Activity Report 2014

Section Software

ATEAMS Project-Team

4. New Software and Platforms

4.1. MicroMachinations

Participant: Riemer Van Rozen [correspondent].

Characterization: A-2-up3, SO-4, SM-2-up3, EM-3, SDL-3-up4, OC-DA-3-CD-3-MS-3-TPM-3.

WWW:

Objective: To create an integrated, live environment for modeling and evolving game economies. This will allow game designers to experiment with different strategies to realize game mechanics. The environment integrates with the SPIN model checker to prove properties (reachability, liveness). A runtime system for executing game economies allows MicroMachinations models to be embedded in actual games.

Users: Game designers

Impact: One of the important problems in game software development is the distance between game design and implementation in software. MicroMachinations has the potential to bridge this gap by providing live design tools that directly modify or create the desired software behaviors.

Competition: None.

Engineering: The front-end of MicroMachinations is built using the Rascal language workbench, including visualization, model checking, debugging and standard IDE features. The runtime library is implemented in C++ and will be evaluated in the context of industrial game design.

Publications: [11]

4.1.1. Novelties

- MMLib was finished to allow the execution of game economies directly within games. This supports "Live programming" of the behavior of games. The library has been used in the development of the real-life game "Johnny Jetstream", designed by IC3DMedia.

4.2. Naked Object Algebras

Participant: Tijs Van Der Storm [correspondent].

Characterization: A5, SO-4, SM-4, EM-4, SDL-4-up5, OC-DA-3-CD-3-MS-3-TPM-3.

WWW: https://github.com/cwi-swat/naked-object-algebras

Objective: Supporting modular and extensible language development.

Users: Programmers, language designers.

Impact: Object Algebras promise a new level of modularity and extensibility in the implementation of recursive data types. The NAO framework lifts this to the implementation of software languages, including the declarative declaration of concrete syntax.

Competition: Language prototyping tools.

Engineering: NAO consists of a few hundred lines of Java code. It has no external dependencies, except ANTLR for parsing.

Publications: [27], [33]

4.2.1. Novelties

- NAO has been used to develop an extensible variant of the QL questionnaire language [33].

4.3. Rascal

Participants: Paul Klint, Jurgen Vinju [correspondent], Tijs Van Der Storm, Pablo Inostroza Valdera, Davy Landman, Bert Lisser, Atze Van Der Ploeg, Vadim Zaytsev, Anastasia Izmaylova, Michael Steindorfer, Jouke Stoel, Ali Afroozeh, Ashim Shahi.

Characterization: A5, SO-4, SM-4, EM-4, SDL-4-up5, OC-DA-3-CD-3-MS-3-TPM-3.

WWW: http://www.rascal-mpl.org

Objective: Provide a completely integrated programming language parametric meta programming language for the construction of any kind of meta program for any kind of programming language: analysis, transformation, generation, visualization.

Users: Researchers in model driven engineering, programming languages, software engineering, software analysis, as well as practitioners that need specialized tools.

Impact: Rascal is making the mechanics of meta programming into a non-issue. We can now focus on the interesting details of the particular fact extraction, model, source analysis, domain analysis as opposed to being distracted by the engineering details. Simple things are easy in Rascal and complex things are manageable, due to the integration, the general type system and high-level programming features.

Competition: There is a plethora of meta programming toolboxes and frameworks available, ranging from plain parser generators to fully integrated environments. Rascal is distinguished because it is a programming language rather than a specification formalism and because it completely integrates different technical domains (syntax definition, term rewriting, relational calculus). For simple tools, Rascal competes with scripting languages and for complex tools it competes context-free general parser generators, with query engines based on relational calculus and with term rewriting and strategic programming languages.

Engineering: Rascal is about 100 kLOC of Java code, designed by a core team of three and with a team of around 8 PhD students and post-docs contributing to its design, implementation and maintenance. The goal is to work towards more bootstrapping and less Java code as the project continues.

Publications: [7], [6], [8], [5], [6]

4.3.1. Novelties

- Improvements of the language-parametric model to represent software projects (M3) [9].
- Performance improvements of the Rascal interpreter throughout.
- Further improvements to the compiler for Rascal, based on new language construct guarded coroutines.
- New language feature: keyword parameters. This will further allow simplificiation of the core language, as well as support better extensibility.
- Significant improvements to the Rascal static type checker.
- Further improvements to the new GLL parser (Iguana).
- Design of a new DSL for describing core banking infrastructure was started (ReBEL). Rascal was also used to develop a state machine DSL for use in embedded devices (Machino).

4.4. IDE Meta-tooling Platform

Participants: Jurgen Vinju [correspondent], Michael Steindorfer.

IMP, the IDE meta tooling platform is an Eclipse plugin developed mainly by the team of Robert M. Fuhrer at IBM TJ Watson Research institute. It is both an abstract layer for Eclipse, allowing rapid development of Eclipse based IDEs for programming languages, and a collection of meta programming tools for generating source code analysis and transformation tools.

Characterization: A5, SO-3, SM4-up5, EM-4, SDL-5, DA-2-CD-2-MS-2-TPM-2

WWW: https://github.com/impulse-org/

Objective: The IDE Meta Tooling Platform (IMP) provides a high-level abstraction over the Eclipse API such that programmers can extend Eclipse with new programming languages or domain specific languages in a few simple steps. IMP also provides a number of standard meta tools such as a parser generator and a domain specific language for formal specifications of configuration parameters.

Users: Designers and implementers of IDEs for programming languages and domain specific languages. Also, designers and implementers of meta programming tools.

Impact: IMP is a popular among meta programmers especially for it provides the right level of abstraction.

Competition: IMP competes with other Eclipse plugins for meta programming (such as Model Driven Engineering tools), but its API is more general and more flexible. IMP is a programmers framework rather than a set of generators.

Engineering: IMP is a long-lived project of many contributors, which is managed as an Eclipse incubation project at `eclipse.org`. Currently we are moving the project to Github to explore more different ways of collaboration.

Publications: [2] [29]

4.4.1. Novelties

- Significant performance improvements to the IMP program database. Performance is now better than equivalent data structure libraries in Scala and Clojure.

4.5. Ensō

Participant: Tijs Van Der Storm [correspondent].

Characterization: A5, SO-4, SM-3-up-4, EM-2-up-4, SDL-4, OC-DA-4-CD-4-MS-4-TPM-4

WWW: http://www.enso-lang.org

Objective: Together with Prof. Dr. William R. Cook of the University of Texas at Austin, and Alex Loh, Tijs van der Storm has been designing and implementing a new programming system, called Ensō. Ensō is theoretically sound and practical reformulation of model-based development. It is based on model-interpretation as opposed to model transformation and code generation. Currently, the system already supports models for schemas (data models), web applications, context-free grammars, diagram editors and security.

Users: All programmers.

Impact: Ensō has the potential to revolutionize the activity of programming. By looking at model driven engineering from a completely fresh perspective, with as key ingredients interpreters and partial evaluation, it may make higher level (domain level) program construction and maintenance as effective as normal programming.

Competition: Ensō competes as a programming paradigm with model driven engineering tools and generic programming and languages that provide syntax macros and language extensions.

Engineering: Ensō is a completely self-hosted system in 7000 lines of code.

