



RESEARCH CENTER
Bordeaux - Sud-Ouest

FIELD

Activity Report 2015

Section New Results

Edition: 2016-03-21

1. CAGIRE Team	4
2. CARDAMOM Team	6
3. CARMEN Team	13
4. CQFD Project-Team	16
5. FLOWERS Project-Team	23
6. GEOSTAT Project-Team	38
7. HIEPACS Project-Team	47
8. LFANT Project-Team	57
9. MAGIQUE-3D Project-Team	59
10. MANAO Project-Team	65
11. MEMPHIS Team	71
12. MNEMOSYNE Project-Team	81
13. Monc Team	83
14. PHOENIX Project-Team	90
15. PLEIADE Team	92
16. POSET Team	94
17. POTIOC Project-Team	96
18. REALOPT Project-Team	106
19. SISTM Project-Team	110
20. STORM Team	111
21. TADAAM Team	115

CAGIRE Team

7. New Results

7.1. DNS of a jet in crossflow

One main achievement of this year is to have done our first DNS computations at third order with the Aerosol software. Two configurations of jet in cross flow have been computed: one with a hole direction aligned with the main flow direction (Fig. 3 -left), and another one with a 90-degree jet skidding (Fig. 3 -right). The first case has been validated by using analytical models of jet trajectory, and has also been compared with experiments made with our experimental bench MAVERIC. The comparison of experiments and DNS showed a good agreement.

The DNS database includes:

- The instantaneous flow at the vertices of the mesh.
- the instantaneous flow at some probes.
- The mean flow.
- The value of the Reynolds stress tensor in all the degrees of freedom.

7.2. Extension of the momentum interpolation method to low Mach Rieman problems

In a previous study [9], the momentum interpolation (MI) method was considered as a guideline to develop a Godunov-like flux scheme called AUSM-IT and able to preserve the acoustic energy at the discrete level for a low-order finite volume approach. This year, the MI method has been successfully extended to the case of low Mach flows featuring discontinuities [8]. The undesirable dispersive effect directly connected to the upwinding of the MI formulation of the face velocity has been corrected (up to second-order errors) by using a central interpolation of momentum in the face velocity definition.

7.3. Main features of highly underexpanded jets

Despite the numerous studies dealing with underexpanded jets, many aspects of their structure were not clearly described, particularly when one seeks for quantitative predictions. Since such flow configuration may be of interest in case of the accidental boring of an aeronautical combustion chamber, an exhaustive review of the main experimental papers dealing with underexpanded jets has been carried out [5]. This study aimed at clarifying the characteristics which were well known, from those where there is clearly a lack of confidence. Curiously enough, such a work has never been done and no exhaustive review was available on such a topic.

7.4. Formulation of a reference EB-RSM model

The Elliptic Blending Reynolds Stress Model (EB-RSM), originally proposed by Manceau & Hanjalic in 2002, has been subject to various modifications by several authors during the last decade, mainly for numerical robustness reasons. We have revisited all these modifications from the theoretical standpoint and investigated in detail their influence on the reproduction of the physical mechanisms at the origin of the influence of the wall on turbulence. Theoretical arguments and comparison with DNS results led to the selection of a recommended formulation for the EB-RSM model [7].

7.5. Development of a new enrichment method

A complex issue in multi-scale simulations is the necessity to *enrich* the solution at the interface between a region described at coarse grain (e.g., using RANS) and a region described at fine grain (e.g., using LES). In order to rapidly generate realistic fluctuations at the beginning of the LES region, we have proposed [4] a method of volumic forcing, the so-called ALF (Anisotropic Linear Forcing). In an overlap region, a time-dependent volume force is introduced into the filtered equations of motion in order to amplify the turbulent fluctuations in order that the LES field satisfy the statistics of the RANS solution, a method that proved simple, efficient and computationally cheap.

7.6. A new criterion to analyse hybrid RANS/LES approaches

Most of the available hybrid RANS/LES methods are completely empirical or based on a formalism which is not applicable in practical application, due to a mismatch between the statistical average and the spatial filtering in inhomogeneous flows. The lack of clear formalism leads to limitations in terms of modeling of the unresolved turbulent motion. We have established a criterion [6] to assess the equivalence between hybrid RANS/LES methods, called *H-equivalence*, that makes it possible to view different hybrid methods as models for the same system of equations: as a consequence, empirical hybrid methods, such as the detached-eddy simulation (DES), can be interpreted as a model for the subfilter stress involved in the *temporally filtered* Navier-Stokes equations, which is an answer to the issue raised above about the formalism underlying such methods.

CARDAMOM Team

7. New Results

7.1. High order discretizations on unstructured meshes

Participants: Héloïse Beaugendre [Corresponding member], Cécile Dobrzynski, Léo Nouveau, Mario Ricchiuto, Quentin Viville.

Our work on high order unstructured discretizations this year has pursued three main avenues:

- We have extended the team's previous work on the consistent residual based approximation of viscous flow equations to the framework of Immersed Boundary Methods (IBM). This is an increasingly popular approach in Computational Fluid Dynamics as it simplifies the mesh generation problem. In our work, we consider a technique based on the addition of a penalty term to the Navier-Stokes equations to account for the wall boundary conditions. To adapt the residual distribution method to the IBM, we developed a new formulation based on a Strang splitting approach in time. This approach, couples in a fully consistent manner an implicit asymptotically exact integration procedure of the penalization ODE, with the explicit residual distribution discretization for the Navier-Stokes equations, based on the method proposed in (Ricchiuto and Abgrall, *J.Comput.Phys* 229, 2010). The ODE integrator provides an operator which is exact up to orders η^2 , with η the penalty parameter assuming values of the order of 10^{-10} . To preserve the accuracy of the spatial discretization in the Navier-Stokes step, we have introduced, in vicinity of the penalised region, a modification of the solution gradient reconstruction required for the evaluation of the viscous fluxes. We have shown formally and numerically that the approach proposed is second order accurate for smooth solutions, and shown its potential when combined with unstructured mesh adaptation strategies w.r.t. the (implicitly described) solid walls. This work has been accepted on *Comp.Meth.Appl.Mech.Eng.* ;
- Another research axis consists in proposing a novel approach that allows to use p-adaptation with continuous finite elements. Under certain conditions, primarily the use of a residual distribution scheme, it is possible to avoid the continuity constraint imposed to the approximate solution, while still retaining the advantages of a method using continuous finite elements. The theoretical material, the complete numerical method and practical results show as a proof of concept that p-adaptation is possible with continuous finite elements. Its extension to penalized Navier-Stokes equations are under progress ;
- We have continued the study of the properties of residual based methods in the time dependent case. We have been able to further characterize one of the variants of the approach proposed in (Ricchiuto and Abgrall, *J.Comput.Phys* 229, 2010) in terms of preservation of the positivity of the density showing this property in practical applications involving the shallow water equations [130]. The impact of the simplified construction leading to these schemes has also been investigated. In particular, we have shown that despite the additional complexity associated to the inversion of the mass matrix, non-linear methods providing monotone solution and yet featuring linear mass matrices can be constructed [142]. These methods, have been shown to have some potential w.r.t. fully diagonal approaches as those used in (Ricchiuto and Abgrall, *J.Comput.Phys* 229, 2010), in terms of error as function of CPU time : the non-diagonal schemes showing error reductions up to one order of magnitude. Current work is devoted to the use of other multistage (defect correction type) and multistep (extrapolated methods) techniques, comparing them to space time approaches.

7.2. Modelling of free surface flows

Participants: Luca Arpaia, Stevan Bellec, Mathieu Colin, Sebastien de Brye, Andrea Filippini, Maria Kazolea, Mario Ricchiuto [Corresponding member].

We have introduced a new systematic method to obtain discrete numerical models for incompressible free-surface flows. The method consists in first discretizing the Euler equations with respect to the horizontal variables, keeping the vertical z variable and time continuous. We have focused so far on (continuous) Galerkin finite element discretizations in the horizontal. We then perform an asymptotic analysis on the resulting semi-discrete system. Our initial result, has led to a new discrete approximation, which we have shown to be consistent with the Boussinesq system known as Peregrine model. We have proven that the method obtained by means of this discrete asymptotic method, has phase and linear shoaling errors far lower than those obtained by discretizing the continuous model directly by means of the Galerkin method. Extensions to other weakly non-linear models have been obtained, and the study of fully nonlinear variants is under way.

We have also investigated the relations between some of the most common weakly nonlinear Boussinesq models. It is known since many years that, for given phase linear shoaling relations, two families of models exist depending on whether the dispersive terms are evaluated using derivatives of the speed, or of the flux (depth times speed). We have shown both analytically and numerically, that, independently on the phase and linear shoaling relations, these two families provide (qualitatively *and* quantitatively) only two distinct behaviours when approaching the nonlinear regime. Models based on velocity derivatives, provide taller more asymmetric waves, all models of the same family produce stunningly similar results, even when the linear relations differ considerably.

To extend our initial work on unstructured solvers for dispersive wave models to the fully nonlinear case we have proposed a new framework to approximate the so-called Green-Naghdi equations [99]. The method proposed, while remaining unsplit in time, is based on a separation of the elliptic and hyperbolic components of the equations. This separation is designed to recover the standard shallow water equations in the hyperbolic step, so that the method can be written as an *algebraic* correction to an existing shallow water code. In particular, in our approach we fix the method used for the elliptic component (a continuous Galerkin method), and couple it to different hyperbolic shallow water solvers. As long as the hyperbolic step is more than second order accurate, the approach proposed allows accuracies comparable to those of a fourth order finite difference method, with a natural potential for h and p - adaptation on unstructured grids. The two-dimensional extension is in the testing phase.

The tools developed have been also used intensively in funded research programs. Within the TANDEM project, several benchmarks relevant to tsunami modelling have been performed and several common publications with the project partners are in preparation. Independently on this activity, this year we used our codes to investigate two case studies. The first is the study of the wave conditions for the old Venetian harbour of Chania in Crete [109]. The study compares fully nonlinear-weakly dispersive COULWAVE code, developed at the University of South California, and TUCWave. The models are used to explore the appearance of resonance, eventually determining the resonant frequencies for the entire basin. The second study concerns the conditions for tidal bore formation in convergent alluvial estuaries [69]. A new set of dimensionless parameters has been introduced to describe the problem, and the code SLOWS has been used to explore the space of these parameters allowing to determine a critical curve allowing to characterize an estuary as "bore forming" or not. Surprising physical behaviours, in terms of dissipation and nonlinearity of the tides, have been highlighted. Part of this work has been accepted on *Estuarine, Coastal and Shelf Science*, with a manuscript on the numerical aspects in review on *Ocean Modelling*.

7.3. Wave energy conversion hydrodynamics

Participants: Umberto Bosi, Mario Ricchiuto [Corresponding member].

We have developed a prototype spectral element solver for a coupled set of differential equations modelling wave propagation (so-called outer domain), and the submerged flow under a floating body (inner domain). Both systems of equations are depth-averaged (Boussinesq type) systems involving some dispersive terms. They are further coupled to a force balance providing a (system of) ODE(s) for the floater. This model constitutes an intermediate fidelity approximation for the hydrodynamics of a wave energy converter. Differently from all industrial state of the art, it is a (fully) nonlinear model. However, its cost is extremely low when compared to full three-dimensional CFD analyses, due to the dimensional reduction brought from the depth averaged modelling. This year we have shown the potential of this Boussinesq-type model to predict the hydrodynamics

of a floater in a simplified case [97], [98] (journal version to appear on *J. Ocean Eng. and Marine Energy*). This result paves the way to the construction of new medium fidelity models to be used in the optimization of converters. This will be achieved in the framework of the MIDWEST project funded this year under the EU OCEANERanet call.

7.4. Two-phase flow numerical simulation with real-gas effects and occurrence of rarefaction shock waves

Participants: Maria Giovanna Rodio, Pietro Marco Congedo [Corresponding member].

We have studied the prediction in numerical simulation of turbulent cavitating flows, which could be strongly influenced by the presence of several empirical coefficients. The aim of this work is to explore the interaction between the cavitation model and turbulence in terms of uncertainty propagation through an unsteady numerical solver, for assessing the robustness and the accuracy of the physical models at different times. Furthermore, the influence of experimental data in the setting of some turbulence and cavitation model coefficients is investigated by means of a Bayesian approach. Finally, the interest is to provide some innovative insights for improving the understanding of these models for cavitating flows.

7.5. Formulation of stochastic methods for CFD

Participants: Maria Giovanna Rodio, Pietro Marco Congedo [Corresponding member].

We have worked on the extension of the Truncate and Encode (TE) approach, previously proposed by some of the authors (Abgrall et al. in *J Comput Phys* 257:19?56, 2014. doi:10.1016/j.jcp.2013.08.006), for taking into account uncertainty in partial differential equations (PDEs). Innovative ingredients are given by an algorithm permitting to recover the multiresolution representation without requiring the fully resolved solution, the possibility to treat a whatever form of pdf and the use of high-order (even non-linear, i.e. data-dependent) reconstruction in the stochastic space. Moreover, the spatial-TE method is introduced, which is a weakly intrusive scheme for uncertainty quantification (UQ), that couples the physical and stochastic spaces by minimizing the computational cost for PDEs. The proposed scheme is particularly attractive when treating moving discontinuities (such as shock waves in compressible flows), even if they appear during the simulations as it is common in unsteady aerodynamics applications. The proposed method is very flexible since it can easily coupled with different deterministic schemes, even with high-resolution features. Flexibility and performances of the present method are demonstrated on various numerical test cases (algebraic functions and ordinary differential equations), including partial differential equations, both linear and non-linear, in presence of randomness. The efficiency of the proposed strategy for solving stochastic linear advection and Burgers equation is shown by comparison with some classical techniques for UQ, namely Monte Carlo or the non-intrusive polynomial chaos methods.

A new scheme for the numerical approximation of a five-equation model taking into account Uncertainty Quantification (UQ) is also presented. In particular, the Discrete Equation Method (DEM) for the discretization of the five-equation model is modified for including a formulation based on the adaptive Semi-Intrusive (aSI) scheme, thus yielding a new intrusive scheme (sDEM) for simulating stochastic two-phase flows. Some reference test-cases are performed in order to demonstrate the convergence properties and the efficiency of the overall scheme. The propagation of initial conditions uncertainties is evaluated in terms of mean and variance of several thermodynamic properties of the two phases.

7.6. Sensitivity analysis, metamodelling, and and robust optimization

Participants: Kunkun Tang, Francesca Fusi, Pietro Marco Congedo [Corresponding member].

We have worked on two different formulations for sensitivity analysis. Moreover, we have proposed a new metamodelling technique and an innovative method for performing robust optimization.

Concerning sensitivity analysis, an anchored analysis of variance (ANOVA) method is proposed to decompose the statistical moments. Compared to the standard ANOVA with mutually orthogonal component functions, the anchored ANOVA, with an arbitrary choice of the anchor point, loses the orthogonality if employing the same measure. However, an advantage of the anchored ANOVA consists in the considerably reduced number of deterministic solver's computations, which renders the uncertainty quantification of real engineering problems much easier. Different from existing methods, the covariance decomposition of the output variance is used to take account of the interactions between non-orthogonal components, yielding an exact variance expansion and thus, with a suitable numerical integration method, provides a strategy that converges. This convergence is verified by studying academic tests. In particular, the sensitivity problem of existing methods to the choice of anchor point is analyzed via the Ishigami case, and we point out that covariance decomposition survives from this issue. Also, with a truncated anchored ANOVA expansion, numerical results prove that the proposed approach is less sensitive to the anchor point. The covariance-based sensitivity indices (SI) are also used, compared to the variance-based SI. Furthermore, we emphasize that the covariance decomposition can be generalized in a straightforward way to decompose higher-order moments. For academic problems, results show the method converges to exact solution regarding both the skewness and kurtosis. Finally, the proposed method is applied on a realistic case, that is, estimating the chemical reactions uncertainties in a hypersonic flow around a space vehicle during an atmospheric reentry.

A sensitivity analysis method is extended in order to compute third and fourth-order statistic moments, i.e. skewness and kurtosis, respectively. It is shown that this decomposition is correlated to a Polynomial Chaos (PC) expansion, permitting to easily compute each term. Then, new sensitivity indexes are proposed basing on the computation of skewness and kurtosis. PC-based numerical technique is used in order to compute the convergence of the sensitivity indexes according to the polynomial order by using the exact solution as the reference one. The interest of the proposed analysis is first depicted by considering several test-functions. In particular, a functional decomposition based on variance, skewness and kurtosis is computed, displaying how sensitivity indexes vary according to the order of the statistical moment. Then, the problem of how reducing the complexity of a stochastic problem is treated by proposing two strategies: i) the reduction of the number of dimensions, the reduction of the order of interaction. In both cases, the impact on the statistics of the reduced function is then assessed. Feasibility of the proposed analysis in a real-case is then demonstrated by presenting a stochastic study about the uncertainty propagation in a challenging engineering simulation: the numerical prediction of a turbine cascade in an Organic Rankine Cycles (ORCs), with the use of complex thermodynamic models and the presence of multiple sources of uncertainty. Basing on high-order statistics decomposition and physical remarks, it is shown how the analysis proposed in this work can help to drive the design process in a real-engineering problem.

For the metamodeling technique, a polynomial dimensional decomposition (PDD) method is proposed for the global sensitivity analysis and uncertainty quantification (UQ) of stochastic systems subject to a moderate to large number of input random variables. Due to the intimate structure between the PDD and the Analysis of Variance (ANOVA) approach, PDD is able to provide a simpler and more direct evaluation of the Sobol' sensitivity indices, when compared to the Polynomial Chaos expansion (PC). Unfortunately, the number of PDD terms grows exponentially with respect to the size of the input random vector, which makes the computational cost of standard methods unaffordable for real engineering applications. In order to address the problem of the curse of dimensionality, this work proposes essentially variance-based adaptive strategies aiming to build a cheap meta-model (i.e. surrogate model) by employing the sparse PDD approach with its coefficients computed by regression. Three levels of adaptivity are carried out in this paper: 1) the truncated dimensionality for ANOVA component functions, 2) the active dimension technique especially for second- and higher-order parameter interactions, and 3) the stepwise regression approach designed to retain only the most influential polynomials in the PDD expansion. During this adaptive procedure featuring stepwise regressions, the surrogate model representation keeps containing few terms, so that the cost to resolve repeatedly the linear systems of the least-squares regression problem is negligible. The size of the finally obtained sparse PDD representation is much smaller than the one of the full expansion, since only significant terms are eventually retained. Consequently, a much less number of calls to the deterministic model is required to compute the final PDD coefficients.

Concerning robust optimization, a strategy is developed to deal with the error affecting the objective functions in uncertainty-based optimization. We refer to the problems where the objective functions are the statistics of a quantity of interest computed by an uncertainty quantification technique that propagates some uncertainties of the input variables through the system under consideration. In real problems, the statistics are computed by a numerical method and therefore they are affected by a certain level of error, depending on the chosen accuracy. The errors on the objective function can be interpreted with the abstraction of a bounding box around the nominal estimation in the objective functions space. In addition, in some cases the uncertainty quantification methods providing the objective functions also supply the possibility of adaptive refinement to reduce the error bounding box. The novel method relies on the exchange of information between the outer loop based on the optimization algorithm and the inner uncertainty quantification loop. In particular, in the inner uncertainty quantification loop, a control is performed to decide whether a refinement of the bounding box for the current design is appropriate or not. In single-objective problems, the current bounding box is compared to the current optimal design. In multi-objective problems, the decision is based on the comparison of the error bounding box of the current design and the current Pareto front. With this strategy, fewer computations are made for clearly dominated solutions and an accurate estimate of the objective function is provided for the interesting, non-dominated solutions. The results presented in this work prove that the proposed method improves the efficiency of the global loop, while preserving the accuracy of the final Pareto front.

7.7. High order mesh generation and mesh adaptation

Participants: Luca Arpaia, Cécile Dobrzynski [Corresponding member], Ghina El Jannoun, Mario Ricchiuto.

This year several new algorithmic improvements have been obtained which will allow to enhance our meshing tools:

- We have enhanced our work on r -adaptation techniques for time dependent equations. These techniques are based on mesh deformations obtained by solving continuous differential equations for the local displacements. These equations are controlled by an error monitor. Several improvements have been made. We have studied in depth the formulation of the coupling of the mesh movement with the flow solver. We have found that for both finite volume and residual distribution methods, a coupling of mesh and solution evolution (by means of an ALE method) provides accuracy enhancements, and is to be preferred to a simpler adapt-project-evolve approach. The method has been fully tested in two space dimensions. The adaptation library has been extended to three dimensions, and benchmarking is under way. We have improved the definition of the error monitor, and we are now able to prescribe directly the local mesh size. For problems with source terms, and in particular problems admitting some important physical invariants as the shallow water equations, we have solved the conflict between the conservation of either mass or the invariant, allowing for the conservation of both quantities up to machine accuracy;
- We upgrade our technique for generating high order curved meshes: starting from a straight mesh with a curved boundary, a new smoothing and untangling approach is proposed to ensure a final valid mesh. The untangling algorithm is a hybrid technique that gathers a local mesh optimization applied on the surface of the domain and a linear elasticity analogy applied in its volume. On the one hand, the local topological optimization consists in simultaneously relocating the vertices and control points of a local patch around the invalid element in order to optimize the quality and validity of the elements inside the patch. The elements' validity problem is formulated as an unconstrained optimization problem using a log-barriers that is solved progressively using the conjugate gradient method. On the other hand, the linear elasticity analogy permits the propagation of the curvature to the volume of the domain hence untangling volume mesh elements.

7.8. Virtual self-healing composite for aeronautic propulsion

Participants: Mathieu Colin [Corresponding member], Xi Lin, Gregory Perrot, Mario Ricchiuto.

As a composite material having excellent properties, Ceramic-matrix composites (CMCs) comprise a ceramic matrix reinforced by a refractory fiber such as silicon carbide fiber. Due to the self-healing process (which consists in filling cracks produced by oxydation by an oxide formed in-situ), CMCs have extremely long lifetimes even under severe mechanical and chemical sollicitations. These time spans make most full-scale experimental investigations impractical : laboratory tests have necessarily to be replaced by predictions based on numerical models. Initial results have already been obtained by in the past with simplified crack averaged models based on simple potential approximations of the flow field of the oxide. In this direction, Xi Lin has developed new asymptotic models by creating a hierarchy inside two different families : the Saint-Venant equations and the thin film equations. The hierarchy is based on the use of several different boundary conditions. The main goal is to obtain more accurate hydrodynamic models accounting for surface tension and viscous effects which may be very important for the oxide evolution.

In parallel, we have made great progress in the coupling of the chemistry module with the structural mechanics solver of the LCTS laboratory in Bordeaux. The first fully coupled simulations of a fatigue test for a so-called mini-composite (on single fibre tow). The simulations have allowed to reproduce the gradual breaking mechanism typical of these materials, allowing to reproduce numerically the delayed rupture observed in practice.

7.9. Numerical simulation of the liquid ablation

Participants: Gérard Gallice, Luc Mieussens [Corresponding member], Simon Peluchon.

During the atmospheric re-entry phase, a space vehicle undergoes a heating due to the friction of the atmospheric gases. Conversion of kinetic energy to thermal energy leads to a sudden increase of the temperature of the solid boundary. This leads to a physical-chemical degradation of the thermal protective system, and to a boundary recession. For some materials, this recession occurs with a melting of the materials into a fluid phase. The numerical simulation of this phenomenon requires to take into account a two-phase flow with a compressible gaz (the air flow) and a weakly compressible liquid (the melted material). Numerically, this problem is strongly stiff.

We have proposed a splitting strategy to simulate compressible two-phase flows using the five equation model. The main idea of the splitting is to separate the acoustic and transport phenomena. The acoustic step is solved in Lagrangian coordinates by using a scheme based on an approximate Riemann solver. On the one hand, since the acoustic time step driven by the fast sound velocity is very restrictive, an implicit treatment of the Lagrangian step is performed. On the other hand, we use an explicit scheme for the transport step driven by the slow material waves. The global scheme resulting from this splitting operator strategy is conservative, positive, and preserves contact discontinuities. Numerical simulations of compressible diphasic flows are presented on 2d-structured grids. The implicit-explicit strategy allows large time steps, which do not depend on the fast acoustic waves.

7.10. Boundary conditions for the Navier-Stokes equations in the transitional regime

Participants: Céline Baranger, Pietro Marco Congedo, Giorgio Martalo, Julien Mathiaud, Luc Mieussens [Corresponding member].

In reentry flows at high altitude, the parietal fluxes along the boundary of a space vehicle are computed by solving the Boltzmann equation of the gas kinetic theory. It depends on the way particles are reflected by the solid wall. In this kind of applications, the reflection is usually supposed to be 80% diffuse (the particle are re-emitted in a random direction in thermal equilibrium with the wall), and 20% specular.

In lower altitude, it is possible to use the compressible Navier-Stokes equations, but the standard boundary conditions do not take into account the specular part. These equations are hence not very accurate in the transition regime (in which the Knudsen number is around 0.01).

By using an asymptotic boundary layer analysis, we have derived boundary conditions for the compressible Navier-Stokes equations that formally make the fluid model a second order approximation of the Boltzmann equation of the kinetic gas theory (with respect to the Knudsen number), and that can take into account the effect of the specular reflections. These boundary conditions include a slip velocity at the wall and a temperature jump, with some coefficients that depend on some auxiliary half-space problems. An existing numerical method has been extended to solve these problems and give numerical values for these coefficients. This makes our boundary conditions practically usable into any Navier-Stokes code.

CARMEN Team

7. New Results

7.1. Inverse Problem

Electrocardiograms simulated by our group with a highly realistic and detailed forward model were used for several inverse-modeling studies [34], [33], [38], [35].

- **Stability analysis of the POD reduced order method for solving the bidomain model in cardiac electrophysiology:** In this work we show the numerical stability of the Proper Orthogonal Decomposition (POD) reduced order method used in cardiac electrophysiology applications. The difficulty of proving the stability comes from the fact that we are interested in the bidomain model, which is a system of degenerate parabolic equations coupled to a system of ODEs representing the cell membrane electrical activity. The proof of the stability of this method is based on an a priori estimate controlling the gap between the reduced order solution and the Galerkin finite element one. We present some numerical simulations confirming the theoretical results. We also combine the POD method with a time splitting scheme allowing a faster solution of the bidomain problem and show numerical results. Finally, we conduct numerical simulation in 2D illustrating the stability of the POD method in its sensitivity to the ionic model parameters. We also perform 3D simulation using a massively parallel code. We show the computational gain using the POD reduced order model. We also show that this method has a better scalability than the full finite element method.
- **In silico assessment of drug effects on human embryonic stem cells derived cardiomyocytes electrical activity:** Computational modeling and simulation is extensively used to investigate diseases in cardiac electrophysiological activity and also drug effects, side effects and interactions. Human embryonic stem cell-derived cardiomyocytes (hESC-CMs) have been recently considered as a promising tool in regenerative medicine: their major role in repairing damaged tissue is due to pluripotency and ability to differentiate. These pluripotent cells are also used in early stages of drugs development. Pharmaceutical companies use the MultiElectrode Array (MEA) device in order to perform many in vitro experiments on hESC-CMs. The goal of our study is to derive a mathematical model and to simulate these in vitro experiments.
- **Sensitivity of the Electrocardiography Inverse Solution to the Torso Conductivity Uncertainties:** Electrocardiography imaging (ECGI) is a new non invasive technology used for heart diagnosis. It allows to construct the electrical potential on the heart surface only from measurement on the body surface and some geometrical informations of the torso. The purpose of this work is twofold: First, we propose a new formulation to calculate the distribution of the electric potential on the heart, from measurements on the torso surface. Second, we study the influence of the errors and uncertainties on the conductivity parameters, on the ECGI solution. We use an optimal control formulation for the mathematical formulation of the problem with a stochastic diffusion equation as a constraint. The discretization is done using stochastic Galerkin method allowing to separate random and deterministic variables. The optimal control problem is solved using a conjugate gradient method where the gradient of the cost function is computed with an ad-joint technique. The efficiency of this approach to solve the inverse problem and the usability to quantify the effect of conductivity uncertainties in the torso are demonstrated through a number of numerical simulations on a 2D geometrical model. Our results show that adding $\pm 50\%$ uncertainties in the fat conductivity does not alter the inverse solution, whereas adding $\pm 50\%$ uncertainties in the lung conductivity affects the reconstructed heart potential by almost 50%.
- **Inverse Localization of Ischemia in a 3D Realistic Geometry: A Level Set Approach:** The reconstruction of cardiac ischemic regions from body surface potential measurements (BSPMs) is usually performed at a single time instant which corresponds to the plateau or resting phase of the cardiac

action potential. Using a different approach, we previously proposed a level set formulation that incorporates the knowledge of the cardiac excitation process in the inverse procedure, thus exploiting the spatio-temporal correlations contained in the BSPMs. In this study, we extend our inverse level-set formulation for the reconstruction of ischemic regions to 3D realistic geometries, and analyze its performance in different noisy scenarios. Our method is benchmarked against zero-order Tikhonov regularization. The inverse reconstruction of the ischemic region is evaluated using the correlation coefficient (CC), the sensitive error ratio (SN), and the specificity error ratio (SP). Our algorithm outperforms zero-order Tikhonov regularization, specially in highly noisy scenarios.

