



RESEARCH CENTER
Nancy - Grand Est

FIELD

Activity Report 2013

Section Software

Edition: 2014-03-19

1. ALGORILLE Project-Team	4
2. ALICE Project-Team	7
3. BIGS Project-Team	9
4. CALVI Project-Team	10
5. CAMUS Team	12
6. CAMEL Project-Team	17
7. CARTE Project-Team	20
8. CASSIS Project-Team	21
9. CORIDA Project-Team	25
10. CORTEX Team	26
11. MADYNES Project-Team	27
12. MAGRIT Project-Team	29
13. MAIA Project-Team	30
14. MASAIE Project-Team (section vide)	32
15. NEUROSYS Team	33
16. ORPAILLEUR Project-Team	34
17. PAREO Project-Team	38
18. PAROLE Project-Team	39
19. SCORE Team	42
20. SÉMAGRAMME Project-Team	43
21. SHACRA Project-Team	46
22. TOSCA Project-Team	47
23. TRIO Team	48
24. VEGAS Project-Team	49
25. VERIDIS Project-Team	51

ALGORILLE Project-Team

5. Software and Platforms

5.1. Introduction

Software is a central part of our output. In the following we present the main tools to which we contribute. We use the [Inria software self-assessment](#) catalog for a classification.

5.2. Implementing parallel models

Several software platforms have served us to implement and promote our ideas in the domain of coarse grained computation and application structuring.

5.2.1. ORWL and P99

Participants: Jens Gustedt, Rodrigo Campos-Catelin, Stéphane Vialle, Mariem Saied.

ORWL is a reference implementation of the Ordered Read-Write Lock tools as described in [4]. The macro definitions and tools for programming in C99 that have been implemented for ORWL have been separated out into a toolbox called P99. ORWL is intended to become opensource, once it will be in a publishable state. P99 is available under a QPL at <http://p99.gforge.inria.fr/>.

Software classification: A-3-up, SO-4, SM-3, EM-3, SDL (P99: 4, ORWL: 2-up), DA-4, CD-4, MS-3, TPM-4

5.2.2. parXXL

Participants: Jens Gustedt, Stéphane Vialle.

ParXXL is a library for large scale computation and communication that executes fine grained algorithms on coarse grained architectures (clusters, grids, mainframes). It has been one of the software bases of the InterCell project and has been proven to be a stable support, there. It is available under a GPLv2 at <http://parxxl.gforge.inria.fr/>. ParXXL is not under active development anymore, but still maintained in the case of bugs or portability problems.

Software classification: A-3, SO-4, SM-3, EM-2, SDL-4, DA-4, CD-4, MS-2, TPM-2

5.3. Distem

Participants: Tomasz Buchert, Emmanuel Jeanvoine, Lucas Nussbaum, Luc Sarzyniec.

Wrekavoc and Distem are distributed system emulators. They enable researchers to evaluate unmodified distributed applications on heterogeneous distributed platforms created from an homogeneous cluster: CPU performance and network characteristics are altered by the emulator.

Wrekavoc was developed until 2010, and we then focused our efforts on **Distem**, that shares the same goals with a different design. Distem is available from <http://distem.gforge.inria.fr/> under GPLv3.

Software classification: A-3-up, SO-4, SM-3-up, EM-3, SDL-4, DA-4, CD-4, MS-4, TPM-4.

5.4. SimGrid

SimGrid is a toolkit for the simulation of distributed applications in heterogeneous distributed environments. The specific goal of the project is to facilitate research in the area of parallel and distributed large scale systems, such as Grids, P2P systems and clouds. Its use cases encompass heuristic evaluation, application prototyping or even real application development and tuning.

5.4.1. Core distribution

Participants: Martin Quinson, Marion Guthmuller, Paul Bédaride, Gabriel Corona, Lucas Nussbaum.

SimGrid has an active user community of more than one hundred members, and is available under GPLv3 from <http://simgrid.gforge.inria.fr/>. One third of the source code is devoted to about 12000 unit tests and 500 full integration tests. These tests are run for each commit for 4 package configurations and on 4 operating systems thanks to the Inria continuous integration platform.

Software classification: A-4-up, SO-4, SM-4, EM-4, SDL-5, DA-4, CD-4, MS-3, TPM-4.

5.4.2. *SimGridMC*

Participants: Martin Quinson, Marion Guthmuller, Gabriel Corona.

SimGridMC is a module of SimGrid that can be used to formally assess any distributed system that can be simulated within SimGrid. It explores all possible message interleavings searching for states violating the provided properties. We recently added the ability to assess liveness properties over arbitrary C 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 arbitrary MPI code written in C.

Software classification: A-3-up, SO-4, SM-3-up, EM-3-up, SDL-5, DA-4, CD-4, MS-4, TPM-4.

5.4.3. *SCHaaS*

Participants: Julien Gossa, Stéphane Genaud, Luke Bertot, Rajni Aron, Étienne Michon.

The *Simulation of Clouds, Hypervisor and IaaS* (SCHaaS) is an extension of SimGrid that can be used to comprehensively simulate clouds, from the hypervisor/system level, to the IaaS/administrator level. The hypervisor level includes models about virtualization overhead and VMs operations like boot, start, suspend, migrate, and network capping. The IaaS level includes models about instances management like image storage and deployment and VM scheduling. This extension allows to fully simulate any cloud infrastructure, whatever the hypervisor or the IaaS manager. This can be used by both cloud administrators to dimension and tune clouds, and cloud users to simulate cloud applications and assess provisioning strategies in term of performances and cost.

Software classification: A-4-up, SO-4, SM-2-up, EM-2-up, SDL-2, DA-4, CD-4, MS-4, TPM-4.

5.5. *Kadeploy*

Participants: Luc Sarzyniec, Emmanuel Jeanvoine, Lucas Nussbaum.

Kadeploy is a scalable, efficient and reliable deployment (provisioning) system for clusters and grids. It provides a set of tools for cloning, configuring (post installation) and managing cluster nodes. It can deploy a 300-nodes cluster in a few minutes, without intervention from the system administrator. It plays a key role on the Grid'5000 testbed, where it allows users to reconfigure the software environment on the nodes, and is also used on a dozen of production clusters both inside and outside INRIA. It is available from <http://kadeploy3.gforge.inria.fr/> under the Cecill license.

Software classification: A-4-up, SO-3, SM-4, EM-4, SDL-4-up, DA-4, CD-4, MS-4, TPM-4.

5.6. *XPFlow*

Participants: Tomasz Buchert, Lucas Nussbaum.

XPFlow is an implementation of a new, workflow-inspired approach to control of experiments involving large-scale computer installations. Such systems pose many difficult problems to researchers due to their complexity, their numerous constituents and scalability problems. The main idea of the approach consists in describing the experiment as a workflow and execute it using achievements of Business Workflow Management (BPM), workflow management techniques and scientific workflows. The website of XPFlow is <http://xpflow.gforge.inria.fr/>. The software itself is not released to the general public yet, a fact that will be addressed during the year 2014.

Software classification: A-2-up, SO-3-up, SM-2-up, EM-3-up, SDL-2-up, DA-4, CD-4, MS-4, TPM-4.

5.7. Grid'5000 testbed

Participants: Luc Sarzyniec, Emmanuel Jeanvoine, Émile Morel, Lucas Nussbaum.

Grid'5000 (<http://www.grid5000.fr>) is a scientific instrument designed to support experiment-driven research in all areas of computer science related to parallel, large-scale or distributed computing and networking. It gathers 10 sites, 25 clusters, 1200 nodes, for a total of 8000 cores. It provides its users with a fully reconfigurable environment (bare metal OS deployment with Kadeploy, network isolation with KaVLAN) and a strong focus on enabling high-quality, reproducible experiments.

The AlGorille team contributes to the design of Grid'5000, to the administration of the local Grid'5000 site in Nancy, and to the design and development of Kadeploy (in close cooperation with the Grid'5000 technical team). The AlGorille engineers also administer *Inria Nancy – Grand Est*'s local production cluster, named *Talc*, leveraging the experience and tools from Grid'5000.

Software classification: A-5, SO-4, SM-4, EM-4, SDL-N/A, DA-4, CD-4, MS-4, TPM-4.

ALICE Project-Team

4. Software and Platforms

4.1. Vorpaline

Participants: Dobrina Boltcheva, Bruno Lévy, Thierry Valentin.

Vorpaline is an automatic surfacic and volumetric mesh generation software, distributed with a commercial license. Vorpaline is based on the main scientific results stemming from projets GoodShape and VORPALINE, funded by the European Research Council, about optimal quantization, centroidal Voronoi diagrams and fast/parallel computation of Voronoi diagrams in high-dimension space. The current version (1.0) provides functionalities such as isotropic/adaptive/anisotropic surface re-meshing, tolerant surface re-meshing, mesh repair and mesh decimation. Next versions will provide functionalities such as constrained surface meshing (2.0), quad-dominant surface meshing (3.0) and hex-dominant volume meshing (4.0).

4.2. IceSL

Participants: Jérémie Dumas, Jean Hergel, Sylvain Lefebvre.

In the new software **IceSL**, we propose to exploit recent advances in GPU and Computer Graphics to accelerate the slicing process of objects modelled via a CSG¹ language. Our target are open source low cost *fused deposition modeling* printers such as RepRaps.

Our approach first inputs a CSG description of a scene which can be composed of both meshes and analytic primitives. During display and slicing the CSG model is converted on the fly into an intermediate representation enabling fast processing on the GPU. Slices can be quickly extracted, and the tool path is prepared through image erosion. The interactive preview of the final geometry uses the exact same code path as the slicer, providing an immediate, accurate visual feedback.

IceSL is the recipient software for our ERC research project “ShapeForge”, led by Sylvain Lefebvre.




Figure 1. Left. A two-colored vase is modeled in IceSL. Right. An early printed result.

4.3. Graphite

Participants: Dobrina Boltcheva, Samuel Hornus, Bruno Lévy, David Lopez, Romain Merland, Jeanne Pellerin, Nicolas Ray.

¹ Constructive Solid Geometry

Graphite is a research platform for computer graphics, 3D modeling and numerical geometry. It comprises all the main research results of our “geometry processing” group. Data structures for cellular complexes, parameterization, multi-resolution analysis and numerical optimization are the main features of the software. Graphite is publicly available since October 2003. It is hosted by Inria GForge since September 2008. Graphite is one of the common software platforms used in the frame of the European Network of Excellence **AIMShape**.

Graphite and its research-plugins are actively developed and extended. The latest version was released on January 2nd, 2014 and has been downloaded 135 times as of January 29.

4.4. GraphiteLifeExplorer

Participant: Samuel Hornus.

GLE is a 3D modeler, developed as a plugin of Graphite, dedicated to molecular biology. It is developed in cooperation with the Fourmentin Guilbert foundation and has recently been renamed "GraphiteLifeExplorer". Biologists need simple spatial modeling tools to help in understanding the role of the relative position of objects in the functioning of the cell. In this context, we develop a tool for easy DNA modeling. The tool generates DNA along any user-given curve, open or closed, allows fine-tuning of atoms position and, most importantly, exports to PDB (the Protein Data Bank file format).

The development of GLE is currently on hold, but it is still downloaded (freely) about twice a day (1600 downloads to date). A paper describing it was published in the broad journal PLOS One [12].

4.5. OpenNL - Open Numerical Library

Participants: Bruno Lévy, Nicolas Ray, Rhaleb Zayer.

OpenNL is a standalone library for numerical optimization, especially well-suited to mesh processing. The API is inspired by the graphics API OpenGL, this makes the learning curve easy for computer graphics practitioners. The included demo program implements our LSCM [5] mesh unwrapping method. It was integrated in **Blender** by Brecht Van Lommel and others to create automatic texture mapping methods. OpenNL is extended with two specialized modules :

- **CGAL parameterization package:** this software library, developed in cooperation with Pierre Alliez and Laurent Saboret, is a **CGAL** package for mesh parameterization.
- **Concurrent Number Cruncher:** this software library extends OpenNL with parallel computing on the GPU, implemented using the CUDA API.

4.6. LibSL

Participants: Anass Lasram, Sylvain Lefebvre.

LibSL is a Simple library for graphics. Sylvain Lefebvre continued development of the LibSL graphics library (under CeCill-C licence, filed at the APP). LibSL is a toolbox for rapid prototyping of computer graphics algorithms, under both OpenGL, DirectX 9/10, Windows and Linux. The library is actively used in both the REVES / Inria Sophia-Antipolis Méditerranée and the ALICE / Inria Nancy Grand-Est teams.

BIGS Project-Team

5. Software and Platforms

5.1. Light diffusion into tissues

We are currently considering the possibility to implement our Matlab algorithms concerning light diffusion into tissues into the Matlab toolbox *Contsid*, developed by the System Identification team of the CRAN (<http://www.iris.cran.uhp-nancy.fr/contsid/>).

5.2. Online data analysis

An R package performing most of the methods of factorial analysis in an online way has been developed by R. Bar and J-M. Monnez. Starting from a simulated data flow, the main goal of the program is to perform online factorial analyses (Principal Component Analyses, Canonical Correlation Analysis, Canonical Discriminant Analysis, Correspondence Analysis). Data are supposed to be independent and identically distributed observations of a random vector (whose distribution is a priori unknown). Defining stochastic approximation processes, the procedure is adaptative in the sense that the results of the analyses are updated recursively each time that a new data is taken into account.

From a theoretical point of view, the i.i.d case has been recently extended to the case of an expectation and/or covariance matrix of the random vector varying with time. We plan to include these improvements into our software.

5.3. Socio-economic index

A R package called *SesIndexCreaoR* has been written by B. Lalloué and J-M. Monnez in order to implement our socio-economic index for health inequalities. The version 1.0 of this package is currently freely available on the website of the Equit'Area project: http://www.equitarea.org/documents/packages_1.0-0/. It contains the functions needed to run the procedure (either integrally or partially) and obtain the corresponding SES index. The user may also create categories of this index with different methods (hierarchical clustering with or without k-nearest neighbors, quantiles, or intervals) and generate automatic reports of the results. Visualization and plotting functions are provided in the package.