Publications: [14], [16], [13]

BONSAI Project-Team

5. New Software and Platforms

5.1. SortMeRNA – Metatranscriptome classification

Software web site: http://bioinfo.lille.inria.fr/RNA/sortmerna

Licence: GPL

Objective: *SortMeRNA* is a tool designed to rapidly filter ribosomal RNA fragments from metatranscriptomic data produced by next-generation sequencers. It is available for download from our website, or through the open web-based platform Galaxy. The development version is also available on GitHub. *SortMeRNA* was first released in October 2012. It is now used in production by Genoscope (French National Center for Sequencing) to process all metatranscriptomic data of the Tara Ocean Expedition, and has been integrated in several other computational pipelines (Qiime developed at University of Colorado at Boulder, MetaMetadb developed at University of Tokyo, Leimena pipeline developed at Wageningen University,...).

SortMeRNA is still under development through a partnership with the Knight lab (University of Colorado at Boulder). Version 2.0 has been released in November 2014, and has extended functionalities. It can now perform sequence alignments to any ribosomal RNA database, which allows the user to study the taxonomic content of a microbial sample. This new version has been presented at the international workshops [12], [11].

5.2. Vidjil – Quantifying lymphocyte rearrangements in high-throughput sequencing data

Software web site: http://bioinfo.lille.inria.fr/vidjil/

Objective: **Vidjil** is a platform for high-throughput V(D)J recombinations analysis. containing three components. The Vidjil *algorithm* process high-throuput sequencing data to extract V(D)J junctions and gather them into clones. Vidjil starts from a set of reads and detects "windows" overlapping the actual CDR3. This is based on an fast and reliable seed-based heuristic and allows to output all sequenced clones. The analysis is extremely fast because, in the first phase, no alignment is performed with database germline sequences [5]. The Vidjil dynamic *browser* is made for the visualization and analysis of clones and their tracking along the time in a "minimal residual disease" setup or in a immunological study. The browser visualize data processed by the Vidjil algorithm or by other V(D)J analysis pipeline and enables to explore further cluterings. Finally, a *patient database* with a server links the browser and the algorithmic part. The goal is that the clinicians will be able to upload, manage and process their runs on a server hosted in their hospital.

In 2014, the development of Vidjil was supported by the SIRIC OncoLille (Marc Duez). We developed the new patient database and added features both on the browser and on the algorithm (multi-system analysis). Several hospital labs in France and in Europe are testing Vidjil. The Lille hospital plans to use Vidjil in 2015 in a pre-production pipeline.

5.3. Norine – A resource for nonribosomal peptides

Software web site: http://bioinfo.lille.inria.fr/norine/

Objective: **Norine** is a public computational resource that contains a database of NRPs with a web interface and dedicated tools, such as a 2D graph viewer and editor for peptides or comparison of NRPs. Norine was created and is maintained by members of BONSAI team, in tight collaboration with members of the ProBioGEM lab, a microbial laboratory of Lille1 University. Since its creation in 2006, Norine has gained an international recognition as the unique database dedicated to non-ribosomal peptides because of its high quality and manually curated annotations, and has been selected by wwPDB as a reference database. It is queried from all around the world by biologists or biochemists. It receives more than 3000 queries per month.

To enhance the Norine resource, we have recently developed a new module, named MyNorine, which is an open interface for biologists and biochemists dedicated to the submission of new non-ribosomal peptides in Norine database. Up to now, peptides were manually inputted and verified before being added in the database, which could potentially lead to human errors. The goal of MyNorine is to help users during the submission of peptides and monomers, by guiding them during all steps. For that, users, all over the world, can create an account on MyNorine. Thus, they contribute to the Norine resource and become curators (author of a peptide entry is mentioned in the corresponding page of Norine). Submitted peptides/monomers are validated, through a workflow process, by Norine team members, to ensure correct and consistent entries.

5.4. miRkwood –microRNAs in plant genomes

Software web site: http://bioinfo.lille.inria.fr/mirkwood/

Objective: **miRkwood** is a web server for the identification of hairpin precursors of both conserved and non-conserved miRNAs in plant genomes. It is able to face the diversity of plant pre-miRNAs and is optimised to take advantage of their distinctive properties: Sequence length, secondary structure, free energy, miRNA conservation, stability of the miRNA/miRNA* duplex, ... Moreover, it offers an intuitive and comprehensive user interface to navigate in the data, as well as many export options to allow the user to conduct further analyses on a local computer. Ongoing work is concerned with integrating small RNA-seq data.

5.5. ProCARs

Software web site: http://bioinfo.lille.inria.fr/procars

Objective: **ProCARs** is a program used to reconstruct ancestral gene orders as CARs (Contiguous Ancestral Regions) with a progressive homology-based method. The method runs from a phylogeny tree, without branch lengths needed, with a marked ancestor and a block file. The method output CARs as sets of ordered contiguous blocks in the targeted ancestor. ProCARs has been developped with Python 2.7.5.

<p style="text-align: center; color: red;">**DOLPHIN Project-Team**</p>

5. New Software and Platforms

5.1. Comparing Continuous Optimizers (Coco) Platform

Participants: Dimo Brockhoff, Arnaud Liefooghe, Thanh-Do Tran.

The Coco Platform (coco.gforge.inria.fr) provides the functionality to automatically benchmark optimization algorithms for unbounded, unconstrained optimization problems in continuous domains. Benchmarking is a vital part of algorithm engineering and a necessary path to recommend algorithms for practical applications. The Coco platform releases algorithm developers and practitioners alike from (re-)writing test functions, logging, and plotting facilities by providing an easy-to-handle interface in several programming languages. The Coco platform has been developed since 2007 and has been used extensively within the "Blackbox Optimization Benchmarking (BBOB)" workshop series since 2009. Overall, 123 algorithms and algorithm variants by contributors from all over the world have been benchmarked with the platform so far and all data is publicly available for the research community (see for example http://coco.gforge.inria.fr/doku.php?id=bbob-2013-algorithms for the submissions to BBOB-2013).

Via the ANR project NumBBO, also Dolphin team members are involved in the development of Coco and it is one of the main purposes of NumBBO to extend the Coco platform towards expensive, large-scale, constrained and multiobjective optimization. In order to facilitate these extensions, a complete overhaul of the platform is currently underway—rewriting the whole functionality from scratch in a single language (ANSI C) which will then be called from all other provided languages (Java, python, Matlab/Octave, R) instead of multiple independent implementations which are highly difficult to maintain. A first release of the new code base is expected in the first half of 2015.

For the rewriting of the source code, we also moved to a publicly available open source code repository at https://code.google.com/p/numbbo/ which, in addition, provides for the first time a bug and feature request tracking system for Coco. As the first Coco-related deliverable of the NumBBO project, the extension towards expensive optimization has been finished this year. Its full functionality will be provided for the first time for the upcoming BBOB special session at the IEEE Congress on Evolutionary Computation (CEC'2015). Coco is also currently extended towards multiobjective optimization in close relation to the PhD topic of Thanh-Do Tran. A first working (but still preliminary) version of the multiobjective Coco has been developed and is expected to be merged with the newly rewritten Coco code just after its first release.