- Inverse problem in electrocardiography via the factorization method of boundary value problems: We present a new mathematical approach for solving the inverse problem in electrocardiography. This approach is based on the factorization of boundary value problems method. In this paper we derive the mathematical equations and test this method on synthetic data generated on realistic heart and torso geometries using the state-of-the-art bidomain model in the heart coupled to the Laplace equation in the torso. We measure the accuracy of the inverse solution using spatial Relative Error (RE) and Correlation Coefficient (CC).
- In the inverse problem in electrocardiology, the goal is to recover electrophysiological activity in the heart without measuring directly on its surface (without using catheter interventions). Note that today the inverse computation is frequently used by solving the quasi-static model. This model doesn't take into account the heart dynamic in time and may result in considerable errors in the reconstruction of the solution on the heart. In [1] we study a 3D numerical inverse problem constrained by the bidomain equations in electrocardiology. The state equations consisting in a coupled reaction-diffusion system modelling the propagation of the intracellular and extracellular electrical potentials, and ionic currents, are extended to further consider the effect of an external bathing medium. Thus, we demonstrate that the novel concept of applying electrophysiological data might be useful to improve noninvasive reconstruction of electrical heart activity. Finally, we present numerical experiments representing the effect of the heart dynamic on the inverse solutions. Moreover in [2], we study the stability result for the conductivities of the approximate bidomain model. The proof is based on the combination of a Carleman estimate obtained in [3] and certain weight energy estimates for parabolic systems.
- The static inverse ECG problem needs to solve the well known ill posed Cauchy problem for the Laplace equation. A new approach investigated in the team uses the method of factorization of boundary value problems. This method, studied for itself, provides in this context the computation of Dirichlet-Neumann operators as solution of a Riccati equation. Results have been presented at IEEE international symposium on biomedical imaging, New York april 16-19, 2015. Further investigations will be lead using more precise numerical methods to solve the Riccati equation. The non-linearity and time dependence of the coupling resistance between cardiac cells (gap junctions) is studied in the Liryc institute and thought to be of importance in the understanding of cardiac arrhythmias. The internship of Nhan Le Thanh was a first step to investigate their numerical simulation.

An Example of calculus on a torso is shown in figure 1

7.2. Cardiac Electromechanics

In [1] we study a coupled elliptic-parabolic system modeling the interaction between the propagation of electric potential and subsequent deformation of the cardiac tissue. The problem consists in a reaction-diffusion system governing the dynamics of ionic quantities, intra and extra-cellular potentials, and the linearized elasticity equations are adopted to describe the motion of an incompressible material. The coupling between muscle contraction, biochemical reactions and electric activity is introduced with a so-called active strain decomposition framework, where the material gradient of deformation is split into an active (electrophysiology-dependent) part and an elastic (passive) one. In this paper we prove existence of weak solutions to the underlying coupled reaction-diffusion system and uniqueness of regular solutions. We close with a numerical example illustrating the convergence of the method and some features of the model.




Figure 1. Mesh of the torso obtained with an IRM picture; this problem is link with the IDAM project

7.3. Cardiac Electrophysiology at the Microscopic Scale

We focused on establishing a microscopic model for cardiac electrophysiology simulations and proving the existence of a solution. We started with writing a mathematical proof allowing from well known physical equations and properties of the cardiac tissue to establish the model. Then, we worked on a variational formulation of the problem, and describing a weak solution of it. The idea is to compute energy estimates and to bound them so that we can extract a convergent sequence of functions in the appropriate Sobolev space. With my PhD advisor, we started to write an article about these two proofs. We also worked on CEPS code to implement some functionalities that will fit my requirements in a near future regarding the simulations we have to design. The main difficulty we identified is, provided we get a well defined geometry and mesh of cardiac cells, to implement the ionic flux between cells. First simulation of a simple "two-cells communication" problem will probably, if the results meet experimental observations, lead to another article. We also attended Imaged Based Biomedical Modelling 2015, a summer course organized by SCI institute (University of Utah), which was designed to give attendees guidelines about visualization and modelling, especially on cardio electrophysiology.

7.4. High order numerical scheme for ionic models

C. Douanla Ionti worked on time numerical schemes like Adams-Bashforth in order to have a high degree of convergence between an exact solution and the approximated solution. This method is a generalisation of Rush-Larshen scheme adapted for electrophysiology cardiac.

CQFD Project-Team

7. New Results

7.1. Control of parallel non-observable queues: asymptotic equivalence and optimality of periodic policies

The following result has been obtained by J. Anselmi (Inria CQFD), T. Nesti and B. Gaujal.

We consider a queueing system composed of a dispatcher that routes deterministically jobs to a set of non-observable queues working in parallel. In this setting, the fundamental problem is which policy should the dispatcher implement to minimize the stationary mean waiting time of the incoming jobs. We present a structural property that holds in the classic scaling of the system where the network demand (arrival rate of jobs) grows proportionally with the number of queues. Assume that each queue of type r is replicated k times and consider the set of policies that are periodic with period $k \sum_r p_r$ and such that exactly p_r jobs are sent in a period to each queue of type r . When $k \rightarrow \infty$, our main result shows that all the policies in this set are equivalent, in the sense that they yield the same mean stationary waiting time, and optimal, in the sense that no other policy having the same aggregate arrival rate to all queues of a given type can do better in minimizing the stationary mean waiting time. This property holds in a strong probabilistic sense. Furthermore, the limiting mean waiting time achieved by our policies is a convex function of the arrival rate in each queue, which facilitates the development of a further optimization aimed at solving the fundamental problem above for large systems.

7.2. Decentralized Proportional Load Balancing

The following result has been obtained by J. Anselmi (Inria CQFD), and N. Walton.

Load balancing is a powerful technique commonly used in communication and computer networks to improve system performance, robustness and fairness. In this paper, we consider a general model capturing the performance of communication and computer networks, and on top of it we propose a decentralized algorithm for balancing load among multiple network paths. The proposed algorithm is inspired by the modus operandi of the processor-sharing queue and on each network entry point operates as follows: every time a unit of load completes its service on a path, it increases by one unit the load of that path and decreases by one unit the load of a path selected at random with probability proportional to the amount of load on each of the available paths. We develop a dynamical system to argue that our load-balancer achieves a desirable network-wide utility optimization.

A paper has been accepted for publication in the SIAM Journal of Applied Mathematics.

7.3. Conditional quantile estimation through optimal quantization

The following result has been obtained by Isabelle Charlier (CQFD member), Davy Paindaveine, and Jérôme Saracco (CQFD member)

We use quantization to construct a nonparametric estimator of conditional quantiles of a scalar response Y given a d -dimensional vector of covariates X . First we focus on the population level and show how optimal quantization of X , which consists in discretizing X by projecting it on an appropriate grid of N points, allows to approximate conditional quantiles of Y given X . We show that this approximation is arbitrarily good as N goes to infinity and provide a rate of convergence for the approximation error. Then we turn to the sample case and define an estimator of conditional quantiles based on quantization ideas. We prove that this estimator is consistent for its fixed- N population counterpart. The results are illustrated on a numerical example. Dominance of our estimators over local constant/linear ones and nearest neighbor ones is demonstrated through extensive simulations in the companion paper Charlier et al. (2014).

7.4. A linear programming formulation for constrained discounted continuous control for piecewise deterministic Markov processes

The following result has been obtained by Oswaldo Costa and François Dufour (CQFD member).

This work deals with the constrained discounted control of piecewise deterministic Markov process (PDMPs) in general Borel spaces. The control variable acts on the jump rate and transition measure, and the goal is to minimize the total expected discounted cost, composed of positive running and boundary costs, while satisfying some constraints also in this form. The basic idea is, by using the special features of the PDMPs, to re-write the problem via an embedded discrete-time Markov chain associated to the PDMP and re-formulate the problem as an infinite dimensional linear programming (LP) problem, via the occupation measures associated to the discrete-time process. It is important to stress however that our new discrete-time problem is not in the same framework of a general constrained discrete-time Markov Decision Process and, due to that, some conditions are required to get the equivalence between the continuous-time problem and the LP formulation. We provide in the sequel sufficient conditions for the solvability of the associated LP problem. We provide some examples to illustrate the obtained results.

7.5. Impulsive control for continuous-time Markov decision processes

The following result has been obtained by Alexey Piunovskiy and François Dufour (CQFD member).

The objective of this work is to study continuous-time Markov decision processes on a general Borel state space with both impulsive and continuous controls for the infinite-time horizon discounted cost. The continuous-time controlled process is shown to be non explosive under appropriate hypotheses. The so-called Bellman equation associated to this control problem is studied. Sufficient conditions ensuring the existence and the uniqueness of a bounded measurable solution to this optimality equation are provided. Moreover, it is shown that the value function of the optimization problem under consideration satisfies this optimality equation. Sufficient conditions are also presented to ensure on one hand the existence of an optimal control strategy and on the other hand the existence of an ε -optimal control strategy. The decomposition of the state space in two disjoint subsets is exhibited where roughly speaking, one should apply a gradual action or an impulsive action correspondingly to get an optimal or ε -optimal strategy. An interesting consequence of our previous results is as follows: the set of strategies that allow interventions at time $t = 0$ and only immediately after natural jumps is a sufficient set for the control problem under consideration.

7.6. Impulsive control for continuous-time Markov decision processes: A Linear Programming Approach

The following result has been obtained by Alexey Piunovskiy and François Dufour (CQFD member).

The objective of this work is to investigate an optimization problem for continuous-time Markov decision processes with both impulsive and continuous controls. We consider the so-called constrained problem where the objective of the controller is to minimize a total expected discounted optimality criterion associated with a cost rate function while keeping other performance criteria of the same form, but associated with different cost rate functions, below some given bounds. Our model allows multiple impulses at the same time moment. The main objective of this work is to study the associated linear program defined on a space of measures including the occupation measures of the controlled process and to provide sufficient conditions to ensure the existence of an optimal control.

7.7. Conditions for the Solvability of the Linear Programming Formulation for Constrained Discounted Markov Decision Processes

The following result has been obtained by François Dufour (CQFD member) and T. Prieto-Rumeau.

This result concerns discrete-time constrained discounted Markov decision processes (MDP) with Borel state and action spaces, compact action sets, and lower semi-continuous cost functions. We introduce a set of hypotheses related to a positive weight function which allow us to consider cost functions that might not be bounded below by a constant, and which imply the solvability of the linear programming formulation of the constrained MDP. In particular, we establish the existence of a constrained optimal stationary policy. Our results are illustrated with an application to a fishery management problem.

7.8. Comparison of Kernel Density Estimators with Assumption on Number of Modes

The following result has been obtained by Gilles Durrieu, Raphaël Coudret and Jérôme Saracco (CQFD member).

A data-driven bandwidth choice for a kernel density estimator called critical bandwidth is investigated. This procedure allows the estimation to have as many modes as assumed for the density to estimate. Both Gaussian and uniform kernels are considered. For the Gaussian kernel, asymptotic results are given. For the uniform kernel, an argument against these properties is mentioned. These theoretical results are illustrated with a simulation study that compares the kernel estimators that rely on critical bandwidth with another one that uses a plug-in method to select its bandwidth. An estimator that consists in estimates of density contour clusters and takes assumptions on number of modes into account is also considered. Finally, the methodology is illustrated using environment monitoring data.

7.9. EEG classification for the detection of mental states

The following result has been obtained by Laurent Vezard, Pierrick Legrand (CQFD member), Marie Chavent (CQFD member), Frederique Faight-Ainseba and Trujillo Leonardo

The objective of the present work is to develop a method that is able to automatically determine mental states of vigilance; i.e., a person's state of alertness. Such a task is relevant to diverse domains, where a person is expected or required to be in a particular state of mind. For instance, pilots and medical staff are expected to be in a highly alert state and the proposed method could help to detect possible deviations from this expected state. This work poses a binary classification problem where the goal is to distinguish between a "relaxed" state and a baseline state ("normal") from the study of electroencephalographic signals (EEG) collected with a small number of electrodes. The EEG of 58 subjects in the two alertness states (116 records) were collected via a cap with 58 electrodes. After a data validation step, 19 subjects were retained for further analysis. A genetic algorithm was used to select a subset of electrodes. Common spatial pattern (CSP) coupled to linear discriminant analysis (LDA) was used to build a decision rule and thus predict the alertness of the subjects. Different subset sizes were investigated and the best compromise between the number of selected electrodes and the quality of the solution was obtained by considering 9 electrodes. Even if the present approach is costly in computation time (GA search), it allows to construct a decision rule that provides an accurate and fast prediction of the alertness state of an unseen individual.

7.10. Modeling and optimization of a launcher integration process

The following result has been obtained by Christophe Nivot (CQFD member), Benoîte De Saporta, François Dufour (CQFD member), Jacques Béhar, Damien Bérard-Bergery and Charles Elegbede.

We deal with the modeling and the optimization of a launcher integration process. The subassemblies go through various types of operations which are split up into workshops. Their operating time is supposed random due to possible breakdowns or staff issues. Storage capacity of output products is limited and costly. Launches have to be performed according to a predetermined schedule, and lateness also costs money. The rate of production of the subassemblies must be decided every year. Therefore, the system can be modeled with a Markov decision process which is suitable for decision optimization and cost minimization. Indeed, one must find a balance between slow production (thus low storage levels and high probability to be late), and fast production (high storage levels but respected schedule).

We propose a model of this integration process based on Markov decision models. We present the simulation we have performed so far and discuss the difficulties of the optimization.

7.11. ClustGeo: Ascendant Hierarchical Clustering (AHC) with geographical constraints

The following result has been obtained by Marie Chavent (CQFD member), Vanessa Kuentz-Simonet, Amaury Labenne and Jerome Saracco (CQFD member).

Hierarchical Ascendant Clustering (HAC) is a well-known method of individual clustering. This method aims to bring together individuals who are similar regarding to variables which describe them. But when individuals are geographical units, the user may wish geographically close individuals to be put in same clusters and that, without too much deteriorating the quality of the partition. The proposed ClustGeo method allows geographical constraints of proximity to be taken into account within the HAC. For that purpose, a new Ward homogeneity criterion based on two different matrices of distances is proposed.

7.12. Approche bayésienne non paramétrique pour la factorisation de matrice binaire à faible rang avec loi de puissance

The following result has been obtained by Adrien Todeschini (CQFD member) and François Caron.

We introduce a low-rank Bayesian nonparametric (BNP) model for bipartite graphs. Recently, Caron (2012) proposed a BNP model where each node is given its own sociability parameter allowing to capture the power-law behavior of real world bipartite graphs. This model can be considered as a rank one nonnegative factorization of the adjacency matrix. Building on the compound random measures recently introduced by Griffin and Leisen (2014), we derive a rank p generalization of this model where each node is associated with a p -dimensional vector of sociability parameters accounting for several latent dimensions. While preserving the desired properties of interpretability, scalability and power-law behavior, our model is more flexible and provides better predictive performance as illustrated on several datasets.

7.13. Compétitions d'apprentissage automatique avec le package R rchallenge

The following result has been obtained by Adrien Todeschini (CQFD member) Robin Genuer.

In machine learning, empirical performance on real data are crucial in the success of a method. Recent years have seen the emergence of a large number of machine learning competitions. These challenges are motivated by industrial (Netflix prize) or academic (HiggsML challenge) applications and put in competition researchers and data scientists to obtain the best performance. We wanted to expose students to this reality by submitting a challenge in the context of the machine learning course. The leaderboard is displayed on an automatically updated web page allowing emulation among students. The history of the results also allows them to visualize their progress through the submissions. In addition, the challenge can continue outside of the supervised sessions promoting independence and exploration of new learning techniques and computer tools. The system we have implemented is available as an R package for reuse by other teachers. Building on the R Markdown and Dropbox tools, it requires no network configuration and can be deployed very easily on a personal computer.

7.14. Novelty Search

The following result has been obtained by Enrique Naredo, Leonardo Trujillo and Pierrick Legrand (CQFD member).

Novelty Search (NS) is a unique approach towards search and optimization, where an explicit objective function is replaced by a measure of solution novelty. However, NS has been mostly used in evolutionary robotics while its usefulness in classic machine learning problems has been unexplored. This work presents a NS-based Genetic Programming (GP) algorithm for supervised classification. Results show that NS can solve real-world classification tasks, validated on real-world benchmarks for binary and multiclass problems. These results are made possible by using a domain-specific behavior descriptor. Two new versions of the NS algorithm are proposed, Probabilistic NS (PNS) and a variant of Minimum Criterion NS (MCNS). The former models the behavior of each solution as a random vector and eliminates all of the original NS parameters while reducing the computational overhead of the NS algorithm. The latter uses a standard objective function to constrain and bias the search towards high performance solutions. The paper also discusses the effects of NS on GP search dynamics and code growth. Results show that NS can be used as a realistic alternative for supervised classification, and for binary problems the NS algorithm exhibits an implicit bloat control ability.

Keywords: Novelty Search, Behavior-based Search, Supervised Classification, Bloat

7.15. Classification of Epileptic states

The following result has been obtained by Emigdio Z. Flores, Leonardo Trujillo and Pierrick Legrand (CQFD member).

The neurological disorder known as Epilepsy is characterized by involuntary recurrent seizures that diminish a patient's quality of life. Automatic seizure detection can help improve a patient's interaction with her/his environment, and while many approaches have been proposed the problem is still not trivially solved. In this work, we present a novel methodology for feature extraction on EEG signals that allows us to perform a highly accurate classification of epileptic states. Specifically, Hölderian regularity and Matching Pursuit are used as the main feature extraction techniques, and are combined with basic statistics to construct the final feature sets. These sets are then delivered to a Random Forests classification algorithm. Furthermore, several versions of the basic problem are tested and statistically validated producing perfect accuracy in most problems and 92% accuracy on the most difficult case. A comparison with recent results in relevant literature using a well known database reveals that our proposal achieves state-of-the-art performance.

Keywords: Epilepsy detection, Hölderian regularity, Matching Pursuit, EEG Classification

7.16. Prediction of expected performance

The following result has been obtained by Yuliana Martinez, Leonardo Trujillo and Pierrick Legrand (CQFD member).

The study of problem difficulty is an open issue in Genetic Programming (GP). The goal of this work is to generate models that predict the expected performance of a GP-based classifier when it is applied to an unseen task. Classification problems are described using domain-specific features, some of which are proposed in this work, and these features are given as input to the predictive models. These models are referred to as predictors of expected performance (PEPs). We extend this approach by using an ensemble of specialized predictors (SPEPs), dividing classification problems into specified groups and choosing the corresponding SPEP. The proposed predictors are trained using 2D synthetic classification problems with balanced datasets. The models are then used to predict the performance of the GP classifier on unseen real-world datasets that are multidimensional and imbalanced. Moreover, as we know, this work is the first to provide a performance prediction of the GP classifier on test data, while previous works focused on predicting training performance. Accurate predictive models are generated by posing a symbolic regression task and solving it with GP. These results are achieved by using highly descriptive features and including a dimensionality reduction stage that simplifies the learning and testing process. The proposed approach could be extended to other classification algorithms and used as the basis of an expert system for algorithm selection.

7.17. Simulation of SPDEs for Excitable Media Using Finite Elements

The following result has been obtained by

M. Boulakia, A. Genadot (CQFD member) and M. Thieullen.

This result concerns the question of the discretization of Stochastic Partial Differential Equations (SPDE's) for excitable media. Working with SPDE's driven by colored noise, we consider a numerical scheme based on finite differences in time (Euler-Maruyama) and finite elements in space. Motivated by biological considerations, we study numerically the emergence of reentrant patterns in excitable systems such as the Barkley or Mitchell-Schaeffer models.

7.18. Conditional quantile estimation through optimal quantization: theoretical aspects

The following result has been obtained by J. Saracco (Inria CQFD) and I. Charlier (Inria CQFD), D. Paindaveine (ULB).

In this work, we use quantization to construct a nonparametric estimator of conditional quantiles of a scalar response Y given a d -dimensional vector of covariates X . First we focus on the population level and show how optimal quantization of X , which consists in discretizing X by projecting it on an appropriate grid of N points, allows to approximate conditional quantiles of Y given X . We show that this approximation is arbitrarily good as N goes to infinity and provide a rate of convergence for the approximation error. Then we turn to the sample case and define an estimator of conditional quantiles based on quantization ideas. We prove that this estimator is consistent for its fixed- N population counterpart. The results are illustrated on a numerical example.

7.19. Conditional quantile estimation based on optimal quantization: From theory to practice

The following result has been obtained by J. Saracco (Inria CQFD) and I. Charlier (Inria CQFD), D. Paindaveine (ULB).

In this work, small-sample properties of a nonparametric estimator of conditional quantiles based on optimal quantization, that was recently introduced(Charlier et al., JSPI, 2015),are investigated. More precisely, (i) the practical implementation of this estimator is discussed (by proposing in particular a method to properly select the corresponding smoothing parameter, namely the number of quantizers) and (ii) its finite-sample performances are compared to those of classical competitors. Monte Carlo studies reveal that the quantization-based estimator competes well in all cases and sometimes dominates its competitors, particularly when the regression function is quite complex. A real data set is also treated. While the main focus is on the case of a univariate covariate, simulations are also conducted in the bivariate case.

7.20. QuantifQuantile: An R Package for Performing Quantile Regression Through Optimal Quantization

The following result has been obtained by J. Saracco (Inria CQFD) and I. Charlier (Inria CQFD), D. Paindaveine (ULB).

In this work, we describe an R package, called QuantifQuantile, that allows to perform quantization-based quantile regression. In quantile regression, various quantiles of a response variable Y are modelled as functions of covariates (rather than its mean). An important application is the construction of reference curves/surfaces and conditional prediction intervals for Y . Recently, a nonparametric quantile regression method based on the concept of optimal quantization was proposed. This method competes very well with k -nearest neighbor, kernel, and spline methods. We describe also the various functions of the package and provide examples.

7.21. Numerical methods for simulation and optimization of piecewise deterministic Markov processes

This book is focused on theoretical and numerical aspects of simulation and optimization for piecewise deterministic Markov processes (PDMP's). PDMP's have been introduced in the literature by M. Davis as a general class of stochastic hybrid models. They form a family of Markov processes involving deterministic motion punctuated by random jumps. The motion of a PDMP includes both continuous and discrete variables. The continuous state variable represents the physical parameters of the system under consideration. The discrete mode characterizes the regimes of operation of the physical system and/or the environment. The process is defined through three local characteristics, namely the flow describing the deterministic trajectory between two consecutive jumps, the intensity function giving the jump rate and the Markov kernel specifying the post-jump location. A suitable choice of the state space and these local characteristics provides stochastic models covering a large number of problems such as engineering systems, operation research, economics, management science, biology, internet traffic, networks and reliability. The class of PDMP's is thus considered and recognized as a powerful modeling tool for complex systems.

However, surprisingly few works are devoted to the development of numerical methods for PDMP's to solve problems of practical importance such as evaluation and optimization of functionals of the process. The main objective of this book consists in presenting mathematical tools recently developed by the authors to address such problems. This book is not only focused on theoretical aspects such as proof of convergence of the approximation procedures but is also concerned with its applicability to practical problems. The approach we are proposing is general enough to be applied to several application domains. In particular, our results are illustrated by examples from the field of reliability.

Our approximation technique is based on the discretization using quantization of the underlying discrete-time Markov chain given by the post-jump locations and jump times of the PDMP. This strategy enables us to address a large class of numerical problems. In particular, in this book we focus, on the one hand, on the computation of expectation of functionals of PDMP's with applications to the evaluation of service times. On the other hand, we are interested in solving optimal control problems with applications to maintenance optimization.

FLOWERS Project-Team

7. New Results

7.1. Robotic And Computational Models Of Human Development and Cognition

7.1.1. Computational Models Of Information-Seeking, Curiosity And Attention in Humans and Animals

Participants: Manuel Lopes, Pierre-Yves Oudeyer [correspondant], Jacqueline Gottlieb, Adrien Baranes, William Schueller, Sebastien Forestier, Nabil Daddaouda, Nicholas Foley.

This project involves a collaboration between the Flowers team and the Cognitive Neuroscience Lab of J. Gottlieb at Columbia Univ. (NY, US) on the understanding and modeling of mechanisms of curiosity and attention that until now have been little explored in neuroscience, computer science and cognitive robotics. It is organized around the study of the hypothesis that information gain could generate intrinsic reward in the brain (living or artificial), controlling attention and exploration independently from material rewards. The project combines expertise about attention and exploration in the brain and a strong methodological framework for conducting experimentations with monkeys and humans (Gottlieb's lab) together with cognitive modeling of curiosity and learning in the Flowers team.

Such a collaboration paves the way towards a central objective, which is now a central strategic objective of the Flowers team: designing and conducting experiments in animals and humans informed by computational/mathematical theories of information seeking, and allowing to test the predictions of these computational theories.

7.1.1.1. Context

Curiosity can be understood as a family of mechanisms that evolved to allow agents to maximize their knowledge of the useful properties of the world - i.e., the regularities that exist in the world - using active, targeted investigations. In other words, we view curiosity as a decision process that maximizes learning (rather than minimizing uncertainty) and assigns value ("interest") to competing tasks based on their epistemic qualities - i.e., their estimated potential allow discovery and learning about the structure of the world.

Because a curiosity-based system acts in conditions of extreme uncertainty (when the distributions of events may be entirely unknown) there is in general no optimal solution to the question of which exploratory action to take [105], [122], [130]. Therefore we hypothesize that, rather than using a single optimization process as it has been the case in most previous theoretical work [90], curiosity is comprised of a family of mechanisms that include simple heuristics related to novelty/surprise and measures of learning progress over longer time scales [16] [74], [112]. These different components are related to the subject's epistemic state (knowledge and beliefs) and may be integrated with fluctuating weights that vary according to the task context. We will quantitatively characterize this dynamic, multi-dimensional system in the framework of Bayesian Reinforcement Learning, as described below.

Because of its reliance on epistemic currencies, curiosity is also very likely to be sensitive to individual differences in personality and cognitive functions. Humans show well-documented individual differences in curiosity and exploratory drives [103], [129], and rats show individual variation in learning styles and novelty seeking behaviors [88], but the basis of these differences is not understood. We postulate that an important component of this variation is related to differences in working memory capacity and executive control which, by affecting the encoding and retention of information, will impact the individual's assessment of learning, novelty and surprise and ultimately, the value they place on these factors [127], [136], [70], [140]. To start understanding these relationships, about which nothing is known, we will search for correlations between curiosity and measures of working memory and executive control in the population of children we test in our tasks, analyzed from the point of view of a computational model based on Bayesian reinforcement learning.

