. LHOTÁK, O. DE MOOR, D. SERENI, G. SITTAMPALAM, J. TIBBLE. *Adding trace matching with free variables to AspectJ*, in "ACM Conference on Object-Oriented Programming, Systems and Languages (OOPSLA)", R. P. GABRIEL (editor), ACM Press, 2005

[48] R. ALLEN, D. GARLAN. *A Formal Basis for Architectural Connection*, in "ACM Transactions on Software Engineering and Methodology", July 1997, vol. 6, n^o 3, pp. 213–49

[49] J. H. ANDREWS. *Process-Algebraic Foundations of Aspect-Oriented Programming*, in "Proceedings of the 3rd International Conference on Metalevel Architectures and Separation of Crosscutting Concerns", Lecture Notes in Computer Science, 2001, vol. 2192, pp. 187–209

[50] M. ARMBRUST, A. FOX, R. GRIFFITH, A. D. JOSEPH, R. KATZ, A. KONWINSKI, G. LEE, D. PATTERSON, A. RABKIN, I. STOICA, M. ZAHARIA. *A view of cloud computing*, in "Communications of the ACM", 2010, vol. 53, n° 4, pp. 50–58

[51] T. H. AUSTIN, C. FLANAGAN. *Multiple facets for dynamic information flow*, in "Proceedings of the 39th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages", New York, USA, POPL '12, ACM, 2012, pp. 165–178, http://doi.acm.org/10.1145/2103656.2103677

[52] A. BELOGLAZOV, R. BUYYA. *Energy efficient resource management in virtualized cloud data centers*, in "in: Proceedings of the 2010 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing, CCGRID'10, IEEE Computer Society", 2010, pp. 826–831

[53] L. D. BENAVIDES NAVARRO, M. SÜDHOLT, W. VANDERPERREN, B. DE FRAINE, D. SUVÉE. *Explicitly distributed AOP using AWED*, in "Aspect-Oriented Software Development (AOSD)", ACM Press, March 2006, pp. 51-62

[54] G. S. BLAIR, G. COULSON, P. ROBIN, M. PAPATHOMAS. *An architecture for next generation middleware*, in "Proceedings of the IFIP International Conference on Distributed Systems Platforms and Open Distributed Processing", Springer-Verlag, 1998

[55] A. BRACCIALIA, A. BROGI, C. CANAL. *A formal approach to component adaptation*, in "Journal of Systems and Software", 2005

[56] R. E. BROWN, E. R. MASANET, B. NORDMAN, W. F. TSCHUDI, A. SHEHABI, J. STANLEY, J. G. KOOMEY, D. A. SARTOR, P. T. CHAN. *Server and Data Center Energy Efficiency: Public Law 109-431*, in "Report to Congress", 06/2008 2008

[57] S. CAPECCHI, I. CASTELLANI, M. DEZANI-CIANCAGLINI, T. REZK. *Session Types for Access and Information Flow Control*, in "CONCUR 2010 - Concurrency Theory, 21th International Conference, CONCUR 2010, Paris, France, August 31-September 3, 2010. Proceedings", P. GASTIN, F. LAROUSSINIE (editors), Lecture Notes in Computer Science, Springer, 2010, vol. 6269, pp. 237–252, http://dx.doi.org/10.1007/978-3-642-15375-4_17

[58] E. M. CLARKE, O. GRUMBERG, D. A. PELED. *Model Checking*, The MIT Press, Cambridge, Massachusetts, 1999

[59] A. COLYER, A. CLEMENT. *Large-scale AOSD for Middleware*, in "Proceedings of the 3rd ACM Int. Conf. on Aspect-Oriented Software Development (AOSD), Lancaster", K. LIEBERHERR (editor), ACM Press, 2004, pp. 56–65

[60] F. DEREMER, H. H. KRON. *Programming-in-the-large versus programming-in-the-small*, in "IEEE Transactions on Software Engineering", 1976, vol. SE-2, n° 2, pp. 80-86

[61] G. DECKER, O. KOPP, F. LEYMANN, M. WESKE. *BPEL4Chor: Extending BPEL for Modeling Choreographies*, in "IEEE International Conference on Web Services (ICWS 2007)", IEEE Computer Society, 2007, pp. 296–303

[62] E. W. DIJKSTRA. *On the role of scientific thought*, in "Selected Writings on Computing: A Personal Perspective", Springer-Verlag, 1974, pp. 60–66, Published in 1982