Inria software self-assessment for Coco: [A-4/5, SO-4, SM-3-up4, EM-3, SDL-4-up5][DA-3, CD-3, MS-3, TPM-2]

5.2. MO-Mine

Participants: Clarisse Dhaenens, Benjamin Fisset, Laetitia Jourdan.

The MO-Mine platform (mo-mine.gforge.inria.fr) aims at providing optimization algorithms, and in particular multi-objective approaches, to deal with data-mining classical tasks (Classification, association rules...). The platform is still in development in the context of an Inria ADT.

$MO - Mine_{clust}$ is the first package of the platform and is dedicated to clustering (unsupervised classification). Indeed, it is well-known that clustering may be seen as a bi-objective optimization problem as the goal is both to minimize distances between data belonging to a same cluster, while maximizing distances between data belonging to different clusters. Several models (objective functions used,...) and engines (optimization algorithms) have been implemented. The framework searches, for a given dataset, the best association of model/engine/parameter without specifiing the number of clusters. $MO - Mine_{clust}$ shows very interesting behavior and shows that the choice of the model and the engine has a great importance in the performance of the method and depends on the dataset to analyze.

Inria software self-assessment for $MO - Mine_{clust}$: [A-1, SO-4, SM-3-up4, EM-2, SDL-1,DA-4,CD-4]

5.3. Platforms

Grid'5000 [0] is a French nation-wide computational grid infrastructure composed of several clusters of processors distributed over 11 sites including Lille. Since 2004, the Dolphin team is the scientific leader of the Grid'5000 site located at Lille. The role of the team is mainly threefold consisting in ensuring (1) the maintenance and evolution of the platform, (2) the local and national coordination of the related activities, and (3) the scientific animation around Grid'5000. Regarding the evolution of the infrastructure, this later has been extended in 2014 with a new powerful storage server hosting mainly the NFS system. In addition, the team has contributed to the 5-year CPER project called "data" proposed by the Inria Lille research center. The project includes a part related to the Grid'5000 platform. From the coordination point of view, Dolphin has participated to the monthly meetings of the national Grid'5000 committee. Finally, regarding the scientific animation, first, Dolphin has organized in December 20[th], 2014, a training [0] around Grid'5000. Second, Dolphin has participated to the program committee of the Grid'5000 spring school [0] held in Lyon in June 2014. Third, the team has also been involved in the final evaluation of the scientific nation-wide Hemera research project [0] related to Grid'5000.

---

[0]https://www.grid5000.fr/
[0]http://www.lifl.fr/~melab/HTML/Journee-G5K-Lille.htm
[0]https://www.grid5000.fr/mediawiki/index.php/Grid5000:School2014
[0]https://www.grid5000.fr/Hemera

DREAMPAL Team

4. New Software and Platforms

4.1. New Software and Platforms

Download page : https://gforge.inria.fr/frs/?group_id=3646

4.1.1. HoMade

HoMade is a softcore processor that we have started developing in 2012. The current version is reflective (i.e., the program it executes is self-modifiable), and statically configurable; dynamically reconfigurable multi-processors are the next steps. Users have to add to it the functionality they need in their applications via IPs. We have also being developing a library of IPs for the most common processor functions (ALU, registers, ...). All the design is in VHDL except for some schematic specifications.

The V5 version of HoMade has been developed in the Spring 2014. It has been used by ∼140 4th-year computer science students at Univ. Lille enrolled in the hardware architecture course (https://sites.google.com/site/tpm1aev/home). The new features of V5 are listed in Section 5.2 .

4.1.2. JHomade

JHomade is a software suite written in JAVA, including compilers and tools for the HoMade processor. It allows us to compile HiHope programs to Homade machine code and load the resulting binaries on FPGA boards. It was first released in 2013. The second version in 2014 includes several new features, like a C-frontend, a binary decoder and a code-generator for VHDL simulation. New features of the HiHope language are described in more detail in Section 5.3 .

4.1.3. Kcheck

Kcheck is a tool for the symbolic execution of programs in arbitrary languages defined in the $\mathbb{K}$ framework (http://k-framework.org), such as C and Java as well as the languages HiHope and Homade machine-code languages developed in out team. It also allow users to formally verify programs against specifications written in Reachability Logic, a specification formalism that can be seen as a language-independent Hoare logic. More information about the theory underlying Kcheck is given in Section 5.5 .

In 2014 we have developed a new and improved version of our tool, in order to keep up with the new modular infrastructure of the $\mathbb{K}$ framework. An online interface has been developed and is available at https://fmse.info.uaic.ro/tools/kcheck/. We have also started (since Nov. 2014) a development in the Coq proof assistant in order to obtained certificates for the program verifications performed by our tool.

<div align="center">

FUN Project-Team

</div>

4. New Software and Platforms

4.1. New ALE module for ASPIRERFID middleware.

Participants: Rim Driss [correspondant], Nathalie Mitton, Ibrahim Amadou, Julien Vandaele.

AspireRFID middleware is a modular OW2 open source RFID middleware. It is compliant with EPC Global standards. This new module integrates the modifications of the new standard release, including new RP and LLRP definitions and fixing bugs. This module has been implemented in the framework of the MIAOU project.

- Version: 1.0. . APP number: IDDN.FR.001.100017.000.S.P.2012.000.10000

4.2. T-SCAN.

Participants: Gabriele Sabatino [correspondant], Nathalie Mitton.

T-Scan is an interface ensuring the translation from a SGTIN tag format to an ONS hostname format according to the EPCGlobal standards. It allows the sending of a DNS request to look up the EPC-IS aides to which the product belongs in order to access the data relative to that product. This module has been implemented in the framework of the TRACAVERRE project.

- Version: 1.0 . March 1st 2014. N IDDN abrégé : 14-440017-000

4.3. GOLIATH 1.0

Participants: Nathalie Mitton [correspondant], Salvatore Guzzo Bonifacio [correspondant].

GOLIATH (Generic Optimized LIghtweight communication stack for Ambient TecHnologies) is a full protocol stack for wireless sensor networks. This module has been implemented in the framework of the ETIPOPS project.

See also the web page https://gforge.inria.fr/projects/goliath/.

4.4. ETINODE-CONTIKI-PORT

Participants: Salvatore Guzzo Bonifacio [correspondant], Roudy Dagher, Nathalie Mitton.

Contiki is an open source embedded OS for Internet of Things (IoT). It is light and portable to different hardware architectures. It embeds communication stacks for IoT Il embarque aussi des piles de communication pour l'internet des objets. This driver allows the running of Contiki OS over Etinode-MSP430. The code dalso allows the use of radio chip and embedded sensors. This module has been implemented in the framework of the ETIPOPS project.

- Version: 1.0 .

4.5. ETINODE-DRIVERS

Participants: Salvatore Guzzo Bonifacio [correspondant], Roudy Dagher, Nathalie Mitton.

These drivers for Etinode-MSP430 control the different embedded sensors and hardware components available on an Etinode-MSP430 node such as gyroscope, accelerometer and barometric sensor. This module has been implemented in the framework of the ETIPOPS project.

- Version: 1.0 .

4.6. FIT IoT-Lab

Participants: Raymond Borenstein, Nathalie Mitton [correspondant], Anne-Sophie Tonneau, Julien Vandaele, Roberto Quilez.