A final premise guiding our research is that essential elements of curiosity are shared by humans and non-human primates. Human beings have a superior capacity for abstract reasoning and building causal models, which is a prerequisite for sophisticated forms of curiosity such as scientific research. However, if the task is adequately simplified, essential elements of curiosity are also found in monkeys [103], [100] and, with adequate characterization, this species can become a useful model system for understanding the neurophysiological mechanisms.

7.1.1.2. Objectives

Our studies have several highly innovative aspects, both with respect to curiosity and to the traditional research field of each member team.

- Linking curiosity with quantitative theories of learning and decision making: While existing investigations examined curiosity in qualitative, descriptive terms, here we propose a novel approach that integrates quantitative behavioral and neuronal measures with computationally defined theories of Bayesian Reinforcement Learning and decision making.
- Linking curiosity in children and monkeys: While existing investigations examined curiosity in humans, here we propose a novel line of research that coordinates its study in humans and non-human primates. This will address key open questions about differences in curiosity between species, and allow access to its cellular mechanisms.
- Neurophysiology of intrinsic motivation: Whereas virtually all the animal studies of learning and decision making focus on operant tasks (where behavior is shaped by experimenter-determined primary rewards) our studies are among the very first to examine behaviors that are intrinsically motivated by the animals' own learning, beliefs or expectations.
- Neurophysiology of learning and attention: While multiple experiments have explored the single-neuron basis of visual attention in monkeys, all of these studies focused on vision and eye movement control. Our studies are the first to examine the links between attention and learning, which are recognized in psychophysical studies but have been neglected in physiological investigations.
- Computer science: biological basis for artificial exploration: While computer science has proposed and tested many algorithms that can guide intrinsically motivated exploration, our studies are the first to test the biological plausibility of these algorithms.
- Developmental psychology: linking curiosity with development: While it has long been appreciated that children learn selectively from some sources but not others, there has been no systematic investigation of the factors that engender curiosity, or how they depend on cognitive traits.

7.1.1.3. Current results

During the first period of the associated team (2013-2015), we laid the operational foundations of the collaboration resulting in several milestone joint journal articles [110], [90], [84][27], new experimental paradigms for the study of curiosity, and organized a major scientific event: the first international interdisciplinary symposium on information seeking, curiosity and attention (web: <https://openlab-flowers.inria.fr/t/first-interdisciplinary-symposium-on-information-seeking-curiosity-and-attention/21>).

In particular, new results in 2015 include:

7.1.1.4. Eye movements reveal epistemic curiosity in human observers

Saccadic (rapid) eye movements are primary means by which humans and non-human primates sample visual information. However, while saccadic decisions are intensively investigated in instrumental contexts where saccades guide subsequent actions, it is largely unknown how they may be influenced by curiosity – the intrinsic desire to learn. While saccades are sensitive to visual novelty and visual surprise, no study has examined their relation to epistemic curiosity – interest in symbolic, semantic information. To investigate this question, we tracked the eye movements of human observers while they read trivia questions and, after a brief delay, were visually given the answer. We showed that higher curiosity was associated with earlier anticipatory orienting of gaze toward the answer location without changes in other metrics of saccades or fixations, and that these influences were distinct from those produced by variations in confidence and surprise. Across

subjects, the enhancement of anticipatory gaze was correlated with measures of trait curiosity from personality questionnaires. Finally, a machine learning algorithm could predict curiosity in a cross-subject manner, relying primarily on statistical features of the gaze position before the answer onset and independently of covariations in confidence or surprise, suggesting potential practical applications for educational technologies, recommender systems and research in cognitive sciences. We published these results in [27], providing full access to the annotated database allowing readers to reproduce the results. Epistemic curiosity produces specific effects on oculomotor anticipation that can be used to read out curiosity states.

7.1.1.5. *Intrinsically motivated oculomotor exploration guided by uncertainty reduction and conditioned reinforcement in non-human primates*

Intelligent animals have a high degree of curiosity – the intrinsic desire to know – but the mechanisms of curiosity are poorly understood. A key open question pertains to the internal valuation systems that drive curiosity. What are the cognitive and emotional factors that motivate animals to seek information when this is not reinforced by instrumental rewards? Using a novel oculomotor paradigm, combined with reinforcement learning (RL) simulations, we show that monkeys are intrinsically motivated to search for and look at reward-predictive cues, and that their intrinsic motivation is shaped by a desire to reduce uncertainty, a desire to obtain conditioned reinforcement from positive cues, and individual variations in decision strategy and the cognitive costs of acquiring information. The results suggest that free-viewing oculomotor behavior reveals cognitive and emotional factors underlying the curiosity driven sampling of information. [84]

7.1.2. **Computational Models Of Speech Development: the Roles of Active Learning, Curiosity and Self-Organization**

Participants: Pierre-Yves Oudeyer [correspondant], Clement Moulin-Frier, Sébastien Forestier.

7.1.2.1. *Special issue on the cognitive nature of speech sounds*

Together with Jean-Luc Schwartz and Kenneth de Jong, Flowers members Clément Moulin-Frier and Pierre-Yves Oudeyer guest-edited a milestone special issue of the Journal of Phonetics focusing on theories of the cognitive nature of speech sounds, and with a special emphasis on presenting and analyzing a rich series of computational models of speech evolution and acquisition developed in the past years internationally, including models developed by the guest-editors. The editorial of this special issue was published in [35] and the special issue is accessible at: <http://www.sciencedirect.com/science/journal/00954470/53>.

7.1.2.2. *The COSMO model: A Bayesian modeling framework for studying speech communication and the emergence of phonological systems*

(Note: this model was developed while C. Moulin-Frier was at GIPSA Lab, and writing was partly achieved while he was at Inria). While the origin of language remains a somewhat mysterious process, understanding how human language takes specific forms appears to be accessible by the experimental method. Languages, despite their wide variety, display obvious regularities. In this paper, we attempt to derive some properties of phonological systems (the sound systems for human languages) from speech communication principles. The article [33] introduces a model of the cognitive architecture of a communicating agent, called COSMO (for “Communicating about Objects using Sensory–Motor Operations”) that allows a probabilistic expression of the main theoretical trends found in the speech production and perception literature. This enables a computational comparison of these theoretical trends, which helps us to identify the conditions that favor the emergence of linguistic codes. It presents realistic simulations of phonological system emergence showing that COSMO is able to predict the main regularities in vowel, stop consonant and syllable systems in human languages.

7.1.2.3. *The role of self-organization, motivation and curiosity in speech development and evolution*

In the article [34], Oudeyer discusses open scientific challenges for understanding development and evolution of speech forms. Based on the analysis of mathematical models of the origins of speech forms, with a focus on their assumptions, the article studies the fundamental question of how speech can be formed out of non-speech, at both developmental and evolutionary scales. In particular, it emphasizes the importance of embodied self-organization, as well as the role of mechanisms of motivation and active curiosity-driven exploration in speech formation. Finally, it discusses an evolutionary-developmental perspective of the origins of speech.

7.1.2.4. Robotic models of the joint development of speech and tool use

A scientific challenge in developmental and social robotics is to model how autonomous organisms can develop and learn open repertoires of skills in high-dimensional sensorimotor spaces, given limited resources of time and energy. This challenge is important both from the fundamental and application perspectives. First, recent work in robotic modeling of development has shown that it could make decisive contributions to improve our understanding of development in human children, within cognitive sciences [90]. Second, these models are key for enabling future robots to learn new skills through lifelong natural interaction with human users, for example in assistive robotics [124].

In recent years, two strands of work have shown significant advances in the scientific community. On the one hand, algorithmic models of active learning and imitation learning combined with adequately designed properties of robotic bodies have allowed robots to learn how to control an initially unknown high-dimensional body (for example locomotion with a soft material body [73]). On the other hand, other algorithmic models have shown how several social learning mechanisms could allow robots to acquire elements of speech and language [79], allowing them to interact with humans. Yet, these two strands of models have so far mostly remained disconnected, where models of sensorimotor learning were too “low-level” to reach capabilities for language, and models of language acquisition assumed strong language specific machinery limiting their flexibility. Preliminary work has been showing that strong connections are underlying mechanisms of hierarchical sensorimotor learning, artificial curiosity, and language acquisition [125].

Recent robotic modeling work in this direction has shown how mechanisms of active curiosity-driven learning could progressively self-organize developmental stages of increasing complexity in vocal skills sharing many properties with the vocal development of infants [113]. Interestingly, these mechanisms were shown to be exactly the same as those that can allow a robot to discover other parts of its body, and how to interact with external physical objects [120].

In such current models, the vocal agents do not associate sounds to meaning, and do not link vocal production to other forms of action. In other models of language acquisition, one assumes that vocal production is mastered, and hand code the meta-knowledge that sounds should be associated to referents or actions [79]. But understanding what kind of algorithmic mechanisms can explain the smooth transition between the learning of vocal sound production and their use as tools to affect the world is still largely an open question.

The goal of this work is to elaborate and study computational models of curiosity-driven learning that allow flexible learning of skill hierarchies, in particular for learning how to use tools and how to engage in social interaction, following those presented in [120], [73], [118], [113]. The aim is to make steps towards addressing the fundamental question of how speech communication is acquired through embodied interaction, and how it is linked to tool discovery and learning.

A first question that we study in this work is the type of mechanisms that could be used for hierarchical skill learning allowing to manage new task spaces and new action spaces, where the action and task spaces initially given to the robot are continuous and high-dimensional and can be encapsulated as primitive actions to affect newly learnt task spaces.

We presented preliminary results on that question in a poster session [89] of the ICDL/Epirob conference in Providence, RI, USA in August 2015. In this work, we rely more specifically on the R-IAC and SAGG-RIAC series of architectures developed in the Flowers team and we develop different ways to extend those architectures to the learning of several task spaces that can be explored in a hierarchical manner. We describe an interactive task to evaluate different hierarchical learning mechanisms, where a robot has to explore its motor space in order to push an object to different locations. The task can be decomposed into two subtasks where the robot can first explore how to make movements with its hand and then integrate this skill to explore the task of pushing an object.

In the Simplest First strategy, the agent explores successively but with a fixed curriculum the different tasks to learn in the good order: from the simplest one (learning hand movements given motor parameters) to the more complex one (pushing a block with hand movements) that need knowledge about the simpler task to be learned.

In the Top-Down Guidance strategy, the module learning the more complex task (pushing a block with hand movements) gives goals (hand movements) to be reached by the lower-level module (learning hand movements given motor parameters) that will explore for a while to reach that goal before switching to a new given goal.

We also compare our architectures to the control ones where the robot learns directly the not decomposed task, with a competence-based intrinsic motivation (goal babbling) or a fully random motor babbling.

The results show a better exploration for the Top-Down Guidance than the Simplest First hierarchical exploration strategy, and that learning intermediate representations is beneficial in this setup.




Figure 8. Left: Poppy Torso in the V-REP simulator pushing a block. Right: Exploration results of the different strategies.

7.1.3. Learning in Adult-Child and Human-Robot Interaction

Participants: Anna-Lisa Vollmer [correspondant], Pierre-Yves Oudeyer.

Learning in Adult-Child and Human-Robot Interaction

7.1.3.1. The Change of ‘Motionese’ Parameters Depending on Children’s Age.

Two adult-child interaction studies were analyzed with the focus on the parental teaching behavior, in particular on *motionese* parameters (modifications of child-directed movement). In the first cross-sectional study, parental action demonstrations to three groups of 8–11, 12–23 and 24–30 month-olds ($N = 84$ parents) were investigated. The youngest group of participants was investigated longitudinally in the second study ($N = 18$ parents). Together the results suggest that some motionese parameters (motion pauses, velocity, acceleration) persist over different ages while other parameters (action length, roundness and pace) occur predominantly in the younger group and seem to be primarily used to attract infants’ attention on the basis of movement. In contrast, parameters appearing to be more in charge of structuring the action by organizing it in motion pauses seem to persist. We discuss the results in terms of facilitative vs. pedagogical learning in a paper currently under review for the Journal of Experimental Child Psychology.

7.1.3.2. An Alternative to Mapping a Word onto a Concept in Language Acquisition: Pragmatic Frames

According to the mapping metaphor, for a child to learn a word, s/he has to map a word onto a concept of an object/event. We are not convinced that associations can explain word learning, because even though children’s attention is on the objects, they do not necessarily remember the connection of the word with the referent. In this theoretical paper, we propose an alternative to the mapping process that is classically assumed as a mechanism for word learning. Our main point holds that word learning is a task, in which children accomplish a goal in cooperation with a partner. In our approach, we follow Bruner’s (1983) idea and further

specify pragmatic frames as learning units that drive language acquisition and cognitive development. These units consist of a sequence of language and actions that are co-constructed with a partner to achieve a joint goal. We elaborate on this alternative, offer some initial parametrizations of the concept and embed it in the current language learning approaches in a paper currently under review for *Frontiers in Psychology*, section *Cognitive Science*.

7.1.3.3. Meta-Analysis of Pragmatic Frames in Human-Robot Interaction for Learning and Teaching: State-of-the-Art and Perspectives

One of the big challenges in robotics today is to learn from inexperienced human users. Despite tremendous research efforts and advances in human-robot interaction (HRI) and robot learning in the past decades, learning interactions with robots remain brittle and rigidly organized, and often are limited to learning only one single task. In this work, we applied the concept of pragmatic frames known from developmental research in humans in a meta-analysis of current approaches on robot learning. This concept offers a new research perspective in HRI as multiple flexible interaction protocols can be used and learned to teach/learn multiple kinds of skills in long-term recurring social interaction. This perspective, thus, emphasizes teaching as a collaborative achievement of teacher and learner. Our meta-analysis focuses on robot learning from a human teacher with respect to the pragmatic frames they (implicitly) use. We show that while the current approaches offer a variety of different learning and teaching behavior, they all employ highly pre-structured, hard-coded pragmatic frames. Compared to natural human-human interaction, interactions are lacking flexibility and expressiveness, and mostly are hardly viable for being realized with truly naive and uninstructed users. We elaborated an outlook on the future research direction with its relevant key challenges that need to be solved for leveraging pragmatic frames for robot learning. These results have been submitted to the *Frontiers in Neurobotics Journal*.

7.1.3.4. Alignment to the Actions of a Robot

Alignment is a phenomenon observed in human conversation: Dialog partners' behavior converges in many respects. Such alignment has been proposed to be automatic and the basis for communicating successfully. Recent research on human-computer dialog promotes a mediated communicative design account of alignment according to which the extent of alignment is influenced by interlocutors' beliefs about each other. Our work aims at adding to these findings in two ways. a) Our work investigates alignment of manual actions, instead of lexical choice. b) Participants interact with the iCub humanoid robot, instead of an artificial computer dialog system. Our results confirm that alignment also takes place in the domain of actions. We were not able to replicate the results of the original study in general in this setting, but in accordance with its findings, participants with a high questionnaire score for emotional stability and participants who are familiar with robots align their actions more to a robot they believe to be basic than to one they believe to be advanced. Regarding alignment over the course of an interaction, the extent of alignment seems to remain constant, when participants believe the robot to be advanced, but it increases over time, when participants believe the robot to be a basic version. These results were published in [38].

7.1.4. Models of Multimodal Concept Acquisition with Non-Negative Matrix Factorization

Participants: Pierre-Yves Oudeyer, Olivier Mangin [correspondant], David Filliat, Louis Ten Bosch.

In the article [32] we introduced MCA-NMF, a computational model of the acquisition of multi-modal concepts by an agent grounded in its environment. More precisely our model finds patterns in multimodal sensor input that characterize associations across modalities (speech utterances, images and motion). We propose this computational model as an answer to the question of how some class of concepts can be learnt. In addition, the model provides a way of defining such a class of plausibly learnable concepts. We detail why the multimodal nature of perception is essential to reduce the ambiguity of learnt concepts as well as to communicate about them through speech. We then present a set of experiments that demonstrate the learning of such concepts from real non-symbolic data consisting of speech sounds, images, and motions. Finally we consider structure in perceptual signals and demonstrate that a detailed knowledge of this structure, named compositional understanding can emerge from, instead of being a prerequisite of, global understanding. An open-source implementation of the MCA-NMF learner as well as scripts and associated experimental data to reproduce the experiments are publicly available.

The python code and datasets allowing to reproduce these experiments and results are available at: <https://github.com/omangin/multimodal>.

7.1.5. Models of Self-organization of lexical conventions: the role of Active Learning in Naming Games

Participants: William Schueller [correspondant], Pierre-Yves Oudeyer.

Our work focuses on the Naming Games framework [135], meant to simulate lexicon evolution in a population from interactions at the individual level. A quite diverse subset of the possible scenarios and algorithms has already been studied, and those do lead to the self-organization of a shared lexicon (understood as associations between meanings and words). However, high values for some parameters (population size, number of possible words and/or meanings that can be referred to) can lead to really slow dynamics. Following the introductory work done in [119], we introduced a new measure of vocabulary evolution based on information theory, as well as various active learning mechanisms in the Naming Games framework allowing the agents to choose what they talk about according to their past. We showed that it improves convergence dynamics in the studied scenarios and parameter ranges. Active learning mechanisms use the confidence an agent has on its own vocabulary (is it already widely used in the population or not?) to choose between exploring new associations (growing vocabulary) or strengthening already existing ones (spreading its own associations to other agents). This was presented at the ICDL/Epirob conference in Providence, RI, USA in August 2015 [59].

A follow-up to this work consisted of changing slightly the base algorithms, allowing agents to select what they want the others to talk about instead of selecting what they would talk about (hearer's choice scenario, the original one being speaker's choice scenario). In the class of algorithms used, with active learning, it leads to faster convergence, with increased robustness to change in parameter values.




Figure 9. Interaction processes in both active scenarios considered in our work. Beforehand, two individuals have been randomly selected among a population, designated as speaker (S) and hearer (H). Speaker's choice: 1. S chooses a topic, 2. S checks its vocabulary to find or invent an associated word, 3. S utters the word, 4. H guesses the intended meaning, 5. S indicates the intended meaning. Hearer's choice: 1. H chooses a topic, 2. H indicates the intended meaning, 3. S checks its vocabulary to find or invent an associated word, 4. S utters the word, 5. H checks its vocabulary for a meaning associated to the uttered word. In both cases, if all meanings match, the interaction is considered a success, otherwise a failure. After the process, both agents can update their vocabularies to take the interaction into account.

All the simulations can be easily rerun using the provided code and explanatory notebooks on <https://github.com/flowersteam/naminggamesal>.

7.2. Life-Long Robot Learning And Development Of Motor And Social Skills

7.2.1. Uncalibrated BCI

Participants: Manuel Lopes [correspondant], Pierre-Yves Oudeyer, Jonathan Grizou, Inaki Iturrate, Luis Montesano.

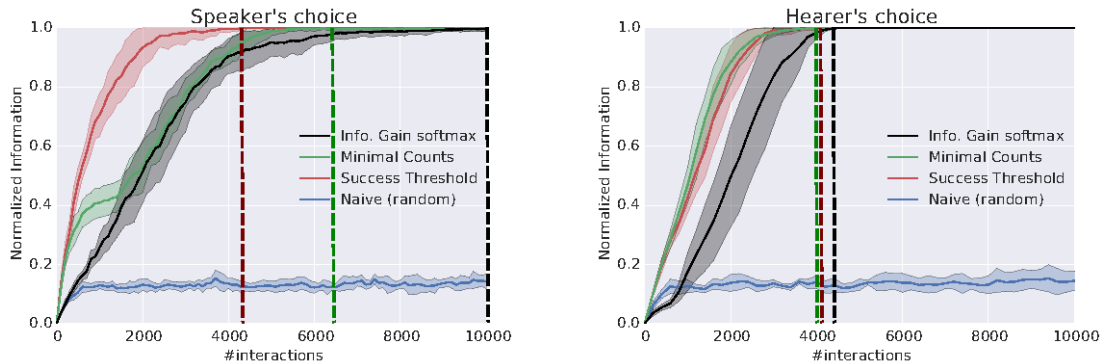


Figure 10. Strategy comparisons in both active scenarios. The measure indicates level of convergence towards a population-wide shared vocabulary (a value of 1 means every agent has the exact same vocabulary). Naive (random, no active learning) strategy converges slowly. Hearer’s choice policy is more efficient for all active learning strategies. Last 5% of information are acquired slower when the speaker is choosing. Vertical lines show full convergence time for each strategy. (number of meanings = number of words = population size = 20, averaged over 8 trials)

We developed an new approach for self-calibration BCI for reaching tasks using error-related potentials. The proposed method exploits task constraints to simultaneously calibrate the decoder and control the device, by using a robust likelihood function and an ad-hoc planner to cope with the large uncertainty resulting from the unknown task and decoder. The method has been evaluated in closed-loop online experiments with 8 users using a previously proposed BCI protocol for reaching tasks over a grid. The results show that it is possible to have a usable BCI control from the beginning of the experiment without any prior calibration. Furthermore, comparisons with simulations and previous results obtained using standard calibration hint that both the quality of recorded signals and the performance of the system were comparable to those obtained with a standard calibration approach. [30]

7.2.2. Learning from Demonstration

Participants: Manuel Lopes, Thibaut Munzer [correspondant], Marc Toussaint, Li Wang Wu, Yoan Mollard, Andrea Baisero, Bilal Piot, Matthieu Geist, Olivier Pietquin.

Learning from Demonstration

7.2.2.1. Relational Activity Processes for Modeling Concurrent Cooperation

In multi-agent domains, human-robot collaboration domains, or single-robot manipulation with multiple end-effectors, the activities of the involved parties are naturally concurrent. Such domains are also naturally relational as they involve multiple objects, multiple agents, and models should generalize over objects and agents. We propose a novel formalization of relational concurrent activity processes that allows us to transfer methods from standard (relational) MDPs, such as Monte-Carlo planning and learning from demonstration, to concurrent cooperation domains. We formally compare the formulation to previous propositional models of concurrent decision making and demonstrate the planning and learning from demonstration methods on a real-world human-robot assembly task.

7.2.2.2. Interactive Learning

In paper [56] we consider that robot programming can be made more efficient, precise and intuitive if we leverage the advantages of complementary approaches such as learning from demonstration, learning from feedback and knowledge transfer. We designed a system that, starting from low-level demonstrations of assembly tasks, is able to extract a high-level relational plan of the task. A graphical user interface (GUI)

allows then the user to iteratively correct the acquired knowledge by refining high-level plans, and low-level geometrical knowledge of the task. A final process allows to reuse high-level task knowledge for similar tasks in a transfer learning fashion. We conducted a user study with 14 participants asked to program assembly tasks of small furniture (chair and bench) to validate this approach. The results showed that this combination of approaches leads to a faster programming phase, more precise than just demonstrations, and more intuitive than just through a GUI.

7.2.2.3. Inverse Reinforcement Learning in Relational Domains

We introduced a first approach to the Inverse Reinforcement Learning (IRL) problem in relational domains. IRL has been used to recover a more compact representation of the expert policy leading to better generalize among different contexts. Relational learning allows one to represent problems with a varying number of objects (potentially infinite), thus providing more generalizable representations of problems and skills. We show how these different formalisms can be combined by modifying an IRL algorithm (Cascaded Supervised IRL) such that it handles relational domains. Our results indicate that we can recover rewards from expert data using only partial knowledge about the dynamics of the environment. We evaluate our algorithm in several tasks and study the impact of several experimental conditions such as: the number of demonstrations, knowledge about the dynamics, transfer among varying dimensions of a problem, and changing dynamics. This was published in [49]

7.2.3. A Unified Model for Regression

Regression is the process of learning relationships between inputs and continuous outputs from example data, which enables predictions for novel inputs. Regression lies at the heart of imitation learning, and value function approximation for reinforcement learning. In [37], we provide a novel perspective on regression, by distinguishing rigorously between the models and representations assumed in regression, and the algorithms used to train the parameters of these models. A rather surprising insight is that many regression algorithms⁰ use very similar models; in fact, we show that the algorithm-specific models are *all* special cases of a “unified model”. This perspective clearly separates between representations and algorithms, and allows for a modular exchange between them, for instance in the context of evolutionary optimization.

7.2.4. Multiple Virtual Guides

In co-manipulation, humans and robots solve manipulation tasks together. Virtual guides are important tools for co-manipulation, as they constrain the movement of the robot to avoid undesirable effects, such as collisions with the environment. Defining virtual guides is often a laborious task requiring expert knowledge. This restricts the usefulness of virtual guides in environments where new tasks may need to be solved, or where multiple tasks need to be solved sequentially, but in an unknown order.

To this end, we have proposed a framework for *multiple probabilistic virtual guides*, and demonstrated a concrete implementation of such guides using kinesthetic teaching and Gaussian mixture models [57], [58]. Our approach enables non-expert users to design virtual guides through demonstration. Also, they may demonstrate novel guides, even if already known guides are active. Finally, users are able to intuitively select the appropriate guide from a set of guides through physical interaction with the robot.

7.2.5. Legible Motion

Participants: Manuel Lopes, Baptiste Busch [correspondant], Jonathan Grizou, Freek Stulp.

⁰Locally Weighted Regression, Receptive Field Weighted Regression, Locally Weighted Projection Regression, Gaussian Mixture Regression, Model Trees, Radial Basis Function Networks, Kernel Ridge Regression, Gaussian Process Regression, Support Vector Regression, Random Features Regularized Least Squares, Sparse Spectrum Gaussian Process Regr., Regression Trees, Extreme Learning Machines.

In a human-robot collaboration context, understanding and anticipating the robot intentions ease the completion of a joint-task. Whereas previous work has sought to explicitly optimize the legibility of behavior, we investigate legibility as a property that arises automatically from general requirements on the efficiency and robustness of joint human-robot task completion. We propose an optimization algorithm, based on policy improvement, that brings out the most legible robot's trajectories during the interaction (cf. Figure 11). The conducted user study highlights that humans become better at predicting sooner the robot's intentions. This leads to faster and more robust overall task completion. This work have been published to IROS 2015[60] and was submitted to the International Journal of Social Robotics under the special issue: Towards a framework for Joint Action.




Figure 11. Illustration of the experimental setup. The robots aims for a button and press it. The human predict the target and is instructed to quickly press the button of the same color when sufficiently confident about this prediction.

7.2.6. Demonstrator of human-robot interface for teaching a collaborative task in the context of assistive robotics

Participants: Pierre Rouanet [correspondant], Yoan Mollard, Thibaut Munzer, Baptiste Busch, Manuel Lopes, Pierre-Yves Oudeyer.

In the context of the Roméo 2 project, we have developed a demonstrator of a human-robot interface designed for non-expert users. It allows them to teach a new collaborative task to a robot through simple and intuitive interactions. It is based on the approach of inverse reinforcement learning in relational domains described above.

The context of the demonstrator is assistive robotics where typically an elderly person wants to teach a robot (we use Baxter in this case) how it can help him to prepare a meal. For instance, the user will show the robot that first he wants the robot to hold its bowl and that he stirs it. Then, the robot should put the bowl on a plate. Then, the user will teach the robot that he wants the robot to grab a glass and put it on the right of the bowl...

7.2.7. Diversity-driven curiosity-driven learning and transfer learning

Participants: Fabien Benureau [correspondant], Pierre-Yves Oudeyer.

7.2.7.1. Diversity-driven selection of exploration strategies in multi-armed bandits

In [40], we considered a scenario where an agent has multiple available strategies to explore an unknown environment. For each new interaction with the environment, the agent must select which exploration strategy to use. We provide a new strategy-agnostic method that treat the situation as a Multi-Armed Bandits problem where the reward signal is the diversity of effects that each strategy produces. We test the method empirically on a simulated planar robotic arm, and establish that the method is both able discriminate between strategies of dissimilar quality, even when the differences are tenuous, and that the resulting performance is competitive with the best fixed mixture of strategies.