CALVI Project-Team

5. Software and Platforms

5.1. SeLaLib

Participants: Aurore Back, Raphaël Blanchard, Edwin Chacon Golcher, Samuel de Santis, Aliou Diouf, Pierre Navaro, Morgane Bergot, Emmanuel Frénod, Virginie Grandgirard, Adnane Hamiaz, Philippe Helluy, Sever Hirstoaga, Michel Mehrenberger, Laurent Navoret, Nhung Pham, Eric Sonnendrücker.

Under the 'Fusion' large scale initiative, we have continued our work in the development of the ADT Selalib (the Semi-Lagrangian Library), now finishing its third year. This library provides building blocks for the development of numerical simulations for the solution of the fundamental equation of plasma physics: the Vlasov equation. In this context we have continued to add new modules improved interfaces and implemented 'continuous integration' software development techniques to improve code robustness and portability. Furthermore, we continue to involve other researchers within France and abroad to aid in the further development of this software product.

One of the aims of the ADT is to provide numerical building blocks for the GYSELA code developed at CEA Cadarache in collaboration with the Calvi project-team. GYSELA is used by physicists for simulating the development of turbulence in magnetic fusion plasmas in particular in view of the ITER project.

This year many developments have incorporated into Selalib: semi-Lagrangian solvers on curvilinear grids, new models, new fully-Eulerian solvers, new linear solvers for the Poisson or quasineutrality equation. More details are given in the corresponding sections, because we always try to test our new algorithms within Selalib.

Selalib is available on the Inria Forge <http://selalib.gforge.inria.fr/>

5.2. CLAC

Participants: Philippe Helluy, Michel Massaro, Thomas Strub.

CLAC is a generic Discontinuous Galerkin solver, written in C/C++, based on the OpenCL and MPI frameworks. CLAC means "Conservation Laws Approximation on many Cores".

It is clear now that a future supercomputer will be made of a collection of thousands of interconnected multicore processors. Globally it appears as a classical distributed memory MIMD machine. But at a lower level, each of the multicore processors is itself made of a shared memory MIMD unit (a few classical CPU cores) and a SIMD unit (a GPU). When designing new algorithms, it is important to adapt them to this architecture. Our philosophy will be to program our algorithms in such a way that they can be run efficiently on this kind of computers. Practically, we will use the MPI library for managing the high level parallelism, while the OpenCL library will efficiently operate the low level parallelism.

We have invested for several years now into scientific computing on GPUs, using the open standard OpenCL (Open Computing Language). We were recently awarded a prize in the international AMD OpenCL innovation challenge thanks to an OpenCL two-dimensional Vlasov-Maxwell solver that fully runs on a GPU. OpenCL is a very interesting tool because it is an open standard now available on almost all brands of multicore processors and GPUs. The same parallel program can run on a GPU or a multicore processor without modification.

CLAC is also a joint project with a Strasbourg small company, AxesSim, which develops software for electromagnetic simulations.

Because of the envisaged applications of CLAC, which may be either academic or commercial, it is necessary to conceive a modular framework. The heart of the library is made of generic parallel algorithms for solving conservation laws. The parallelism can be both fine-grained (oriented towards GPUs and multicore processors) and coarse-grained (oriented towards GPU clusters). The separate modules allow managing the meshes and some specific applications. In this way, it is possible to isolate parts that should be protected for trade secret reasons. The open source part of CLAC will be made freely available on the web later on. We have made an APP deposit of the first version of CLAC in October 2012. The versioning of CLAC project is also registered in the Inria Forge <http://clac.gforge.inria.fr>.

CAMUS Team

5. Software and Platforms

5.1. PolyLib

Participant: Vincent Loechner.

PolyLib ¹ is a C library of polyhedral functions, that can manipulate unions of rational polyhedra of any dimension, through the following operations: intersection, difference, union, convex hull, simplify, image and preimage. It was the first to provide an implementation of the computation of parametric vertices of a parametric polyhedron, and the computation of an Ehrhart polynomial (expressing the number of integer points contained in a parametric polytope) based on an interpolation method.

It is used by an important community of researchers (in France and the rest of the world) in the area of compilation and optimization using the polyhedral model. Vincent Loechner is the maintainer of this software. It is distributed under GNU General Public License version 3 or later, and it has a Debian package maintained by Serge Guelton (Parkas Projet, Inria Paris - Rocquencourt).

5.2. ZPolyTrans

Participant: Vincent Loechner.

ZPolyTrans ² is a C library and a set of executables, that permits to compute the integer transformation of a union of parametric \mathbb{Z} -polyhedra (the intersection between lattices and parametric polyhedra), as a union of parametric \mathbb{Z} -polyhedra. The number of integer points of the result can also be computed. It is build upon PolyLib and Barvinok library. This work is based on some theoretical results obtained by Rachid Seghir and Vincent Loechner [29].

It allows for example to compute the number of solutions of a Presburger formula by eliminating existential integer variables, or to compute the number of different data accessed by some array accesses contained in an affine parametric loop nest.

The authors of this software are Rachid Seghir (Univ. Batna, Algeria) and Vincent Loechner. It is distributed under GNU General Public License version 3 or later.

5.3. NLR

Participants: Alain Ketterlin, Philippe Clauss.

We have developed a program implementing our loop-nest recognition algorithm, detailed in [7]. This standalone, filter-like application takes as input a raw trace and builds a sequence of loop nests that, when executed, reproduce the trace. It is also able to predict forthcoming values at an arbitrary distance in the future. Its simple, text-based input format makes it applicable to all kinds of data. These data can take the form of simple numeric values, or have more elaborate structure, and can include symbols. The program is written in standard ANSI C. The code can also be used as a library.

We have used this code to evaluate the compression potential of loop nest recognition on memory address traces, with very good results. We have also shown that the predictive power of our model is competitive with other models on average.

¹<http://icps.u-strasbg.fr/PolyLib>

²<http://ZPolyTrans.gforge.inria.fr>

The software is available upon request to anybody interested in trying to apply loop nest recognition. It has been distributed to a dozen of colleagues around the world. In particular, it has been used by Andres Charif-Rubial for his PhD work (Université de Versailles Saint-Quentin en Yvelines), and is now included in a released tool called MAQAO (<http://www.maqao.org>). Our code is also used by Jean-Thomas ACQUAVIVA, at Commissariat à l'Énergie Atomique, for work on compressing instruction traces. These colleagues have slightly modified the code we gave them. We plan to release a stable version incorporating most of their changes in the near future. We also plan to change the license to avoid such drifts in the future.

5.4. Binary files decompiler

Participant: Alain Ketterlin.

Our research on efficient memory profiling has led us to develop a sophisticated decompiler. This tool analyzes x86-64 binary programs and libraries, and extracts various structured representations of the code. It works on a routine per routine basis, and first builds a loop hierarchy to characterize the overall structure of the algorithm. It then puts the code into Static Single Assignment (SSA) form to highlight the fine-grain data-flow between registers and memory. Building on these, it performs the following analyzes:

- All memory addresses are expressed as symbolic expressions involving specific versions of register contents, as well as loop counters. Loop counter definitions are recovered by resolving linearly incremented registers and memory cells, i.e., registers that act as induction variables.
- Most conditional branches are also expressed symbolically (with registers, memory contents, and loop counters). This captures the control-flow of the program, but also helps in defining what amounts to loop “trip-counts”, even though our model is slightly more general, because it can represent any kind of iterative structure.

This tool embodies several passes that, as far as we know, do not exist in any existing similar tool. For instance, it is able to track data-flow through stack slots in most cases. It has been specially designed to extract a representation that can be useful in looking for parallel (or parallelizable) loops [27]. It is the basis of several of our studies.

Because binary program decompilation is especially useful to reduce the cost of memory profiling, our current implementation is based on the Pin binary instrumenter. It uses Pin's API to analyze binary code, and directly interfaces with the upper layers we have developed (e.g., program skeletonization, or minimal profiling). However, we have been careful to clearly decouple the various layers, and to not use any specific mechanism in designing the binary analysis component. Therefore, we believe that it could be ported with minimal effort, by using a binary file format extractor and a suitable binary code parser. It is also designed to abstract away the detailed instruction set, and should be easy to port (even though we have no practical experience in doing so).

We feel that such a tool could be useful to other researchers, because it makes binary code available under abstractions that have been traditionally available for source code only. If sufficient interest emerges, e.g., from the embedded systems community, or from researchers working on WCET, or from teams working on software security, we are willing to distribute and/or to help make it available under other environments.

5.5. Parwiz: a dynamic dependency analyser

Participant: Alain Ketterlin.

We have developed a dynamic dependence analyzer. Such a tool consumes the trace of memory (or object) accesses, and uses the program structure to list all the data dependences appearing during execution. Data dependences in turn are central to the search for parallel sections of code, with the search for parallel loops being only a particular case of the general problem. Most current works of these questions are either specific to a particular analysis (e.g., computing dependence densities to select code portions for thread-level speculation), or restricted to particular forms of parallelism (e.g., typically to fully parallel loops). Our tool tries to generalize existing approaches, and focuses on the program structures to provide helpful feedback

either to a user (as some kind of “smart profiler”), or to a compiler (for feedback-directed compilation). For example, the tool is able to produce a dependence schema for a complete loop nest (instead of just a loop). It also targets irregular parallelism, for example analyzing a loop execution to estimate the expected gain of parallelization strategies like inspector-executor.

We have developed this tool in relation to our minimal profiling research project. However, the tool itself has been kept independent of our profiling infrastructure, getting data from it via a well-defined trace format. This intentional design decision has been motivated by our work on distinct execution environments: first on our usual x86-64 benchmark programs, and second on less regular, more often written in Java, real-world applications. The latter type of applications is likely the one that will most benefit from such tools, because their intrinsic execution environment does not offer enough structure to allow effective static analysis techniques. Parallelization efforts in this context will most likely rely on code annotations, or specific programming language constructs. Programmers will therefore need tools to help them choose between various constructs. Our tool has this ambition. We already have a working tool-chain for C/C++/Fortran programs (or any binary program). We are in the process of developing the necessary infrastructure to connect the dynamic dependence profiler to instrumented Java programs. Other managed execution environments could be targeted as well, e.g., Microsoft’s .Net architecture, but we have no time and/or workforce to devote to such time-consuming engineering efforts.

5.6. APOLLO software and LLVM

Participants: Aravind Sukumaran-Rajam, Juan Manuel Martinez Caamaño, Willy Wolff, Alexandra Jimborean, Jean-François Dollinger, Philippe Clauss.

We are developing a new framework called APOLLO (Automatic speculative POLyhedral Loop Optimizer) whose main concepts are based on our previous framework VMAD. However, several important implementation issues are now handled differently in order to improve the performance and usability of the framework, and also to open its evolution to new interesting perspectives. Thus VMAD played the role of a prototype which is now being re-written as a sustainable tool named APOLLO. As VMAD, APOLLO is dedicated to automatic, dynamic and speculative parallelization of loop nests that cannot be handled efficiently at compile-time. It is composed of a static part consisting of specific passes in the LLVM compiler suite, plus a modified Clang frontend, and a dynamic part consisting of a runtime system. It is described in more details in subsection 6.1 .

Aravind Sukumaran-Rajam (PhD student), Juan Manuel Martinez Caamaño (PhD student), Jean-François Dollinger (PhD student), Willy Wolff (Master student) and Philippe Clauss are the main contributors of APOLLO. It will soon be distributed.

5.7. IBB source-to-source compiler

Participants: Imen Fassi, Philippe Clauss.

We have developed a multiloop-compiler called IBB for Iterate-But-Better. IBB translates any C program containing multiloop-loops into an equivalent C program in which all multiloop-loops are replaced with equivalent for-loops. The resulting source code can then be compiled using any C compiler to produce executable code. IBB will soon be distributed.

5.8. Polyhedral prover

Participants: Nicolas Magaud, Julien Narboux, Éric Violard [correspondant].

We are currently developing a formal proof of program transformations based on the polyhedral model. We use the CompCert verified compiler [28] as a framework. This tool is written in the specification language of Coq.

5.9. CLooG

Participant: Cédric Bastoul.

CLoog³ is a free software and library to generate code (or an abstract syntax tree of a code) for scanning Z-polyhedra. That is, it finds a code (e.g. in C, FORTRAN...) that reaches each integral point of one or more parameterized polyhedra. CLoog has been originally written to solve the code generation problem for optimizing compilers based on the polyhedral model. Nevertheless it is used now in various area e.g. to build control automata for high-level synthesis or to find the best polynomial approximation of a function. CLoog may help in any situation where scanning polyhedra matters. While the user has full control on generated code quality, CLoog is designed to avoid control overhead and to produce a very effective code. CLoog is widely used (including by GCC and LLVM compilers), disseminated (it is installed by default by the main Linux distributions) and considered as the state of the art in polyhedral code generation.

5.10. OpenScop

Participant: Cédric Bastoul.

OpenScop⁴ is an open specification that defines a file format and a set of data structures to represent a static control part (SCoP for short), i.e., a program part that can be represented in the polyhedral model. The goal of OpenScop is to provide a common interface to the different polyhedral compilation tools in order to simplify their interaction. To help the tool developers to adopt this specification, OpenScop comes with an example library (under 3-clause BSD license) that provides an implementation of the most important functionalities necessary to work with OpenScop.

5.11. Clan

Participant: Cédric Bastoul.

Clan⁵ is a free software and library which translates some particular parts of high level programs written in C, C++, C# or Java into a polyhedral representation called OpenScop. This representation may be manipulated by other tools to, e.g., achieve complex analyses or program restructurations (for optimization, parallelization or any other kind of manipulation). It has been created to avoid tedious and error-prone input file writing for polyhedral tools (such as CLoog, LeTSeE, Candl etc.). Using Clan, the user has to deal with source codes based on C grammar only (as C, C++, C# or Java). Clan is notably the frontend of the two major high-level compilers Pluto and PoCC.