FIT IoT-LAB is a very large scale open testbed that features over 2700 wireless sensor nodes and more than 200 robots spread across six different sites in France. Nodes are either fixed or mobile and can be allocated in various topologies throughout all sites. A variety of wireless sensors are available, with different processor architectures (MSP430, STM32 and Cortex-A8) and different wireless chips (802.15.4 PHY at 800 MHz or 2.4 GHz). In addition, "open nodes" can receive custom wireless sensors for inclusion in IoT-LAB testbed. This platform is completely open and can be used by any one wishing to run experiment on wireless sensors and robots.

The Lille site displays 3 subsets of the platforms:

- Euratechnologies : this site features 256 WSN430 sensor nodes operating in the 2.4GHz band. 64 nodes are mobile, embedded on mobile trains.

- Haute Borne : this site features 256 M3 sensor nodes operating in the 2.4GHz band and 64 mobile robots (32 turtlebots and 32 wifibots) completely remotely programmable.

- Opennodes : this site will feature (opening beginning 2015) 64 hardware open slots to allow any one to plug his own hardware and benefits from the platform debugging and monitoring tools.

<p style="text-align:center; color:red">**LINKS Team**</p>

4. New Software and Platforms

4.1. QuiX-Tool Suite

Participants: Joachim Niehren [correspondant], Denis Debarbieux, Tom Sebastian.

The QuiX-Tool Suite provides tools to process XML streams and documents. The QuiX-Tool Suite is based on early algorithms: query answers are delivered as soon as possible and in all practical cases at the earliest time point. The QuiX-Tool Suite provides an implementation of the main XML standart over streams. XPath, XSLT, XQuery and XProc are W3C standarts while Schematron is an ISO one. The QuiX-Tool suite is developed in the Inria transfer project QuiXProc in cooperation with Innovimax. It includes among the others existing tools such as FXP and QuixPath, along with new tools, namely X-Fun. Both, a free and a professional version are available. The ownership of QuiX-Tool Suite is shared between Inria and Innovimax. The main application of QuiX-Tool Suite is its usage in QuiXProc, an professional implementation of the W3C pipeline language XProc owned by Innovimax.

The QuiXPath language is a large fragment of XPath with full support for the XML data model. The QuiXPath library provides a compiler from QuiXPath to FXP, which is a library for querying XML streams with a fragment of temporal logic.

The X-Fun language is a functional language for defining transformations between XML data trees, while providing shredding instructions. X-Fun can be understood as an extension of Frisch's XStream language with output shredding, while pattern matching is replaced by tree navigation with XPath expressions. The QuiX-Tool suite includes QuiXSLT, which is a compiler from XSLT into a fragment of X-Fun, which can be considered as the core of XSLT. It also provides QuiXSchematron, which is a compiler from Schematron to X-Fun, and QuiXQuery, which is a compiler from XQuery to X-Fun.

QuixPath now covers 100 per cent of the XPathMark, a W3C benchmark for the language Xpath (querying XML trees). In particular, it includes aggregation operators, joins and arithmetics operations.

See also the web page https://project.inria.fr/quix-tool-suite/.

- Version: QuixPath v2.0.3
- Version: X-Fun v0.5.0
- Version: QuiXSLT v0.5.0
- Version: QuiXSchematron v1.0.2

4.2. SmartHal

Participants: Joachim Niehren [correspondant], Antoine Mbaye Ndione, Guillaume Bagan.

SmartHal is a better tool for querying the HAL bibliography database, while is based on Haltool queries. The idea is that a Haltool query returns an XML document that can be queried further. In order to do so, SmartHal provides a new query language. Its queries are conjunctions of Haltool queries (for a list of laboratories or authors) with expressive Boolean queries by which answers of Haltool queries can be refined. These Boolean refinement queries are automatically translated to XQuery and executed by Saxon. A java application for extraction from the command line is available. On top of this, we have build a tool for producing the citation lists for the evaluation report of the LIFL, which can be easily adapter to other Labs.

See also the web page http://smarthal.lille.inria.fr/.

This year, Smathal has been adapted for querying the version 3.0 of Hal. Moreover, maintenance and optimization has been proceded all over the year.

- Version: SmartHal v1.0.0

MAGNET Team

5. New Software and Platforms

5.1. CoRTex

Participants: Pascal Denis [correspondent], David Chatel.

CoRTex is a LGPL-licensed Python library for Noun Phrase coreference resolution in natural language texts. This library contains implementations of various state-of-the-art coreference resolution algorithms, including those developed in our research. In addition, it provides a set of APIs and utilities for text pre-processing, reading the main annotation formats (ACE, CoNLL and MUC), and performing evaluation based on the main evaluation metrics (MUC, B-CUBED, and CEAF). As such, CoRTex provides benchmarks for researchers working on coreference resolution, but it is also of interest for developers who want to integrate a coreference resolution within a larger platform. This project is hosted on Inria gforge: https://gforge.inria.fr/projects/cortex/.

MEPHYSTO Team

5. New Software and Platforms

5.1. Platforms

5.1.1. Modulef

The numerical method to approximate the constitutive laws for rubber elasticity derived from polymer physics (as used in [15], [25]) are implemented in the Inria software Modulef (joint work of M. Vidrascu, projetc-team REO, and A. Gloria).

It is based on

- algorithms from stochastic geometry to generate suitable polymer networks;
- Delaunay tessellation algorithms to deal with steric effects (courtesy of the Inria project-team GAMMA2);
- the introduction of 1-dimensional finite elements for the polymer-chains in Modulef.

5.1.2. CMA-ES

To solve the inverse problem for the reconstruction of an explicit constitutive law from in silico experiments in [25], we relied on the Covariance Matrix Adaptation Evolution Strategy developed in the project-team TAO.

https://www.lri.fr/~hansen/cmaesintro.html

5.1.3. FreeFEM++

The numerical methods proposed in [14] for the approximation of homogenized coefficients were implemented in FreeFEM++, a user-friendly PDE-solver.

http://www.freefem.org/ff++/

<p style="text-align:center; color:red; font-weight:bold">MINT Project-Team</p>

5. New Software and Platforms

5.1. LibGINA

Participant: Laurent Grisoni [correspondant].

This year we used it with Ankama SME for a 3D videogame installation (La mine), done in collaboration with Idées-3com and LightUp. The library architecture has been rethought in order to provide ease of use and genericity.

Current version: version 1.1

Software characterization: A-2 SO-3 SM-2-up EM-3 SDL-3 OC-DA4-CD4-MS2-TPM4

5.2. 3D interaction using mobile phone

Participants: Samuel Degrande [correspondant], Laurent Grisoni.

This work has been achieved in the context of the Idées-3com I-lab. In this context a module, that allows to use any android based smartphone to control an Explorer module for navigation and interaction with VRML-based content. This module was used as a basis by Idées-3com in their commercial product this year.