[63] R. DOUENCE, P. FRADET, M. SÜDHOLT. *A framework for the detection and resolution of aspect interactions*, in "Proceedings of the ACM SIGPLAN/SIGSOFT Conference on Generative Programming and Component Engineering (GPCE'02)", Lecture Notes in Computer Science, Springer-Verlag, October 2002, vol. 2487, pp. 173–188, http://hal.inria.fr/inria-00072153

[64] R. DOUENCE, P. FRADET, M. SÜDHOLT. *Trace-Based Aspects*, in "Aspect-Oriented Software Development", M. AKŞIT, S. CLARKE, T. ELRAD, R. E. FILMAN (editors), Addison-Wesley, 2004, pp. 201-218

[65] R. DOUENCE, O. MOTELET, M. SÜDHOLT. *A formal definition of crosscuts*, in "Proceedings of the 3rd International Conference on Metalevel Architectures and Separation of Crosscutting Concerns", Lecture Notes in Computer Science, Springer-Verlag, 2001, vol. 2192, pp. 170–186

[66] R. DOUENCE, D. LE BOTLAN, J. NOYÉ, M. SÜDHOLT. *Concurrent Aspects*, in "Proc. of the Int. ACM Conf. on Generative Programming and Component Engineering (GPCE)", ACM Press, October 2006, pp. 79-88

[67] H. FOSTER, S. UCHITEL, J. MAGEE, J. KRAMER. *Model-based Verification of Web Service Compositions*, in "Proceedings of the 18th IEEE Int. Conf. on Automated Software Engineering (ASE'03)", IEEE Computer Society, 2003, pp. 152–163

[68] I. FOSTER, Y. ZHAO, I. RAICU, S. LU. *Cloud computing and grid computing 360-degree compared*, in "Grid Computing Environments Workshop, 2008. GCE'08", Ieee, 2008, pp. 1–10

[69] A. FUGGETTA, G. P. PICCO, G. VIGNA. *Understanding Code Mobility*, in "IEEE Transactions on Software Engineering", May 1998, vol. 24, n[o] 5, pp. 342–361

[70] GREENPEACE. *Make IT green: Cloud computing and its contribution to climate change*, Greenpeace International, March 2010

[71] K. HONDA, N. YOSHIDA, M. CARBONE. *Multiparty asynchronous session types*, in "Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008, San Francisco, California, USA, January 7-12, 2008", G. C. NECULA, P. WADLER (editors), ACM, 2008, pp. 273–284, http://www.doc.ic.ac.uk/~yoshida/multiparty/multiparty.pdf

[72] G. KICZALES, E. HILSDALE, J. HUGUNIN, M. KERSTEN, J. PALM, W. G. GRISWOLD. *An Overview of AspectJ*, in "ECOOP 2001 — Object-Oriented Programming 15th European Conference, Budapest Hungary", Berlin, J. L. KNUDSEN (editor), Lecture Notes in Computer Science, Springer-Verlag, June 2001, vol. 2072, pp. 327–353, http://www.eclipse.org/aspectj/

[73] G. KICZALES. *Aspect Oriented Programming*, in "Proc. of the Int. Workshop on Composability Issues in Object-Orientation (CIOO'96) at ECOOP", July 1996, Selected paper published by dpunkt press, Heidelberg, Germany

[74] G. KICZALES, J. DES RIVIERES, DANIEL G. BOBROW. *The Art of the Meta-Object Protocol*, MIT Press, Cambridge (MA), USA, 1991

[75] J. KIENZLE, R. GUERRAOUI. *AOP - Does It Make Sense? The Case of Concurrency and Failures*, in "16th European Conference on Object-Oriented Programming (ECOOP'2002)", Malaga, Spain, B. MAGNUSSON (editor), Lecture Notes in Computer Science, Springer-Verlag, 2002

[76] T. LEDOUX. *OpenCorba: a Reflective Open Broker*, in "ACM Meta-Level Architectures and Reflection, Second International Conference, Reflection'99", Saint-Malo, France, P. COINTE (editor), Lecture Notes in Computer Science, Springer-Verlag, July 1999, vol. 1616, pp. 197–214

[77] M. MCILROY. *Mass produced software components*, in "Mass produced software components", Garmish, Germany, P. NAUR, B. RANDELL (editors), NATO Science Committee, October 1968, pp. 138-155

[78] N. MEDVIDOVIC, R. N. TAYLOR. *A Classification and Comparison Framework for Software Architecture Description Languages*, in "IEEE Transactions on Software Engineering", January 2000, vol. 26, n⁰ 1, pp. 70-93