5.12. Candl

Participant: Cédric Bastoul.

Candl⁶ is a free software and a library devoted to data dependences computation. It has been developed to be a basic bloc of our optimizing compilation tool chain in the polyhedral model. From a polyhedral representation of a static control part of a program, it is able to compute exactly the set of statement instances in dependence relation. Hence, its output is useful to build program transformations respecting the original program semantics. This tool has been designed to be robust and precise. It implements some usual techniques for data dependence removal, as array privatization or array expansion, offers simplified abstractions like dependence vectors and performs violation dependence analysis. Candl is notably the dependence analyzer of the two major high-level compilers Pluto and PoCC.

5.13. Clay

Participant: Cédric Bastoul.

³<http://www.cloog.org>

⁴<http://icps.u-strasbg.fr/~bastoul/development/openscop>

⁵<http://icps.u-strasbg.fr/~bastoul/development/clang>

⁶<http://icps.u-strasbg.fr/~bastoul/development/candl>

Clay⁷ is a free software and library devoted to semi-automatic optimization using the polyhedral model. It can input a high-level program or its polyhedral representation and transform it according to a transformation script. Classic loop transformations primitives are provided. Clay is able to check for the legality of the complete sequence of transformation and to suggest corrections to the user if the original semantics is not preserved. Clay is still experimental at this report redaction time but is already used during advanced compilation labs at Paris-Sud University and is one of the foundations of our ongoing work on simplifying code manipulation by programmers.

⁷<http://icps.u-strasbg.fr/~bastoul/development/clay>

CAMEL Project-Team

5. Software and Platforms

5.1. Introduction

A major part of the research done in the CAMEL team is published within software. On the one hand, this enables everyone to check that the algorithms we develop are really efficient in practice; on the other hand, this gives other researchers — and us of course — basic software components on which they — and we — can build other applications.

5.2. GNU MPFR

Participant: Paul Zimmermann [contact].

GNU MPFR is one of the main pieces of software developed by the CAMEL team. Since end 2006, with the departure of Vincent Lefèvre to ENS Lyon, it has become a joint project between CAMEL and the ARÉNAIRE project-team (now AriC, INRIA Grenoble - Rhône-Alpes). GNU MPFR is a library for computing with arbitrary precision floating-point numbers, together with well-defined semantics, and is distributed under the LGPL license. All arithmetic operations are performed according to a rounding mode provided by the user, and all results are guaranteed correct to the last bit, according to the given rounding mode.

Several software systems use GNU MPFR, for example: the GCC and GFORTRAN compilers; the SAGE computer algebra system; the KDE calculator Abakus by Michael Pyne; CGAL (Computational Geometry Algorithms Library) developed by the Geometrica project-team (INRIA Sophia Antipolis - Méditerranée); Gappa, by Guillaume Melquiond; Sollya, by Sylvain Chevillard, Mioara Joldeş and Christoph Lauter; Genius Math Tool and the GEL language, by Jiri Lebl; Giac/Xcas, a free computer algebra system, by Bernard Parisse; the iRRAM exact arithmetic implementation from Norbert Müller (University of Trier, Germany); the Magma computational algebra system; and the Wcalc calculator by Kyle Wheeler.

The main development in 2013 is the release of version 3.1.2 (the “canard à l’orange” release) in March. This version fixes a few bugs from previous version.

5.3. GNU MPC

Participant: Paul Zimmermann [contact].

GNU MPC is a floating-point library for complex numbers, which is developed on top of the GNU MPFR library, and distributed under the LGPL license. It is co-written with Andreas Enge (LFANT project-team, INRIA Bordeaux - Sud-Ouest). A complex floating-point number is represented by $x + iy$, where x and y are real floating-point numbers, represented using the GNU MPFR library. The GNU MPC library provides correct rounding on both the real part x and the imaginary part y of any result. GNU MPC is used in particular in the TRIP celestial mechanics system developed at IMCCE (*Institut de Mécanique Céleste et de Calcul des Éphémérides*), and by the Magma and Sage computational number theory systems.

No new version of GNU MPC was released in 2013, which confirms the status of mature library.

5.4. GMP-ECM

Participants: Cyril Bouvier, Paul Zimmermann [contact].

GMP-ECM is a program to factor integers using the Elliptic Curve Method. Its efficiency comes both from the use of the GMP library, and from the implementation of state-of-the-art algorithms. GMP-ECM contains a library (LIBECM) in addition to the binary program (ECM). The binary program is distributed under GPL, while the library is distributed under LGPL, to allow its integration into other non-GPL software. The Magma computational number theory software and the SAGE computer algebra system both use LIBECM.

In February 2013, a new version 6.4.4 was released. Apart from bug fixes, this new release provides some improvements (better integration of the GPU code, of the new *-batch* option, ...).

In September 2013, a new record prime of 83 digits was found by R. Propper using GMP-ECM.

5.5. Finite Fields

Participants: Pierrick Gaudry, Emmanuel Thomé [contact], Luc Sanselme.

$\text{mp}\mathbb{F}_q$ is (yet another) library for computing in finite fields. The purpose of $\text{mp}\mathbb{F}_q$ is not to provide a software layer for accessing finite fields determined at runtime within a computer algebra system like Magma, but rather to give a very efficient, optimized code for computing in finite fields precisely known at *compile time*. $\text{mp}\mathbb{F}_q$ can adapt to finite fields of any characteristic and any extension degree. However, one of the targets being the use in cryptology, $\text{mp}\mathbb{F}_q$ somehow focuses on prime fields and on fields of characteristic two.

When it was first written in 2007, $\text{mp}\mathbb{F}_q$ established reference marks for fast elliptic curve cryptography: the authors improved over the fastest examples of key-sharing software in genus 1 and 2, both over binary fields and prime fields. A stream of academic works followed the idea behind $\text{mp}\mathbb{F}_q$ and improved over such timings, notably by Scott, Aranha, Longa, Bos, Hisil, Costello.

The library's purpose being the *generation* of code rather than its execution, the working core of $\text{mp}\mathbb{F}_q$ consists of roughly 18,000 lines of Perl code, which generate most of the C code. $\text{mp}\mathbb{F}_q$ is distributed at <http://mpfq.gforge.inria.fr/>.

In 2013, version 1.1 of $\text{mp}\mathbb{F}_q$ has been released. This new release includes new assembly code by Luc Sanselme providing optimized arithmetic over fields whose characteristic fits in a number of bits which fit within half-word boundaries.

In 2013, Hamza Jeljeli collaborated with Bastien Vialla from LIRMM, Montpellier to integrate experimental code based on RNS arithmetic (Residue Number System), intending to provide back-end functionality for the linear algebra code in CADO-NFS. This feature set is still experimental.

5.6. gf2x

Participants: Pierrick Gaudry, Emmanuel Thomé [contact], Paul Zimmermann.

GF2X is a software library for polynomial multiplication over the binary field, developed together with Richard Brent (Australian National University, Canberra, Australia). It holds state-of-the-art implementation of fast algorithms for this task, employing different algorithms in order to achieve efficiency from small to large operand sizes (Karatsuba and Toom-Cook variants, and eventually Schönhage's or Cantor's FFT-like algorithms). GF2X takes advantage of specific processors instruction (SSE, PCLMULQDQ).

The current version of GF2X is 1.1, released in May 2012 under the GNU GPL. Since 2009, GF2X can be used as an auxiliary package for the widespread software library NTL, as of version 5.5.

In 2013, the development version of GF2X has been updated to incorporate detection of Intel Haswell micro-processors, which provide much improved performance for the PCLMULQDQ instruction (this instruction is of utmost importance for GF2X).

An LGPL-licensed portion of GF2X is also part of the CADO-NFS software package.

5.7. CADO-NFS

Participants: Cyril Bouvier, Jérémie Detrey, Alain Filbois, Pierrick Gaudry, Alexander Kruppa, Emmanuel Thomé [contact], Paul Zimmermann.

CADO-NFS is a program to factor integers using the Number Field Sieve algorithm (NFS), originally developed in the context of the ANR-CADO project (November 2006 to January 2010).

NFS is a complex algorithm which contains a large number of sub-algorithms. The implementation of all of them is now complete, but still leaves some places to be improved. Compared to existing implementations, the CADO-NFS implementation is already a reasonable player. Several factorizations have been completed using our implementations.

Since 2009, the source repository of CADO-NFS is publicly available for download, and is referenced from the software page at <http://cado-nfs.gforge.inria.fr/>. A major new release, CADO-NFS 2.0, was published in November 2013. The client/server framework was completely rewritten to allow the use of CADO-NFS routinely on clusters of 100 to 1000 nodes.

More and more people use CADO-NFS to perform medium to large factorizations. Also in 2013 some researchers in the field wrote some papers where they study the implementation and default parameters of CADO-NFS. This is very useful feedback from the scientific community.

5.8. Belenios

Participants: Pierrick Gaudry, Stéphane Glondou [contact].

In collaboration with the CASSIS team, we develop an open-source private and verifiable electronic voting protocol, named BELENIOS. Our system is an evolution of an existing system, Helios, developed by Ben Adida, and used e.g. by UCL and the IACR association in real elections. The main differences with Helios are the following ones:

- In Helios, the ballot box publishes the encrypted ballots together with their corresponding voters. This raises a privacy issue in the sense that whether someone voted or not shall not necessarily be publicized on the web. Publishing this information is in particular forbidden by CNIL's recommendation. BELENIOS no longer publishes voters' identities, still guaranteeing correctness of the tally.
- Helios is verifiable except that one has to trust that the ballot box will not add ballots. The addition of ballots is particularly hard to detect as soon as the list of voters is not public. We have therefore introduced an additional authority that provides credentials that the ballot box can verify but not forge.

This new version has been implemented by Stéphane Glondou ¹. and has been tested in July 2013 in a mock election in the teams CASSIS and CAMEL.

¹<http://belenios.gforge.inria.fr/>

CARTE Project-Team

5. Software and Platforms

5.1. Morphus/MMDEX

MMDEX is a virus detector based on morphological analysis. It is composed of our own disassembler tool, on a graph transformer and a specific tree-automaton implementation. The tool is used in the EU-Fiware project and by some other partners (e.g. DAVFI project).

Written in C, 20k lines.

APP License, IDDN.FR.001.300033.000.R.P.2009.000.10000, 2009.

5.2. TraceSurfer

TraceSurfer is a self-modifying code analyzer coming with an IDA add-on. It works as a wave-builder. In the analysis of self-modifying programs, one basic task is indeed to separate parts of the code which are self-modifying into successive layers, called waves. TraceSurfer extracts waves from traces of program executions. Doing so drastically simplifies program verification.

Written in C, 5k lines.

<http://code.google.com/p/tartetatintools/>

5.3. CROCUS

CROCUS is a program interpretation synthetizer. Given a first order program (possibly written in OCAML), it outputs a quasi-interpretation based on max, addition and product. It is based on a random algorithm. The interpretation is actually a certificate for the program's complexity. Users are non academics (some artists).

Written in Java, 5k lines.

CASSIS Project-Team

5. Software and Platforms

5.1. Protocol Verification Tools

Participants: Stéphane Glondu, Pierre-Cyrille Héam, Olga Kouchnarenko, Steve Kremer, Michaël Rusinowitch, Mathieu Turuani, Laurent Vigneron.

5.1.1. AVISPA

Cassis has been involved in the European project AVISPA, which has resulted in the distribution of a tool for automated verification of security protocols, named AVISPA Tool. It is freely available on the web ¹ and it is well supported. The AVISPA Tool compares favourably to related systems in scope, effectiveness, and performance, by (i) providing a modular and expressive formal language for specifying security protocols and properties, and (ii) integrating 4 back-ends that implement automatic analysis techniques ranging from *protocol falsification* (by finding an attack on the input protocol) to *abstraction-based verification* methods for both finite and infinite numbers of sessions.

5.1.2. CL-AtSe

We develop, as a back-end of AVISPA, *CL-AtSe*, a Constraint Logic based Attack Searcher for cryptographic protocols. The *CL-AtSe* approach to verification consists in a symbolic state exploration of the protocol execution, for a bounded number of sessions. This necessary restriction (for decidability, see [77]) allows *CL-AtSe* to be correct and complete. Each protocol step is represented by a constraint on the protocol state, used to check for reachability of the next state. *CL-AtSe* includes a proper handling of sets, choice points, specification of any attack states through a language for expressing e.g. secrecy, authentication, fairness, or non-abuse freeness, advanced protocol simplifications and optimizations to reduce the problem complexity, and protocol analysis modulo the algebraic properties of cryptographic operators such as XOR (exclusive or) and Exp (modular exponentiation).

CL-AtSe has been successfully used [65] to analyse France Telecom R&D, Siemens AG, IETF, or Gemalto protocols in funded projects. It is also employed by external users, e.g., from the AVISPA's community. Moreover, *CL-AtSe* achieves very good analysis times, comparable and sometimes better than state-of-the-art tools in the domain (see [82] for tool details and precise benchmarks).

CL-Atse has been enhanced in various ways. In particular, the tool fully supports Aslan semantics introduced in [63], including Horn Clauses (for intruder-independent deductions, e.g. for credential management), and LTL-based security properties. Also, bug information and correction are processed through a bugzilla server, and online analysis and orchestration are available on our team server (<https://cassis.loria.fr>). CL-Atse supports negative constraints on the intruder's knowledge [66]. This extension of CL-Atse allows us to reduce drastically the orchestrator's processing times. It has also been used to model e.g. separation of duties and non-disclosure policies. We have also extended the syntax and semantics of ASLan to better model lists of undefined length, directly inside messages. CL-AtSe tool now supports membership predicates, deletion operators and so on for managing these lists, and offers a first reference implementation for other tools in Avantssar. In particular, the ASLan translator has been updated by our partners.