Current version: version 1.0

Software characterization: A-2 SO-3 SM-2-up EM-2-up SDL-3 OC-DA4-CD4-MS2-TPM4

5.3. tIO (tactile input & output)

Participants: Marc-Antoine Dupre, Matthieu Falce, Nicolas Roussel [correspondant], Takashi Miyaki.

tIO is a library designed to facilitate the implementation of doubly tactile interaction techniques (tactile input coupled with tactile feedback) based on the STIMTAC technology. Supporting all current STIMTAC prototypes, it makes it easy to move the system pointer of the host computer according to motions detected on them and adapt their vibration amplitude based on the color of the pointed pixel or the nature of the pointed object. The library includes a set of demo applications that illustrate these two different approaches and makes it easy to "augment" existing `Qt` applications with tactile feedback. It also makes it possible to supplement or substitute tactile feedback with basic auditory feedback synthesized using `portaudio` (friction level is linearly mapped to the frequency of a sine wave). This not only facilitates the development and documentation of tactile-enhanced applications but also makes it easier to demonstrate them to a large audience.

Software characterization: A2, SO3-up, SM-2, EM2, SDL1.

5.4. libpointing

Participants: Géry Casiez [correspondant], Damien Marchal, Nicolas Roussel, Izzatbek Mukhanov.

Libpointing is a software toolkit that provides direct access to HID pointing devices and supports the design and evaluation of pointing transfer functions [2]. The toolkit provides resolution and frequency information for the available pointing and display devices and makes it easy to choose between them at run-time through the use of URIs. It allows to bypass the system's transfer functions to receive raw asynchronous events from one or more pointing devices. It replicates as faithfully as possible the transfer functions used by Microsoft Windows, Apple OS X and Xorg (the X.Org Foundation server). Running on these three platforms, it makes it possible to compare the replicated functions to the genuine ones as well as custom ones. The toolkit is written in C++ with Python and Java bindings available. It is publicly available under the GPLv2 license.

Izzatbek Mukhanov was recruited in October 2014 for two years as an engineer (IJD) to support the development and deployment of the library.

Web site: http://libpointing.org/

Software characterization: A3, SO3, SM-2, EM2, SDL4

5.5. PIRVI platform

Participants: Fabrice Aubert [correspondant], Damien Marchal.

MINT participates to the PIRVI platform (Framework for Computer Human Animation, Virtual Reality and Images, which aims at promoting research achieved by participant research teams (6 research teams, among which MINT), as well as encouraging collaborations with regional economical tissue on the knowledge fields covered within the associated research teams. The PIRVI allows these research teams to share a Virtual-Reality Room and various mid-size research equipments : multitouch tables, cameras (depth, infrared, ...), interactive devices (force-feedback, multitouch, smartphones...), a configurable multitouch wall.

MODAL Project-Team

5. New Software and Platforms

5.1. BlockCluster package for co-clustering

Participants: Serge Iovleff, Vincent Kubicki.

BlockCluster is an R package on top of the coclust C++ library. Maintenance of the CRAN package (http://cran.r-project.org/web/packages/blockcluster/index.html) and user assistance on the forum have been ensured.

5.2. clere package for high dimensional regression

Participants: Christophe Biernacki, Loïc Yengo, Julien Jacques.

The clere package for R proposes variable clustering in high dimensional linear regression. It is available on CRAN (http://cran.r-project.org/web/packages/clere/index.html) and now submitted to an international journal dedicated to software [52].

5.3. Clustericat package for correlated categorical variable

Participants: Christophe Biernacki, Matthieu Marbac-Lourdelle, Vincent Vandewalle.

Clustericat is an R package for model-based clustering of categorical data. In this package, the Conditional Correlated Model (CCM), published in 2014 [24], takes into account the main conditional dependencies between variables through extreme dependence situations (independence and deterministic dependence). Clustericat performs the model selection and provides the best model according to the BIC criterion and the maximum likelihood estimates. It is available online on Rforge (https://r-forge.r-project.org/R/?group_id=1803).

5.4. CoModes package for correlated categorical variables

Participants: Christophe Biernacki, Matthieu Marbac-Lourdelle, Vincent Vandewalle.

CoModes is another R package for model-based clustering of categorical data. In this package, the Conditional Modes Model (CMM), submitted for publication in 2014 [49], takes into account the main conditional dependencies between variables through particular modality crossings (so-called modes). CoModes performs the model selection and provides the best model according to the exact integrated likelihood criterion and the maximum likelihood estimates. It is available online on Rforge (https://r-forge.r-project.org/R/?group_id=1809).

5.5. CorReg package for correlated variables in regression

Participants: Christophe Biernacki, Clément Théry.

The main idea of the CorReg package is to consider some form of sub-regression models, some variables defining others. We can then remove temporarily some of the variables to overcome ill-conditioned matrices inherent in linear regression and then reinject the deleted information, based on the structure that links the variables. The final model therefore takes into account all the variables but without suffering from the consequences of correlations between variables or high dimension. The CorReg package is now available on CRAN (http://cran.r-project.org/web/packages/CorReg/index.html) and graphical functionalities have been added in 2014. It has been presented to a conference [32] and is currently written as a research paper [51]. It is a joint work with Gaétan Loridant.

5.6. HDPenReg package for penalized regressions in high dimension

Participants: Quentin Grimonprez, Serge Iovleff.

HDPenReg is an R-package based on a C++ code dedicated to the estimation of regression model with l1-penalization. It is now available on CRAN (http://cran.r-project.org/web/packages/HDPenReg/index.html). More cross-validation options were added. Maintenance in 2014 concerned bugs correction and documentation updates.

5.7. FunFEM package for functional data

Participant: Julien Jacques.

FunFEM package for R proposes a clustering tool for functional data. The model-based algorithm clusters the functional data into discriminative subspaces. It is available on CRAN (http://cran.r-project.org/web/packages/funFEM/index.html).

5.8. FunHDDC package for functional data

Participant: Julien Jacques.

FunHDDC package for R proposes a clustering tool for functional data. The model-based clustering algorithm considers that functional data live in cluster-specific subspaces. It is available on CRAN (http://cran.r-project.org/web/packages/funHDDC/index.html).

5.9. Galaxy-Modal platform

Participants: Samuel Blanck, Guillemette Marot.

Galaxy is an open, web-based platform for data intensive biomedical research. Galaxy features user friendly interface, workflow management, sharing functionalities and is widely used in the biologist community. The MPAgenomics R package developed by MODAL has been integrated into Galaxy, and the Galaxy-Modal instance has been publicly deployed thanks to the IFB-cloud infrastruture (http://cloud.france-bioinformatique.fr/). An APP repository with Galaxy-Modal source code has been created (reference : Galaxy - MPAgenomics)

5.10. metaMA package for meta-analysis of microarray data

Participant: Guillemette Marot.

metaMA is a specialised software for microarrays. It is an R package which combines either p-values or modified effect sizes from different studies to find differentially expressed genes. The main competitor of metaMA is geneMeta. Compared to geneMeta, metaMA offers an improvement for small sample size datasets since the corresponding modelling is based on shrinkage approaches.

This software is routinely used by biologists from INRA, Jouy en Josas (it has been included in a local analysis pipeline) but its diffusion on the CRAN (http://cran.r-project.org/web/packages/metaMA/index.html) makes it available to a wider community, as attested by the citations of publications related to the methods implemented in the software.

Maintenance in 2014 concerned documentation updates and users assistance.

5.11. metaRNASeq package for meta-analysis of RNA-Seq data

Participant: Guillemette Marot.