7.2.7.2. Behavioral Diversity Generation in Autonomous Exploration Through Reuse of Past Experience

The production of behavioral diversity—producing a diversity of effects—is an essential strategy for robots exploring the world when facing situations where interaction possibilities are unknown or non-obvious. It allows to discover new aspects of the environment that cannot be inferred or deduced from available knowledge. However, creating behavioral diversity in situations where it is the most crucial, i.e. new and unknown ones, is far from trivial. In particular in large and redundant sensorimotor spaces, some effects can typically only be produced by a few number of specific motor commands. We introduced a method to create behavioral diversity by re-enacting past experiences, along with a measure that quantifies this diversity. We showed that our method is robust to morphological and representation changes, that it can learn how to interact with an object by reusing experience from another and how scaffolding behaviors can emerge by simply switching the attention of the robot to different parts of the environment. Finally, we showed that the method can robustly use simulated experiences and crude cognitive models to provide behavioural diversity in the real world. This result are under review.

7.3. Autonomous And Social Perceptual Learning

Participants: David Filliat [correspondant], Freek Stulp, Celine Craye, Yuxin Chen, Clement Masson, Adrien Matricon.

7.3.1. Incremental Learning of Object-Based Visual Saliency

Searching for objects in an indoor environment can be drastically improved if a task-specific visual saliency is available. We describe a method to learn such an object-based visual saliency in an intrinsically motivated way using an environment exploration mechanism. We first define saliency in a geometrical manner and use this definition to discover salient elements given an attentive but costly observation of the environment. These elements are used to train a fast classifier that predicts salient objects given large-scale visual features. In order to get a better and faster learning, we use intrinsic motivation to drive our observation selection, based on uncertainty and novelty detection. Our approach has been tested on RGB-D images, is real-time, and outperforms several state-of-the-art methods in the case of indoor object detection. We published these results in two conferences [43], [42]

7.3.2. Cross-situational noun and adjective learning in an interactive scenario

Learning word meanings during natural interaction with a human faces noise and ambiguity that can be solved by analysing regularities across different situations. We propose a model of this cross-situational learning capacity and apply it to learning nouns and adjectives from noisy and ambiguous speeches and continuous visual input. This model uses two different strategy: a statistical filtering to remove noise in the speech part and the Non Negative Matrix Factorization algorithm to discover word-meaning in the visual domain. We present experiments on learning object names and color names showing the performance of the model in real interactions with humans, dealing in particular with strong noise in the speech recognition. We published these results in a conference paper [41]

7.3.3. Learning representation with gated auto-encoders

We investigated algorithms that would be able to learn relevant visual or multi-modal features from data recorded while the robot performed some task. Representation learning is a currently very active research field, mainly focusing on deep-learning, which investigates how to compute more meaningful features from the raw high dimensional input data, providing a more abstract representation from which it should be easier to make decision or deduction (e.g classification, prediction, control, reinforcement learning). In the context of robotics, it is notably interesting to apply representation learning in a temporal and multi-modal approach exploiting vision and proprioception so as to be able to find feature that are relevant for building models of the robot itself and of its actions and their effect on the environment. Among the many existing approaches, we decided to explore the use of gated auto-encoders, a particular kind of neural networks including multiplicative connections, as they seem well adapted to this problem. Preliminary experimentations have been carried out with gated auto-encoders to learn transformations between two images. We observed that Gated Auto-Encoders (GAE) can successfully find compact representations of simple transformations such as translations, rotation or scaling between two small images. This is however not directly scalable to realistic images such as ones acquired by a robot's camera because of the number of parameters, memory size and computational power it would require (unless drastically downsampling the image which induces sensible loss of information). In addition, the transformation taking an image to the next one can be the combination of transformations due to the movement of several object in the field of view, composed with the global movement of the camera. This induces the existence of an exponential number of possible transformations to model, for which the basic GAE architecture is not suited. To tackle both issue, we are developing a convolutional architectures inspired from Convolutional Neural Networks (CNNs) that provide different modelisations for different parts of the image, which might be usefull to model combinations of transformations. Our Convolutional Gated Auto-Encoder is designed to perform generic feature learning in an unsupervised way (while most CNNs are trained in a supervised fasion) and we are currently testing it on realistic image sequences. We plan to extend this architecture to find relations between modalities as, for instance, proprioceptive information and its evolution could be used to predict the next visual features. Similarly, proprioceptive information could be used as a supervising signal to learn visual features.

7.3.4. Learning models by minimizing complexity

In machine learning, it is commonly assumed that simpler models have better chances at generalizing to new, unseen data. Following this principle, we developped an algorithm relying on minimization of a given complexity measure to build a collection of models which jointly explain the observation of the training datapoints. The resulting collection is composed of as few models as possible, each using as few dimensions as possible and each as regular as possible. As of now, each model is a multivariate polynomial, with the complexity of a polynomial of degree N in d variables being $N*d+1$. The complexity of the collection is the sum of the complexity of all its models. The algorithm starts by associating each datapoint to a local model of complexity 1 (degree 0, no variables), then models are iteratively merged into models of higher complexity, as long as those merges don't increase the complexity of the collection and as long as the resulting models stay within a certain distance of their associated datapoints. We applied this algorithm to the problem of inverse dynamics, which we studied in simulation. For a given robot, torques needed to compensate gravity at equilibrium are entirely determined by the values of its joint angles. As it is common that robots actually perform only low-dimensional tasks, and do not explore their full state space during normal operation, we would like the complexity of our models to mirror the structure of the task. When the task was expressed in the joint space, we got satisfying results on that point, and got good predictions for unseen datapoints. When the task was expressend in end-effector position, it turned out to be impossible to learn the underlying manifolds because a given end-effector position can correspond to various joint configurations, and thus to various torques, making it impossible to predict those torques from the end-effector position alone. We are currently working on applying this model to data generated by an exploration algorithm on a robot arm manipulating objects.

7.4. Applications for Robotic myoelectric prostheses: co-adaptation algorithms and design of a 3D printed robotic arm prosthesis

Participants: Pierre-Yves Oudeyer [correspondant], Manuel Lopes, Joel Ortiz, Mathilde Couraud, Aymar de Rugy, Daniel Cattaert, Florent Paclet.

Together with the Hybrid team at INCIA, CNRS, the Flowers team continued to work on establishing the foundations of a long-term project related to the design and study of myoelectric robotic prosthesis. The ultimate goal of this project is to enable an amputee to produce natural movements with a robotic prosthetic arm (open-source, cheap, easily reconfigurable, and that can learn the particularities/preferences of each user). This will be achieved by 1) using the natural mapping between neural (muscle) activity and limb movements in healthy users, 2) developing a low-cost, modular robotic prosthetic arm and 3) enabling the user and the prosthesis to co-adapt to each other, using machine learning and error signals from the brain, with incremental learning algorithms inspired from the field of developmental and human-robot interaction. In particular, in 2015 two lines of work were achieved, concerning two important scientific challenges, and in the context of a PEPS CNRS project:

First, an experimental setup was designed to allow fast prototyping of 3D printed robotic prostheses (internship of Joel Ortiz). This work was based on the use of the Poppy open-source modular platform, and resulted in a functional prototype. Several video demonstrations are available at: <https://forum.poppy-project.org/t/real-time-control-of-a-prosthetic-robotic-arm-poppy-with-muscle-activities/1656>.

Second, first versions of co-adaptation algorithms were designed, implemented and tested with human subjects, based on the combination of advanced models of the arm biomechanics and incremental learning algorithms (internship of Mathilde Couraud). An article is under preparation.

7.5. Applications for Educational Technologies

7.5.1. KidLearn

Participants: Manuel Lopes [correspondant], Pierre-Yves Oudeyer, Didier Roy, Benjamin Clement.

Kidlearn is a research project studying how machine learning can be applied to intelligent tutoring systems. It aims at developing methodologies and software which adaptively personalize sequences of learning activities to the particularities of each individual student. Our systems aim at proposing to the student the right activity at the right time, maximizing concurrently his learning progress and its motivation. In addition to contributing to the efficiency of learning and motivation, the approach is also made to reduce the time needed to design ITS systems.

Intelligent Tutoring System (ITS) are computer environments designed to guide students in their learning. Through the proposal of different activities, it provides teaching experience, guidance and feedback to improve learning. The FLOWERS team has developed several computational models of artificial curiosity and intrinsic motivation based on research on psychology that might have a great impact for ITS. Results showed that activities with intermediate levels of complexity, neither too easy nor too difficult but just a little more difficult than the current level, provide better teaching experiences. The system is based on the combination of three approaches. First, it leverages Flowers team's recent models of computational models of artificial curiosity and intrinsic motivation based on research in psychology and neuroscience. Second, it uses state-of-the-art Multi-Arm Bandit (MAB) techniques to efficiently manage the exploration/exploitation challenge of this optimization process. Third, it leverages expert knowledge to constrain and bootstrap initial exploration of the MAB, while requiring only coarse guidance information of the expert and allowing the system to deal with didactic gaps in its knowledge. In 2014, we have run a second pilot experiment in elementary schools of Région Aquitaine, where 7-8 year old kids could learn elements of mathematics thanks to an educational software that presented the right exercises at the right time to maximize learning progress. [29]

7.5.2. Poppy System

Participants: Matthieu Lapeyre [correspondant], Nicolas Rabault, Pierre Rouanet, Pierre-Yves Oudeyer.

In the Poppy project we are working on the Poppy System which is a new modular and open-source robotic architecture. It is design to help people create and build custom robots. It permits, in a similar approach as Lego, building robots or smart objects using standardized elements.

Poppy System is an unified system where each essential robotic components (actuators, sensors, ...) is an independant module, connected with other through standardized interfaces.

- Unified mechanical interfaces which simplifies the assembly process and the design of 3D printable parts.
- Unified communication between elements using the same connector and bus for each module.
- Unified software makes it easy to program each module independantly.

The current Poppy robots (Humanoid, Torso, Ergo) will be updated using this novel architecture.

Our ambition is to create an ecosystem around this system so communities can develop custom modules, following the Poppy System standards, which can be compatible with all other Poppy robots.

7.5.3. Poppy Education

Participants: Pierre-Yves Oudeyer [correspondant], Didier Roy, Théo Segonds, Stéphanie Noirpoudre, Marie Demangeat, Thibault Desprez, Matthieu Lapeyre, Pierre Rouanet, Nicolas Rabault.

Poppy Education aims to create, evaluate and disseminate pedagogical kits “turnkey solutions” complete, open-source and low cost, for teaching computer science and robotics. It is designed to help young people to take ownership with concepts and technologies of the digital world, and provide the tools they need to allow them to become actors of this world, with a considerable socio-economic potential. It is carried out in collaboration with teachers and several official french structures (French National Education, Highschools, engineers schools, ...). For secondary education and higher education, scientific literacy centers, Fablabs.

The Poppy robotic platform used in the project is free hardware and software, printed in 3D, and is intended primarily for:

- learning of computer science and robotics,
- introduction to digital manufacturing (3D printing ...)
- initiation to the integration of IT in physical objects in humanoid robotics, mechatronics.
- artistic activities.

Educational sectors covered by the project are mainly: Enseignement d’exploration ICN en seconde, enseignement ISN en terminale S et bientôt en 1ère , filière STI2D, MPS seconde.

Users and their needs are placed at the center of this project. The pedagogical tools of the project are being created directly with them and evaluated in real life by experiments. Poppy Education is based on the robotic platform poppy, from which it is possible to construct different robots, including:

- Poppy Humanoid is a robust and complete robotics platform designed for genuine experiments in the real world and can be adapted to specific user needs.
- Poppy Torso is a variant of Poppy Humanoid. It is a torso humanoid robot that can be easily installed on a table.
- Poppy Ergo Jr is a robotic arm. Solid and inexpensive, it is perfect to be used in class. Poppy robots are easy to program. Different options are possible based on students level and teaching objectives :
- Pixl is a board who manage power and communication between a raspberry pi and robotis XL320 low cost motors. We use this bord for all our low cost robots.
- Python. Directly from a web browser, using Ipython notebooks (an interactive terminal, in a web interface for the Python Programming Language).
- Snap. The visual programming system Snap, which is a variant of Scratch. Its features allow a thorough introduction of IT.
- C ++, Java, Matlab, Ruby, Javascript, etc. thanks to a REST API that allows you to send commands and receive information from the robot with simple HTTP requests.

Poppy Humanoid, Torso and Ergo robots can be simulated with the free simulator V-REP. It is possible in the classroom to work on the simulated model and then allow students to run their program on the physical robot.

Experimentations have began to be setup in 10 high-schools of Region Aquitaine, and 3 university level institutions: Lycée Camille Jullian (Bordeaux), Lycée Victor Louis (Talence), Lycée Saint Genès (Talence), Lycée François Mauriac (Bordeaux), Lycée Jean Moulin (Langon), Lycée des Graves (Gradignan), Lycée Sud Medoc (Le Taillan Medoc), Lycée Alfred Kastler (Talence), Lycée Raoul Follereau (Nevers), Aérocampus Auqitaine, ENSEIRB/IPB, ENSAM Talence.

7.5.4. IniRobot : Education and Thymio II Project (partnership with EPFL)

Participants: Didier Roy [correspondant], Pierre-Yves Oudeyer.

IniRobot Project consists to produce and diffuse a pedagogical kit for teachers and animators, to help to train them directly or by the way of external structures. The aim of the kit is to initiate children to computer science and robotics. The kit provides a micro-world for learning, and takes an enquiry-based educational approach, where kids are led to construct their understanding through practicing an active investigation methodology within teams. It is based on the use of the Thymio II robotic platform. More details about this projects were published in RIE 2015 [50], which presents the detailed pedagogical objectives and a first measure of results showing that children acquired several robotics-related concepts. See also <https://dm1r.inria.fr/inirobot> or <http://www.inirobot.fr>. The project is carried out in main collaboration with the LSRO Laboratory from EPFL (Lausanne) and others collaborations with French National Education/Rectorat d' Aquitaine.

Deployment: After 16 months of activity, IniRobot is used by about 900 adults and 8000 children in 35 cities of France. Example of action in university: MEEF teacher training for the hope of Aquitaine. Example of action in school: training of all Gironde Pedagogical ICT Advisors, covering nearly 1000 schools. Example of action in the extracurricular time: training 82 facilitators TAP cities of Talence, Pessac, Lille, ..., CDC Gates of inter-seas. Example of national action: Training of the digital mediators of the 8 Inria centers.

GEOSTAT Project-Team

7. New Results

7.1. Super-resolution, multiscale data fusion and complex dynamics in Earth Observation and Universe Sciences

Participants: Hussein Yahia, Nicolas Brodu, Guillaume Attuel, Sylvain Bontemps, Nicola Schneider, Camila Artana, Dharmendra Singh, Joel Sudre, Véronique Garçon, Christine Provost, Anass El Aouni, Oriol Pont, Khalid Daoudi, Ayoub Tamim, Akankhsa Garg, Frédéric Frappart, Luc Bourrel.

In these thematics the following research is started or continued:

- Super-resolution and data fusion in Earth Observation. Important results obtained in validation either in ocean dynamics or partial pCO_2 pressures in ocean/atmosphere exchanges, coastal upwelling.
- Development of a new super-resolution model for multispectral images, demonstration on MODIS (NASA) and Sentinel-2 (ESA) data.
- Adaptive optics.
- Starting of a strong collaboration with Labroatoire d'Astrophysique de Bordeaux on the dynamics of galactic clouds.
- Supervised classification of ground terrain through multispectral imagery (with OPTIC associated team).
- Anomaly detection in SAR images, application to flood monitoring in Equator.
- Starting of a project on dune monitoring.

Publications: [21], [19], [29], [22], [25], [18] [A. Tamim's PhD HAL link](#), [IEEE TGRS article on AO \[21\]](#).

7.2. Characterization of underlying stochastic dynamic of the cardiac muscle under fibrillation: singularity analysis and modeling

Participants: Guillaume Attuel, Binbin Xu, Oriol Pont, Hussein Yahia.

Signals of heart electrical activity obtained through invasive measurements show properties not compatible with the purely excitable nature of cellular dynamics. We have developed a synaptic perturbation model of that dynamics showing good properties. Perturbations propagate an inter-cell desynchronization formally like diffusion-coupled chaotic maps. The model enters the universality class of directed propagative fronts of the type random branching or directed polymer in a disordered medium in 1+1D and pinning-depinning contact lines in 1+2D. In the continuum limit, the universality class is supposed to be the one of KPZ (Kardar Parisi Zhang) or VM (Voter Model). This is a change in paradigm for the description of cardiac dynamics. We make use of this hypothesis to characterize precisely the state of the substratum through appropriate signal analysis, with the goal of being able to distinguish between different states or types in the pathology. We are involved in a technological transfert on this activity since summer 2015.

Publication: [26].

7.3. Classification of Cardiac Arrhythmia in vitro based on Multivariate Complexity Analysis

Participant: Binbin Xu.

Background: The animal models (in vitro or in vivo) provide an excellent tool to study heart diseases, among them the arrhythmia remains one of the most active research subjects. It can be induced or treated by drugs, electrical stimulation, hypothermia etc. **Problems:** However, the inducing or treating effects in cardiac culture often happened long after the initial applications or in some relatively short time windows. So, it is necessary to capture and classify rapidly the signal change. Human-assisted monitoring is time-consuming and less efficient. An automatic classification method for real-time use would be useful and necessary. **Methods:** Since electrocardiological signals are features by repetitive or similar patterns reflecting the intrinsic information about the patient (or culture), analyzing these patterns could help not only to monitor the status's change but also to evaluate/explore the physiologic control mechanisms. Methods based on complexity analysis are of considerable interest in this case. **Aims:** Compare different complexity analysis methods in order to find the most appropriate ones to discriminate the normal cardiac signals from arrhythmic ones acquired from a cardiac cell culture in vitro. The selected features are then used by a SVM classifier.

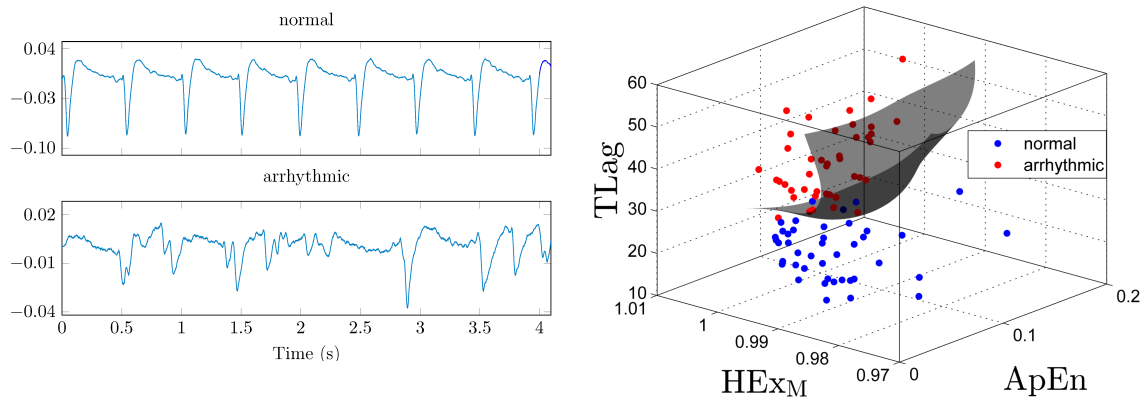


Figure 1. Left: Electrical field potential (EFP) of in vitro cardiac culture on Multielectrodes Array; Right: SVM classification of normal and arrhythmic EFP signals.

Results: Among the six complexity analysis methods, Time Lagging (TLag) method allowed obtaining the best discrimination index (normal vs. arrhythmic, p -value, $9e-23$). The proposed Modified Hurst Exponent (HEX_M) showed better performance than original Hurst Exponent with well-improved p -value (from 0.019 to $2e-9$). The Approximate Entropy (ApEn), Sample Entropy (SampEn) and Detrended Fluctuation Analysis gave good discrimination ratio but with larger p -values (at order 10^{-3}). Combination of TLag, HEX_M and ApEn can provide a more robust classifier and allow monitoring and classifying in an automatic way the electrical activities' changes in the cardiac cultures.

Publication: [28].

7.4. Classification of Cardiac Arrhythmia in vitro based on Multivariate Complexity Analysis

Participant: Binbin Xu.

Physiological signals are temporal series containing a lot of information, and their analysis (either for diagnosis or evolution monitoring) necessitates tools that take into account their intrinsic characteristics, notably in terms of unpredictability and high number of parameters. Methodologies coming from chaotic and nonlinear dynamical systems contain some useful building blocks in that perspective, and allow a qualitative link with phenomenological and bio-inspired models. The objective of this work is to introduce some methods in nonlinear dynamics useful for the processing of these types of signals. An application of these tools is

illustrated in the processing of potential electrical fields acquired from in vitro culture cells on newborn rats. Both normal (regular contraction of cells) and arrhythmic (disordered contractions) cases are contemplated.




Figure 2. CPE period bifurcation diagram. Cells are stimulated by M1 electrode during 5 minutes with an impulsion train of $200\mu\text{m}$ and frequency 100Hz. Three particular phenomena in cell behaviour: **A** ($t = 1$ hour) chaotic state, **B** ($t = 3$ hour) and **C** ($t = 3.5$ hour) period doubling phase, **D** ($t = 4$ hour) regular and stable rhythm.

The bifurcation diagram is an example of a tool that can be used in the temporal analysis of an experimental system.

Publication: [32].

7.5. Nonlinear trend removal should be carefully performed in heart rate variability analysis

Participants: Binbin Xu, Oriol Pont, Hussein Yahia, Rémi Dubois.

Background : In Heart rate variability analysis, the rate-rate time series suffer often from aperiodic non-stationarity, presence of ectopic beats etc. It would be hard to extract helpful information from the original signals. **Problem :** Trend removal methods are commonly practiced to reduce the influence of the low frequency and aperiodic non-stationary in RR data. This can unfortunately affect the signal and make the analysis on detrended data less appropriate. **Objective :** Investigate the detrending effect (linear & nonlinear) in temporal / nonlinear analysis of heart rate variability of long-term RR data (in normal sinus rhythm, atrial fibrillation, congestive heart failure and ventricular premature arrhythmia conditions). **Methods :** Temporal method : standard measure SDNN; Nonlinear methods : multi-scale Fractal Dimension (FD), Detrended Fluctuation Analysis (DFA) & Sample Entropy (SampEn) analysis.

Results : The linear detrending affects little the global characteristics of the RR data, either in temporal analysis or in nonlinear complexity analysis. After linear detrending, the SDNNs are just slightly shifted and all distributions are well preserved. The cross-scale complexity remained almost the same as the ones for original RR data or correlated. Nonlinear detrending changed not only the SDNNs distribution, but also the order among different types of RR data. After this processing, the SDNN became indistinguishable between SDNN for normal sinus rhythm and ventricular premature beats. Different RR data has different complexity signature. Nonlinear detrending made the all RR data to be similar, in terms of complexity. It is thus impossible

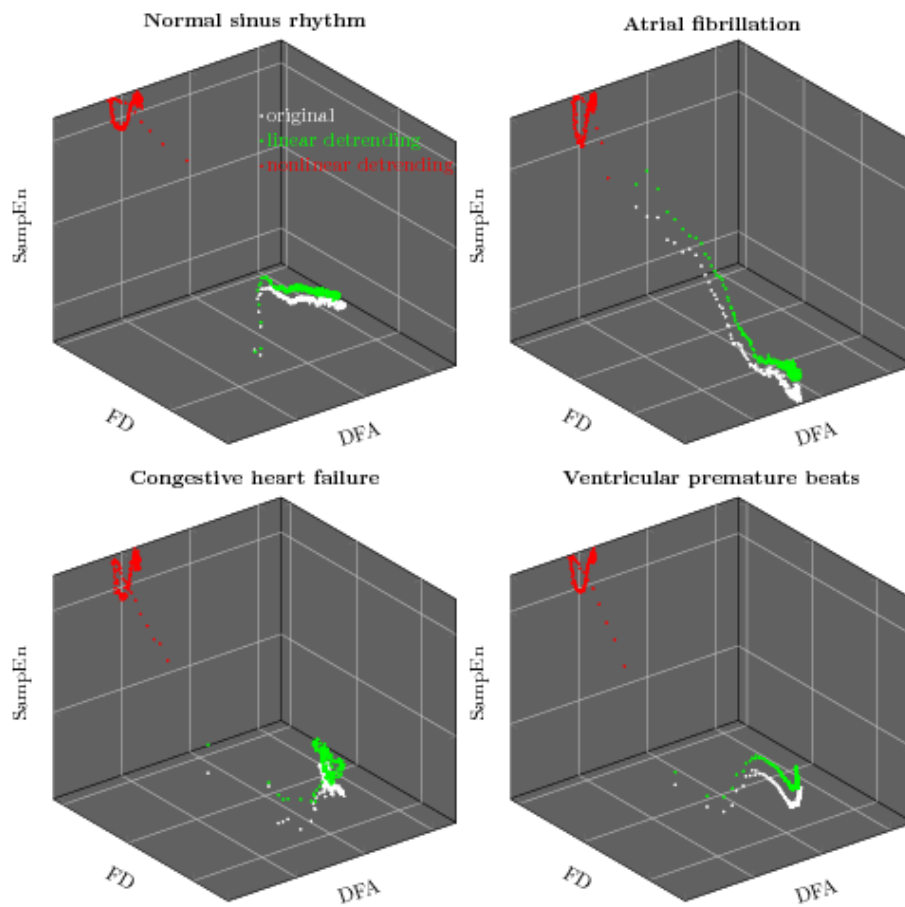


Figure 3. Complexity Space based on FD, DFA & SampEn, for RR data in NSR, AF, CHF & VPB conditions.

to distinguish them. The FD showed that nonlinearly detrended RR data has a dimension close to 2, the exponent from DFA is close to zero and SampEn is larger than 1.5 – these complexity values are very close to those for random signal. **Conclusions** : Pre-processing by linear detrending can be performed on RR data, which has little influence on the corresponding analysis. Nonlinear detrending could be harmful and it is not advisable to use this type of pre-processing. Exceptions do exist, but only combined with other appropriate techniques to avoid complete change of the signal's intrinsic dynamics.

One submitted publication.

7.6. Quantification of Heart's Recover by Multiscale Complexity Analysis of Heart Rate: a Validation Study

Participants: Binbin Xu, Hussein Yahia, Rémi Dubois.

Background : Heart rate analysis is the common analysis of heart's function. After the drug treatment of cardiac arrhythmia, the heart rate looks like the same as the group with normal sinus rhythm.




Figure 4. Heart rate RR time series in three cases : (1) arrhythmia, pre-treatment; (2) arrhythmia suppressed, post-treatment; (3) normal healthy group.

Problem : However, the visibly "same dynamics" for post-treatment & normal group does not reflect the true intrinsic dynamics of the heart. **Methods** : Using multi-scale complexity analysis to quantify and qualify the heart rate's dynamics.

Results : Though the analysis shown in time domain that the dynamics of post-treatment and normal group looked similar. Their dynamics is completely different : (1) for normal heart rate, the multiscale fractal dimension is almost linearly decreased – invariance; (2) for arrhythmic heart rate before and after treatment, they converged to a certain value. All these suggested that item after the drug treatments, the heart's function is not still fully restored and more recovering time is needed. The multiscale complexity analysis can be used to quantify the heart function's recovering and optimize the post-treatment.

One submitted publication.




Figure 5. Analysis of heart rate variation by multiscale (coarse-graining) fractal dimension

7.7. Quantification and Its Approximate Solution of Action Potential in Neuron Models by Anharmonicity Analysis

Participants: Binbin Xu, Hussein Yahia, Rémi Dubois.