¹<http://www.avispa-project.org>

5.1.3. *Akiss*

We develop the *Akiss* (Active Knowledge in Security Protocols) tool for verifying indistinguishability properties in cryptographic protocols. Indistinguishability properties are essential in formal verification of cryptographic protocols. They are needed to model anonymity properties, strong versions of confidentiality and resistance against offline guessing attacks, which can be conveniently modeled using process equivalences. *Akiss* implements a procedure to verify equivalence properties for a bounded number of sessions of cryptographic protocols. As in the applied pi-calculus, the protocol specification language is parametrized by a first-order sorted term signature and an equational theory which allows formalization of algebraic properties of cryptographic primitives. *Akiss* is able to verify trace equivalence for determinate cryptographic protocols. On determinate protocols, trace equivalence coincides with observational equivalence which can therefore be automatically verified for such processes. When protocols are not determinate *Akiss* can be used for both under- and over-approximations of trace equivalence, which proved successful on several examples. The procedure can handle a large set of cryptographic primitives, namely those that can be modeled by an optimally reducing convergent rewrite system.

The underlying procedure is based on a fully abstract modelling of the traces of a bounded number of sessions of the protocols into first-order Horn clauses on which a dedicated resolution procedure is used to decide equivalence properties. Although termination of the resolution procedure has not been proved, the procedure has been effectively tested on examples, some of which are outside the scope of other existing tools, including checking anonymity in several electronic voting protocols.

Recent developments include the possibility for checking everlasting indistinguishability properties. This feature was added when analyzing everlasting privacy properties in electronic voting protocols. We are currently working on a generalization of the procedure to allow associative-commutative operators and in particular a re-design of the resolution procedure for allowing analysis of protocols that use exclusive or. Expected case studies for this development include unlinkability in RFID protocols.

The *Akiss* tool is freely available at <https://github.com/ciobaca/akiss>.

5.1.4. *Belenios*

In collaboration with the Caramel team, we develop an open-source private and verifiable electronic voting protocol, named *Belenios*. Our system is an evolution of an existing system, Helios, developed by Ben Adida, and used e.g. by UCL and the IACR association in real elections. The main differences with Helios are the following ones:

- In Helios, the ballot box publishes the encrypted ballots together with their corresponding voters. This raises a privacy issue in the sense that whether someone voted or not shall not necessarily be publicized on the web. Publishing this information is in particular forbidden by the CNIL's recommendations. *Belenios* no longer publishes voters' identities, still guaranteeing the correctness of the tally.
- Helios is verifiable except that one has to trust that the ballot box will not add ballots. The addition of ballots is particularly hard to detect as soon as the list of voters is not public. We have therefore introduced an additional authority that provides credentials that the ballot box can verify but not forge.

This new version has been implemented by Stéphane Glondu and has been tested in July 2013 in a mock election in the teams Cassis and Caramel.

In a first step, *Belenios* has been implemented as an extension of existing Helios system. However, the existing software development of Helios is large and its security becomes difficult to assess. We have therefore re-implemented entirely the code of the bulletin box, yielding a now independent software ².

²<http://belenios.gforge.inria.fr/>

In Helios as well as *Belenios*, votes are encrypted using the public key of the election. To ensure privacy, the corresponding decryption key is not known to anyone. Instead, several authorities detain a share of it. For robustness reasons (and as recommended by the CNIL), it is important to be able to decrypt even if some of the authorities are missing. We have implemented the threshold decryption scheme that we have proposed [40]. This implementation is currently available only within the Helios system and we plan to integrate it to *Belenios* in the next months.

5.2. Testing Tools

Participants: Fabrice Bouquet, Frédéric Dadeau, Kalou Cabrera.

5.2.1. Hydra

Hydra is an Eclipse-like platform, based on Plug-ins architecture. Plug-ins can be of five kinds: *parser* is used to analyze source files and build an intermediate format representation of the source; *translator* is used to translate from a format to another or to a specific file; *service* denotes the application itself, i.e. the interface with the user; *library* denotes an internal service that can be used by a service, or by other libraries; *tool* encapsulates an external tool. The following services have been developed so far:

- BZPAnimator: performs the animation of a BZP model (a B-like intermediate format);
- Angluin: makes it possible to perform a machine learning algorithm (à la Angluin) in order to extract an abstraction of a system behavior;
- UML2SMT: aims at extracting first order logic formulas from the UML Diagrams and OCL code of a UML/OCL model to check them with a SMT solver.

These services involve various libraries (sometimes reusing each other), and rely on several *tool* plug-ins that are: SMTProver (encapsulating Z3 solver), PrologTools (encapsulating CLPS-B solver), Grappa (encapsulating a graph library). We are currently working on transferring the existing work on test generation from B abstract machines, JML, and statecharts using constraint solving techniques.

5.2.2. jMuHLPSL

jMuHLPSL [9] is a mutant generator tool that takes as input a verified HLPSL protocol, and computes mutants of this protocol by applying systematic mutation operators on its contents. The mutated protocol then has to be analyzed by a dedicated protocol analysis tool (here, the AVISPA tool-set). Three verdicts may then arise. The protocol can still be *safe*, after the mutation, this means that the protocol is not sensitive to the realistic “fault” represented by the considered mutation. This information can be used to inform the protocol designers of the robustness of the protocol w.r.t. potential implementation choices, etc. The protocol can also become *incoherent*, meaning that the mutation introduced a functional failure that prevents the protocol from being executed entirely (one of the participants remains blocked in a given non-final state). The protocol can finally become *unsafe* when the mutation introduces a security flaw that can be exploited by an attacker. In this case, the AVISPA tool-set is able to compute an attack-trace, that represents a test case for the implementation of the protocol. If the attack can be replayed entirely, then the protocol is not safe. If the attack can not be replayed then the implementation does not contain the error introduced in the original protocol.

The tool is written in Java, and it is freely available at: <http://members.femto-st.fr/sites/femto-st.fr/frederic-dadeau/files/content/pub/jMuHLPSL.jar>.

5.3. Collaborative Tools

Participant: Abdessamad Imine.

The collaborative tools allow us to manage collaborative works on shared documents using flexible access control models. These tools have been developed in order to validate and evaluate our approach on combining collaborative edition with optimistic access control.

5.3.1. P2PEdit

This prototype is implemented in Java and supports the collaborative editing of HTML pages and it is deployed on P2P JXTA platform ³. In our prototype, a user can create a HTML page from scratch by opening a new collaboration group. Other users (peers) may join the group to participate in HTML page editing, as they may leave this group at any time. Each user can dynamically add and remove different authorizations for accessing to the shared document according the contribution and the competence of users participating in the group. Using JXTA platform, users exchange their operations in real-time in order to support WYSIWIS (What You See Is What I See) principle. Furthermore, the shared HTML document and its authorization policy are replicated at the local memory of each user. To deal with latency and dynamic access changes, an optimistic access control technique is used where enforcement of authorizations is retroactive.

5.3.2. P2PCalendar

To extend our collaboration and access control models to mobile devices, we implemented a shared calendar on iPhone OS which is decentralized and scalable (i.e. it can be used over both P2P and ad-hoc networks). This application aims to make a collaborative calendar where users can simultaneously modify events (or appointments) and control access on events. The access rights are determined by the owner of an event. The owner decides who is allowed to access the event and what privileges they have. Likewise to our previous tool, the calendar and its authorization policy are replicated at every mobile device.

5.4. Other Tools

Several software tools described in previous sections are using tools that we have developed in the past. For instance BZ-TT uses the set constraints solver CLPS. Note that the development of the SMT prover haRVey has been stopped. The successor of haRVey is called veriT and is developed by David Déharbe (UFRN Natal, Brasil) and Pascal Fontaine (Veridis team). We have also developed, as a second back-end of AVISPA, TA4SP (Tree Automata based on Automatic Approximations for the Analysis of Security Protocols), an automata based tool dedicated to the validation of security protocols for an unbounded number of sessions.

³<http://www.sun.com/software/jxta/>

CORIDA Project-Team

4. Software and Platforms

4.1. Simulation of viscous fluid-structure interactions

Participants: Takeo Takahashi [correspondant], Jean-François Scheid, Jérôme Lohéac.

A number of numerical codes for the simulation for fluids and fluid-structure problems has been developed by the team. These codes are mainly written in MATLAB Software with the use of C++ functions in order to improve the sparse array process of MATLAB. We have focused our attention on 3D simulations which require large CPU time resources as well as large memory storage. In order to solve the 3D Navier-Stokes equations which model the viscous fluid, we have implemented an efficient 3D Stokes sparse solver for MATLAB and a 3D characteristics method to deal with the nonlinearity of Navier-Stokes equations. This year, we have also started to unify our 2D fluid-structure codes (fluid alone, fluid with rigid bodies and fluid with fishes).

Another code has been developed in the case of self-propelled deformable object moving into viscous fluid. Our aim is to build a deformable ball which could swim in a viscous fluid. In order to do this we have started a collaboration with a team from the CRAN (Research Centre for Automatic Control). This software solves numerically 3D Stokes equations using finite elements methods. The source code is written for use with MATLAB thanks to a C++ library developed by ALICE.

- Version: v0.5
- Programming language: MATLABc++

4.2. Fish locomotion in perfect fluids with potential flow

Participants: Alexandre Munnier [correspondant], Marc Fuentes, Bruno Pinçon.

SOLEIL is a Matlab suite to simulate the self-propelled swimming motion of a single 3D swimmer immersed in a potential flow. The swimmer is modeled as a shape-changing body whose deformations can be either prescribed as a function of time (simulation of the direct swimming problem) or computed in such a way that the swimmer reaches a prescribed location (control problem). For given deformations, the hydrodynamical forces exerted by the fluid on the swimmer are expressed as solutions of 2D integral equations on the swimmer's surface, numerically solved by means of a collocation method.

SOLEIL is free, distributed under licence GPL v3. More details are available on the project web page <http://soleil.gforge.inria.fr/>.

The next step of SOLEIL (under progress) is to take into account a fluid whose flow is governed by Stokes equations.

- Version: 0.1
- Programming language: Matlab/C++

4.3. SUSHI3D : SimUlations of Structures in Hydrodynamic Interactions

Participants: Marc Fuentes, Jean-François Scheid, Jérémy Sinoir, Takéo Takahashi, Rhaleb Zayer.

SUSHI3D is a 3D solver for numerical simulations of Fluid/Structures Interactions. The Navier-Stokes equations are coupled with the dynamics of immersed bodies which can be either rigid or deformable. The deformable body case is handled and designed for fish-swimming. The numerical method used to solve the full differential system is based on a Lagrange-Galerkin method with finite elements.

- Version: 1.0
- Programming language: Matlab/C++

CORTEX Team

5. Software and Platforms

5.1. Spiking neural networks simulation

Participants: Dominique Martinez, Yann Boniface.

A spiking neuron is usually modeled as a differential equation describing the evolution over time of its membrane potential. Each time the voltage reaches a given threshold, a spike is sent to other neurons depending on the connectivity. A spiking neural network is then described as a system of coupled differential equations. For the simulation of such a network we have written two simulation engines: (i) Mvaspike based on an event-driven approach and (ii) sirene based on a time-driven approach.

- Mvaspike: The event-driven simulation engine was developed in C++ and is available on <http://mvaspike.gforge.inria.fr>. Mvaspike is a general event-driven purpose tool aimed at modeling and simulating large, complex networks of biological neural networks. It allows to achieve good performance in the simulation phase while maintaining a high level of flexibility and programmability in the modeling phase. A large class of spiking neurons can be used ranging from standard leaky integrate-and-fire neurons to more abstract neurons, e.g. defined as complex finite state machines.
- Sirene: The time-driven simulator engine was written in C and is available on <http://sirene.gforge.inria.fr>. It has been developed for the simulation of biologically detailed models of neurons —such as conductance-based neurons— and synapses. Its high flexibility allows the user to implement easily any type of neuronal or synaptic model and use the appropriate numerical integration routine (e.g. Runge-Kutta at given order).

5.2. CLONES: Closed-Loop Neural Simulations

Participant: Thomas Voegtlin.

The goal of this work is to provide an easy-to-use framework for closed-loop simulations, where interactions between the brain and body of an agent are simulated.

We developed an interface between the Sofa physics engine, (<http://www.sofa-framework.org>) and the Brian neural simulator (<http://www.briansimulator.org>). The interface consists in a Sofa plugin and a Python module for Brian. Sofa and Brian use different system processes, and communicate via shared memory. Synchronization between processes is achieved through semaphores.

As a demonstration of this interface, a physical model of undulatory locomotion in the nematode *c. elegans* was implemented, based on the PhD work of Jordan H. Boyle.

MADYNES Project-Team

5. Software and Platforms

5.1. SecSIP

Participants: Abdelkader Lahmadi [contact], Olivier Festor.

*SecSip*¹ is developed by the team to defend SIP-based (The Session Initiation Protocol) services from known vulnerabilities. It presents a proactive point of defense between a SIP-based network of devices (servers, proxies, user agents) and the open Internet. Therefore, all SIP traffic is inspected and analyzed against authored Veto specification before it is forwarded to these devices. When initializing, the SecSIP runtime starts loading and parsing authored VeTo blocks to identify different variables, event patterns, operations and actions from each rule. Veto is a generic declarative language for attack patterns specification. SecSIP implements an input and output layer, to capture, inject, send and receive SIP packets from and to the network. Intercepted packets are moved to the SIP Packet parser module. The main function of this module is to extract different fields within a SIP message and trigger events specified within the definition blocks. During each execution cycle when a SIP message arrives, the SecSIP runtime uses a data flow acyclic graph network to find definition matching rules and triggers defined events. The paired events in each operator node are propagated over the graph until a pattern is satisfied. When the pattern is satisfied, the respective rule is fired and the set of actions is executed.

5.2. NDPMon

Participants: Isabelle Chrisment, Olivier Festor [contact].