This is joint work with Andrea Rau (INRA, Jouy-en-Josas). metaRNASeq is a specialised software for RNA-seq experiments. It is an R package which is an adaptation of the metaMA package presented previously. Both implement the same kind of methods but specificities of the two types of technologies require some adaptations to each one. metaRNASeq is now available on CRAN (http://cran.r-project.org/web/packages/metaRNASeq/index.html).

5.12. MixCluster package for correlated mixed variables

Participants: Christophe Biernacki, Matthieu Marbac-Lourdelle, Vincent Vandewalle.

MixCluster is an R package for model-based clustering of mixed data (continuous, binary, integer). In this package, the model, submitted for publication in 2014 [48], takes into account the main conditional dependencies between variables through Gaussian copula. Mixcluster performs the model selection and provides the best model according to Bayesian approaches. It is available online on Rforge (https://r-forge.r-project.org/R/?group_id=1939).

5.13. MIXMOD and Rmixmod package for mixed data

Participants: Vincent Kubicki, Christophe Biernacki, Serge Iovleff.

MIXMOD (MIXture MODelling) is an important software for the MODAL team since it concerns its main topics: model-based supervised, unsupervised and semi-supervised classification for various data situations. MIXMOD is now a well-distributed software with over 250 downloads/month are recorded for several years. MIXMOD is written in C++ (more than 10 000 lines) and distributed under GNU General Public License. Several other institutions participate in the MIXMOD development since several years: CNRS, Inria Saclay-Île de France, Université de Franche-Comté, Université Lille 1. The software already benefits from several APP depositions and an R package (Rmixmod) has been associated to MIXMOD in 2012. In 2014, it has led to publication in an international journal dedicated to software [23].

5.14. MixtComp package for full mixed data

Participants: Vincent Kubicki, Christophe Biernacki, Serge Iovleff.

MixtComp (Mixture Computation) is another important software fot the MODAL team since it concerns model-based clustering for mixed data. Main difference with the MIXMOD/Rmixmod software is that MixtComp's architecture is able to integrate easily and quickly all new univariate models, under the conditional independence assumption, that will be sequentially available from researches of the MODAL team or others. Currently, central architecture of MixtComp is built and three models (Gaussian, multinomial, Poisson) are implemented with ability to natively manage missing data (completely or by interval). MixtComp stands both as a C++ library and an R package. The code is currently developed internally, and has been field-tested through two contracted partnerships.

5.15. MPAgenomics package for multi-patient analysis of genomic markers

Participants: Quentin Grimonprez, Guillemette Marot, Alain Celisse.

MPAgenomics is an R package for multi-patients analysis of genomics markers. It enables to study several copy number and SNP data profiles at the same time. It offers wrappers from commonly used packages to offer a pipeline for beginners in R. It also proposes a special way of choosing some crucial parameters to change some default values which were not adapted in the original packages. For multi-patients analysis, it wraps some penalized regression methods implemented in HDPenReg.

MPAgenomics is now available on CRAN (http://cran.r-project.org/web/packages/MPAgenomics/index.html). New segmentation methods were added to MPAgenomics. Maintenance in 2014 concerned bugs correction, documentation updates and code factorization.

5.16. RankCluster package to cluster ranking data

Participants: Christophe Biernacki, Quentin Grimonprez, Julien Jacques.

Rankcluster package for R proposes a clustering tool for ranking data. Multivariate and partial rankings can be also taken into account. It is available on CRAN (http://cran.r-project.org/web/packages/Rankcluster/index.html).

Rankcluster now supports tied ranking data. Maintenance in 2014 concerned bugs correction, documentation updates and addition of parallelism.

5.17. rtkpp package: STK++ Integration To R Using Rcpp

Participant: Serge Iovleff.

rtkpp is the integration of the library STK++ (see 5.18) into R. It is using Rcpp. Some functionalities of the Clustering project provided by the library are available in the R environment as R functions. The rtkpp package includes the header files from the STK++ library (currently version 0.8.2). Thus users do not need to install STK++ itself in order to use it. rtkpp is licensed under the GNU GPL version 2 or later and available on CRAN (http://cran.r-project.org/web/packages/rtkpp/index.html).

5.18. STK++ release 0.8.4: The Statistical ToolKit

Participant: Serge Iovleff.

STK++ is a versatile, fast, reliable and elegant collection of C++ classes for statistics, clustering, linear algebra, arrays (with an API Eigen-like), regression, dimension reduction, etc. STK++ is licensed under the GNU LGPL version 2 or later. See: http://www.stkpp.org/

<p style="text-align:center">**NON-A Project-Team**</p>

5. New Software and Platforms

5.1. SLIM

Multi-robots cooperation can be found as an application in many domains of science and technology: manufacturing, medical robotics, personal assistance, military/security and spatial robots. The market of robots is quickly developing and its capacity is continuously growing. Concerning cooperation of mobile multi-robots, 3 key issues have to be studied: Localization, path planning and robust control, for which Non-A team has worked and proposed new algorithms. Due to the ADT SLIM, we implement our algorithms (localization, path planning and robust control) and integrate them into ROS (Robotic Operating System) as a package, named SLIM.

5.2. Blimp

Scientific research and development on the control of autonomous airship have shown a significant growth in recent years. New applications appear in the areas such as freight carrier, advertising, monitoring, surveillance, transportation, military and scientific research. The control of autonomous airship is a very important problem for the aerial robots research.

The development of Blimp by Non-A is used for experimentation and demonstration of controlling algorithms. The blimp is required to provide some environment information and status of itself, such as surveillance video of surrounding environment, gesture of blimp, altitude of blimp. With these basic information, one could localize blimp with certain algorithm (visual SLAM for example) or implement one controller in order to improve the stability and maneuverability of blimp.

<p align="center" style="color:red">RMOD Project-Team</p>

5. New Software and Platforms

5.1. Moose5.0

Participants: Stéphane Ducasse [correspondant], Muhammad Bhatti, Andre Calvante Hora, Nicolas Anquetil, Anne Etien, Guillaume Larcheveque, Tudor Gîrba [University of Bern].

Web: http://www.moosetechnology.org/

The platform. Moose is a language-independent environment for reverse- and re-engineering complex software systems. Moose provides a set of services including a common meta-model, metrics evaluation and visualization, a model repository, and generic GUI support for querying, browsing and grouping. The development of Moose began at the Software Composition Group in 1997, and is currently contributed to and used by researchers in at least seven European universities. Moose offers an extensible meta-described metamodel, a query engine, a metric engine and several visualizations. Moose is currently in its fourth major release and comprises 55,000 lines of code in 700 classes.

The RMoD team is currently the main maintainer of the Moose platform. There are 200 publications (journal, international conferences, PhD theses) based on execution or use of the Moose environment.

The first version running on top of Pharo (Moose 4.0) was released in June 2010. The current focus is Moose 5.0, in late beta testing as of December 2014. A major development of 2014 is that tools and frameworks built for Moose are being integrated into Pharo4 as the default development tools.

Here is the self-assessment of the team effort following the grid given at http://www.inria.fr/institut/organisation/instances/commission-d-evaluation.