Action potential (AP) plays an important role to initiate and maintain the cell-cell communication. The nerve impulses are extensively studied but the action potential is less investigated as in other types of cells (for example, cardiac action potential). The AP can tell more about the state of the cell. It reflects the physical / chemical intracellular exchanges. Any changes in the cell would change the form/geometry of AP, or a more relevant term the *harmony*. The intrinsic changes would modify the harmony of the impulses train. The broken harmony (form/geometry change) of the impulse train means that there would be some problems in the cell. This provides an indirect way to study the intrinsic dynamics the cells.

In the work of P. Hanusse proposed a very interesting signal analyzing approach by anharmonicity, especially for signals with nonlinear oscillations properties exhibited in many physical / biological systems. This is exactly the case for neuron impulses trains. The principle is to describing the signal with their harmonic behaviors by solving the nonlinear phase equation. The obtained phase is thus used to reconstruct a solution of the original signal. The key notion is the nonlinear trigonometry that they developed. According to this approach, for any periodic signal $x(t) = x_0 + x_1 \cos(\phi(t))$, its phase can be obtained by the proposed general solution $t(\phi) = \phi + \sum_{k=1}^n a_k \text{hpsin}_1(\phi - p_k, r_k) - b_k \text{hpcos}_1(\phi - p_k, r_k)$ which can be used to reconstruct the original signal $x(t)$.

There is no practical implementation in their papers. Here we propose a first order solution of the original analytical equation. It can be used to quantify the harmonicity of the action potential.

$$\text{hpsin}_1 = \frac{-i(-\ln(1 - e^{it}r) + \ln(1 - re^{-it}))}{2r}, \quad \text{hpcos}_1 = -\frac{\ln(1 - e^{it}r) + \ln(1 - re^{-it})}{2r}$$

The phase of a signal can be solved as : $\phi(t) = t - t_0 + a_1 \text{hpsin}_1(t - t_1, r) - a_2 \text{hpsin}_1(t - t_2, r)$, so the signal can be reconstructed as $x(t) = \cos(\phi(t))$. The related parameters can be obtained by regression or nonlinear optimization methods. In consequence, all AP can be quantified by the anharmonicity parameter r .




Figure 6. Quantification of action potential in FitzHugh-Nagumo (FHN) model by anharmonicity analysis

FitzHugh-Nagumo (FHN) model has been one of the basic models to study the action potential's dynamics. It's derived from the Hodgkin-Huxley model and is physiologically correct. We take here dynamics of APs from FHN model to illustrate the anharmonicity analysis. As shown in the following figure, anharmonicity analysis is far more efficient than Fourier representation. Even with 8 Fourier terms, the signal is ill-represented with Gibbs phenomena. In contrast, 3 anharmonic terms exhibit a quasi-identical fit. These AP are quantified with an anharmonicity $r = 0.7384$.

More development of anharmonicity analysis on AP is ongoing in order to provide a more efficient way to generalize the practices and for better solutions. We believe that anharmonicity analysis can help to quantify/qualify the AP in a different yet efficient way than conventional analysis.

One submitted publication.

7.8. Image Reconstruction from Highly Corrupted Gradients

Participants: Hicham Badri, Hussein Yahia, Driss Aboutajdine.

Surface-from-Gradients (SfG) is an important step in many imaging applications. It consists in reconstructing an image/surface from corrupted gradient fields, which results in an ill-posed problem. We propose to use sparsity to regularize the problem. The first approach uses sparsity in the gradient field together with a robust norm on the data-fitting term and was presented at CVPR 2014. The new approach uses a non-local regularization that manipulates non-local similar patches of the corrupted gradient and forcing them to be low-rank. The two approaches significantly outperform previous optimization-based SfG methods on both synthetic and real data.

One submitted publication.

7.9. Fast Image Edge-Aware Processing

Participants: Hicham Badri, Hussein Yahia, Driss Aboutajdine.

We present a framework for fast edge-aware processing of images and videos. This is an extension of our previous SIGGRAPH Asia 2013 paper. The proposed approach uses non-convex sparsity on the gradients of the latent smooth image to better preserve sharp edges. We develop tools based on first order proximal estimation for fast processing. We also propose fast and efficient numerical solutions based on separable filters estimation, which enables our method to perform fast high-quality smoothing on large-scale images. Extensive experiments show that the proposed method produces high-quality smoothing compared to state-of-the-art methods, while being fast and simple to implement.

Publication: [15].

7.10. Low-Rankness Transfer for Realistic Denoising

Participants: Hicham Badri, Hussein Yahia.

Image restoration is a very challenging task in low-level vision and is extensively used in many imaging applications. Sparsity in various forms (dictionary learning, low-rank estimation,...) has shown to be the key for successful image denoising. However, the standard noise model used to validate the results is mainly Gaussian and uniform, with known standard deviation. Unfortunately, these assumptions do not hold for real camera noise. Instead of using sparsity to model the singular values of non-local clean similar patches, we use a learning model that trains a mapping between the noisy and ground-truth clean singular values. The training is performed on real camera noise, contrary to previous methods. Experiments show that the proposed method significantly outperforms previous denoising works on real non-uniform noise and does not require estimating the standard deviation of the corruption.

One publication accepted with minor revision at IEEE Transactions on Image Processing, publication date: 2016.

7.11. Turbulent Flow Estimation

Participants: Hicham Badri, Hussein Yahia.

We use singularity exponents (SE) to regularize the problem of turbulent flow estimation under the assumption that the brightness constancy constraint holds also for (SE). We also use weighted filtering (Lucas-Kanade's solution) and sparsity on the data-fitting term to improve robustness to outliers. The proposed motion estimation is built on a Gaussian pyramid and uses the theory of warping for a better estimation of large displacements. Experiments on synthetic data show that the proposed method outperforms sophisticated methods while being simple.

Work in progress.

7.12. Pathological voice classification

Participants: Khalid Daoudi, Nicolas Brodu.

Based on our GCI detection algorithm, we redefined the classical pitch perturbation measures that are widely used in voice quality assessment. We showed that our perturbation measures yield significantly better performance in pathological voice classification than classical measures. We also showed that some matching pursuit features can allow good performances in discrimination between pathological voice categories.

Publications: [31], [30].

7.13. Emotion detection: project with Batvoice start-up

Participants: Khalid Daoudi, Nicolas Brodu.

Geostat has been granted in 2015 a Carnot-Inria contact to fund a 1 year engineer to develop a prototype of a speech emotion detection system.

7.14. Heartbeat signal analysis: Proof of Concept with IHU LIRYC

Participants: Hussein Yahia, Guillaume Attuel, Oriol Pont, Binbin Xu.

Geostat has been granted in 2015 a fund from Inria DGT to conduct pre-clinical validation from patient database acquired by IHU LIRYC.

HIEPACS Project-Team

7. New Results

7.1. High-performance computing on next generation architectures

7.1.1. *Soft error sensitivity of PCG and reliability of detection mechanisms*

Soft errors can be defined as failures arising from several electricity fluctuations, cosmic particle effects on chip or any other unexpected problem while computations are in progress. If computational environment grows up to exascale, the rate of these types of error is likely to increase. These bit-flips may have a strong impact on iterative methods, that might diverge or converge to an unexpected final accuracy. Consequently, soft errors deserve to be examined in details especially in the perspective of extreme scale computing platforms. In this work, we investigate the combination of different numerical techniques to tackle the challenge of the detection. The first ingredient relies on checksum mechanisms, that are applied to secure the sparse matrix vector (SpMV) products. However, the checksum equalities are only valid in exact arithmetic while calculation are performed in finite precision. Another possibility is to monitor the residual deviation between the true and computed residual. Exploiting finite precision analysis of the round-off provides us with an upper bound on the residual norm deviation that can be used. Through intensive numerical experiments and statistical analysis we shown how round-off error analysis for the residual norm deviation can be an efficient and robust soft error detection criterion alternative to checksum approaches. This methodology has also be applied to other variants of CG, namely the pipelined and chronopolus/gear versions.

This research effort was conducted in collaboration with colleagues S. Cools and W. Vanroose from the Applied Mathematics Group of Antwerp university within the framework of the **EXA2CT** project. In this context, we also studied the impact of soft errors on a variant of the algorithm designed in their group (so-called pipelined CG). This study allowed to highlight some numerical instability in the baseline version of this variant of CG in the presence of round-off errors and we jointly proposed a correction of it that led a new both scalable and stable variant (see Section 7.2.5).

We have also designed an self-recovering CG algorithm which detects large magnituded faults with ABFT and smoothes low and average magnituded faults with deviation-based criteria.

7.1.2. *Resilience of parallel sparse hybrid solvers*

As the computational power of high performance computing (HPC) systems continues to increase by using a huge number of CPU cores or specialized processing units, extreme-scale applications are increasingly prone to faults. Consequently, the HPC community has proposed many contributions to design resilient HPC applications. These contributions may be system-oriented, theoretical or numerical. In this study we consider an actual fully-featured parallel sparse hybrid (direct/iterative) linear solver, **MaPHyS**, and we propose numerical remedies to design a resilient version of the solver. The solver being hybrid, we focus in this study on the iterative solution step, which is often the dominant step in practice. We furthermore assume that a separate mechanism ensures fault detection and that a system layer provides support for setting back the environment (processes, ...) in a running state. The present manuscript therefore focuses on (and only on) strategies for recovering lost data *after* the fault has been detected (a separate concern beyond the scope of this study), *once* the system is restored (another separate concern not studied here). The numerical remedies we propose are twofold. Whenever possible, we exploit the natural data redundancy between processes from the solver to perform exact recovery through clever copies over processes. Otherwise, data that has been lost and no longer available on any process is recovered through a so-called interpolation-restart mechanism. This mechanism is derived from our earlier studies by carefully taking into account the properties of the target hybrid solver. These numerical remedies have been implemented in the **MaPHyS** parallel solver so that we can assess their efficiency on a large number of processing units (up to 12,288 CPU cores) for solving large-scale real-life problems.

These contributions will be presented at the international conference HiPC [42].

7.1.3. Hierarchical DAG scheduling for hybrid distributed systems

Accelerator-enhanced computing platforms have drawn a lot of attention due to their massive peak computational capacity. Despite significant advances in the programming interfaces to such hybrid architectures, traditional programming paradigms struggle mapping the resulting multi-dimensional heterogeneity and the expression of algorithm parallelism, resulting in sub-optimal effective performance. Task-based programming paradigms have the capability to alleviate some of the programming challenges on distributed hybrid many-core architectures. In this work we take this concept a step further by showing that the potential of task-based programming paradigms can be greatly increased with minimal modification of the underlying runtime combined with the right algorithmic changes. We propose two novel recursive algorithmic variants for one-sided factorizations and describe the changes to the PaRSEC task-scheduling runtime to build a framework where the task granularity is dynamically adjusted to adapt the degree of available parallelism and kernel efficiency according to runtime conditions. Based on an extensive set of results we show that, with one-sided factorizations, i.e. Cholesky and QR, a carefully written algorithm, supported by an adaptive tasks-based runtime, is capable of reaching a degree of performance and scalability never achieved before in distributed hybrid environments.

These contributions will be presented at the international conference IPDPS 2015 [34] in Hyderabad.

7.1.3.1. Comparison of Static and Dynamic Resource Allocation Strategies for Matrix Multiplication

The tremendous increase in the size and heterogeneity of supercomputers makes it very difficult to predict the performance of a scheduling algorithm. In this context, relying on purely static scheduling and resource allocation strategies, that make scheduling and allocation decisions based on the dependency graph and the platform description, is expected to lead to large and unpredictable makespans whenever the behavior of the platform does not match the predictions. For this reason, the common practice in most runtime libraries is to rely on purely dynamic scheduling strategies, that make short-sighted scheduling decisions at runtime based on the estimations of the duration of the different tasks on the different available resources and on the state of the machine. In this work, we considered the special case of Matrix Multiplication, for which a number of static allocation algorithms to minimize the amount of communications have been proposed. Through a set of extensive simulations, we analyzed the behavior of static, dynamic, and hybrid strategies, and we assessed the possible benefits of introducing more static knowledge and allocation decisions in runtime libraries. These contributions have been presented at the international conference SBAC-PAD 2015.

7.1.3.2. Scheduling Trees of Malleable Tasks for Sparse Linear Algebra

Scientific workloads are often described as directed acyclic task graphs. In this paper, we focus on the multifrontal factorization of sparse matrices, whose task graph is structured as a tree of parallel tasks. Among the existing models for parallel tasks, the concept of *malleable* tasks is especially powerful as it allows each task to be processed on a time-varying number of processors. Following the model advocated by Prasanna and Musicus for matrix computations, we considered malleable tasks whose speedup is p^α , where p is the fractional share of processors on which a task executes, and α ($0 < \alpha \leq 1$) is a parameter which does not depend on the task. We first motivated the relevance of this model for our application with actual experiments on multicore platforms. Then, we studied the optimal allocation proposed by Prasanna and Musicus for makespan minimization using optimal control theory. We largely simplified their proofs by resorting only to pure scheduling arguments. Building on the insight gained thanks to these new proofs, we extended the study to distributed multicore platforms. There, a task cannot be distributed among several distributed nodes. In such a distributed setting (homogeneous or heterogeneous), we proved the NP-completeness of the corresponding scheduling problem, and proposed some approximation algorithms. We finally assessed the relevance of our approach by simulations on realistic trees. We showed that the average performance gain of our allocations with respect to existing solutions (that are thus unaware of the actual speedup functions) is up to 16% for $\alpha = 0.9$ (the value observed in the real experiments). These contributions have been presented at the international conference Europar 2015.

7.1.3.3. Task-based multifrontal QR solver for GPU-accelerated multicore architectures

Recent studies have shown the potential of task-based programming paradigms for implementing robust, scalable sparse direct solvers for modern computing platforms. Yet, designing task flows that efficiently exploit heterogeneous architectures remains highly challenging. In this work we first tackled the issue of data partitioning using a method suited for heterogeneous platforms. On the one hand, we designed task of sufficiently large granularity to obtain a good acceleration factor on GPU. On the other hand, we limited the size in order to both fit the GPU memory constraints and generate enough parallelism in the task graph. Secondly we handled the task scheduling with a strategy capable of taking into account workload and architecture heterogeneity at a reduced cost. Finally we proposed an original evaluation of the performance obtained in our solver on a test set of matrices. We showed that the proposed approach allows for processing extremely large input problems on GPU-accelerated platforms and that the overall performance is competitive with equivalent state of the art solvers designed and optimized for GPU-only use. These contributions have been presented at the international conference HiPC 2015 where they received the best paper award.

7.1.3.4. Fast and Accurate Simulation of Multithreaded Sparse Linear Algebra Solvers

The ever growing complexity and scale of parallel architectures imposes to rewrite classical monolithic HPC scientific applications and libraries as their portability and performance optimization only comes at a prohibitive cost. There is thus a recent and general trend in using instead a modular approach where numerical algorithms are written at a high level independently of the hardware architecture as Directed Acyclic Graphs (DAG) of tasks. A task-based runtime system then dynamically schedules the resulting DAG on the different computing resources, automatically taking care of data movement and taking into account the possible speed heterogeneity and variability. Evaluating the performance of such complex and dynamic systems is extremely challenging especially for irregular codes. In this work, we explained how we crafted a faithful simulation, both in terms of performance and memory usage, of the behavior of `qr_mumps`, a fully-featured sparse linear algebra library, on multi-core architectures. In our approach, the target high-end machines are calibrated only once to derive sound performance models. These models can then be used at will to quickly predict and study in a reproducible way the performance of such irregular and resource-demanding applications using solely a commodity laptop. These contributions have been presented at the international conference ICPADS 2015.

7.2. High performance solvers for large linear algebra problems

7.2.1. Divide and conquer symmetric tridiagonal eigensolver for multicore architectures

Computing eigenpairs of a symmetric matrix is a problem arising in many industrial applications, including quantum physics and finite-elements computation for automobiles. A classical approach is to reduce the matrix to tridiagonal form before computing eigenpairs of the tridiagonal matrix. Then, a back-transformation allows one to obtain the final solution. Parallelism issues of the reduction stage have already been tackled in different shared-memory libraries. In this work, we focus on solving the tridiagonal eigenproblem, and we describe a novel implementation of the Divide and Conquer algorithm. The algorithm is expressed as a sequential task-flow, scheduled in an out-of-order fashion by a dynamic runtime which allows the programmer to play with tasks granularity. The resulting implementation is between two and five times faster than the equivalent routine from the INTEL MKL library, and outperforms the best MRRR implementation for many matrices. These contributions have been presented at the international conference IPDPS 2015 [32] in Hyderabad.

7.2.2. Blocking strategy optimizations for sparse direct linear solver on heterogeneous architectures

Solving sparse linear systems is a problem that arises in many scientific applications, and sparse direct solvers are a time consuming and key kernel to those applications or more advanced solvers such as hybrid direct-iterative solvers. That is why optimizing their performance on modern architectures is a crucial problem. The preprocessing steps of sparse direct solvers: ordering and symbolic factorization, are two major steps that lead to a reduced amount of computation and memory, and to a better task granularity to reach a good level of performance when using BLAS kernels. With the advent of GPUs, the granularity of the symbolic factorization

became more important than ever. In this work, we present a reordering strategy that increases the block granularity. This strategy relies on the symbolic factorization to refine the ordering produced by tools such as METIS or *Scotch*, and does not impact the number of operations required to solve the problem. We integrated this algorithm in the *PaStiX* solver and show a reduction of the number of off-diagonal blocks by two to three on a large spectrum of matrices. This improvement leads to an efficiency on GPUs raised by up to 40%. These contributions have been presented at the Sparse Days [51] in Saint-Girons.

7.2.3. *On the use of \mathcal{H} -Matrix Arithmetic in PaStiX: a Preliminary Study*

The objective is to investigate innovative lowrank approximations based on \mathcal{H} -matrix variants for direct solver and Schur complements. The intent is to improve scalability of those components involved in preconditioners and hybrid solvers by reducing the computational and memory costs of the dense calculation. The quality of hybrid ordering algorithms combining topdown (such as nested dissection) and bottomup (such as minimum degree) ordering techniques in the context of sparse linear solvers will be investigated.

In this work, we describe a preliminary fast direct solver using HODLR library to compress large blocks appearing in the symbolic structure of the *PaStiX* sparse direct solver. We present our general strategy before analyzing the practical gains in terms of memory and floating point operations with respect to a theoretical study of the problem. Finally, we discuss ways to enhance the overall performance of the solver.

Some contributions have already been presented at the Workshop on Fast Solvers [52] in Toulouse. This work is a joint effort between Professor Darve's group at Stanford and the Inria HiePACS team within *FASTLA*.

7.2.4. *Data sparse techniques for parallel hybrid solvers*

In this work we describe how data sparse techniques exploiting \mathcal{H} -matrix calculations can be implemented in a parallel hybrid sparse linear solver based on an algebraic non overlapping domain decomposition approach.

Various graph-based clustering techniques to approximate the local Schur complements are investigated, with the aim of optimally complying with the interface structure of the local interfaces of the subdomains. We consider strong-hierarchical (sH) matrix arithmetic as efficient means for obtaining low rank approximations in terms of workload distribution as well as memory consumption. We also show how sH-arithmetic can be utilized to form an effective global preconditioner for the iterative phase of the hybrid solver. Numerical and parallel experiments are presented to evaluate the advantages and drawbacks of the different variants.

This work is a joint effort between Professor Darve's group at Stanford and the Inria HiePACS team within *FASTLA*. Some intermediate progresses have already been presented [38], [37]

7.2.5. *Analysis of the rounding error accumulation in Conjugate Gradient to improve the maximal attainable accuracy of pipelined CG*

Pipelined Krylov solvers typically offer better scalability in the strong scaling limit compared to standard Krylov methods. The synchronization bottleneck is mitigated by overlapping time-consuming global communications with useful computations in the algorithm. However, to achieve this communication hiding strategy, pipelined methods feature multiple recurrence relations on additional auxiliary variables to update the guess for the solution. This paper aims to study the influence of rounding errors on the convergence of the pipelined Conjugate Gradient method. It is analyzed why rounding effects have a significantly larger impact on the maximal attainable accuracy of the pipelined CG algorithm compared to the traditional CG method. Based on a rounding error model, we then propose an automated residual replacement strategy to reduce the effect of rounding errors on the final iterative solution. The resulting pipelined CG method with residual replacement improves the maximal attainable accuracy of pipelined CG while maintaining its efficient parallel performance.

This research effort was conducted in collaboration with colleagues S. Cools and W. Vanroose from the Applied Mathematics Group of Antwerp university within the framework of the *EXA2CT* project.

7.3. High performance Fast Multipole Method for N-body problems

7.3.1. Task-based Fast Multipole Method

Last year we have worked primarily on developing an efficient fast multipole method for heterogeneous architecture. Some of the accomplishments for this year include:

1. We have finalized the Uniform FMM (ufmm) based on polynomial interpolations combined with a hierarchical (data sparse) representation of a kernel matrix. The algorithm is close to the Black Box FMM by Fong and Darve developed with Chebyshev polynomials, however it uses an interpolation scheme based on an equispaced grid, which allows the use of FFT and consequently reduce both running time and memory footprint but has implications on accuracy and stability. The theory behind the Uniform FMM kernel is explained in a research report [63] along with numerical benchmarks on artificial test cases and presented in [44]. This new kernel was extended to be used for dislocation kernel.
2. Concerning the Group-Tree approach, we have shown in past studies its advantages of the task-based FMM and how the group-tree is well suited for runtime systems. In fact, it improves the locality, but it also reduces the number of dependencies which is an important asset to decrease the runtime overhead. These prospective task-based FMM can solve problems on heterogeneous architecture as presented in [36]. Therefore, we have continued this work and created a robust group-tree that has been included in ScalFMM and which is now available to the community. This data structure is generic and can be used with the different ScalFMM kernels. Moreover, we have extended our work and implemented a distributed task-based FMM above StarPU. The description of the data structure and some experimental studies will be presented in February 2016 during PhD defense of B. Bérenger.
3. With the advent of complex modern architectures, the low-level paradigms long sufficient to build high performance computing (HPC) numerical codes have met their limits. Achieving efficiency, ensuring portability, while preserving programming tractability on such hardware prompted the HPC community to design new, higher level paradigms. Indeed, several robust runtime systems proposed recently have shown the benefit of task-based parallelism models in terms of performance portability on complex platforms, on top of which full-featured numerical libraries have been ported successfully. However, the common weakness of these projects is to deeply tie applications to specific expert-only runtime system APIs. The OPENMP specification, which aims at providing a common parallel programming means for shared-memory platforms, appears a good candidate to address this issue thanks to the latest task-based constructs introduced as part of its revision 4.0. The goal of this joint work with STORM team is to assess the effectiveness and limits of this support for designing a high-performance numerical library like ScalFMM library, which implements state-of-the-art fast multipole methods (FMM) algorithms and that we have considerably re-designed with respect to the most advanced features provided by OPENMP 4.0. We show that OPENMP 4.0 allows for significant performance improvements over previous OPENMP revisions on recent multicore processors. We furthermore propose extensions to the OPENMP 4.0 standard and show how they could enhance FMM performance. To assess our statement, we have implemented this support within the KLANG-OMP source-to-source compiler that translates OPENMP directives into calls to the StarPU task-based runtime system. This study shows that we can take advantage of the advanced capabilities of a fully-featured runtime system without resorting to a specific, native runtime port, hence bridging the gap between the OPENMP standard and the very high performance that was so far reserved to expert-only runtime system APIs.

7.3.2. Time-domain boundary element method

The Time-domain Boundary Element Method (TD-BEM) has not been widely studied but represents an interesting alternative to its frequency counterpart. Usually based on inefficient Sparse Matrix Vector-product (SpMV), we investigate other approaches in order to increase the sequential flop-rate.

The TD-BEM formulation we is naturally expressed using sparse-matrix vector product (SpMV). We describe how the Flop-rate can be improved using a so-called multi-vectors/vector product, and we provide an efficient implementation of this operation using vectorization. We have extended our TD-BEM solver to support NVidia GPUs, and we have looked at different blocking schemes and their respective implementations. We have created a new blocking storage which matches our operators and allows to obtain a high Flop-rate. In addition, we provide a balancing heuristic to divide the work between the CPUs and the GPUs dynamically. The results have been published in [20], and our solver is now able to work on distributed heterogeneous nodes.

Our TD-BEM solver is efficient, but it still has a quadratic complexity which might become a problem for large problems. This high complexity motivates the study of an FMM based TD-BEM solver with the objective of being more competitive as the problem size increases. Therefore, we have implemented an FMM-based solver but while the complexity should be lower than the matrix approach, it remains unclear from which problem size. Moreover, we show in [PhD defense of B. Béranger] different results and point-out that the memory cost is much more expensive for the FMM approach compare to the matrix one. The method has been discussed in [43] among other ScalFMM applications.

All the implementations should be in high quality in the Software Engineering sense since the resulting library is going to be used by industrial applications.

This work is developed in the framework of Béranger Bramas's PhD and contributes to the EADS-ASTRIUM, Inria, Conseil Régional initiative.

7.3.3. Randomized algorithms for covariance matrices

7.3.3.1. Covariance kernel matrices

Random projection based Low Rank Approximation (LRA) algorithms such as the randomized SVD produce approximate matrix factorizations in quadratic instead of cubic time in N (N being the matrix size). This complexity can be further improved if fast matrix multiplication is available. A paper explaining our recent advances in fast randomized LRA of covariance kernel matrices using FMM is available as a research report [63] and presented in [44]. In particular, the fast multipole acceleration of the randomized SVD allowed for generating Gaussian random fields on arbitrary grids in linear running time and memory requirements. The code is available in the open source C++ project FMR: <https://gforge.inria.fr/projects/fmr>, it relies heavily on the ScalFMM library for data structures and fast matrix multiplication.

7.3.3.2. New applications: Data Assimilation and Taxonomy

Many applications like data assimilation (e.g. Kalman Filtering or variational approaches) or biology (e.g. taxonomy) involve covariance matrices that are only known in algebraic form, as opposed to kernel matrices that can be explicitly build given a kernel function. In a joint project (called FastMDS) with Alain Franc (INRA, Inria PLEIADE) addressing fast methods for the classification of biological species (taxonomy) our randomized SVD algorithm was used in order to accelerate a MultiDimensionalScaling (MDS) algorithm. The MDS is a widely used method in machine learning and data analysis that aim at visualizing the information contained in a distance matrix. Our MDS algorithm is applied to DNA sequences coming from various sources (e.g. Leman's lake), it consists in forming an euclidian image of the sample by taking the square root of a covariance matrix computed from the distance matrix. The randomized SVD approach lead to promising results, since it allowed to treat up to 100.000 samples in a few seconds. Since the covariance matrix still needs to be loaded in memory, storage might become problematic for larger samples. Therefore we are now considering matrix-free methods in order to decrease the memory requirements but also hierarchical algorithms in order to compute the MDS in near-linear time. The following methods are currently under investigation:

- Random column selection based LRA methods such as the Nystrom method or blocked variant of the Nystrom method (BBF, see Wang, Darve, Mahoney).
- Random projection based LRA powered by general H2-methods.

All these techniques are considered since they apply well, when the relevant information is spread uniformly among the data, just like in our data sets.

7.4. Efficient algorithmic for load balancing and code coupling in complex simulations

7.4.1. Dynamic load balancing for massively parallel coupled codes

In the field of scientific computing, load balancing is a major issue that determines the performance of parallel applications. Nowadays, simulations of real-life problems are becoming more and more complex, involving numerous coupled codes, representing different models. In this context, reaching high performance can be a great challenge. In the PhD of Maria Predari (started in October 2013), we develop new graph partitioning techniques, called co-partitioning, that address the problem of load balancing for two coupled codes: the key idea is to perform a *coupling-aware* partitioning, instead of partitioning these codes independently, as it is usually done. However, our co-partitioning technique requires to use graph partitioning with *fixed vertices*, that raises serious issues with state-of-the-art software, that are classically based on the well-known recursive bisection paradigm (RB). Indeed, the RB method often fails to produce partitions of good quality. To overcome this issue, we propose a *new* direct *k*-way greedy graph growing algorithm, called KGGGP, that overcomes this issue and succeeds to produce partition with better quality than RB while respecting the constraint of fixed vertices. Experimental results compare KGGGP against state-of-the-art methods for graphs available from the popular DIMACS'10 collection. This work will be presented in the 24th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP 2016).