The Neighbor Discovery Protocol Monitor (**NDPMon**) is an IPv6 implementation of the well-known ArpWatch tool. NDPMon monitors the pairing between IPv6 and Ethernet addresses (NDP activities: new station, changed Ethernet address, flip flop...). NDPMon also detects attacks on the NDP protocol, as defined in RFC 3756 (bogon, fake Router Advertisements...). New attacks based on the Neighbor Discovery Protocol and Address Auto-configuration (RFC 2461 and RFC 2462) have been identified and integrated in the tool. An XML file describes the default behavior of the network, with the authorized routers and prefixes, and a second XML document containing the neighbors database is used. This second file can be filled during a learning phase. All NDP activities are logged in the syslog utility, and so the attacks, but these ones are also reported by mail to the administrator. Finally, NDPMon can detect stack vulnerabilities, like the assignment of an Ethernet broadcast address on an interface.

NDPMon comes along with a WEB interface acting as a GUI to display the informations gathered by the tool, and give an overview of all alerts and reports. Thanks to color codes, the WEB interface makes possible for the administrator to have an history of what happened on his network and identify quickly problems. All the XML files used or produced by the daemon (neighbor cache, configuration file and alerts list) are translated in HTML via XSL for better readability. A statistic module is also integrated and gives informations about the discovery of the nodes and their type (MAC manufacturer distribution ...).

The software package and its source code is freely distributed under an opensource license (LGPL). It is implemented in C, and is available through a SourceForge project at <http://ndpmon.sf.net>. An open source community is now established for the tool which has distributions for several Operating Systems (Linux, FreeBSD, OpenBSD, NetBSD and Mac OS X). It is also integrated in FreeBSD ports². Binary distributions are also available for .deb and .rpm based Linux flavors.

¹<http://secsip.gforge.inria.fr/doku.php>

²<http://www.freebsd.org/cgi/cvsweb.cgi/ports/net-mgmt/ndpmon/>

5.3. AA4MM

Participants: Laurent Ciarletta [contact], Yannick Presse.

Vincent Chevrier (MAIA team, contact), and Benjamin Camus and Julien Vaubourg (MAIA team, LORIA) are contributors for this software.

AA4MM (Agents and Artefacts for Multi-modeling and Multi-simulation) is a framework for coupling existing and heterogeneous models and simulators in order to model and simulate complex systems. The first implementation of the AA4MM meta-model was proposed in Julien Siebert's PhD [49] and written in Java. This version is currently being put into APP (Agence pour la protection des programmes).

This year, we have used this software in a strategic action with EDF R&D in the context of the simulation of smart-grids. Julien Vaubourg started a PhD on this project that is co-directed by Laurent Ciarletta and Vincent Chevrier.

5.4. MASDYNE

Participant: Laurent Ciarletta [contact].

This work was undertaken in a joint PhD Thesis between MAIA and Madynes Team. Vincent Chevrier (MAIA team, LORIA) has been director and co-advisor of this PhD and is correspondent for this software, which has been used by Tomas Navarrete (MAIA team, LORIA). Other contributors to this software were: Julien Siebert, Tom Leclerc, François Klein, Christophe Torin, Marcel Lamenu, Guillaume Favre and Amir Toly.

MASDYNE (Multi-Agent Simulator of DYnamic Networks usErs) is a multi-agent simulator for modeling and simulating users behaviors in mobile ad hoc network. This software is part of joint work with the MAIA team, as part as a modeling and simulation of ubiquitous networks effort.

MAGRIT Project-Team

5. Software and Platforms

5.1. Software and Platforms

Our software efforts are integrated in a library called RALib which contains our research development on image processing, registration (2D and 3D) and visualization. This library is licensed by the APP (French agency for software protection).

The visualization module is called QGLSG: it enables the visualization of images, 2D and 3D objects under a consistent perspective projection. It is based on Qt ¹ and OpenScenegraph ² libraries. The QGLSG library integrates innovative features such as online camera distortion correction, and invisible objects that can be incorporated in a scene so that virtual objects can cast shadows on real objects, and occlusion between virtual and real objects are easier to handle. The library was also ported to Mac OS and Windows and a full doxygen documentation was written.

¹<http://qt.digia.com>

²<http://www.openscenegraph.org/projects/osg>

MAIA Project-Team

5. Software and Platforms

5.1. AA4MM

Participants: Vincent Chevrier [correspondant], Benjamin Camus, Julien Vaubourg.

Laurent Ciarletta (Madyne team, LORIA) is a collaborator and correspondant for this software. Yannick Presse (Madyne team, LORIA) is collaborator for this software.

AA4MM (Agents and Artefacts for Multi-modeling and Multi-simulation) is a framework for coupling existing and heterogeneous models and simulators in order to model and simulate complex systems. The first implementation of the AA4MM meta-model was proposed in Julien Siebert's PhD [65] and written in Java. A newer version with more coupling models is currently submitted to the APP (Agence pour la protection des programmes).

This year, we used this software in a strategic action with EDF R&D in the context of the simulation of smart-grids.

5.2. MASDYNE

Participants: Vincent Chevrier [correspondant], Tomas Navarrete.

This work was undertaken in the PhD Thesis of Julien Siebert, a joint thesis between MAIA and Madyne Team. Laurent Ciarletta (Madyne team, LORIA) has been co-advisor of this PhD and correspondant for this software.

Other contributors to this software were: Tom Leclerc, François Klein, Christophe Torin, Marcel Lamenu, Guillaume Favre and Amir Toly.

MASDYNE (Multi-Agent Simulator of DYnamic Networks usErs) is a multi-agent simulator for modeling and simulating users behaviors in mobile ad hoc network. This software is part of joint work with MADYNES team, on modeling and simulation of ubiquitous networks. It has been updated by Tomas Navarrete with new functionalities for the simulation of scenarii.

5.3. FiatLux

Participant: Nazim Fatès [correspondant].

FiatLux is a discrete dynamical systems simulator that allows the user to experiment with various models (for example 1D and 2D cellular automatas, moving agents on cellular automatas) and to perturb them. Its main feature is to allow users to change the type of updating, for example from a deterministic parallel updating to an asynchronous random updating. FiatLux has a Graphical User Interface and can also be launched in a batch mode for the experiments that require statistics.

In 2013, FiatLux was officially registered by the Agence pour la protection des programmes (APP). A new release is available under the CeCILL licence on the FiatLux website : fiatlux.loria.fr

5.4. Cart-o-matic

Participants: Olivier Simonin [correspondant], François Charpillet, Antoine Bautin, Nicolas Beaufort.

Philippe Lucidarme (Université d'Angers, LISA) is a collaborator and the coordinator of the Cart-o-matic project.

Cart-o-matic is a software platform for (multi-)robot exploration and mapping tasks. It has been developed by Maia members and LISA (Univ. Angers) members during the robotics ANR/DGA Carotte challenge (2009-2012). This platform is composed of three softwares tools which are protected by software copyrights (through the Agence pour la Protection des Programmes): Slam-o-matic a SLAM algorithm developed by LISA members, Plan-o-matic a robot trajectory planning algorithm developed by Maia and LISA members, and Expl-o-matic a distributed multi-agent strategy for multi-robot exploration developed by Maia members (which is based on algorithms proposed in the PhD Thesis of Antoine Bautin). Cf. illustration at [Cart-o-matic](#). The purchase of Cart-o-matic by some robotics companies is underway.

MASAIE Project-Team (section vide)

NEUROSYS Team

5. Software and Platforms

5.1. Software and Platform

5.1.1. Visualization

- The NeuralFieldSimulator¹ computes numerically activity in two-dimensional neural fields by solving integral-differential equations involving transmission delays and visualizes the spatio-temporal activity. The tool includes a GUI that allows the user to choose field parameters. It is written in Python, open-source and is aimed to be promoted to become a major graphical visualization tool in the domain of neural field theory.
- AnaesthesiaSimulator² simulates the activity of networks of spiking neurons subject to specific receptor dynamics. The tool is a platform to test effects of anaesthetics on neural activity and is still in its first stage of development. The neural activity is planned to be visualized in a 2D and 3D-plot evolving in time. It is written in Python, open-source and involves heavily the simulation package BRIAN³.

5.1.2. Platforms

OpenViBE⁴ is a C++ open-source software devoted to the design, test and use of Brain-Computer Interfaces. The OpenViBE platform consists of a set of software modules that can be integrated easily and efficiently to design BCI applications. Key features of the platform are its modularity, high-performance, portability, its multiple-users facilities and its connection with high-end/Virtual Reality displays. The designer tool of the platform enables to build complete scenarios based on existing software modules using a dedicated graphical language and a simple Graphical User Interface (GUI). This software is available on the Inria Forge⁵ under the terms of the LGPL-V2 license. The development of OpenVibe is done in association with other Inria research teams (Hybrid, Athena, Potioc) for the national Inria project: ADT OpenViBE-NT. Neurosys is in charge of machine learning techniques and the interoperability with other tools such as Matlab, BCI2000, or TOBI.

5.1.3. Others

The package DEvariants⁶ includes Matlab routines which implements new variants of the Differential Evolution (an evolutionary algorithm) strategies. The novelty lies in the selection process where we proposed to use a multinomial law to recombine the individuals/vectors. Compared to the standard strategies, our variants allow a faster convergence and a better avoidance of local minima. The different variants are provided with a test sample of functions, the DeJong benchmark. The audience is any scientific user familiar with evolutionary optimization.

¹<https://gforge.inria.fr/projects/nfsimulator/>

²<https://gforge.inria.fr/projects/anasim/>

³<http://briansimulator.org/>

⁴<http://openvibe.inria.fr/>

⁵<https://gforge.inria.fr/projects/openvibe/>

⁶<https://sites.google.com/site/laurebuhry/publications/optimization-algorithms>

ORPAILLEUR Project-Team

5. Software and Platforms

5.1. Generic Symbolic KDD Systems

5.1.1. *The Coron Platform*

Participants: Jérémie Bourseau [contact person], Aleksey Buzmakov, Victor Codocedo, Adrien Coulet, Amedeo Napoli, Yannick Toussaint.

Keywords: data mining, frequent itemset, closed itemset, generator, association rule, rare itemset

The Coron platform [117], [101] is a KDD toolkit organized around three main components: (1) Coron-base, (2) AssRuleX, and (3) pre- and post-processing modules. The software was registered at the “Agence pour la Protection des Programmes” (APP) and is freely available (see <http://coron.loria.fr>). The Coron-base component includes a complete collection of data mining algorithms for extracting itemsets such as frequent itemsets, closed itemsets, generators and rare itemsets. In this collection we can find APriori, Close, Pascal, Eclat, Charm, and, as well, original algorithms such as ZART, Snow, Touch, and Talky-G. AssRuleX generates different sets of association rules (from itemsets), such as minimal non-redundant association rules, generic basis, and informative basis. In addition, the Coron system supports the whole life-cycle of a data mining task and proposes modules for cleaning the input dataset, and for reducing its size if necessary. The Coron toolkit is developed in Java, is operational, and was already used in several research projects.

5.1.2. *Orion: Skycube Computation Software*

Participant: Chedy Raïssi [contact person].

Keywords: skyline, skycube

This program implements the algorithms described in a research paper published at VLDB 2010 [111]. The software provides a list of four algorithms discussed in the paper in order to compute skycubes. This is the most efficient –in term of space usage and runtime– implementation for skycube computation (see <https://github.com/leander256/Orion>).

5.2. Stochastic systems for knowledge discovery and simulation

5.2.1. *The CarottAge system*

Participants: Florence Le Ber, Jean-François Mari [contact person].

Keywords: Hidden Markov Models, stochastic process

The system CarottAge is based on Hidden Markov Models of second order and provides a non supervised temporal clustering algorithm for data mining and a synthetic representation of temporal and spatial data. CarottAge is currently used by INRA researchers interested in mining the changes in territories related to the loss of biodiversity (projects ANR BiodivAgrim and ACI Ecoger) and/or water contamination. CarottAge is also used for mining hydromorphological data. Actually a comparison was performed with three other algorithms classically used for the delineation of river continuum and CarottAge proved to give very interesting results for that purpose [102].

CarottAge is freely available under GPL license (see <http://www.loria.fr/~jfmari/App/>).

5.2.2. *The ARPENAge system*

Participants: Florence Le Ber, Jean-François Mari [contact person].

Keywords: Hidden Markov Models, stochastic process

ARPEntAge¹ (for *Analyse de Régularités dans les Paysages: Environnement, Territoires, Agronomie* is a software based on stochastic models (HMM2 and Markov Field) for analyzing spatio-temporal data-bases [107]. ARPEntAge is built on top of the CarottAge system to fully take into account the spatial dimension of input sequences. It takes as input an array of discrete data in which the columns contain the annual land-uses and the rows are regularly spaced locations of the studied landscape. It performs a Time-Space clustering of a landscape based on its time dynamic Land Uses (LUS). Displaying tools and the generation of Time-dominant shape files have also been defined.

ARPEntAge is freely available (GPL license) and is currently used by INRA researchers interested in mining the changes in territories related to the loss of biodiversity (projects ANR BiodivAgrim and ACI Ecoger) and/or water contamination. In these practical applications, CarottAge and ARPEntAge aim at building a partition –called the hidden partition– in which the inherent noise of the data is withdrawn as much as possible. The estimation of the model parameters is performed by training algorithms based on the Expectation Maximization and Mean Field theories. The ARPEntAge system takes into account: (i) the various shapes of the territories that are not represented by square matrices of pixels, (ii) the use of pixels of different size with composite attributes representing the agricultural pieces and their attributes, (iii) the irregular neighborhood relation between those pixels, (iv) the use of shape files to facilitate the interaction with GIS (geographical information system).

ARPEntAge and CarottAge were used for mining decision rules in a territory showing environmental issues. They provide a way of visualizing the impact of farmers decision rules in the landscape and revealing new extra hidden decision rules [116].