[79] M. MERNIK, J. HEERING, A. M. SLOANE. *When and How to Develop Domain-Specific Languages*, in "ACM Computing Surveys", December 2005, vol. 37, n⁰ 4, pp. 316-344

[80] L. MIKHAJLOV, E. SEKERINSKI. *A study of the fragile base class*, in "A study of the fragile base class", Brussels, Belgium, E. JUL (editor), Lecture Notes in Computer Science, July 1998, vol. 1445, pp. 355-382

[81] D. H. NGUYEN, M. SÜDHOLT. *VPA-based aspects: better support for AOP over protocols*, in "4th IEEE International Conference on Software Engineering and Formal Methods (SEFM'06)", IEEE Computer Society Press, September 2006

[82] O. NIERSTRASZ. *Regular Types for Active Objects*, in "Object-Oriented Software Composition", O. NIERSTRASZ, D. TSICHRITZIS (editors), Prentice Hall, 1995, chap. 4, pp. 99–121

[83] M. NISHIZAWA, S. CHIBA, M. TATSUBORI. *Remote Pointcut - A Language Construct for Distributed AOP*, in "Proceedings of the 3rd ACM Int. Conf. on Aspect-Oriented Software Development (AOSD), Lancaster", ACM Press, 2004

[84] D. L. PARNAS. *On the criteria for decomposing systems into modules*, in "Communications of the ACM", December 1972, vol. 15, n⁰ 12, pp. 1053-1058

[85] S. PEARSON. *Toward Accountability in the Cloud*, in "Internet Computing, IEEE", July-Aug. 2011, vol. 15, n⁰ 4, pp. 64-69, http://dx.doi.org/10.1109/MIC.2011.98

[86] F. PLASIL, S. VISNOVSKY. *Behavior Protocols for Software Components*, in "Transactions on Software Engineering", January 2002, vol. 28, n⁰ 9

[87] F. PUNTIGAM. *Coordination Requirements Expressed in Types for Active Objects*, in "ECOOP'97—Object-Oriented Programming", M. AKŞIT, S. MATSUOKA (editors), Lecture Notes in Computer Science, Springer-Verlag, 1997, vol. 1241, pp. 367–388

[88] N. SHARMA, S. BARKER, D. IRWIN, P. SHENOY. *Blink: managing server clusters on intermittent power*, in "SIGARCH Comput. Archit. News", March 2011, vol. 39, pp. 185–198, http://dx.doi.org/10.1145/1961295.1950389

[89] M. SHAW, D. GARLAN. *Software Architecture: Perspectives on an Emerging Discipline*, Prentice-Hall, 1996

[90] B. C. SMITH. *Reflection and Semantics in LISP*, Xerox Palto Alto Research Center, Palo Alto, 1984, n⁰ P84-00030

[91] S. SOARES, E. LAUREANO, P. BORBA. *Implementing distribution and persistence aspects with AspectJ* , in "Proceedings of the 17th ACM conference on Object-oriented programming, systems, languages, and applications (OOPSLA-02)", C. NORRIS, J. J. B. FENWICK (editors), ACM SIGPLAN Notices, ACM Press, November 4–8 2002, vol. 37, 11, pp. 174–190

[92] S. SUNDARESWARAN. *Ensuring Distributed Accountability for Data Sharing in the Cloud*, in "Dependable and Secure Computing", 2012, vol. 9, http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6165313

[93] R. J. WALKER, K. VIGGERS. *Implementing Protocols via Declarative Event Patterns*, in "Proceedings of the ACM SIGSOFT International Symposium on Foundations of Software Engineering (FSE-12)", ACM Press, 2004, pp. 159 - 169

[94] M. WAND, G. KICZALES, C. DUTCHYN. *A Semantics for Advice and Dynamic Join Points in Aspect-Oriented Programming*, in "ACM Transactions on Programming Languages and Systems (TOPLAS)", 2004, vol. 26, n⁰ 5, pp. 890–910

[95] D. M. YELLIN, R. E. STROM. *Protocol specifications and component adaptors*, in "ACM Transactions of Programming Languages and Systems", March 1997, vol. 19, n⁰ 2, pp. 292–333

[96] Q. ZHANG, L. CHENG, R. BOUTABA. *Cloud computing: state-of-the-art and research challenges*, in "Journal of internet services and applications", 2010, vol. 1, n⁰ 1, pp. 7–18

[97] A. VAN DEURSEN, P. KLINT, J. VISSER. *Domain-Specific Languages: An Annotated Bibliography*, in "ACM SIGPLAN Notices", June 2000, vol. 35, n⁰ 6, pp. 26-36