7.5. Application Domains

7.5.1. Material physics

7.5.1.1. Molecular Vibrational Spectroscopy

Quantum chemistry eigenvalue problem is a big challenge in recent research. Here we are interested in solving eigenvalue problems coming from the molecular vibrational analysis. These problems are challenging because the size of the vibrational Hamiltonian matrix to be diagonalized is exponentially increasing with the size of the molecule we are studying. So, for molecules bigger than 10 atoms the actual existent algorithms suffer from a curse of dimensionality or computational time. We propose a new variational algorithm (namely residue-based adaptive vibrational configuration interaction) intended for the resolution of the vibrational Schrödinger equation. The main advantage of this approach is to efficiently reduce the dimension of the active space generated into the configuration interaction (CI) process. This adaptive algorithm is developed with the use of three correlated conditions i.e. a suitable starting space ; a criterion for convergence, and a procedure to expand the approximate space. The speed of the algorithm was increased with the use of a posteriori error estimator (residue) to select the most relevant direction to increase the space. Two examples have been selected for benchmark. In the case of Formalde hydemolecule (H_2CO) with a dimension space of 6, we mainly study the performance of RA-VCI algorithm: comparison with the variation-perturbation method, choice of the initial space, residual contributions. For Acetonitrile molecule (CH_3CN) with dimension space of 12 the active space computed by our algorithm is divided by 20 compared to the computations done by Avila et. al using the same potential energy surface. This work was presented in [54], [53].

7.5.1.2. Dislocations

7.5.1.2.1. Direct evaluation of the anisotropic elastic force field

The anisotropic elastic force field created by dislocations is not explicitly given, in fact it is only known in integral form using Green's or Stroh's formalism. The approach considered in OptiDis is based on Stroh's formalism, i.e. we compute the stress field using tensorial angular functions known as Stroh matrices. A benefit of using Stroh's formalism is that it only requires the evaluation of a single line integral for the force field and no integration for the stress field, while Green's formalism involve double and single line integral respectively. The evaluation of Stroh matrices in arbitrary directions is not affordable, therefore spherical harmonic expansions were considered in order to approximate the stress field efficiently. Until now the integration of the stress field on target dislocations was performed numerically using simple quadratures,

although the quadrature size required to evaluate the force field at a given precision may explode as segments get closer and computation may become untractable. In order to avoid this behaviour, we developed semi-analytical expressions of the force field based on the analytic integration of the expansions of the stress field (in spherical harmonics). This new method is an adaptation of Aubry et al. approach to Stroh's formalism, in the sense that it also provides optimized recursive formulae to efficiently evaluate these semi-analytic expressions. Numerous verifications and further improvements of the expressions are required before implementing it inside OptiDis.

7.5.1.2.2. Parallel dislocation dynamics simulation

We have focused on the improvements of our hybrid MPI-OpenMP parallelism of the OptiDis code. More precisely, we have continued the development of parallel algorithm to add/remove element in the cache-conscious data structure. This data structured combined with an octree manages efficiently large set of data (segments and nodes) during all the steps of the algorithm. Moreover, we have tuned and improved our hybrid MPI-OpenMP parallelism to run simulations with large number of radiation induced defects forming our dislocation network. To obtain a good scalability, we have introduced a better load balancing at thread level as well as process level. By combining efficient data structure and hybrid parallelism we obtained a speedup of 112 on 160 cores for a simulation of half a million of segments.

All this work was developped in the Phd of A. Etchevery.

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

7.5.2.1. MHD instabilities edge localized modes

The last contribution of Xavier Lacoste's thesis deals with the integration of our work in **JOEK**, a production controlled plasma fusion simulation code from CEA Cadarache. We described a generic finite element oriented distributed matrix assembly and solver management API. The goal of this API is to optimize and simplify the construction of a distributed matrix which, given as an input to **PaStiX**, can improve the memory scaling of the application. Experiments exhibit that using this API we could reduce the memory consumption by moving to a distributed matrix input and improve the performance of the factorized matrix assembly by reducing the volume of communication. All this study is related to **PaStiX** integration inside **JOEK** but the same API could be used to produce a distributed assembly for another solver or/and another finite elements based simulation code.

7.5.2.2. Turbulence of plasma particules inside a tokamak

Concerning the **GSELA** global non-linear electrostatic code, the efforts during the period have concentrated on predicting memory requirement and on the gyroaverage operator.

The Gysela program uses a mesh of 5 dimensions of the phase space (3 dimensions in configuration space and 2 dimensions in velocity space). On the large cases, the memory consumption already reaches the limit of the available memory on the supercomputers used in production (Tier-1 and Tier-0 typically). Furthermore, to implement the next features of Gysela (e.g. adding kinetic electrons in addition to ions), the needs of memory will dramatically increase, the main unknown will represents hundreds of TB. In this context, two tools were created to analyze and decrease the memory consumption. The first one is a tool that plots the memory consumption of the code during a run. This tool helps the developer to localize where the memory peak is located. The second tool is a prediction tool to compute the peak memory in offline mode (for production use mainly). A post processing stage combined with some specific traces generated on purpose during runtime allow the analysis of the memory consumption. Low-level primitives are called to generate these traces and to model memory consumption : they are included in the libMTM library (Modeling and Tracing Memory). Thanks to this work on memory consumption modeling, we have decreased the memory peak of the Gysela code up to 50 % on a large case using 32,768 cores and memory scalability improvement has been shown using these tools up to 65k cores.

The main unknown of the Gysela is a distribution function that represents either the density of the guiding centers, either the density of the particles in a tokamak (depending of the location in the code). The switch between these two representations is done thanks to the gyroaverage operator. In the actual version of Gysela, the computation of this operator is achieved thanks to the so-called Padé approximation. In order to improve the precision of the gyroaveraging, a new implementation based on interpolation methods has been done (mainly by researchers from the Inria Tonus project-team and IPP Garching). We have performed the integration of this new implementation in Gysela and also some parallel benchmarks. However, the new gyroaverage operator is approximatively 10 times slower than the original one. Investigations and optimizations on this operator are still a work in progress.

This work has been carried on in the framework of Fabien Rozar's PhD in collaboration with CEA Cadarache (defended in November 2015). A new PhD (Nicolas Bouzat) has started in October 2015 and the scientific objectives of this work will be first to consolidate the parallel version of the gyroaverage operator, in particular by designing a complete MPI+OpenMP parallel version, and then to design new numerical methods for the gyroaverage, source and collision operators to deal with new physics in Gysela. The objective is to tackle kinetic electron configurations for more realistic simulations.

7.5.2.3. *SN Cartesian solver for nuclear core simulation*

High-fidelity nuclear power plant core simulations require solving the Boltzmann transport equation. In discrete ordinate methods, the most computationally demanding operation of this equation is the sweep operation. Considering the evolution of computer architectures, we propose in this work, as a first step toward heterogeneous distributed architectures, a hybrid parallel implementation of the sweep operation on top of the generic task-based runtime system: **PaRSEC**. Such an implementation targets three nested levels of parallelism: message passing, multi-threading, and vectorization. A theoretical performance model was designed to validate the approach and help the tuning of the multiple parameters involved in such an approach. The proposed parallel implementation of the Sweep achieves a sustained performance of 6.1 Tflop/s, corresponding to 33.9% of the peak performance of the targeted supercomputer. This implementation compares favorably with state-of-art solvers such as PARTISN; and it can therefore serve as a building block for a massively parallel version of the neutron transport solver DOMINO developed at EDF.

The main contribution has been presented at the international conference IPDPS 2015 [31] in Hyderabad.

7.5.2.4. *3D aerodynamics for unsteady problems with moving bodies*

In the first part of our research work concerning the parallel aerodynamic code FLUSEPA, a first OpenMP-MPI version based on the previous one has been developed. By using an hybrid approach based on a domain decomposition, we achieved a faster version of the code and the temporal adaptive method used without bodies in relative motion has been tested successfully for real complex 3D-cases using up to 400 cores. Moreover, an asynchronous strategy for computing bodies in relative motion and mesh intersections has been developed and has been used for actual 3D-cases. A journal article (for JCP) to sum-up this part of the work is under redaction and a presentation at ISC at the "2nd International Workshop on High Performance Computing Simulation in Energy/Transport Domains" on July 2015 is scheduled.

This intermediate version exhibited synchronization problems for the aerodynamic solver due to the time integration used by the code. To tackle this issue, a task-based version over the runtime system **StarPU** is currently under development and evaluation. This year was mainly devoted to the realisation of this version. Task generation function have been designed in order to maximize asynchronism in execution. Those functions respect the data pattern access of the code and led to the refactorization of the actual kernels. A task-based version is now available for the aerodynamic solver and is available for both shared and distributed memory. This work has been presented as a poster during the SIAM CSE' 15 conference and at the Parallel CFD' 15 and HPCSET' 15 conferences.

The next steps will be to validate the correction of this task-based version and to work on the performance of this new version on actual cases. Later, the task description should be extended to the motion and intersection operations.

This work is carried on in the framework of Jean-Marie Cousteyen's PhD in collaboration with Airbus Defence and Space.

7.5.2.5. Spectral recycling strategies for the solution of nonlinear eigenproblems in thermoacoustics

In this work we consider the numerical solution of large nonlinear eigenvalue problems that arise in thermoacoustic simulations involved in the stability analysis of large combustion devices. We briefly introduce the physical modeling that leads to a nonlinear eigenvalue problem that is solved using a nonlinear fixed point iteration scheme. Each step of this nonlinear method requires the solution of a complex non-Hermitian linear eigenvalue problem. We review a set of state of the art eigensolvers and discuss strategies to recycle spectral information from one nonlinear step to the next. More precisely, we consider the Jacobi-Davidson algorithm, the Implicitly Restarted Arnoldi method, the Krylov-Schur solver and its block-variant, as well as the subspace iteration method with Chebyshev acceleration. On a small test example we study the relevance of the different approaches and illustrate on a large industrial test case the performance of the parallel solvers best suited to recycle spectral information for large scale thermoacoustic stability analysis.

The results of this work conducted in collaboration with S. Moreau (Sherbrooke University) and Y. Saad (University of Minnesota Twin-cities) are detailed in [22]

7.5.2.6. A conservative 2-D advection model towards large-scale parallel calculation

To exploit the possibilities of parallel computers, we designed a large-scale bidimensional atmospheric advection model named Pangolin. As the basis for a future chemistry-transport model, a finite-volume approach for advection was chosen to ensure mass preservation and to ease parallelization. To overcome the pole restriction on time steps for a regular latitude–longitude grid, Pangolin uses a quasi-area-preserving reduced latitude–longitude grid. The features of the regular grid are exploited to reduce the memory footprint and enable effective parallel performances. In addition, a custom domain decomposition algorithm is presented. To assess the validity of the advection scheme, its results are compared with state-of-the-art models on algebraic test cases. Finally, parallel performances are shown in terms of strong scaling and confirm the efficient scalability up to a few hundred cores

The results of this work are detailed in [21].

LFANT Project-Team

6. New Results

6.1. Class groups and other invariants of number fields

Participants: Karim Belabas, Jean-Paul Cerri, Henri Cohen, Pinar Kılıçer, Pierre Lezowski.

Ohno and Nakagawa have proved relations between the counting functions of certain cubic fields. These relations may be viewed as complements to the Scholz reflection principle, and Ohno and Nakagawa deduced them as consequences of 'extra functional equations' involving the Shintani zeta functions associated to the prehomogeneous vector space of binary cubic forms. The paper [14] by Henri Cohen, Simon Rubinstein-Salzedo and Frank Thorne proves an identity relating certain degree fields with Galois groups D and F , respectively, for any odd prime, giving another proof of the Ohno–Nakagawa relation between the counting functions of certain cubic fields.

Pinar Kılıçer and Marco Streng have solved a variant of the class number 1 problem for quartic CM fields with a geometric motivation [27]; the question is whether a certain class group is trivial, which corresponds to a genus 2 curve with that complex multiplication being defined over a real-quadratic number field (instead of an extension). Using classical techniques provides a bound on the discriminant of such fields, which they refine taking ramification into account to obtain a practically useful bound. A carefully crafted enumeration algorithm finishes the proof.

In the article [28], P. Lezowski studies the Euclidean properties of matrix algebras $M_n(R)$ over commutative rings R . In particular, he shows that for any integer $n > 1$, $M_n(R)$ is a left and right Euclidean ring if and only if R is principal. The proof is constructive and allows to better understand the Euclidean order types of matrix algebras. Similar ideas are also applied to prove k -stage Euclidean properties of $M_n(R)$, linking them with Bézout property for the ring R . The article [28] has been submitted to *Journal of Algebra*.

The article by Aurel Page on the computation of arithmetic Kleinian groups has appeared [21].

6.2. Complex L -functions and certified arithmetic

Participants: Bill Allombert, Karim Belabas, Henri Cohen, Fredrik Johansson.

Fredrik Johansson's paper [23] has been published and presented at the 22nd IEEE Symposium on Computer Arithmetic (ARITH22), Lyon, France. This paper describes a new implementation of the elementary transcendental functions \exp , \sin , \cos , \log and atan for variable precision up to approximately 4096 bits. Compared to the MPFR library, it achieves a maximum speedup ranging from a factor 3 for \cos to 30 for atan .

Bill Allombert, Karim Belabas, Henri Cohen and Pascal Molin (Paris 7) have implemented a new framework in PARI/GP to compute and manipulate complex L -functions and the associated ϑ and Λ functions, exporting 25 new functions to the GP computer algebra system.

6.3. Elliptic curve and Abelian varieties cryptology

Participants: Jean-Marc Couveignes, Andreas Enge, Enea Milio, Damien Robert.

In [29] David Lubicz and Damien Robert explain how to improve the arithmetic of Abelian and Kummer varieties. The speed of the arithmetic is a crucial factor in the performance of cryptosystems based on abelian varieties. Depending on the cryptographic application, the speed record holders are elliptic curves (in the Edwards model) or the Kummer surface of an hyperelliptic curves of genus 2 (in the level 2 theta model). One drawback of the Kummer surface is that only scalar multiplications are available, which may be a problem in certain cryptographic protocols. The previous known models to work on the Jacobian rather than the Kummer surface (Mumford coordinates or the theta model of level 4) are too slow and not competitive with elliptic curves. This paper explains how to use geometric properties (like projective normality) to speed up the arithmetic. In particular it introduces a novel addition algorithm on Kummer varieties (compatible addition), and uses it to speed up multi-exponentiations in Kummer varieties and to obtain new models of abelian surfaces in which the scalar multiplication is as fast as on the Kummer surface. This paper was written last year but heavily revised in 2015 and has been accepted (up to minor revisions) in the journal *Finite Fields and Their Applications*.

The paper [19] by David Lubicz and Damien Robert about computing certain isogenies in quasi optimal time has been published in the *LMS Journal of Computation and Mathematics* and the paper [18] by the same authors about optimal pairing computation on abelian varieties has been published in the *Journal of Symbolic Computation*. This paper expands the article [15] by Romain Cosset and Damien Robert about the computation of (ℓ, ℓ) -isogenies in dimension 2 which has been published in *Mathematics of Computation*.

Enea Milio has published one of the main results of his PhD thesis [20]. He has generalised the work of Régis Dupont for computing modular polynomials in dimension 2 to new invariants. He describes an algorithm to compute modular polynomials for invariants derived from theta constants and proves under some heuristics that this algorithm is quasi-linear in its output size. Some properties of the modular polynomials defined from quotients of theta constants are analysed and experiments with an implementation are related.

The paper [16] by Jean-Marc Couveignes and Tony Ezome explaining how to efficiently evaluate functions, including Weil functions and canonical theta functions, on Jacobian varieties and their quotients has been published in the *LMS Journal of Computation and Mathematics*. This paper also describes a quasi-optimal algorithm to compute (l, l) -isogenies between Jacobians of genus two curves, using a compact representation and differential characterisation of isogenies.

In [26], Sorina Ionica and Emmanuel Thomé look at the structure of isogeny graphs of genus 2 Jacobians with maximal real multiplication. They generalise a result of Kohel's describing the structure of the endomorphism rings of the isogeny graph of elliptic curves. Their setting considers genus 2 jacobians with complex multiplication, with the assumptions that the real multiplication subring is maximal and has class number 1. Over finite fields, they derive a depth first search algorithm for computing endomorphism rings locally at prime numbers, if the real multiplication is maximal.

6.4. Cryptology with quadratic fields

Participant: Guilhem Castagnos.

In [22] Guilhem Castagnos and Fabien Laguillaumie design a linearly homomorphic encryption scheme the security of which relies on the hardness of the decisional Diffie-Hellman problem. The approach requires some special features of the underlying group. In particular, its order is unknown and it contains a subgroup in which the discrete logarithm problem is tractable. Therefore, their instantiation holds in the class group of a non-maximal order of an imaginary quadratic field. Its algebraic structure makes it possible to obtain such a linearly homomorphic scheme in which the message space is the whole set of integers modulo a prime p and which supports an unbounded number of additions modulo p from the ciphertexts. A notable difference with previous work is that, for the first time, the security does not depend on the hardness of the factorisation of integers. As a consequence, under some conditions, the prime p can be scaled to fit the application needs. This paper has been presented at the cryptographer's track at the RSA Conference 2015.

MAGIQUE-3D Project-Team

6. New Results

6.1. Seismic Imaging and Inverse Problems

6.1.1. *hp-adaptive simulation and inversion of magnetotelluric measurements*

Participants: Hélène Barucq, Julien Alvarez Aramberri, David Pardo.

The magnetotelluric (MT) method is a passive exploration technique that aims at estimating the resistivity distribution of the Earth's subsurface, and therefore at providing an image of it. This process is divided into two different steps. The first one consists in recording the data. In a second step, recorded measurements are analyzed by employing numerical methods. In this work, we provide a rigorous mathematical setting in the context of the Finite Element Method (FEM) that helps to understand the MT problem and its inversion process. In order to recover a map of the subsurface based on 2D MT measurements, we employ for the first time in MTs a multigoal oriented self adaptive hp-Finite Element Method (FEM). We accurately solve both the full formulation as well as a secondary field formulation where the primary field is given by the solution of a 1D layered medium. To truncate the computational domain, we design a Perfectly Matched Layer (PML) that automatically adapts to high-contrast material properties that appear within the subsurface and on the air-ground interface. For the inversion process, we develop a first step of a Dimensionally Adaptive Method (DAM) by considering the dimension of the problem as a variable in the inversion. Additionally, this dissertation supplies a rigorous numerical analysis for the forward and inverse problems. Regarding the forward modelization, we perform a frequency sensitivity analysis, we study the effect of the source, the convergence of the hp-adaptivity, or the effect of the PML in the computation of the electromagnetic fields and impedance. As far as the inversion is concerned, we study the impact of the selected variable for the inversion process, the different information that each mode provides, and the gains of the DAM approach.

6.1.2. *Ultrasonic imaging of complex media*

Participants: Hélène Barucq, Juliette Chabassier, Marc Duruflé, Julien Diaz, Sébastien Tordeux, Ha Howard Faucher.

In 2015 we have begun a collaborating project with I2M (Physics Acoustics Department of Bordeaux 1 University). We aim at modeling and simulating efficiently the propagation of acoustic waves and later elastodynamic waves in highly heterogeneous media, the final goal is to use topological gradient imaging techniques. Classical techniques as finite elements can be too costly, we propose to design more efficient numerical techniques that exploit the fact that the wavelength is big with respect to the heterogeneities. For instance, we will use numerical upscaling, multiscale homogenization or asymptotic methods. A funding has been obtained for a PhD and a post doctoral position, that have both started in 2015. Our first step is to design a laboratory experiment and a simulation code in order to challenge the limits of the newly derived models and quantify their validity.

6.1.3. *Impedance transmission conditions for the electric potential across a highly conductive casing*

Participants: Hélène Barucq, Aralar Erdozain, David Pardo, Victor Péron.

In this study we present Impedance Transmission Conditions (ITCs) for the electric potential in the framework of borehole through-casing resistivity measurements. Such ITCs substitute the part of the domain corresponding to a highly conductive casing. The naturally small thickness of the casing makes it ideal for exhibiting ITCs. We numerically observe the delivered order of accuracy.

6.1.4. *An efficient truncated SVD of large matrices based on the low-rank approximation for inverse geophysical problems*

Participant: Sébastien Tordeux.

We have proposed a new algorithm to compute a truncated singular value decomposition of the Born matrix based on a low-rank arithmetic. Theoretical background to the low-rank SVD method has been investigated: the Born matrix of an acoustic problem can be approximated by a low-rank approximation derived thanks to a kernel independent multipole expansion. The new algorithm to compute T-SVD approximation consists of four steps, and they are described in detail. The largest singular values and their left and right singular vectors can be approximated numerically without performing any operation with the full matrix. The low-rank approximation is computed due to a dynamic panel strategy of cross approximation technique.

6.1.5. *Handling clusters with a task-based runtime system: application to Geophysics*

Participants: Emmanuel Agullo, Hélène Barucq, Lionel Boillot, George Bosilca, Julien Diaz.

The extreme complexity of hardware platforms makes them harder and harder to program. To fully exploit such machines, the High Performance Computing community often uses a MPI + X (X being pthreads, OpenMP, Cuda ...) programming models. We propose to use an alternative solution consisting of programming at a higher level of abstractions by describing a scientific, high performance computing application as a sequence of tasks whose execution is delegated to a runtime system. We compared MPI-based version and task-based version on Geophysics simulations, especially on the DIVA code of Total. Our previous studies demonstrated the task-based paradigm superiority on shared memory architectures (CPU or MIC), we are now working on distributed and heterogeneous architectures (CPUs+MICs) and, according to our preliminary results, the performances are still better than the MPI-version.

This work has been presented to the conferences PRACEdays [60], Rice Oil&Gas [43] and PASC [37].

6.2. Mathematical modeling of multi-physics involving wave equations

6.2.1. *Elasto-acoustic coupling*

Participants: Hélène Barucq, Lionel Boillot, Henri Calandra, Julien Diaz, Simon Ettouati, Conrad Hillairet, Elvira Shishenina.

In the framework of her Master thesis, Elvira Shishenina developed a Discontinuous Galerkin Method for the elastoacoustic coupling in time domain. The proposed solution methodology in general and can be applied to any kind of fluxes. We have implemented and validated in Elasticus a centered flux version and an upwind flux version in two dimensions. The time discretization is achieved thanks to Runge Kutta schemes of second and fourth orders.

In frequency domain, Conrad Hillairet developed a 3D elasto-coupling IPDG scheme, in the framework of his Master thesis. It has been implemented and validated in Hou10ni. Moreover, the code is able to handle p -adaptivity and we have proposed a strategy in order to determine the order of the cell as a function of the size of the cell and of the physical parameters. The results of this work have been presented to the Siam Conference on Geosciences in Stanford [39] and to the XXIV Congress on Differential Equations and Applications in Cadiz [32].

Finally, we have considered elastoacoustic coupling with curved interfaces and we have proposed a solution methodology based on Finite Element techniques, which allows for a flexible coupling between the fluid and the solid domain by using non-conforming meshes and curved elements. Since characteristic waves travel at different speeds through different media, specific levels of granularity for the mesh discretization are required on each domain, making impractical a possible conforming coupling in between. Advantageously, physical domains may be independently discretized in our framework due to the non-conforming feature. Consequently, an important increase in computational efficiency may be achieved compared to other implementations based on non-conforming techniques, namely by reducing the total number of degrees of freedom. Differently from other non-conforming approaches proposed so far, our technique is relatively simpler and requires only a

geometrical adjustment at the coupling interface at a preprocessing stage, so that no extra computations are necessary during the time evolution of the simulation. On the other hand, as an advantage of using curvilinear elements, the geometry of the coupling interface between the two media of interest is faithfully represented up to the order of the scheme used. In other words, higher order schemes are in consonance with higher order approximations of the geometry. Concerning the time discretization, we analyzed both explicit and implicit schemes. These schemes are energy conserving and, for the explicit case, the stability is guaranteed by a CFL condition.

This work, which has been achieved in collaboration with Angel Rodriguez Rozas, former post-doc of the team, was published in *Journal of Computational Physics* [27].

6.2.2. *Atmospheric boundary conditions for helioseismology*

Participants: H el ene Barucq, Juliette Chabassier, Marc Durufl e, Victor P eron.

The sun does not have a clear boundary like a solid ball, but it has an atmosphere which can be modeled as an exponential decay of the density. We have studied the replacement of this atmosphere by an equivalent boundary condition in order to avoid meshing the atmosphere. When we assume that the exponential decay is large enough, asymptotic modeling can be performed with respect to this large parameter. Equivalent boundary conditions have been obtained for order 1, 2 and 3, and they substantially improve Dirichlet condition (order 0) for low frequencies. However for high frequencies, these conditions are no longer relevant. We have developed a first-order absorbing boundary condition adapted to an exponential decay of the density, this last condition provides good results for the tested range of frequency. These conditions have been used by the team of Laurent Gizon (Max Planck institute) to retrieve experimental dispersion curves, so called ‘‘power spectrum’’.

6.2.3. *Absorbing Boundary Conditions for 3D elastic TTI modeling*

Participants: H el ene Barucq, Lionel Boillot, Julien Diaz.

We propose stable low-order Absorbing Boundary Conditions (ABC) for elastic TTI modeling. Their derivation is justified in elliptic TTI media but it turns out that they are directly usable to non-elliptic TTI configurations. Numerical experiments are performed by using a new elastic tensor source formula which generates P-waves only in an elliptic TTI medium. Numerical results have been performed in 3D to illustrate the performance of the ABCs.

This work has been presented to the conferences PANACM [38] and SEG [33].

6.2.4. *The airfoil equation on near disjoint intervals : Approximate models and polynomial solutions*

Participants: Leandro Farina, Marcos Ferreira, Victor P eron.

In [26], the airfoil equation is considered over two disjoint intervals. Assuming the distance between the intervals is small an approximate solution is found and relationships between this approximation and the solution of the classical airfoil equation are obtained. Numerical results show the convergence of the approximation to the solution of the original problem. Polynomial solutions for an approximate model are obtained and a spectral method for the generalized airfoil equation on near disjoint intervals is proposed.

6.2.5. *Finite element subproblem method*

Participants: Patrick Dular, Christophe Geuzaine, Laurent Kr ahenb uhl, Victor P eron.

In [25], progressive refinements of inductors are done with a subproblem method, from their wire or filament representations with Biot-Savart models up to their volume finite-element models, from statics to dynamics. The reaction fields of additional magnetic and/or conducting regions are also considered. Accuracy improvements are efficiently obtained for local fields and global quantities, i.e., inductances, resistances, Joule losses, and forces.

6.2.6. Asymptotic study for Stokes-Brinkman model with Jump embedded transmission conditions

Participants: Philippe Angot, Gilles Carbou, Victor Péron.

In [18], one considers the coupling of a Brinkman model and Stokes equations with jump embedded transmission conditions. Assuming that the viscosity in the porous region is very small, we derive a Wentzel-Kramers-Brillouin (WKB) expansion in power series of the square root of this small parameter for the velocity and the pressure which are solution of the transmission problem. This WKB expansion is justified rigorously by proving uniform errors estimates.

6.2.7. On the solution of the Laplace equation in 3-D domains with cracks and elliptical edges

Participants: Victor Péron, Samuel Shannon, Zohar Yosibash.