5.3. KDD in Systems Biology

5.3.1. IntelliGO online

The IntelliGO measure computes semantic similarity between terms from a structured vocabulary (Gene Ontology: GO) and uses these values for computing functional similarity between genes annotated by sets of GO terms [83]. The IntelliGO measure is available on line (<http://plateforme-mbi.loria.fr/intelligo/>) to be used evaluation purposes. It is possible to compute the functional similarity between two genes, the intra-set similarity value in a given set of genes, and the inter-set similarity value for two given sets of genes.

5.3.2. WAFObI : KNIME nodes for relational mining of biological data

KNIME (for “Konstanz Information Miner”) is an open-source visual programming environment for data integration, processing, and analysis. KNIME includes a rich library of data manipulation tools (import, export) and several mining algorithms which operate on a single data matrix (decision trees, clustering, frequent itemsets, association rules...). The KNIME platform aims at facilitating the data mining experiment settings as many tests are required for tuning the mining algorithms. The evaluation of the mining results is also an important issue and its configuration is made easier.

Various KNIME nodes were developed for supporting relational data mining using the ALEPH program (<http://www.comlab.ox.ac.uk/oucl/research/areas/machlearn/Aleph/aleph.pl>). These nodes include a data preparation node for defining a set of first-order predicates from a set of relation schemes and then a set of facts from the corresponding data tables (learning set). A specific node allows to configure and run the ALEPH program to build a set of rules. Subsequent nodes allow to test the first-order rules on a test set and to perform configurable cross validations.

5.3.3. MOdel-driven Data Integration for Mining (MODIM)

Participants: Marie-Dominique Devignes [contact person], Malika Smaïl-Tabbone.

¹<http://www.loria.fr/~jfmari/App/>

The MODIM software (MOdel-driven Data Integration for Mining) is a user-friendly data integration tool which can be summarized along three functions: (i) building a data model taking into account mining requirements and existing resources; (ii) specifying a workflow for collecting data, leading to the specification of wrappers for populating a target database; (iii) defining views on the data model for identified mining scenarios. A version of the software was declared through Inria APP procedure in December, 2010.

Although MODIM is domain independent, it was used so far for biological data integration in various internal research studies. MODIM was also used for organizing data about non ribosomal peptide syntheses. The sources can be downloaded at <https://gforge.inria.fr/projects/modim/>.

5.4. Knowledge-Based Systems and Semantic Web Systems

5.4.1. *The Kasimir System for Decision Knowledge Management*

Participants: Nicolas Jay, Jean Lieber [contact person], Amedeo Napoli, Thomas Meilender.

Keywords: classification-based reasoning, case-based reasoning, decision knowledge management, knowledge edition, knowledge base maintenance, semantic portal

The objective of the Kasimir system is decision support and knowledge management for the treatment of cancer. A number of modules have been developed within the Kasimir system for editing treatment protocols, visualization, and maintenance. Kasimir is developed within a semantic portal, based on OWL. KatexOWL (Kasimir Toolkit for Exploiting OWL Ontologies, <http://katexowl.loria.fr>) is developed in a generic way and is applied to Kasimir. In particular, the user interface EdHibou of KatexOWL is used for querying the protocols represented within the Kasimir system (see [17] where an extension of Kasimir for multi-viewpoint case-based reasoning is presented).

Cabamaka (case base mining for adaptation knowledge acquisition) is a module of the Kasimir system. This system performs case base mining for adaptation knowledge acquisition and provides information units to be used for building adaptation rules. Actually, the mining process in Cabamaka is based on a frequent close itemset extraction module from the Coron platform (see §5.1.1).

The Oncologik system [12] is a collaborative editing tool aiming at facilitating the management of medical guidelines (<http://www.oncologik.fr/>). Based on a semantic wiki, it allows the acquisition of formalized decision knowledge. Oncologik also includes a graphical decision tree editor called KcatoS.

5.4.2. *Taaable: a system for retrieving and creating new cooking recipes by adaptation*

Participants: Valmi Dufour-Lussier, Emmanuelle Gaillard, Laura Infante Blanco, Florence Le Ber, Jean Lieber, Amedeo Napoli, Emmanuel Nauer [contact person].

Keywords: knowledge acquisition, ontology engineering, semantic annotation, case-based reasoning, hierarchical classification, text mining

Taaable [69] is a system whose objectives are to retrieve textual cooking recipes and to adapt these retrieved recipes whenever needed. Suppose that someone is looking for a “leek pie” but has only an “onion pie” recipe: how can the onion pie recipe be adapted?

The Taaable system combines principles, methods, and technologies such as case-based reasoning (CBR), ontology engineering, text mining, text annotation, knowledge representation, and hierarchical classification. Ontologies for representing knowledge about the cooking domain, and a terminological base for binding texts and ontology concepts, were built from textual web resources. These resources are used by an annotation process for building a formal representation of textual recipes. A CBR engine considers each recipe as a case, and uses domain knowledge for reasoning, especially for adapting an existing recipe w.r.t. constraints provided by the user, holding on ingredients and dish types.

The Taaable system is available on line since 2008 at <http://taaable.fr>. A new version of Taaable was implemented using Tuurbine, a generic ontology-guided CBR engine based on semantic web technologies (see Section 5.4.3). BeGood (see Section 5.4.4), a generic system for managing non-regression tests on knowledge bases, is also plugged for acquiring test sets. When the Taaable system returns answers to a query, the user may evaluate the relevance of the answers. Currently, user feedback is collected using BeGood and will be used in the future to run tests when the knowledge exploited by the CBR system evolves. The objective is to ensure that the knowledge base evolution does not affect the quality of answers given by the CBR system.

5.4.3. *Tuurbine: a generic ontology guided case-based inference engine*

Participants: Laura Infante Blanco, Jean Lieber, Emmanuel Nauer [contact person].

Keywords: case-based reasoning, inference engine, knowledge representation, ontology engineering, semantic web

The experience acquired since 5 years with the Taaable system conducted to the creation of a generic case-based reasoning system, whose reasoning procedure is based on a domain ontology. This new system, called Tuurbine (<http://tuurbine.loria.fr/>), takes into account the retrieval step, the case base organization, but also an adaptation procedure which is not addressed by other generic case-based reasoning tools. Moreover, Tuurbine is built over semantic web standards allowing to be connected to the web of data. The domain knowledge is represented in an RDF store, which can be interfaced with a semantic wiki, for collaborative edition and management of the knowledge involved in the reasoning system (cases, ontology, adaptation rules). The development of Tuurbine was supported by an Inria ADT funding until October 2013.

5.4.4. *BeGood: a generic system for managing non-regression tests on knowledge-bases*

Participants: Laura Infante Blanco, Emmanuel Nauer [contact person].

Keywords: tests, non-regression, knowledge evolution

BeGood [67] is a system allowing to define test plans, independent of any application domain, and usable for testing any system answering queries by providing results in the form of sets of strings. BeGood provides all the features usually found in test systems, such as tests, associated queries, assertions, and expected result sets, test plans (sets of tests) and test reports. The system is able to evaluate the impact of a system modification by running again test plans and by evaluating the assertions which define whether a test fails or succeeds. The main components of BeGood are (1) the “test database” that stores every test artifacts, (2) the “remote query evaluator” which evaluates test queries, (3) the “assertion engine” which evaluates assertions over the expected and effective query result sets. and finally (4) the “REST API” which offers the test functionalities as web services.

BeGood is available under a AGPL license on github². BeGood is used to manage the non-regression of the Taaable system (see Section 5.4.2) when the knowledge base used by the CBR system is modified.

5.4.5. *Revisor: a library of revision operators and revision-based adaptation operators*

Participants: Valmi Dufour-Lussier, Alice Hermann, Florence Le Ber, Jean Lieber [contact person], Emmanuel Nauer, Gabin Personeni.

Keywords: belief revision, adaptation, revision-based adaptation, case-based reasoning, inference engines, knowledge representation

Revisor is a library of inference engines dedicated to belief revision and to revision-based adaptation for case-based reasoning [60]. It is open source, under a GPL license and available on the web (<http://revisor.loria.fr/>). It gathers several engines developed during the previous years, for various knowledge representation formalisms (propositional logic—with or without the use of adaptation knowledge [65]—conjunction of linear constraints, and qualitative algebras [3]). Some of these engines are already used in the Taaable system. Current developments on Revisor aim at defining new engines in other formalisms.

²<https://github.com/kolflow/begood>

PAREO Project-Team

5. Software and Platforms

5.1. ATerm

Participant: Pierre-Etienne Moreau [correspondant].

ATerm (short for Annotated Term) is an abstract data type designed for the exchange of tree-like data structures between distributed applications.

The ATerm library forms a comprehensive procedural interface which enables creation and manipulation of ATerms in C and Java. The ATerm implementation is based on maximal subterm sharing and automatic garbage collection.

We are involved (with the CWI) in the implementation of the Java version, as well as in the garbage collector of the C version. The Java version of the ATerm library is used in particular by *Tom*.

The ATerm library is documented, maintained, and available at the following address: <http://www.meta-environment.org/Meta-Environment/ATerms>.

5.2. Tom

Participants: Jean-Christophe Bach, Christophe Calvès, Horatiu Cirstea, Pierre-Etienne Moreau [correspondant].

Since 2002, we have developed a new system called *Tom* [31], presented in [17], [18]. This system consists of a pattern matching compiler which is particularly well-suited for programming various transformations on trees/terms and XML documents. Its design follows our experiments on the efficient compilation of rule-based systems [29]. The main originality of this system is to be language and data-structure independent. This means that the *Tom* technology can be used in a C, C++ or Java environment. The tool can be seen as a Yacc-like compiler translating patterns into executable pattern matching automata. Similarly to Yacc, when a match is found, the corresponding semantic action (a sequence of instructions written in the chosen underlying language) is triggered and executed. *Tom* supports sophisticated matching theories such as associative matching with neutral element (also known as list-matching). This kind of matching theory is particularly well-suited to perform list or XML based transformations for example.

In addition to the notion of *rule*, *Tom* offers a sophisticated way of controlling their application: a strategy language. Based on a clear semantics, this language allows to define classical traversal strategies such as *innermost*, *outermost*, *etc.*. Moreover, *Tom* provides an extension of pattern matching, called *anti-pattern matching*. This corresponds to a natural way to specify *complements* (*i.e.* what should not be there to fire a rule). *Tom* also supports the definition of cyclic graph data-structures, as well as matching algorithms and rewriting rules for term-graphs.

Tom is documented, maintained, and available at <http://tom.loria.fr> as well as at <http://gforge.inria.fr/projects/tom>.

PAROLE Project-Team

5. Software and Platforms

5.1. WinSnoori

WinSnoori is a speech analysis software that we have been developing for 15 years. It is intended to facilitate the work of the scientist in automatic speech recognition, phonetics or speech signal processing. Basic functions of WinSnoori enable several types of spectrograms to be calculated and the fine edition of speech signals (cut, paste, and a number of filters) as the spectrogram allows the acoustical consequences of all the modifications to be evaluated. Beside this set of basic functions, there are various functionalities to annotate phonetically or orthographically speech files, to extract fundamental frequency, to pilot the Klatt synthesizer and to utilize PSOLA resynthesis.

The current version of WinSnoori is available on <http://www.winsnoori.fr>.

5.2. JSnoori

JSnoori is written in Java and uses signal processing algorithms developed within WinSnoori software with the double objective of being a platform independent signal visualization and manipulation tool, and also for designing exercises for learning the prosody of a foreign language. JSnoori thus focused the calculation of F0, the forced alignment of non native English uttered by French speakers and the correction of prosody parameters (F0, rhythm and energy). Since phonetic segmentations and annotations play a central role in the derivation of diagnosis concerning the realization of prosody by learners, several tools have been incorporated to segment and annotate speech. In particular, a complete phonetic keyboard is available, several kinds of annotation can be used (phonemes, syllables and words) and forced alignment can exploit variants to cope with non native accents. In addition, JSnoori offers real time F0 calculation which can be useful from a pedagogical point of view.

5.3. Xarticulators

Xarticulators software is intended to delineate contours of speech articulators in X-ray images, to construct articulatory models and to synthesize speech from X-ray films. This software provide tools to track contours automatically, semi-automatically or by hand, to make the visibility of contours easier, to add anatomical landmarks to speech articulators and to synchronize images together with the sound.

It also enables the construction of adaptable linear articulatory models from the X-ray images.

This year we particularly worked on the possibility of synthesizing speech from X-ray images. We thus substantially improved algorithms used to compute the centerline of the vocal tract in order to segment the vocal tract into elementary tubes approximating the propagation of a one-dimensional wave. We also developed time patterns used to synthesize sequences of voiceless consonants and vowels (VCV). In addition we also added the possibility of processing digitized manual delineation results made on sheet of papers in the seventies.

5.4. SUBWEB

We published in 2007 a method which allows to align sub-titles comparable corpora [94]. In 2009, we proposed an alignment web tool based on the developed algorithm. It allows to: upload a source and a target files, obtain an alignment at a sub-title level with a verbose option, and a graphical representation of the course of the algorithm. This work has been supported by CPER/TALC/SUBWEB ².

²<http://wikitalc.loria.fr/dokuwiki/doku.php?id=operations:subweb>

5.5. ANTS

The aim of the Automatic News Transcription System (ANTS) is to transcribe radio or TV shows. ANTS is composed of several stages. The first processing steps aim at splitting the audio stream into homogeneous segments of a manageable size and at identifying the segment characteristics in order to allow the use of specific algorithms or models according to the nature of the segment. This includes broad-band/narrow-band speech segmentation, speech/music classification, speaker segmentation and clustering, detection of silences/breathing segments and generally speaker gender classification.

Each segment is then decoded using a large vocabulary continuous speech recognition engine, either the Julius engine or the Sphinx engine. The Julius engine operates in two passes: in the first pass, a frame-synchronous beam search algorithm is applied on a tree-structured lexicon assigned with bigram language model probabilities. The output of this pass is a word-lattice. In the second pass, a stack decoding algorithm using a trigram language model gives the N-best recognition sentences. The Sphinx engine processes the speech input segment in a single forward pass using a trigram language model.