- **(A5)** Audience : 5 – Moose is used by several research groups, a consulting company, and some companies using it in ad-hoc ways.
- **(SO4)** Software originality : 4 – Moose aggregates the last results of several research groups.
- **(SM4)** Software Maturity : 4 – Moose is developed since 1996 and got two main redesign phases.
- **(EM4)** Evolution and Maintenance : 4 – Moose will be used as a foundation of our Synectique start up so its maintenance is planned.
- **(SDL4)** Software Distribution and Licensing : 4 – Moose is licensed under BSD
- **(OC)** Own Contribution : (Design/Architecture)DA-4, (Coding/Debugging)-4, (Maintenance/Support)-4, (Team/Project Management)-4

5.2. Pharo3.0

Participants: Marcus Denker [correspondant], Damien Cassou, Stéphane Ducasse, Esteban Lorenzano, Damien Pollet, Igor Stasenko, Camillo Bruni, Camille Teruel, Clément Bera.

Web: http://www.pharo.org/

The platform. Pharo is an open-source Smalltalk-inspired language and environment. It provides a platform for innovative development both in industry and research. By providing a stable and small core system, excellent developer tools, and maintained releases, Pharo's goal is to be a platform to build and deploy mission critical applications, while at the same time continue to evolve.

The first stable version, Pharo 1.0, was released in 2010. We are now releasing one new version of Pharo every year, with Pharo3 released in spring 2014. Pharo4 has seen already over 400 incremental updates and is scheduled for early 2015. It should be noted that Pharo, even though already used outside of research, still continues to improve radically.

In November 2012 RMoD launched the Pharo Consortium (http://consortium.pharo.org/) and the Pharo Association (http://association.pharo.org). The consortium has now 14 industrial members, 3 sponsors and 10 academic partners.

RMoD is the main maintainer and coordinator of Pharo.

Here is the self-assessment of the team effort following the grid given at http://www.inria.fr/institut/organisation/instances/commission-d-evaluation.

- **(A5)** Audience: 5 – Used in many universities for teaching, more than 25 companies.
- **(SO3)** Software originality : 3 – Pharo offers a classical basis for some aspects (UI). It includes new frameworks and concepts compared to other Smalltalk implementations.
- **(SM4)** Software Maturity: 4 – Bug tracker, continuous integration, large test suites are in place.
- **(EM4)** Evolution and Maintenance: 4 – Active user group, consortium and association has been set up.
- **(SDL4)** Software Distribution and Licensing: 4 – Pharo is licensed under MIT.
- **(OC5)** Own Contribution: (Design/Architecture) DA-5, (Coding/Debugging) CD-5, (Maintenance/Support) MS-5, (Team/Project Management) TPM-5

5.3. Pillar 0.17

Pillar is a markup syntax and associated tools to write and generate documentation and books. Pillar is currently used to write several books and other documentation. Two platforms have already been created on top of Pillar: PillarHub and Marina. http://www.smalltalkhub.com/#!/~Pier/Pillar

<h1 style="text-align:center;color:red;">SEQUEL Project-Team</h1>

5. New Software and Platforms

5.1. Computer Games

Participant: Rémi Coulom.

- ***Crazy Stone*** is a top-level Go-playing program that has been developed by Rémi Coulom since 2005. Crazy Stone won several major international Go tournaments in the past. In 2013, a new version was released in Japan. This new version won the 6th edition of the UEC Cup (the most important international computer-Go tournament). It also won the first edition of the Denseisen, by winning a 4-stone handicap game against 9-dan professional player Yoshio Ishida. It is distributed as a commercial product by *Unbalance Corporation* (Japan). 6-month work in 2013. URL: http://remi.coulom.free.fr/CrazyStone/

- ***Kifu Snap*** is an Android image-recognition app. It can automatically recognize a Go board from a picture, and analyze it with Crazy Stone. It was released on Google Play in November, 2013. 6-month work in 2013. URL: http://remi.coulom.free.fr/kifu-snap/

5.2. Function optimization

Participant: Philippe Preux.

5.2.1. yaStoSOO

We have worked on the efficient implementation of the StoSOO algorithm in order to have a software that can be used for real to optimize real functions, and to be able to experiment with the algorithm, and assess its practical usefulness. This led to yaStoSOO, an implementation in C available on the web at http://www.grappa.univ-lille3.fr/~ppreux/software/StoSOO/. The code is distributed under the GPL licence.

Thanks to this implementation, we were able to compete in the CEC'2014 competition on Real-Parameter Single Objective optimization at which we ranked honorably (10th out of 17 competitor algorithms). More experimental work is under-way.

SPIRALS Team

5. New Software and Platforms

5.1. APISENSE®

Participants: Clive Ferret-Canape, Julien Duribreux, María Gómez Lacruz, Nicolas Haderer, Christophe Ribeiro, Romain Rouvoy [correspondant], Antoine Veuiller.

In 2014, new developments have been made on our APISENSE® distributed crowdsensing platform. APISENSE® now builds on a distributed infrastructure hosted in the Cloud that can better cope with scalability issues in the number of experiments, users, and volume of data to be collected in the wild. Data collected by participants can be exposed to applications and stakeholders via an Open Data API, which provides the ability to build realtime web applications from crowdsourced datasets. The APISENSE® mobile app, named BEE, can be freely downloaded from the Google Play Store. APISENSE® is part of the results of the PhD thesis of Nicolas Haderer [12] that was defended in November 2014. In 2014, APISENSE® has also been at the core of an industrial transfer action that aims at creating a spin-off company. The project is managed by Christophe Ribeiro and Romain Rouvoy. The project has been accepted (so-called qualification) in 2014 by the Inria investment fund IT-Translation. The project in supported by Direction Transfert & Innovation which will fund in 2015 the 1-year engineer contract of Christophe Ribeiro for maturing the project.

APISENSE® is a distributed platform dedicated to crowdsensing activities. Crowdsensing intends to leverage mobile devices to seamlessly collect valuable dataset for different categories of stakeholders. APISENSE® intends to be used in a wide variety of scientific and industrial domains, including network quality monitoring, social behavior analysis, epidemy predictions, emergency crisis support, open maps initiatives, debugging of applications in the wide. APISENSE® is composed of BEE.HIVE delivered as a *Platform-as-a-Service* (PaaS) to the stakeholders who can pilot and customize their own crowdsensing environment [108], and *Bee.mob* supporting participants with a mobile application to control the sensors to be shared with the rest of the world [96], [97]. The platform is used by the Metroscope consortium, an Internet scientific observatory initiative supported by Inria.

APISENSE® is at the core of the Inria ADT Focus CrowdLab project (see Section 8.2).

Web site: http://www.apisense.fr. Registered with the APP (*Agence pour la Protection des Programmes*) under reference IDDN.FR.001.080006.000.S.P.2013.000.10000 is pending. License: Proprietary.

5.2. FraSCAti

Participants: Philippe Merle [correspondant], Fawaz Paraiso, Romain Rouvoy, Lionel Seinturier.

The novelty of 2014 consists in the development of the SOCLOUD platform for distributed multi-cloud systems. This platform has been defined in the context of the PhD thesis of Fawaz Paraiso [15] that was defended in June 2014. SOCLOUD is built on top of our existing FRASCATI platform. SOCLOUD enables to deploy, execute and manage an application that spans on several different cloud systems.