An explicit asymptotic solution to the elasticity system in a three-dimensional domain in the vicinity of an elliptical crack front, or for an elliptical sharp V-notch is still unavailable. Towards its derivation we first consider in [30] the explicit asymptotic solutions of the Laplace equation in the vicinity of an elliptical singular edge in a three-dimensional domain. Both homogeneous Dirichlet and Neumann boundary conditions on the surfaces intersecting at the elliptical edge are considered. The dual singular solution is also provided to be used in a future study to extract the edges flux intensity functions by the quasi-dual function method. We show that just as for the circular edge case, the solution in the vicinity of an elliptical edge is composed of three series, with eigenfunctions being functions of two coordinates.

In [29] the singular solution of the Laplace equation with a straight-crack is represented by a series of eigenpairs, shadows and their associated edge flux intensity functions (EFIFs). We address the computation of the EFIFs associated with the integer eigenvalues by the quasi dual function method (QDFM). The QDFM is based on the dual eigenpairs and shadows, and we show that the dual shadows associated with the integer eigenvalues contain logarithmic terms. These are then used with the QDFM to extract EFIFs from p-version finite element solutions. Numerical examples are provided.

6.3. Supercomputing for Helmholtz problems

6.3.1. High order methods for Helmholtz problems in highly heterogeneous media

Participants: Théophile Chaumont-Frelet, Henri Calandra, H el ene Barucq, Christian Gout.

The numerical solution of Helmholtz problems set in highly heterogeneous media is a tricky task. Classical high order discretizations fail to handle such propagation media, because they are not able to capture any of the scales of the velocity parameter. Indeed, they are build upon coarse meshes and therefore, if the velocity parameter is taken to be constant in each cell (through averaging, or local homogenization strategy), scale information is (at least partially) lost. We propose to overcome this difficulty by introducing a multiscale medium approximation strategy. The velocity parameter is not assumed to be constant on each cell, but on a submesh of each cell. If the submeshes are designed properly, the medium approximation method is equivalent to a quadrature formula, adapted to the medium. In particular, we show that this methodology has roughly the same computational cost as the classical finite element method. This new solution methodology has been presented in a paper under revision. We have performed a mathematical analysis of the multiscale medium approximation techniques to higher order discretization. First, we show that the heterogeneous Helmholtz problem is well-posed and derive stability estimates with respect to the right hand side, and with respect to variations of the velocity parameter, justifying the use of medium approximation. Those results are obtained assuming the velocity parameter is monotonous and that the propagation medium is closed by first order absorbing boundary conditions. However, these hypothesis are not mandatory to discretize the problem. Second, we turn to the analysis of finite element schemes with subcell variations of the velocity. In particular, we show that even if the solution can be rough inside each cell because of velocity jumps, we are able to extend the asymptotic error estimates obtained in [93] to heterogeneous media with non-matching mesh in case of elements of order $1 \leq p \leq 3$. Third, we investigate numerically the stability of the scheme when the

frequency is increasing to figure out optimal meshing conditions. We show that in simple media, the optimal homogeneous pre-asymptotic error estimates are still valid. However, in more complex cases, it looks like this condition is not sufficient anymore. Apart from showing that the homogeneous results are not always applicable to the heterogeneous Helmholtz equation, we are not able to give a clear answer to the question. Finally, we are able to conclude that high order methods are actually interesting: in our examples, $p = 4$ discretizations always yield a smaller linear system than lower order discretizations for the same precision.

6.3.2. *Hybridizable Discontinuous Galerkin method for the elastic Helmholtz equations*

Participants: Marie Bonnasse-Gahot, Henri Calandra, Julien Diaz, Stéphane Lanteri.

In the framework of the PhD thesis of Marie Bonnasse-Gahot, we have proposed an hybridizable discontinuous Galerkin method for solving the anisotropic elastodynamics wave equations in harmonic domain, in two and three dimensions. The method was implemented in Hou10ni and in the platform of Total. We have analyzed the performance of the proposed method in 2D on simple test case and compared it to classical DG methods. We have shown that the HDG method provides a more accurate solution for less computational cost provided that the order is high enough. We have illustrated the usefulness of the p -adaptivity in 2D, which allows to reach the accuracy of a global method of degree p for the costs of a global method of degree $p - 1$ or $p - 2$. This feature is already implemented in the 3D code. We now have to determine an accuracy criteria for assigning an order to a given cell, similar to the criteria we proposed in 2D.

For the numerical analysis of the scheme, we have shown that the HDG method could be rewritten as an upwind fluxes DG method and one of our perspectives is to use this equivalence in order to perform a dispersion analysis following the work of Ainsworth, Monk and Muniz [64].

We have shown that HDG could be used for 2D simulation on geophysical benchmark, and we will now implement the method in a Reverse Time Migration software, the ultimate goal being to couple HDG method with a full waveform inversion solver. In order to tackle more realistic test cases in 3D, it will be mandatory to improve the linear solver and we are now considering the use of an hybrid solver such as Maphys developed by the Inria team-project HIEPACS.

The results of this work have been presented at the ‘‘SIAM Conference on Geosciences’’ [48] and at the ‘‘Oil and Gas HPC Workshop’’ [49].

6.4. Hybrid time discretizations of high-order

6.4.1. *High-order symmetric multistep schemes for wave equation*

Participants: Juliette Chabassier, Marc Durufl e, Guillaume Marty.

We have studied high-order symmetric multistep schemes for the second-order formulation $y'' = f(t, y)$ during the internship of Guillaume Marty. The stability condition (CFL) can be optimized for explicit schemes since they have free parameters. However, this optimization procedure is not easy since the optimum is reached for forbidden values (values for which the high-order accuracy is no longer obtained). We have proposed acceptable values of free parameters for schemes of order 4, 6 and 8. These schemes have been tested for the wave equation, they suffer from a lack of robustness with respect to rounding numerical errors. The stability of implicit schemes has also been explored. For fourth-order schemes, a family of energy-conserving schemes has been obtained. However, we have not found unconditionally stable high-order schemes, which is well-known for the first-order formulation as Dahlquist’s barrier. It seems that for the second-order formulation, this barrier holds and only second-order accurate schemes are unconditionally stable. Implicit high-order schemes have a maximum CFL of $\sqrt{6}$, the same CFL as the standard θ -scheme with $\theta = \frac{1}{12}$. As a result, the implicit version of these schemes does not have a practical interest.

6.4.2. *High order conservative explicit and implicit schemes for wave equations.*

Participants: Juliette Chabassier, S ebastien Imperiale.

In 2015 we have studied the space/time convergence of a family of high order conservative explicit and implicit schemes for wave equations. An original proof of convergence has been proposed and provides an understanding of the lack of convergence of some schemes when the time step approaches its greatest admissible value for stability (CFL condition). An article is being written and will be submitted soon.

6.4.3. Multi-level explicit local time-stepping methods for second-order wave equations

Participants: Julien Diaz, Marcus Grote.

Local mesh refinement severely impedes the efficiency of explicit time-stepping methods for numerical wave propagation. Local time-stepping (LTS) methods overcome the bottleneck due to a few small elements by allowing smaller time-steps precisely where those elements are located. Yet when the region of local mesh refinement itself contains a sub-region of even smaller elements, any local time-step again will be overly restricted. To remedy the repeated bottleneck caused by hierarchical mesh refinement, multi-level local time-stepping methods are proposed, which permit the use of the appropriate time-step at every level of mesh refinement. Based on the LTS methods from Diaz and Grote [82], these multi-level LTS methods are explicit, yield arbitrarily high accuracy and conserve the energy.

The method was published in *Computer Methods in Applied Mechanics and Engineering* [24].

MANAO Project-Team

7. New Results

7.1. Analysis and Simulation

7.1.1. Parametrization of BRDFs

Opaque materials are represented in computer graphics by Bi-directional Reflectance Distribution Functions (BRDF), which are 4D functions of light and view direction. Dealing with such a high dimensionality is problematic for the modeling and rendering of material appearance. The choice of a BRDF parametrization greatly simplifies this task by identifying the axis where most variations occur in common opaque materials. The 4D parametrization of Rusinkiewicz [86] is classically used in graphics, in particular because of its direct connection to micro-facet theory. Alternative parametrization by Neumann et al. [71] and Stark et al. [91] have been proposed, but are restricted to 2D parametrizations, and hence a restricted class of materials.

We have extended the work of Neumann et al. [71] and Stark et al. [91] to a pair of 4D BRDF parameterizations with explicit changes of variables. Revealing some of their mathematical properties and relationships to Rusinkiewicz' parametrization allows us to better understand their benefits and drawbacks for representing measured BRDFs. Our preliminary study suggests that the alternative parametrization inspired by Stark et al. [91] is superior, and should thus be considered in future work involving BRDFs.

7.1.2. New BRDF Model and Diffraction Effects identification


Finding the appropriate BRDF model, with meaningful physical parameters, that can represent accurately measured data remains a challenging task. In [20], we show that two different physical phenomena are present in measured reflectance: reflection and diffraction. Taking both into account, we present a reflectance model [24] that is compact and a very good approximation (cf. Figure 8) of measured reflectance. Designers can act on model parameters, related to surface properties, to create new materials.

7.1.3. Statistical analysis of BRDFs

On the one hand, a BRDF is a complex 4D function, which should ensure reciprocity and energy conservation laws. On the other hand, when computing radiance reaching the eye from a surface point, the view direction is held fixed. In this respect, we are only interested in a 2D BRDF slice that acts as a filter on the local environment lighting. In [21], our goal is to understand the statistical properties of such a filter as a function of viewing elevation. To this end, we have conducted a study of measured BRDFs where we have computed statistical moments for each viewing angle. We show that some moments are correlated together across dimensions and orders, while some others are close to zero and may safely be discarded. Our study opens the way to novel applications such as moment-based manipulation of measured BRDFs, material estimation and image-based material editing. It also puts empirical and physically-based material models in a new perspective, by revealing their effect as view-dependent filters.

7.1.4. Importance Sampling of Real Light Sources

Realistic images can be rendered by simulating light transport with Monte Carlo methods. The possibility to use realistic light sources for synthesizing images greatly contributes to their physical realism. Among existing models, the ones based on light fields are attractive due to their ability to capture faithfully the far-field and near-field effects as well as their possibility of being acquired directly. Since acquired light sources have arbitrary frequencies and possibly high dimensions (4D), using such light sources for realistic rendering leads to performance problems. We have investigated [12] how to balance the accuracy of the representation and the efficiency of the simulation (cf. Figure 9). The work relies on generating high quality samples from the input light sources for unbiased Monte Carlo estimation [67]. This is a foundation work that has led to new sampling techniques for physically-based rendering with light field light sources. The results show that physically accurate rendering with realistic light sources can be achieved in real time.



- (a) Car model (Path tracing with adaptive sampling, 128 to 16384 samples per pixel), with smoother red-metallic-paint body, chrome wheels, white-marble floor. Inside the car: pickled-oak-260, specular-white-phenolic
- (b) Kitchen model (Energy Redistribution Path Tracing, 5h), with colonial-maple-223 cupboards, chrome tap, sink and handles, nickel kitchen wall, alumina-oxide oven door, white-marble and black-obsidian tiles, brass bowls, aluminium glasses.

Figure 8. Example scenes using our new BRDF model [24] to represent different measured materials.



Figure 9. Our new light importance sampling technique estimates direct lighting interactively (7-9 fps) with only 200 samples per pixel that are distributed among the different images of the light field luminaire. The car headlights are represented by the same light field composed of 11×9 images (256×256 pixels).

7.1.5. Exact Relations between Wave and Ray Aberrations

The aberrations of an optical system can be described in terms of the wave aberrations, defined as the departure from the ideal spherical wavefront; or the ray aberrations, which are in turn the deviations from the paraxial ray intersections measured in the image plane. The classical connection between the two descriptions is an approximation, the error of which has, so far, not been quantified analytically.

We derive [13] exact analytical equations for computing the wavefront surface, the aberrated ray directions, and the transverse ray aberrations in terms of the wave aberrations (a.k.a., Optical Path Difference) and the reference sphere. We introduce precise conditions for a function to be an OPD function, show that every such function has an associated wavefront, and study the error arising from the classical approximation. We establish strict conditions for the error to be small. We illustrate our results with numerical simulations. Our results show that large numerical apertures and OPD functions with strong gradients yield larger approximation errors.

7.2. From Acquisition to Display

7.2.1. Lytro Microscope




Figure 10. Light field microscopy with a consumer light field camera. Light fields can be beneficial for microscopic applications because they provide 3D information on a sample. Access to the technology, has, however, been limited by the need for custom-building the device. Our work enables an easy entry-level experimentation with the technology. (left) Light field microscope with a consumer light field camera. (right) Example specimens photographed with our system.

We explore [22] the use of inexpensive consumer light-field camera technology for the purpose of light-field microscopy. Our experiments are based on the Lytro (first generation) camera. Unfortunately, the optical systems of the Lytro and those of microscopes are not compatible, leading to a loss of light-field information due to angular and spatial vignetting when directly recording microscopic pictures. We therefore consider an

adaptation of the Lytro optical system. We demonstrate that using the Lytro directly as an ocular replacement, leads to unacceptable spatial vignetting. However, we also found a setting that allows the use of the Lytro camera in a virtual imaging mode which prevents the information loss to a large extent. We analyze the new virtual imaging mode and use it in two different setups for implementing light-field microscopy using a Lytro camera. As a practical result, we show that the camera can be used for low magnification work, as e.g. common in quality control, surface characterization, etc. (cf. Figure 10) We achieve a maximum spatial resolution of about 6.25 micrometers, albeit at a limited SNR for the side views.

7.3. Editing and Modeling

7.3.1. MatCap Decomposition for Dynamic Appearance Manipulation




Figure 11. Our method decomposes a MatCap into a representation that permits dynamic appearance manipulation via image filters and transforms. (a) An input MatCap applied to a sculpted head model (with a lookup based on screen-space normals). (b) The low- & high-frequency (akin to diffuse & specular) components of our representation stored in dual paraboloid maps. (c) A rotation of our representation orients lighting toward the top-left direction. (d) Color changes applied to each component. (e) A rougher-looking material obtained by blurring, warping and decreasing the intensity of the high-frequency component.

In sculpting software, MatCaps are often used by artists as a simple and efficient way to design appearance. Similar to LitSpheres, they convey material appearance into a single image of a sphere, which can be easily transferred to an individual 3D object. Their main purpose is to capture plausible material appearance without having to specify lighting and material separately. However, this also restricts their usability, since material or lighting cannot later be modified independently. Manipulations as simple as rotating lighting with respect to the view are not possible. In [23], we show how to decompose a MatCap into a new representation that permits dynamic appearance manipulation. We consider that the material of the depicted sphere acts as a filter in the image, and we introduce an algorithm that estimates a few relevant filter parameters interactively. We show that these parameters are sufficient to convert the input MatCap into our new representation, which enables real-time appearance manipulations through simple image re-filtering operations. This includes lighting rotations, the painting of additional reflections, material variations, selective color changes and silhouette effects that mimic Fresnel or asperity scattering (cf. Figure 11).

7.3.2. Dynamic On-Mesh Procedural Generation

In collaboration with Technicolor, we developed a method to generate procedural models with global structures, such as growth plants, on existing surfaces at interactive time [18]. Our approach extends shape grammars to enable context-sensitive procedural generation on the GPU. To this end, we unified the representation of external contexts as texture maps, which can be spatially varying parameters controlling the grammar expansion through very fast texture fetches (e.g., a density map). External contexts also include the shape of the



Figure 12. Given a base mesh and a procedural grammar of ivy growth, our GPU-based marching rule generated the ivy geometry on-the-fly in parallel with interactive performance. In this example, the grammar expansion is guided through a user-friendly painting interface.

underlying surface itself that we represent as a texture atlas of geometry images. Extrusion along the surface is then performed by a marching rule working in texture space using indirection pointers. We also introduce a lightweight deformation mechanism of the generated geometry maintaining a C1 continuity between the terminal primitives while taking into account the shape and trajectory variations. Our method is entirely implemented on the GPU and it allows to dynamically generate highly detailed models on surfaces at interactive time (cf. Figure 12). Finally, by combining marching rules and generic contexts, users can easily guide the growing process by directly painting on the surface with a live feedback of the generated model. This provides friendly editing in production environments.

7.3.3. Boolean on general 3D meshes

Computing Boolean operations (Booleans) of 3D polyhedra/meshes is a basic and essential task in many domains, such as computational geometry, computer-aided design, and constructive solid geometry. Booleans are challenging to compute when dealing with meshes, because of topological changes, geometric degeneracies, etc. Most prior art techniques either suffer from robustness issues, deal with a restricted class of input/output meshes, or provide only approximate results.

We overcome these limitations and introduced an exact and robust approach performing on general surface meshes (closed and orientable) [11]. Our method is based on a few geometric and topological predicates that allow to handle all input/output cases considered as degenerate in existing solutions, such as voids, non-manifold, disconnected, and unbounded meshes, and to robustly deal with special input configurations. Our experimentation showed that our more general approach is also more robust and more efficient than Maya's implementation ($\times 3$), CGAL's robust Nef polyhedra ($\times 5$), and recent plane-based approaches.

During this work, we also developed a complete benchmark intended to validate Boolean algorithms under relevant and challenging scenarios, and we successfully ascertain both our algorithm and implementation with it.

7.3.4. Extending MLS surfaces

Moving least squares (MLS) surface approximation is a popular tool for the processing and reconstruction of non-structured and noisy point clouds. We introduce [14] a new variant improving the approximation

quality when the underlying surface is assumed to be locally developable, which is often the case in point clouds coming from the acquisition of manufactured objects. Our approach follows Levin's classical MLS procedure: the point cloud is locally approximated by a bivariate quadratic polynomial height-field defined in a local tangent frame. The a priori developability knowledge is introduced by constraining the fitted polynomials to have a zero-Gaussian curvature leading to the actual fit of so-called parabolic cylinders. When the local developability assumption cannot be made unambiguously, our fitted parabolic cylinders seamlessly degenerate to linear approximations. We show that our novel MLS kernel reconstructs more locally-developable surfaces than previous MLS methods while being faithful to the data.

MEMPHIS Team

7. New Results

7.1. Plastic impact of iron on aluminium

A new model for plasticity has been developed this year. An iron projectile is impacting an aluminium plate immersed in air. The initial horizontal velocity of the iron is $1000m.s^{-1}$. The computation is performed on a 2000×1600 mesh with 144 processors. Homogeneous Neumann conditions are imposed on the left and right borders and embedded on the top and bottom.

The results are presented in Fig 6 with a schlieren image (bottom) and the von Mises criteria (top) at different time steps. A log scale is used and the minimum value is fixed to 10^9 . We can see that the plate is strongly deformed and form at the end a filament. The projectile is flattened but not as much as in the literature because the yield plastic limit is higher. We see a longitudinal wave propagating in the plate followed by a shear wave that causes the plasticity of the material.

7.2. Air-helium shock-bubble interaction

A three-dimensional hyperelastic model has been developed. It can deal with multi-fluid and solids. Here we show the propagation in air of a Mach 1.22 shock through an helium bubble. The computation is performed on a $1000 \times 400 \times 400$ mesh and lasts for 50h on 300 processors. The zero iso-value of the level set function and schlieren on the horizontal plane through the center of the bubble are presented at different times on Fig. 7 .

7.3. Particles flowing in a fluid

A new type of algorithm is designed to enable contacts efficiently between particles immersed in a fluid by adding a short range repulsive force. The algorithm is derived from the multi geometric deformable model introduced for image segmentation. It can handle multiple deforming bodies and avoid collision using a short range repulsive force depending on the distance to the closest interface. The main advantages of this method is it requires only five fields (three label maps and two distance functions) and one level set function to capture an arbitrary number of cells and it can, at the same time, deal with collisions.

7.4. Inertial Sea Wave Energy Converter (ISWEC)

The ISWEC is a floater and was design by Wave For Energy (<http://www.waveforenergy.com>) to extract the energy of typical waves in the Mediterranean Sea. The energy is extracted using a mechanical system based on a gyroscope activated by the motion of the floater generated by sea waves. This is a complex system coupling Fluid/Fluid/structure interfaces, computation of the rigid motion of the floater and computation of the power extraction. The problem is solved using in-house numerical solver (NaSCar) developed in MEMPHIS team. The interfaces are tracked using level se functions. The bi fluid interface is computed using Continuous Surface Force method (CSF), the motion of the floater imposed by penalization is computed using the forces and the torques exerted by the flow, and finally this motion activates the gyroscope for power extraction. The gyroscope model was developed by the Politecnico di Torino. Figure 1 shows a numerical simulation of the iswec (see <http://www.math.u-bordeaux1.fr/~mbergman/> for a movie)

7.5. Flow with many particles

A version of the code NaSCar has been developed to simulated the flow around particules with high volume fraction (see figure 10). The standard central lubrication forces are used to computed the interaction between spherical particules. Ongoing project will deal with non spherical particules (Lisl Weynans and PhD Baptiste Lambert).




Figure 6. Impact of Iron on Aluminium TC2. Schlieren image and von Mises criteria at $t = 0.03\text{ms}$, 0.06ms , 0.13ms , 0.26ms , 0.53ms and $t = 1.04\text{ms}$ from left to right, top to bottom




Figure 7. Interaction of a Mach 1.22 shock propagating in air through an helium bubble (TC1). Pictures at $t = 62\mu\text{s}$, $110\mu\text{s}$, $163\mu\text{s}$, $264\mu\text{s}$, $471\mu\text{s}$, $735\mu\text{s}$. From left to right, top to bottom.




Figure 8. Simulation of 30 rigid bodies of different radii ($R = 0.05$ or $R = 0.025$) falling under gravity. The colors indicate the values of the first label map from dark blue for the first body to dark orange for the 30th body and red for the fluid that is the 31st object.




Figure 9. Simulation of 30 rigid bodies of different radii ($R = 0.05$ or $R = 0.025$) falling under gravity. The colors indicate the vorticity levels from dark blue for -200 and dark red for 200.




Figure 10. Rheology of a Couette flow with solid particles.

7.6. Overset method

With Valeol (Cifre PhD Claire Morel) we are developing an Overlapping grids approach coupling a background cartesian grid with a body fitted grid around wind turbine blades. This method allows us to push away the limit of the numerical simulation on octree grids when small boundary layers play an important role. The generation of overset grids is based on level set functions (post doc AMIES by Franck Luddens). This method has been implemented for two dimensional test cases using the Schwartz method to deal with the domain decomposition. In the same time, we are also building the global operators onto the whole domain, i.e. without Schwartz iterations (PhD Federico Tesser).

7.7. Electrostrictive materials: modelling and simulation

In this work, a result of the collaboration between physicists, chemists (Annie Colin and Philippe Poulin, CRPP Bordeaux) and applied mathematicians, we deal with mathematical modelling and simulation of electrostrictive materials. These kinds of materials are composed by a polymeric matrix with carbon nanotubes embedded in and this structure gives them interesting electrical properties. Their dielectric constant varies as a function of the mechanical deformation. Housed in a capacitor, they show variable capacity when subjected to vibration and they can generate potential differences from mechanical deformations. Because of their composition, their structure involves different physical scales, from the small nanotube dimension, through the scale of nanotube clusters, to the large dimension of the sample. Our purpose is to provide physicists and chemists with a tool to test *in silico* several material configurations and to have a deeper insight into the features of these materials, developing numerical models which can predict their steady and unsteady behaviour. We propose to model the physical problem by reducing the nanotubes to dipoles and solving a Gauss equation for the electrical potential equation informed of the presence of nanotubes with zero electrical field conditions on the centres of the nanotubes. We started considering the steady problem, which is interesting for the purpose

of understanding the basic electrical properties of different nanotube configurations and of designing of the material. In order to discretize and simulate the mathematical problem, we chose parallel linear octree-based adaptive meshes and we developed an original hybrid Finite Volume/Finite Difference second-order scheme for 2D and 3D elliptic problems on this kind of mesh. A convergence analysis of the numerical scheme has been developed and validating test cases have been performed. Good qualitative agreement between numerical and real experiments has been observed for the steady model. In the future we aim to quantitatively compare the numerical results and the real material behaviour, to model the unsteady problem and to deal with electrical consequences of mechanical deformations.

7.8. Development of a sharp cartesian method for the simulation of flows with high density ratios

We have developed a sharp cartesian method for the simulation of incompressible flows with high density and viscosity ratios, like air-water interfaces. This method is inspired from a second-order cartesian method for elliptic problems with immersed interfaces (Cisternino-Weynans 2012). A classical predictor-corrector algorithm is used to solve the fluid equations, in a non-incremental version, which means that the guess value for the pressure is zero. This choice avoids instability issues due to the discontinuous pressure values when the interface moves. We take into account the viscous forces by regularizing the density and viscosity values. This approach allows for a more straightforward and robust treatment, and has been proven to provide satisfactory accuracies for high Reynolds numbers. To compute the pressure, it is necessary to solve an elliptic problem. This elliptic problem with discontinuous values across an interface is solved with the second order method cited above. The originality of this method lies in the use of unknowns located at the interface. These interface unknowns are used to discretize the flux jump conditions and the elliptic operator accurately enough to get a second order convergence in maximum norm.





Figure 11. Electrostrictive materials: Mesh and convergence results for the solver.




Figure 12. Electrostrictive materials: 3D configuration with nanotubes




Figure 13. Horizontal et vertical velocity profile of an air bubble rising into water.

MNEMOSYNE Project-Team

7. New Results

7.1. Overview

Though our view is systemic, our daily research activities are also concerned with the design, at a given scale of description, of models of neuronal structures, each concerned with a specific learning paradigm. Of course, a major challenge is to integrate these elements in a systemic view, i.e. to put a specific emphasis on the way each neuronal structure communicates with the rest of the system and to highlight how its learning paradigms interact with other memory systems.

Among the numerous loops involving the brain, the body and the environment, a basic grid of description corresponds to distinguish “perception aspects of loops”, the goal of which is to extract from the inner and outer world sensory invariants helpful to identify and evaluate the current state and to make predictions from previous learning, and “action aspects of loops”, the goal of which is to rely on this sensory and emotional information to decide, plan and trigger actions for the benefit of the body.

This year, our team was engaged on the following topics: Concerning perception aspects of loops, we published original models of the amygdala and of the hippocampus and considered their role in pavlovian conditioning and their evaluation as classical models in machine learning. Concerning action aspects of loops, in addition to a critical analysis of the current views of the interactions between the prefrontal cortex and the basal ganglia [15], we have proposed an original model for the formation of habits and have also studied related theoretical problems in machine learning, for data representation. Finally, we also report here more methodological achievements, corresponding to the design of algorithmic ersatz of cerebral subsystems.

7.2. Pavlovian conditioning

Within perception aspects of loops, pavlovian conditioning is a very interesting learning paradigm to study in a systemic view because it is tightly related to other learning paradigms like episodic and semantic memory. This year, we have published papers presenting the biological basis of models of two fundamental structures in pavlovian conditioning, the amygdala [1] and the hippocampus [2]. We have also evaluated their most critical features, when considered as models of machine learning, namely their architecture and implementations at both rate and spiking levels [10] and their robustness to interference [11].

7.3. The formation of habits

Concerning action aspects of loops, we made important extensions to a model of basal ganglia that we developed recently [54] in interaction with another team in our neuroscience lab. In addition to extending the bio-plausibility of this model with an implementation at the spiking level, we have also developed this year a new theoretical framework that provides a novel explanation for the formation and the expression of habits in the cortex of primates by considering the basal ganglia as an implicit supervisor. This has been achieved with a model of basal ganglia running both at the rate and spiking levels. This framework predicts that Hebbian learning and reinforcement learning can be explicitly dissociated by inactivating the output of the basal ganglia during learning and later tested in normal conditions. Experimental results in the monkey confirmed this prediction and explain how a behavioral decision results from both the cooperation (acquisition) and competition (expression) of two distinct but entangled memory systems.

7.4. Beyond symbolic models

Using a biologically plausible model, we have been investigating some external and internal factors related to the stimulus representation that might affect the decision making and action selection [41]. We used a computational model of the cerebral structure Basal Ganglia, inspired and replicated from a model designed in previous studies [54]. One of the questions we attempt to address is to what extent the physical properties of the stimulus affect the decision to overcome the impact of reward associated to the stimuli.