Further processing passes are usually run in order to apply unsupervised adaptation processes on the feature computations (VTLN: vocal tract length normalization) and/or on the model parameters (MLLR: maximum likelihood linear regression), or to use speaker adaptive training (SAT) based models. Moreover decoding results of both systems can be efficiently combined for improved decoding performance.

The latest version which relies on a perl script exploits the multiple CPUs available on a computer to reduce the processing time, and runs on both a stand alone linux machine and on the cluster.

5.6. CoALT

CoALT (Comparing Automatic Labeling Tools) compares two automatic labelers or two speech-text alignment tools, ranks them and displays statistics about their differences. The main feature of our software is that a user can define its own criteria for evaluating and comparing two speech- text alignment tools. With CoALT, a user can give more importance to either phoneme labels or phoneme boundaries because the CoALT elastic comparison algorithm takes into account time boundaries. Moreover, by providing a set of phonetic rules, a user can define the allowed discrepancies between the automatic labeling result and the hand-labeling one.

5.7. TTS SoJA

TTS SoJA (Speech synthesis platform in Java) is a software for text-to-speech synthesis. The aim of this software is to provide a toolkit to test some steps of natural language processing and to provide a whole system of TTS based on non uniform unit selection algorithm. The software performs all steps from the text to the speech signal. Moreover, it provides a set of tools to elaborate a corpus for a TTS system (transcription alignment, ...). Currently, the corpus contains 1800 sentences (about 3 hours of speech) recorded by a female speaker.

Most of the modules are developed in Java. Some modules are in C. The platform is designed to make easy the addition of new modules. The software runs under Windows and Linux (tested on Mandriva, Ubuntu). It can be launch with a graphical user interface or directly integrated in a Java code or by following the client-server paradigm.

The software license should easily allow associations of impaired people to use the software. A demo web site has been built: <http://soja-tts.loria.fr>

5.8. JCorpusRecorder

JCorpusRecorder is a software for the recording of audio corpora. It provides a easy tool to record with a microphone. The audio input gain is controlled during the recording. From a list of sentences, the output is a set of wav files automatically renamed with textual information given in input (nationality, speaker language, gender...). An easy syntactic tagging allows displaying a textual/visual/audio context of the sentence to pronounce. This software is suitable for recording sentences with information to guide the speaker. The sentences can be presented randomly.

The software is now developed in Java (since 2013). It is currently used for the recording of sentences in several projects (including IFCASL).

5.9. VisArtico

VisArtico is intended to visualize articulatory data acquired using an articulograph [97]. It is intended for researchers that need to visualize data acquired from the articulograph with no excessive processing. It is well adapted to the data acquired using the AG500 and AG501 (developed by Carstens Medizinelektronik GmbH), and the articulograph NDI Wave, developed by Northern Digital Inc.

The software allows displaying the positions of the sensors that are simultaneously animated with the speech signal. It is possible to display the tongue contour and the lips contour. The software helps to find the midsagittal plane of the speaker and find the palate contour. In addition, VisArtico allows labeling phonetically the articulatory data.

All this information is very useful to researchers working in the field of speech production, as phoneticians for instance. VisArtico provides several possible views: (1) temporal view, (2) 3D spatial view and (3) 2D midsagittal view. In the temporal view, it is possible to display different articulatory trajectories in addition to the acoustic signal and eventually labels. The midsagittal view can display the tongue contour, the jaw, the lips and the palate.

VisArtico provides several tools to help to improve the quality of interpreting the data. It is a cross-platform software as it is developed in JAVA and does not need any additional external library or framework. It was tested and worked on Windows, Mac OS, and Linux. It should work on any system having JAVA installed. VisArtico is freely distributed via a dedicated website <http://visartico.loria.fr>.

5.10. FASST

The Flexible Audio Source Separation Toolbox (FASST) is a toolbox for audio source separation (<http://bass-db.gforge.inria.fr/fasst/>). It aims to become the reference software for research and applications of audio source separation. Its unique feature is the possibility for users to specify easily a suitable algorithm for their use case thanks to the general modeling and estimation framework. Besides, it forms the basis of most of our current research in audio source separation, some of which may be incorporated into future versions of the software.

SCORE Team

4. Software and Platforms

4.1. Rivage

Participants: Claudia-Lavinia Ignat, Stéphane Martin [contact].

Rivage (Real-time Vector graphic Group Editor) is a real-time collaborative graphical editor. Several users can edit at the same time and in real-time a graphical document, user changes being immediately seen by the other users. The editor relies on a peer-to-peer architecture where users can join and leave the group at any time. Each user has a copy of the shared document and user changes on the document copies are merged in real-time by using a CRDT (Commutative Replicated Data Type) algorithm. The code is available at <https://github.com/stephanemartin/rivage>.

4.2. Replication Benchmark

Participants: Pascal Urso [contact], Mehdi Ahmed-Nacer, Stéphane Martin, Gérald Oster.

The Replication Benchmark is a performance evaluation framework for optimistic replication mechanisms used in collaborative applications. It contains a library of implementation of several CRDT (Commutative Replicated Data Type) and OT (Operational Transformation) algorithms for different data types: text, set, trees. The framework is able to evaluate the performance of comparable algorithms on different corpus of events traces. These events traces can be produced randomly according to different parameters, can be extracted from real real-time editing session that have been recorded, or can be automatically extracted from distributed version control repositories such as the one produced with Git. Performances of the algorithms are measured in term of execution time, memory footprint and merge result quality (compared to manual merge history stored in git repositories). The source code of this evaluation framework is available at <https://github.com/score-team/replication-benchmark>.

4.3. BeGood

Participant: G r me Canals.

BeGood is a generic system for managing non-regression tests on knowledge-bases. BeGood allows to define test plans in order to monitor the evolution of knowledge-bases. Any system answering queries by providing results in the form of set of strings can be tested with BeGood. BeGood has been developed following a REST architecture and is independent of any application domain. BeGood is a part of the Kolflow infrastructure and is available at <https://github.com/kolflow>.

SÉMAGRAMME Project-Team

5. Software and Platforms

5.1. Leopard

Participants: Bruno Guillaume [correspondent], Guy Perrier, Tatiana Ekeinhor.

5.1.1. Software description

Leopard is a parser for natural languages which is based on the formalism of Interaction Grammars [40]. It uses a parsing principle, called “electrostatic parsing” which consists in neutralizing opposite polarities. A positive polarity corresponds to an available linguistic feature and a negative one to an expected feature.

Parsing a sentence with an Interaction Grammar consists in first selecting a lexical entry for each of its words. A lexical entry is an underspecified syntactic tree, a tree description in other words. Then, all selected tree descriptions are combined by partial superposition guided by the aim of neutralizing polarities: two opposite polarities are neutralized by merging their support nodes. Parsing succeeds if the process ends with a minimal and neutral tree. As IGs are based on polarities and under-specified trees, Leopard uses some specific and non-trivial data-structures and algorithms.

The electrostatic principle has been intensively considered in Leopard. The theoretical problem of parsing IGs is NP-complete; the nondeterminism usually associated to NP-completeness is present at two levels: when a description for each word is selected from the lexicon, and when a choice of which nodes to merge is made. Polarities have shown their efficiency in pruning the search tree:

- In the first step (tagging the words of the sentence with tree descriptions), we forget the structure of descriptions, and only keep the bag of their features. In this case, parsing inside the formalism is greatly simplified because composition rules reduce to the neutralization of a negative feature-value pair $f \leftarrow v$ by a dual positive feature-value pair $f \rightarrow v$. As a consequence, parsing reduces to a counting of positive and negative polarities present in the selected tagging for every pair (f, v) : every positive occurrence counts for +1 and every negative occurrence for -1, the sum must be 0.
- Again in the tagging step, original methods were developed to filter out bad taggings. Each unsaturated polarity p in the grammar induces constraints on the set of contexts in which it can be used: the unsaturated polarity p must find a *companion* (i.e. a tree description able to saturated it); and the set of companions for the polarity p can be computed statically from the grammar. Each lexical selection which contains an unsaturated polarity without one of its companions can be safely removed.
- In the next step (node-merging phase), polarities are used to cut off parsing branches when their trees contain too many non neutral polarities.

5.1.2. Current state of the implementation

Leopard is presented and documented at <http://leopard.loria.fr>; an online demonstration page can be found at <http://leopard.loria.fr/demo>.

It is open-source (under the CECILL License <http://www.cecill.info>) and it is developed using the InriaGforge platform (<http://gforge.inria.fr/projects/semagramme/>)

The main features of current software are:

- automatic parsing of a sentence or a set of sentences,
- dependency and parse-tree representation of sentences,
- interactive parsing (the user chooses the couple of nodes to merge),
- visualization of grammars produced by XMG-2 or of sets of description trees associated to some word in the linguistic resources.

One of the difficulties with symbolic parsing is that several solutions can be produced for a single sentence and we want to be able to rank them. Tatiana Ekeinhor, during her second year Master Internship (from February to June 2013), implemented a ranker based on statistical techniques. Using the Sequoia TreeBank as a training corpus, she obtained an improvement of the system compared to the handcrafted rules.

5.2. ACG Development Toolkit

Participants: Sylvain Pogodalla [correspondent], Philippe de Groot.

In order to support the theoretical work on ACG, we have been developing a support system. The objectives of such a system are twofold:

1. To make possible to implement and experiment grammars the modeling of linguistic phenomena.
2. To make possible to implement and experiment results related to the ACG formalisms. Such results can concern parsing algorithms, type extensions, language extensions, etc.

The ACG Development toolkit development effort is part of the POLYMNIE project (see Section 7.2.1.1). It will support the experimentation and evaluation parts of the project.

The current version of the ACG development toolkit prototype¹ issues from a first release published in October 2008. Further releases have been published before the ESSLLI 2009 course on ACG. It focuses on providing facilities to develop grammars. To this end, the type system currently implemented is the linear core system plus the (non-linear) intuitionistic implication, and a special attention has been paid to type error management. As a major limitation, this version only considers transformation from abstract terms to object terms, and not the other way around.

The prototype now enables the transformation from the object terms to the abstract terms. The parsing algorithm follows [43]'s method which is being implemented for second-order ACGs. It is based on a translation of ACG grammars into Datalog programs and is well-suited to fine-grained optimization.

However, since we're interested not only by recognizability (hence whether some fact is provable) but also by the parsing structure (hence the proof), the Datalog solver has been adapted to produce not only yes/no answer to queries, but also all the proofs of the answers to the queries. The next steps concern optimization and efficiency. Note however that in the general case, the decidability of translating an object term to an abstract one is still an open problem.

5.3. Grew

Participants: Bruno Guillaume [correspondent], Guy Perrier.

Graph rewriting, Interface syntaxe-sémantique

Grew is a Graph Rewriting tools dedicated to applications in NLP. It is freely-available (from the page <http://grew.loria.fr>) and it is developed using the InriaGforge platform (<http://gforge.inria.fr/projects/semagramme/>)

We list below some of the major specificities of the GREW software.

- Graph structures can use a build-in notion of feature structures.
- The left-hand side of a rule is described by a graph called a pattern; injective graph morphisms are used in the pattern matching algorithm.
- Negative pattern can be used for a finer control on the left-hand side of rules.
- The right-hand side of rules is described by a sequence of atomic commands that describe how the graph should be modified during the rule application.
- Rules can be parametrized by lexical information.
- Filters can be used at the output of each module to control the structure produced are well-formed.

¹ Available at <http://acg.gforge.inria.fr> with a CeCILL license.

- Subset of rules are grouped in modules; the full rewriting process being a sequence of module applications.
- The Grew software has support both for confluent and non-confluent modules; when a non-confluent modules is used, all normal forms are returned and then ambiguity is handled in a natural way.
- Grew can be used on Corpus mode with statistics about rules usage or with an a Graphical User Interface which can show all intermediate graphs used during the rewriting process (useful either to debug rewriting system or for demonstrations).

The Grew software was used for several kind of applications manipulating syntactic and/or semantic graph representations. It was used to build DMRS semantic representation from syntactic dependency trees in the French TreeBank [51].

More recently, it was used in the project “Deep Syntax Annotation of the Sequoia French Treebank”. First, it was used as a pre-annotation tool and; second, it is used to detect ill-formed structures that don’t fit the annotation guide requirement.

5.4. Other developments

Participants: Bruno Guillaume [correspondent], Maxime Amblard [correspondent].

Concordancer, Dependencies, Graphical tools Other peripheral developments of the team are available either as web service or as downloadable code:

- A concordancer named CONDOR which is usable online: <http://condor.loria.fr>. With Condor, it is possible to search for all inflexions (given by a lexicon) of some lemma; it is possible to search for a couple of lemmas to find collocations.
- A program (named DEP2PICT) to build graphical representations (PNG, SVG or PDF) of dependency structures. It is presented in <http://dep2pict.loria.fr>; it is usable online <http://dep2pict.loria.fr/demo>.
- a management chain of the transcriptions of interviews for the SLAMproject. including the production of a full anonymized randomized version of the resources.
- A program which use Distagger and propose different analyze of the repartition of disfluencies.