FRASCATI is a service-oriented component-based middleware platform implementing OASIS *Service Component Architecture* (SCA) specifications. The main originality of FRASCATI is to bring FRACTAL-based reflectivity to SCA, *i.e.*, any FRASCATI software component is equipped with both the SOA capabilities brought by SCA and the reflective capabilities (*i.e.*, introspection and reconfiguration) brought by FRACTAL. Various micro-benchmarks have shown that FRASCATI reflectivity is achieved without hindering its performance relative to the de facto reference SCA implementation, *i.e.*, Apache Tuscany. Non-functional concerns (logging, transaction, security, etc.), so-called intents in SCA terms, are also programmed as FRASCATI components and are (un)woven on business components dynamically at runtime, this is based on aspect-oriented

concepts defined in FAC [110]. FRASCATI supports various implementation technologies (SCA Composite, Java, WS-BPEL, Spring Framework, OSGi, Fractal ADL, native C library, Apache Velocity templates, and seven scripting languages as BeanShell, FScript, Groovy, JavaScript, JRuby, Jython, XQuery) for programming services or integrating legacy code, various binding protocols (SOAP, REST, JSON-RPC, UPnP, HTTP servlets, Java RMI, JMS, JGroups) and interface definition languages (WSDL, Java, WADL) for interoperating with existing services. FRASCATI provides management tools like standalone, Web-based, and JMX-based graphical consoles and a dedicated scripting language for reconfiguring SCA applications. The whole FRASCATI platform is itself built as a set of reflective SCA components.

Inria Evaluation Committee Criteria for Software Self-Assessment: A-4-up, SO-4, SM-4-up, EM-3-up, SDL-4-up, DA-4, CD-4, MS-4, TPM-4. FRASCATI is a project of the OW2 consortium for open-source middleware. Web site: http://frascati.ow2.org. 292 Kloc (mainly Java). Registered with the APP (Agence pour la Protection des Programmes) under reference FR.001.050017.000.S.P.2010.000.10000. License: LGPL. Embedded into several industrial software systems: EasySOA, Petals Link EasyViper, EasyBPEL, EasyESB, OW2 PEtALS, OW2 Scarbo. Various demonstrators built during funded projects: ANR SCOrWare, FP7 SOA4All, ANR ITEmIS, ANR SALTY, ANR SocEDA, FUI Macchiato, FUI EasySOA, ADT Galaxy and ADT Adapt. Main publications: [117], [116], [103], [104], [93], [92].

5.3. PowerAPI

Participants: Maxime Colmant, Loïc Huertas, Adel Noureddine, Romain Rouvoy [correspondant].

In 2014, new developments have been made on our POWERAPI library for monitoring energy in software systems. POWERAPI now includes an accurate power model, which supports both DFVS, hyper threads and turbo boost features of modern processors. This model has been assessed on acknowledged benchmarks (PARSEC, SPEC CPU, SPECjbb) and is used as a basis to estimate the power consumption of applications running in virtualised environments. Finally, POWERAPI has evolved towards a modular toolkit that can be used to build software-defined power meters supporting a wide range of input sources (*e.g.*, hardware performance counters, RAPL, PowerSpy). POWERAPI is part of the results of the PhD thesis of Adel Noureddine [14] that was defended in March 2014.

POWERAPI is a Scala-based library for monitoring energy in software systems. It is based on a modular and asynchronous event-driven architecture using the Akka library. POWERAPI differs from existing energy process-level monitoring tool in its pure software, fully customizable and modular aspect which let users precisely define what they want to monitor, without plugging any external device. POWERAPI offers an API which can be used to express requests about energy spent by a process, following its hardware resource utilization (in terms of CPU, memory, disk, network, etc.). Its applications cover energy-driven benchmarking [105], [88], [86], [87], energy hotspots and bugs detection [106], [107], and real-time distributed system monitoring.

POWERAPI is at the core of the Inria ADT eSurgeon project (see Section 8.2).

Web site: http://www.powerapi.org. Registered with the APP (Agence pour la Protection des Programmes) under reference IDDN.FR.001.400015.000.S.P.2012.000.10000. License: AGPL.

5.4. Saloon

Participants: Laurence Duchien, Clément Quinton, Daniel Romero Acero, Lionel Seinturier [correspondant].

SALOON is a framework for the selection and configuration of Cloud providers according to application requirements. The framework enables the specification of such requirements by defining ontologies. Each ontology provides a unified vision of provider offers in terms of frameworks, databases, languages, application servers and computational resources (i.e., memory, storage and CPU frequency). Furthermore, each provider is related to a Feature Model (FM) with attributes and cardinalities, which captures its capabilities. By combining the ontology and FMs, the framework is able to match application requirements with provider capabilities and select a suitable one. Specific scripts to the selected provider are generated in order to enable its configuration.

SALOON is the result of the PhD thesis of Clément Quinton [16] that was defended in October 2014. SALOON is partially developed in the context of the FP7 PaaSage project (see Section 8.3).

Registered with the APP (Agence pour la Protection des Programmes) under reference IDDN.FR.001.300002.000.S.P.2014.000.10800.

5.5. Spoon

Participants: Martin Monperrus [correspondant], Gérard Paligot, Nicolas Petitprez.

In 2014, SPOON has been at the core of an industrial transfer action that aims at creating a spin-off company. The project is managed by Nicolas Petitprez and Martin Monperrus. The project has been accepted (so-called qualification) in 2014 by the Inria investment fund IT-Translation. The project in supported by Direction Transfert & Innovation which will fund in 2015 the 1-year engineer contract of Nicolas Petitprez for maturing the project. As an open source project Spoon has attracted new contributors in 2014. The Spoon development team is now composed of 8 active members, including 4 that are not at all related to Inria. Second, Spoon now supports analyzing and transforming Java 7 code, which is the now the dominant version of Java. Third, Spoon is the technical foundation of five important papers published in 2014. To sum up, year 2014 was a major year for warming up the Spoon project. Thanks to the support of Inria through the ADT, year 2015 is expected to be as vibrant and rich.

SPOON is a library for analyzing and transforming Java source code [76] [109]. SPOON provides a core API and associated tools for static analysis and generative programming within the Java 5+ environment. SPOON must be seen as a basis to ensure Software Quality through code validation and generation. It can be used in the software development process during the validation phase, as well as for engineering and re-engineering software. The first key point of SPOON is to provide a well-typed and comprehensive AST API which is designed to facilitate analysis and transformation work for programmers. Scanners and processors allow the programmer to implement various program traversal strategies on the Java program. Also, the program representation is built with a well-known and well-tested open source Java compiler: the Eclipse JDT compiler, which ensures the support of the latest Java features. The second key point of SPOON is to provide a pure Java API to specify program transformations using a well-typed generative programming technique (called Spoon Templates). By using well-typed templates, SPOON makes programming of transformations easier and safer for the end-user programmers.

SPOON is at the core of the Inria ADT Spoon3R project (see Section 8.1).

Web site: http://spoon.gforge.inria.fr. Registered with the APP (Agence pour la Protection des Programmes) under reference IDDN.FR.001.070037.000.S.P.2007.000.10600. License: CeCILL-C.