SHACRA Project-Team

5. Software and Platforms

5.1. SOFA

SOFA <http://www.sofa-framework.org> is an open-source software framework targeted at interactive computational (medical) simulation. The idea of SOFA was initiated by members of the SHACRA team, and strongly supported by Inria through a development program that we lead. SOFA facilitates collaborations between specialists from various domains, by decomposing complex simulators into components designed independently. Each component encapsulates one of the key aspects of a simulation, such as the degrees of freedom, the forces and constraints, the differential equations, the linear solvers, the collision detection algorithms or the interaction devices. The simulated objects can be represented using several models, each of them optimized for a different task such as the computation of internal forces, collision detection, haptics or visual display. These models are synchronized during the simulation using a mapping mechanism. CPU and GPU implementations can be transparently combined to exploit the computational power of modern hardware architectures. Thanks to this flexible yet efficient architecture, SOFA can be used as a test-bed to compare models and algorithms, or as a basis for the development of complex, high-performance simulators. As proof of its success, SOFA has been downloaded nearly 150,000 times, and is used today by many research groups around the world, as well as a number of companies. The mailing list used to exchange with the community includes several hundreds of researchers, from about 50 different institutions. SOFA is at the heart of a number of research projects, including cardiac electro-physiology modeling, interventional radiology planning and guidance, planning for cryosurgery and deep brain stimulation, robotics, percutaneous procedures, laparoscopic surgery, non-rigid registration, etc. SOFA is the only software developed by our team, but practically speaking it is a collection of plugins (each one aimed at a specific application) organized around a common core that provides a large number of functionalities. As mentioned previously, SOFA is currently used by a number of companies (Siemens Corporate Research, Digital Trainers, Epona Medical, Moog, SenseGraphics, etc.) and also provides the key technology on which our newly created start-up (InSimo) is relying. We strongly believe that today SOFA has become a reference for academic research, and is increasingly gaining recognition for product prototyping and development. The best illustration of this worldwide positioning is the role of SOFA in the challenge set by the HelpMeSee foundation to win the contract for the development of a very ambitious and high-risk project on cataract surgery simulation.

We also gave a 4 hours workshop on SOFA at MMVR/NextMed conference in february 2013 in San Diego. This workshop was done in collaboration with the swedish company SenseGraphics. The topic was to demonstrate the setup of a dental surgery simulation in Sofa, and use SenseGraphics visual tools for the rendering. The attendees feedback was beyond our expectations, with an unexpected interest in new SOFA features like the SofaPython plugin. Still about SOFA, like last year we gave in october a 3 days training session in Montpellier for about twenty SOFA beginners (mostly engineers). These are new engineers of the three teams involved in SOFA development, and employees of companies using SOFA in their business. Last, a "SOFA Day" in november in prelude of the Vriphys conference gave us a unique opportunity to meet SOFA users from various research institutes or companies, and exchange about the future improvements and development of the engine. We use these occasions to share and discuss with SOFA users, to refine the roadmap and stay tuned with our audience.

TOSCA Project-Team

5. Software and Platforms

5.1. SDM

Participant: Mireille Bossy [correspondant].

The computation of the wind at small scale and the estimation of its uncertainties is of particular importance for applications such as wind energy resource estimation. To this aim, starting in 2005, we have developed a new method based on the combination of an existing Numerical Weather Prediction model providing a coarse prediction, and a Lagrangian Stochastic Model for turbulent flows. This Stochastic Downscaling Method (SDM) requires a specific modelling of the turbulence closure, and involves various simulation techniques whose combination is totally original (such as Poisson solvers, optimal transportation mass algorithm, original Euler scheme for confined Langevin stochastic processes, and stochastic particle methods).

In 2013, the SDM code became the kernel of the wind farm modelling of the Fundacion Inria Chile. In France, its development is pursuing through the collaborative Modéol project on the evaluation of wind potential.

This is a joint work with Antoine Rousseau from the project-team MOISE.

- Version: 2.0

5.2. CarbonQuant

Participants: Mireille Bossy [correspondant], Selim Karia.

CarbonQuant is a simulator project of CO₂ allowances prices on a EU-ETS type market, by an indifference price approach.

It aims to demonstrate the high potentiality of stochastic control solvers, to quantify sensibilities of a carbon market with respect to its design.

Starting in September 2011, CarbonQuant is an ADT ¹ Inria.

See also the web page <http://carbonvalue.gforge.inria.fr>, from where CarbonQuant can be now downloaded for various architectures.

- Version: 2.0

¹Technology Development Action

TRIO Team

5. Software and Platforms

5.1. ANR Open-PEOPLE platform

Participants: Anis Koubaa, Olivier Zendra.

The aim of Open-PEOPLE is to provide a platform for estimating and optimizing the power and energy consumption of systems. The Open-PEOPLE project formally started in April 2009.

In 2013, work in TRIO on this platform was minimal, because of the lack of development resources.

We performed software updates on the servers, ensuring the continuity of access (which is especially important since other partners of the former ANR Open-PEOPLE project still use this platform and actively develop on it). We fixed a few bugs on the platform.

In late 2013, we started working, in the context of Anis Koubaa's student project, on adding a new functionality to help develop energy combustion models, namely the automatic extraction of mathematical laws, derived from the measurements (cloud of points) coming from the experimental hardware platform.

5.2. VITRAIL

Participants: Pierre Caserta, Romarik Jodin, Olivier Zendra.

The aim of the VITRAIL operation is to provide tools for the advanced and immersive visualization of programs. Some of this work has been done with the University of Montréal, the University of Montpellier and to a lesser extent the Pareo team of Inria Nancy Grand Est.

Last years, in VITRAIL, we had developed software to instrument and trace Java programs at the bytecode level. We then had developed an analysis tool able to exploit these traces to compute relevant software metrics.

In 2013, we were able to restart developments on the VITRAIL platform.

We first explored a Linux port, toward which we progressed but were stopped by the fact we relied on Ogre3D, a library that calls some Windows specific APIs. We have identified those, and we believe that an OpenGL port would be a sensible path to OS portability.

We also ported our VITRAIL Vizualizer software to a new type of display, namely 3 interactive whiteboards placed so as to form a 3-side box.

Finally, we also successfully implemented a first prototype of interaction through a head-tracking system relying on two cameras. First experiments gave promising results.

VEGAS Project-Team

4. Software and Platforms

4.1. QI: Quadrics Intersection

QI stands for “Quadrics Intersection”. QI is the first exact, robust, efficient and usable implementation of an algorithm for parameterizing the intersection of two arbitrary quadrics, given in implicit form, with integer coefficients. This implementation is based on the parameterization method described in [7], [10] and represents the first complete and robust solution to what is perhaps the most basic problem of solid modeling by implicit curved surfaces.

QI is written in C++ and builds upon the LiDIA computational number theory library [29] bundled with the GMP multi-precision integer arithmetic [28]. QI can routinely compute parameterizations of quadrics having coefficients with up to 50 digits in less than 100 milliseconds on an average PC; see [10] for detailed benchmarks.

Our implementation consists of roughly 18,000 lines of source code. QI has being registered at the Agence pour la Protection des Programmes (APP). It is distributed under the free for non-commercial use Inria license and will be distributed under the QPL license in the next release. The implementation can also be queried via a web interface [30].

Since its official first release in June 2004, QI has been downloaded six times a month on average and it has been included in the geometric library EXACUS developed at the Max-Planck-Institut für Informatik (Saarbrücken, Germany). QI is also used in a broad range of applications; for instance, it is used in photochemistry for studying the interactions between potential energy surfaces, in computer vision for computing the image of conics seen by a catadioptric camera with a paraboloidal mirror, and in mathematics for computing flows of hypersurfaces of revolution based on constant-volume average curvature.

4.2. Isotop: Topology and Geometry of Planar Algebraic Curves

ISOTOP is a Maple software for computing the topology of an algebraic plane curve, that is, for computing an arrangement of polylines isotopic to the input curve. This problem is a necessary key step for computing arrangements of algebraic curves and has also applications for curve plotting. This software has been developed since 2007 in collaboration with F. Rouillier from Inria Paris - Rocquencourt. It is based on the method described in [4] which incorporates several improvements over previous methods. In particular, our approach does not require generic position.

Isotop is registered at the APP (June 15th 2011) with reference IDDN.FR.001.240007.000.S.P.2011.000.10000. This version is competitive with other implementations (such as ALCIX and INSULATE developed at MPII Saarbrücken, Germany and TOP developed at Santander Univ., Spain). It performs similarly for small-degree curves and performs significantly better for higher degrees, in particular when the curves are not in generic position.

We are currently working on an improved version integrating our new bivariate polynomial solver.

4.3. CGAL: Computational Geometry Algorithms Library

Born as a European project, CGAL (<http://www.cgal.org>) has become the standard library for computational geometry. It offers easy access to efficient and reliable geometric algorithms in the form of a C++ library. CGAL is used in various areas needing geometric computation, such as: computer graphics, scientific visualization, computer aided design and modeling, geographic information systems, molecular biology, medical imaging, robotics and motion planning, mesh generation, numerical methods...

In computational geometry, many problems lead to standard, though difficult, algebraic questions such as computing the real roots of a system of equations, computing the sign of a polynomial at the roots of a system, or determining the dimension of a set of solutions. We want to make state-of-the-art algebraic software more accessible to the computational geometry community, in particular, through the computational geometric library CGAL. On this line, we contributed a model of the *Univariate Algebraic Kernel* concept for algebraic computations [32] (see Sections 8.2.2 and 8.4). This CGAL package improves, for instance, the efficiency of the computation of arrangements of polynomial functions in CGAL [34]. We are currently developing a model of the *Bivariate Algebraic Kernel* based on a new bivariate polynomial solver.

4.4. Fast_polynomial: fast polynomial evaluation software

The library *fast_polynomial*¹ provides fast evaluation and composition of polynomials over several types of data. It is interfaced for the computer algebra system *Sage* and its algorithms are documented². This software is meant to be a first step toward a certified numerical software to compute the topology of algebraic curves and surfaces. It can also be useful as is and is submitted for integration in the computer algebra system *Sage*.

This software is focused on *fast online computation*, *multivariate evaluation*, *modularity*, and *efficiency*.

Fast online computation. The library is optimized for the evaluation of a polynomial on several point arguments given one after the other. The main motivation is numerical path tracking of algebraic curves, where a given polynomial criterion must be evaluated several thousands of times on different values arising along the path.

Multivariate evaluation. The library provides specialized fast evaluation of multivariate polynomials with several schemes, specialized for different types such as *mpz* big ints, *boost* intervals with hardware precision, *mpfi* intervals with any given precision, etc.

Modularity. The evaluation scheme can be easily changed and adapted to the user needs. Moreover, the code is designed to easily extend the library with specialization over new C++ objects.

Efficiency. The library uses several tools and methods to provide high efficiency. First, the code uses templates, such that after the compilation of a polynomial for a specific type, the evaluation performance is equivalent to low-level evaluation. Locality is also taken into account: the memory footprint is minimized, such that an evaluation using the classical Hörner scheme will use $O(1)$ temporary objects and divide and conquer schemes will use $O(\log n)$ temporary objects, where n is the degree of the polynomial. Finally, divide and conquer schemes can be evaluated in parallel, using a number of threads provided by the user.

¹http://trac.sagemath.org/sage_trac/ticket/13358

²<http://arxiv.org/abs/1307.5655>

VERIDIS Project-Team

5. Software and Platforms

5.1. The veriT solver

Participants: David Déharbe, Pablo Dobal, Haniel Barbosa, Pascal Fontaine [correspondent].

The veriT solver is an SMT (Satisfiability Modulo Theories) solver developed in cooperation with David Déharbe from the Federal University of Rio Grande do Norte in Natal, Brazil. The solver can handle large quantifier-free formulas containing uninterpreted predicates and functions, and arithmetic over integers and reals. It features a very efficient decision procedure for difference logic, as well as a simplex-based reasoner for full linear arithmetic. It also has some support for user-defined theories, quantifiers, and lambda-expressions. This allows users to easily express properties about concepts involving sets, relations, etc. The prover can produce an explicit proof trace when it is used as a decision procedure for quantifier-free formulas with uninterpreted symbols and arithmetic. To support the development of the tool, a regression platform using Inria's grid infrastructure is used; it allows us to extensively test the solver on thousands of benchmarks in a few minutes. The veriT solver is available as open source under the BSD license at the [veriT Web site](#).

Efforts in 2013 have been focused on efficiency, and more specifically on arithmetic. A preliminary prototype integrating the solver [Redlog](#) for non-linear arithmetic has been stabilized. First results are encouraging; this prepares the ground for the starting ANR project SMARt (Satisfiability Modulo Arithmetic Theories), involving both sites of the VeriDis team (veriT being developed in Nancy and Redlog being designed in Saarbrücken), as well as Systerel as an industrial partner.

In late 2013, Haniel Barbosa joined the team as a PhD student. He will work on theoretical and practical aspects of handling quantifiers in SMT frameworks, which is currently an important challenge for SMT, and he will implement his techniques in veriT.

We target applications where validation of formulas is crucial, such as the validation of TLA^+ and B specifications, and work together with the developers of the respective verification platforms to make veriT even more useful in practice. The solver is available as a plugin for the Rodin platform for discharging proof obligations generated in Event-B [39]; on a large repository of industrial and academic cases, this SMT-based plugin decreased by 75% the number of proof obligations requiring human interactions, compared to the original B prover.

5.2. The TLA+ proof system

Participants: Bhargav Bhatt, Stephan Merz [correspondent], Hernán Vanzetto.

TLAPS, the TLA^+ proof system, is a platform for developing and mechanically verifying proofs about TLA^+ specifications. It is developed at the Joint MSR-Inria Centre. The TLA^+ proof language is hierarchical and explicit. TLAPS consists of a *proof manager* that interprets the proof language and generates a collection of proof obligations that are sent to *backend verifiers* that include theorem provers, proof assistants, SMT solvers, and decision procedures.

The current version 1.2.1 of TLAPS was released in September 2013, it is distributed under a BSD-like license at <http://tla.msr-inria.inria.fr/tlaps/content/Home.html>. The prover currently handles the non-temporal part of TLA^+ and can be used to prove safety, but not liveness properties. Its backends include a tableau prover for first-order logic, an encoding of TLA^+ in the proof assistant Isabelle, and a backend for interfacing with SMT solvers. The SMT backend, developed in Nancy, has been further improved in 2013 and is now considered by users as the most useful backend prover for system verification. During his internship in the summer of 2013, Bhargav Bhatt helped design and implement a standard library of TLA^+ theorems about functions, sequences, and finite sets that is now part of the TLAPS distribution. Development of support for temporal reasoning in TLAPS has started in late 2013.