7.5. Algorithmic ersatz of cerebral subsystems

As far as the systemic modeling and simulation of high-level brain functions are concerned (e.g., sensory-motor behavior, action selection and planning, perceptual categorization), we need to confront biologically plausible models at different scales of description with functional models that are not constrained by biological facts but still reproduce the expected functional response. This is mandatory to benchmark bio-physical models with respect to their equivalent in classical machine learning, in order to evaluate the degree of naiveness of their performances and also to build feasible simulation in which detailed biological models can interact with less plausible modules in order to be evaluated in realistic numerical situations.

This year, a set of formalism such as the Friston free-energy minimization general principle, deep-learning and related architectures, and more specific formalisms such as harmonic control or adaptive-subspace self-organizing maps have been studied and reviewed. The next step is to write a review, with the challenge of proposing an unifying view of those, and at a more concrete level, to propose the integration of a relevant subset of the related algorithms as a easily usable toolbox. This can be particularly useful to design global models of cognitive functions, even if biologically-inspired models are not yet available for all their components.

Preliminary key points regarding numerical experimentations have been published in this methodological paper [16].

Monc Team

7. New Results

7.1. Axis 1: Tumor modeling for patient-specific simulations

7.1.1. Lung metastasis

Patient specific simulation of tumor growth, response to the treatment and relapse of a lung metastasis: a clinical case [10], [1]

Team participants: Thierry Colin, Julien Jouganous, François Cornelis (Hôpital Pellegrin), Olivier Saut

Other participant: Jean Palussière (Bergonié Institute)

In this work, a parametrization strategy based on reduced order methods is presented for tumor growth PDE models. This is applied to a new simple spatial model for lung metastasis including angiogenesis. The goal is to help clinicians monitoring tumors and eventually predicting its evolution or response to a particular kind of treatment. To illustrate the whole approach, a clinical case including the natural history of the lesion, the response to a chemotherapy and the relapse before a radiofrequency ablation is presented.

Nenuphar

Team participants: Thierry Colin, Julien Jouganous, Marie Martin, Olivier Saut

This work concerns the development of *Nenuphar* which is a software devoting to the evaluation and the surveillance of the tumor aggressiveness.

7.1.2. Take into account the drug resistance

Modeling and analysis of tumor heterogeneity during treatments resistance: GIST liver metastases case

Team participants: Thierry Colin, François Cornelis, Guillaume Lefebvre, Clair Poignard, Olivier Saut

This work deals with tumor heterogeneity analysis and modeling during treatments resistances. A patient-dependent PDEs model, that takes into account two kinds of treatments, is presented. It qualitatively and quantitatively reproduces the different stage during the tumor growth undergoing treatments. In order to overcome a numerical instability linked to the type of modeling, a new numerical scheme is built. Then, an image synthesis method is developed to enable a better comparison between the numerical results and the clinical data. Finally, a robust criteria that quantifies the tumor heterogeneity from the clinical data and from the synthesis images, is built.

Mathematical study and asymptotic analysis of a model for tumour drug resistance [19]

Team participants: Thierry Colin, Thomas Michel, Clair Poignard

In this work we study a partial differential equations model for tumour growth taking into account drug resistance. It is well known that angiogenesis, the process of creation of new blood vessels from existing ones, is induced by tumour cells to get the amount of nutrients and oxygen needed to continue their proliferation when the tumour has reached a critical size. Angiogenesis is therefore a target for therapy. The model we study takes into account two kinds of treatments: a cytotoxic treatment and a treatment which is both cytotoxic and anti-angiogenic. It is based on mass-balance equations on cells densities coupled with a diffusion equation for the nutrients and oxygen concentration. In a first part we prove that the model is well-posed if the initial tumour is compactly supported in the domain, which is the case for tumour metastases. The proof states that the tumour remains compactly supported in a finite time. In the model, we also consider the presence of a necrotic compartment composed of dead cells. Since some tumours can present necrosis while other do not, we want a model which can reproduce these two different cases. The second part of this work is devoted to an asymptotic analysis which proves that the absence of necrosis is the limit case of our model when the necrosis is immediately evacuated.

7.1.3. Motility phenotype

TMOD-03 * Motility controls growth and progression patterns of glioblastoma multiforme [13]

Team participants: Olivier Saut, Thierry Colin

Other participants: Hassan Fathallah, Elizabeth Scribner

Purpose: Glioblastoma multiforme (GBM) is a malignant brain tumor with poor prognosis and high morbidity due to its invasiveness. Hypoxia-driven motility (HM) and concentration-driven motility (CM) are two mechanisms of GBM invasion in the brain. The use of anti-angiogenic drugs has uncovered new progression patterns of GBM associated with significant differences in overall survival times. Here, we test the hypotheses that the types and rates of GBM motility predict its progression pattern and the patients' survival times. **Methods:** We applied a mathematical model of GBM growth and invasion in humans to simulate a clinical trial and study the effects of the rate and mechanism of motility on the patterns of progression and on survival times. **Results:** The motility phenotype appears to determine the progression pattern as well as the survival time of a patient treated by anti-angiogenesis. Highly-dispersive tumors are associated with the longest survival times ($p < 0.001$) and with progression by Expanding FLAIR. Moderately-Dispersive tumors are associated with short survival times and with progression by Expanding FLAIR + Necrosis. Tumors with HM are associated with the shortest survival times and with progression by Expanding Necrosis. The survival times of the latter are similar to non-responders. This investigation also uncovered the HM-CM principle: the aggressive HM-dependent phenotype surfaces only when the rate of CM is low in both untreated and bevacizumab-treated GBM. **Conclusions:** Finding that the motility phenotype is a fundamental property that controls progression and survival times, has biological, clinical and therapeutic implications.

7.2. Axis 2: Bio-physical modeling for personalized therapies

7.2.1. Electroporation

Non-Linear Steady-State Electrical Current Modeling for the Electroporation of Biological Tissue [8]

Team participants: Clair Poignard, Michael Leguebe

Other participants: Marie Breton, Lluís M. Mir (Vectorology and Anticancer Therapies), Francois Buret, Riccardo Scorretti, Damien Voyer, Laurent Krähenbühl (Ampère Laboratory (Lyon) participants), Ronan Perrussel (LAPLACE - Laboratoire Plasma et Conversion d'Énergie, Toulouse)

We propose a non-linear steady-state model of irreversible electroporation in a biological tissue. The non-linear problem is solved using a modified fixed point iteration. The unknown parameters are experimentally estimated from the observation of the necrosis on a potato tissue for different applied voltages. A variability study of the parameters involved in the model is performed.

A second-order Cartesian method for the simulation of electroporation cell models [12]

Team participants: Clair Poignard, Michael Leguebe

Other participant: Lizl Weynans (Memphis team, Inria)

In this work, we present a new finite differences method to simulate electroporation models, like the model of Neu and Krassowska or the recent model of Kaviani et al. These models are based on the evolution of the electric potential in a cell embedded in a conducting medium. The main feature lies in the transmission of the voltage potential across the cell membrane: the jump of the potential is proportional to the normal flux thanks to the well-known Kirchhoff law. An adapted scheme is thus necessary to accurately simulate the voltage potential in the whole cell, notably at the membrane separating the cell from the outer medium. We present a second-order finite differences scheme in the spirit of the method introduced by Cisternino and Weynans for elliptic problems with immersed interfaces. This is a Cartesian grid method based on the accurate discretization of the fluxes at the interface, through the use of additional interface unknowns. The main novelty of our present work lies in the fact that the jump of the potential is proportional to the flux, and therefore is not explicitly known. The original use of interface unknowns makes it possible to discretize the transmission conditions with enough accuracy to obtain a second-order spatial convergence. We prove the second-order spatial convergence in the stationary linear one-dimensional case, and the first-order temporal convergence for the dynamical non-linear model in one dimension. We then perform numerical experiments in two dimensions that corroborate these results.

Cell membrane permeabilization by 12-ns electric pulses: Not a purely dielectric, but a charge-dependent phenomenon [15]

Team participants: Clair Poignard, Michael Leguèbe

Other participants: Aude Silve (KIT - Karlsruhe Institute of Technology), Isabelle Leray, Lluís M. Mir (Université Paris Sud)

Electric pulses of a few nanoseconds in duration can induce reversible permeabilization of cell membrane and cell death. Whether these effects are caused by ionic or purely dielectric phenomena is still discussed. We address this question by studying the impact of conductivity of the pulsing buffer on the effect of pulses of 12 ns and 3.2 MV/m on the DC-3F mammalian cell line. When pulses were applied in a high-conductivity medium (1.5 S/m), cells experienced both reversible electroporation and cell death. On the contrary, no effect was observed in the low-conductivity medium (0.1 S/m). Possible artifacts due to differences in viscosity, temperature increase or electrochemical reactions were excluded. The influence of conductivity reported here suggests that charges still play a role, even for 12-ns pulses. All theoretical models agree with this experimental observation, since all suggest that only high-conductivity medium can induce a transmembrane voltage high enough to induce pore creation, in turn. However, most models fail to describe why pulse accumulation is experimentally required to observe biological effects. They mostly show no increase of permeabilization with accumulation of pulses. Currently, only one model properly describes pulse accumulation by modeling diffusion of the altered membrane regions.

7.2.2. Cell protrusion

Free boundary problem for cell protrusion formations: theoretical and numerical aspects [20]

Team participants: Olivier Gallinato, Clair Poignard

Other participants: Masahito Ohta (Tokyo University of Sciences), Takashi Suzuki (Osaka University)

In this work, we derive a free boundary problem for cell protrusion formation in which the cell membrane is precisely described thanks to a level-set function, whose motion is due to specific signalling pathways. The model consists in Laplace equation with Dirichlet condition inside the cell coupled to Laplace equation with Neumann condition in the outer domain. The motion of the interface is due to the gradient of the inner quantity. We prove the well-posedness of our free boundary problem under a sign condition on the datum similarly to the Taylor criterion in water waves. We also propose an accurate numerical scheme to solve the problem and we exhibit the main biological features that can be accounted for by the model. Even though simplistic from the modeling point of view, we claim that this work provides the theoretical and numerical grounds for single cell migration modeling. In particular, specific chemical reactions that occurred at the cell membrane could be precisely described in forthcoming works.

7.3. Axis 3: Quantitative cancer modeling for biological and preclinical studies

7.3.1. Modelling of metastasis development

Computational Modelling of Metastasis Development in Renal Cell Carcinoma [2]

Team participants: Etienne Baratchart, Sébastien Benzekry, Thierry Colin, Olivier Saut

Other participants: Andreas Bikfalvi, Lindsay S. Cooley, Raphaël Pineau, Wilfried Souleyreau (LAMC - Laboratoire Angiogenèse et Micro-environnement des Cancers), Emeline J Ribot (RMSB - Résonance magnétique des systèmes biologiques)

To improve our understanding of the biology of the metastatic colonization process, we conducted a modelling study based on multi-modal data from an orthotopic murine experimental system of metastatic renal cell carcinoma. The standard theory of metastatic colonization usually assumes that secondary tumours, once established at a distant site, grow independently from each other and from the primary tumour. Using a mathematical model describing the metastatic population dynamics under this assumption, we challenged the theory against our data that included: 1) dynamics of primary tumour cells in the kidney and metastatic cells in the lungs, retrieved by green fluorescent protein tracking, and 2) magnetic resonance images (MRI) informing on the number and size of macroscopic lesions. While the model could fit the primary tumour and total metastatic burden, the predicted size distribution was not in agreement with the MRI observations. Moreover, the model was incompatible with the growth rates of individual metastatic tumours. To explain the observed metastatic patterns, we hypothesised that metastatic foci derived from one or a few cells could aggregate, resulting in a similar total mass but a smaller number of metastases. This was indeed observed in our data and led us to investigate the effect of spatial interactions on the dynamics of the global metastatic burden. We derived a novel mathematical model for spatial tumour growth, where the intra-tumour increase in pressure is responsible for the slowdown of the growth rate. The model could fit the growth of lung metastasis visualized by magnetic resonance imaging. As a non-trivial outcome from this analysis, the model predicted that the net growth of two neighbouring tumour lesions that enter in contact is considerably impaired (of $31\% \pm 1.5\%$, mean \pm standard deviation), as compared to the growth of two independent tumours. Together, our results have implications for theories of metastatic development and suggest that global dynamics of metastasis development is dependent on spatial interactions between metastatic lesions.

Modeling spontaneous metastasis following surgery: an in vivo-in silico approach [6]

Team participant: Sebastien Benzekry

Other participants: Amanda Tracz, Michalis Matri, Ryan Corbelli, Dominique Barbolosis, John Ebos (Buffalo University)

Rapid improvements in the detection and tracking of early-stage tumor progression aim to guide decisions regarding cancer treatments as well as predict metastatic recurrence in patients following surgery. Mathematical models may have the potential to further assist in estimating metastatic risk, particularly when paired with in vivo tumor data that faithfully represent all stages of disease progression. Herein we describe mathematical analysis that uses data from mouse models of spontaneous metastasis developing after surgical removal of orthotopically implanted primary tumors. Both presurgical (primary tumor) and postsurgical (metastatic) growth was quantified using bioluminescence and was then used to generate a mathematical formalism based on general laws of the disease (i.e. dissemination and growth). The model was able to fit and predict pre-/post-surgical data at the level of the individual as well as the population. Our approach also enabled retrospective analysis of clinical data describing the probability of metastatic relapse as a function of primary tumor size. In these data-based models, inter-individual variability was quantified by a key parameter of intrinsic metastatic potential. Critically, our analysis identified a highly nonlinear relationship between primary tumor size and postsurgical survival, suggesting possible threshold limits for the utility of tumor size as a predictor of metastatic recurrence. These findings represent a novel use of clinically relevant models to assess the impact of surgery on metastatic potential and may guide optimal timing of treatments in neoadjuvant (presurgical) and adjuvant (postsurgical) settings to maximize patient benefit.

Migration and orientation of endothelial cells on micropatterned polymers: A simple model based on classical mechanics [11]**Team participants:** Thierry Colin, Clair Poignard, Olivier Saut**Other participants:** Julie Joie, Marie-Christine Durrieu (IMB - Institut de Mathématiques de Bordeaux), Yifeng Lei (French Institute of Health and Medical Research, Paris)

Understanding the endothelial cell migration on micropatterned polymers, as well as the cell orientation is a critical issue in tissue engineering, since it is the preliminary step towards cell polarization and that possibly leads to the blood vessel formation. In this work, we derive a simple agent-based model to describe the migration and the orientation of endothelial cells seeded on bioactive micropatterned polymers. The aim of the modeling is to provide a simple model that corroborates quantitatively the experiments, without considering the complex phenomena inherent to cell migration. Our model is obtained thanks to a classical mechanics approach based on experimental observations. Even though its simplicity, it provides numerical results that are quantitatively in accordance with the experimental data, and thus our approach can be seen as a preliminary way towards a simple modeling of cell migration.

7.3.2. Tumor-host crosstalk**Host age is a systemic regulator of gene expression impacting cancer progression [3]****Team participant:** Sebastien Benzekry**Other participants:** Afshin Beheshti, Lili Ma, Philip Hahnfeldt, Lynn Hlatky (CCSB - Center of Cancer and Systems Biology), J. Tyson McDonald (University of Houston), Michael Peluso (Cancer Risk Factor Branch, Molecular Biology Laboratory)

Aging is the major determinant of cancer incidence, which, in turn, is likely dictated in large part by processes that influence the progression of early subclinical (occult) cancers. However, there is little understanding of how aging informs changes in aggregate host signaling that favor cancer progression. In this study, we provide direct evidence that aging can serve as an organizing axis to define cancer progression-modulating processes. As a model system to explore this concept, we employed adolescent (68 days), young adult (143 days), middle-aged (551 days), and old (736 days) C57BL/6 mice as syngeneic hosts for engraftment of Lewis lung cancer to identify signaling and functional processes varying with host age. Older hosts exhibited dysregulated angiogenesis, metabolism, and apoptosis, all of which are associated with cancer progression. $TGF\beta 1$, a central player in these systemic processes, was downregulated consistently in older hosts. Our findings directly supported the conclusion of a strong host age dependence in determining the host tumor control dynamic. Furthermore, our results offer initial mechanism-based insights into how aging modulates tumor progression in ways that may be actionable for therapy or prevention.

Capturing the Driving Role of Tumor-Host Crosstalk in a Dynamical Model of Tumor Growth [4]**Team participant:** Sebastien Benzekry**Other participants:** Afshin Beheshti, Philip Hahnfeldt, Lynn Hlatky (CCSB - Center of Cancer and Systems Biology)

In 1999, Hahnfeldt et al. proposed a mathematical model for tumor growth as dictated by reciprocal communications between tumor and its associated vasculature, introducing the idea that a tumor is supported by a dynamic, rather than a static, carrying capacity. In this original work, the carrying capacity was equated with the variable tumor vascular support resulting from the net effect of tumor-derived angiogenesis stimulators and inhibitors. This dynamic carrying capacity model was further abstracted and developed in our recent publication to depict the more general situation where there is an interaction between the tumor and its supportive host tissue; in that case, as a function of host aging. This allowed us to predict a range of host changes that may be occurring with age that impact tumor dynamics. More generally, the basic formalism described here can be (and has been), extended to the therapeutic context using additional optimization criteria. The model depends on three parameters: one for the tumor cell proliferation kinetics, one for the stimulation of the stromal support, and one for its inhibition, as well as two initial conditions. We describe here the numerical method to estimate these parameters from longitudinal tumor volume measurements.

7.3.3. *Metronomic oncology*

Metronomic Reloaded: Theoretical Models Bringing Chemotherapy into the Era of Precision Medicine [5]

Team participant: Sebastien Benzekry

Other participants: Eddy Pasquier, Dominique Barbolosi, Joseph Ciccolini, Nicolas André (CRO2 - Centre de recherches en oncologie biologique et oncopharmacologie), Bruno Lacarelle (Clinical Pharmacokinetics), Fabrice Barlési (Service d'Oncologie Multidisciplinaire et d'Innovations Thérapeutiques)

Oncology has benefited from an increasingly growing number of groundbreaking innovations over the last decade. Targeted therapies, biotherapies, and the most recent immunotherapies all contribute to increase the number of therapeutic options for cancer patients. Consequently, substantial improvements in clinical outcomes for some disease with dismal prognosis such as lung carcinoma or melanoma have been achieved. Of note, the latest innovations in targeted therapies or biotherapies do not preclude the use of standard cytotoxic agents, mostly used in combination. Importantly, and despite the rise of bioguided (a.k.a. precision) medicine, the administration of chemotherapeutic agents still relies on the maximum tolerated drug (MTD) paradigm, a concept inherited from theories conceptualized nearly half a century ago. Alternative dosing schedules such as metronomic regimens, based upon the repeated and regular administration of low doses of chemotherapeutic drugs, have emerged as possible strategies to improve response rates while reducing toxicities. The recent changes in paradigm in the way we theorize cancer biology and evolution, metastatic spreading and tumor ecology, alongside the recent advances in the field of immunotherapy, have considerably strengthened the interest for metronomic approaches. This work aims at reviewing the recent evolutions in the field of theoretical biology of cancer and computational oncology, with a focus on the consequences these changes have on the way we administer chemotherapy. In particular, a step towards developing adaptive dosing should help to further optimize the efficacy of metronomic therapy. There is a rising trend to establish personalized medicine in oncology. Developing extensive bio-guided strategies for decision-making in the choice of drugs to be administered is now a common practice at the bedside. Similarly, developing extensive model-guided strategies for decision-making in refining dosing and scheduling should be undertaken to achieve precision medicine in oncology.

7.3.4. *Protein-protein interaction networks*

Design principles for cancer therapy guided by changes in complexity of protein-protein interaction networks [7]

Team participant: Sebastien Benzekry

Other participants: Jack A Tuszyński (Alberta University), Edward Rietman, Giannoula Lakka Klement (Newman-Lakka Institute)

The ever-increasing expanse of online bioinformatics data is enabling new ways to, not only explore the visualization of these data, but also to apply novel mathematical methods to extract meaningful information for clinically relevant analysis of pathways and treatment decisions. One of the methods used for computing topological characteristics of a space at different spatial resolutions is persistent homology. This concept can also be applied to network theory, and more specifically to protein-protein interaction networks, where the number of rings in an individual cancer network represents a measure of complexity. Results: We observed a linear correlation of $R = -0.55$ between persistent homology and 5-year survival of patients with a variety of cancers. This relationship was used to predict the proteins within a protein-protein interaction network with the most impact on cancer progression. By re-computing the persistent homology after computationally removing an individual node (protein) from the protein-protein interaction network, we were able to evaluate whether such an inhibition would lead to improvement in patient survival. The power of this approach lied in its ability to identify the effects of inhibition of multiple proteins and in the ability to expose whether the effect of a single inhibition may be amplified by inhibition of other proteins. More importantly, we illustrate specific examples of persistent homology calculations, which correctly predict the survival benefit observed effects in clinical trials using inhibitors of the identified molecular target. Conclusions: We propose that computational approaches such as persistent homology may be used in the future for selection of molecular therapies in clinic.

The technique uses a mathematical algorithm to evaluate the node (protein) whose inhibition has the highest potential to reduce network complexity. The greater the drop in persistent homology, the greater reduction in network complexity, and thus a larger potential for survival benefit. We hope that the use of advanced mathematics in medicine will provide timely information about the best drug combination for patients, and avoid the expense associated with an unsuccessful clinical trial, where drug(s) did not show a survival benefit.

7.4. Other new results

Superconvergent Cartesian Methods for Poisson type Equations in 2D-domains [21]

Team participants: Olivier Gallinato, Clair Poinard

In this work, we present three superconvergent Finite Difference methods on Cartesian grids for Poisson type equations with Dirichlet, Neumann or Robin conditions. Our methods are based on finite differences and high-order discretizations of the Laplace operator, to reach the superconvergence properties, in the sense that the first-order (and possibly the second-order) derivatives of the numerical solution are computed at the same order as the solution itself. We exhibit the numerical conditions that have to be fulfilled by the schemes to get such superconvergences and extensively illustrate our purpose by numerical simulations. We conclude by applying our method to a free boundary problem for cell protrusion formation recently proposed by the authors and colleagues. Note that quasistatic Stefan-like problem can be accurately solved by our methods.

Adaptive radiotherapy in routine: The radiation oncologist's point of view [14]

Team participant: Olivier Saut

Other participants: Bénédicte Henriques de Figueiredo, Adeline Petit, Paul Sargos, Guy Kantor, Claudia Pouypoudat, Christina Zacharitou, Mikael Antoine (Institut Bergonié, radiology department)

Adaptive radiotherapy is defined as all processes leading to the modification of a treatment plan on the basis of patient-specific variations observed during the course of a treatment. This concept is currently of particular relevance due to the development of onboard volumetric imaging systems, which allow for daily viewing of variations in both tumour and organs at risk in terms of position, shape or volume. However, its application in routine clinical practice is limited due to the demanding nature of the processes involved (re-delineation and replanning) and increased dependence on available human resources. Even if "online" strategies, based on deformable image registration (DIR) algorithms, could lead to a reduction in both work and calculation time, for the moment their use is limited to the research field due to uncertainties surrounding the validity of results gathered. Other strategies without DIR can be used as "offline" or "hybrid offline-online" strategies that seem to offer a compromise between time consumption and therapeutic gain for the patient.

PHOENIX Project-Team

7. New Results

7.1. Tablet-Based Activity Schedule in Mainstream Environment for Children with Autism and Children with ID

Including children with Autism Spectrum Disorders (ASD) in mainstreamed environments creates a need for new interventions whose efficacy must be assessed in situ. We present a tablet-based application for activity schedules that has been designed following a participatory design approach involving mainstream teachers, special-education teachers and school aides. This application addresses two domains of activities: classroom routines and verbal communications. We assessed the efficiency of our application with two overlapping user-studies in mainstream inclusion, sharing a group of children with ASD. The first experiment involved 10 children with ASD, where 5 children were equipped with our tablet-based application and 5 were not equipped. We show that (1) the use of the application is rapidly self-initiated (after two months for almost all the participants) and that (2) the tablet-supported routines are better performed after three months of intervention. The second experiment involved 10 children equipped with our application; it shared the data collected for the 5 children with ASD and compared them with data collected for 5 children with Intellectual Disabilities – ID. We show that (1) children with ID are not autonomous in the use of the application at the end of the intervention; (2) both groups exhibited the same benefits on classroom routines; and, (3) children with ID improve significantly less their performance on verbal communication routines. These results are discussed in relation with our design principles. Importantly, the inclusion of a group with another neurodevelopmental condition provided insights about the applicability of these principles beyond the target population of children with ASD.

7.2. Age and active navigation effects on episodic memory: A virtual reality study

We investigated the navigation-related age effects on learning, proactive interference semantic clustering, recognition hits, and false recognitions in a naturalistic situation using a virtual apartment-based task. We also examined the neuropsychological correlates (executive functioning [EF] and episodic memory) of navigation-related age effects on memory. Younger and older adults either actively navigated or passively followed the computer-guided tour of an apartment. The results indicated that active navigation increased recognition hits compared with passive navigation, but it did not influence other memory measures (learning, proactive interference, and semantic clustering) to a similar extent in either age group. Furthermore, active navigation helped to reduce false recognitions in younger adults but increased those made by older adults. This differential effect of active navigation for younger and older adults was accounted for by EF score. Like for the subject-performed task effects, the effects from the navigation manipulation were well accounted for by item-specific/relational processing distinction, and they were also consistent with a source monitoring deficit in older adults.

7.3. Constraining application behaviour by generating languages

Writing a platform for reactive applications which enforces operational constraints is difficult, and has been approached in various ways. In this experience report, we detail an approach using an embedded DSL which can be used to specify the structure and permissions of a program in a given application domain. Once the developer has specified which components an application will consist of, and which permissions each one needs, the specification itself evaluates to a new, tailored, language. The final implementation of the application is then written in this specialised environment where precisely the API calls associated with the permissions which have been granted, are made available. Our prototype platform targets the domain of mobile computing ,

and is implemented using Racket. It demonstrates resource access control (e.g., camera, address book, etc.) and tries to prevent leaking of private data. Racket is shown to be an extremely effective platform for designing new programming languages and their run-time libraries. We demonstrate that this approach allows reuse of an inter-component communication layer, is convenient for the application developer because it provides high-level building blocks to structure the application, and provides increased control to the platform owner, preventing certain classes of errors by the developer.

7.4. A Unifying Notification System To Scale Up Assistive Services

Aging creates needs for assistive technology to support all activities of daily living (meal preparation, dressing, social participation, stove monitoring, *etc.*). These needs are mostly addressed by a silo-based approach that requires a new assistive service (*e.g.*, a reminder system, a pill prompter) to be acquired for every activity to be supported. In practice, these services manifest their silo-based nature in their user interactions, and more specifically, in the heterogeneity of their notification system. This heterogeneity incurs a cognitive cost that prevents scaling up assistive services and compromises adoption by older adults. We present an approach to scaling up the combination of technology-based, assistive services by proposing a unifying notification system. To do so, (1) we propose a decomposition of assistive services to expose their needs in notification; (2) we introduce a notification framework, allowing heterogeneous assistive services to homogeneously notify users; (3) we present how this notification framework is carried out in practice for an assisted living platform. We successfully applied our approach to a range of existing and new assistive services. We used our notification framework to implement an assistive platform that combines a variety of assistive services. This platform has been deployed and used 24/7 in the home of 15 older adults for up to 6 months. This study provides empirical evidence of the effectiveness and learnability of the notification system of our platform, irrespective of the cognitive and sensory resources of the user. Additional results show that our assisted living platform achieved high user acceptance and satisfaction.

7.5. Orchestrating Masses of Sensors: A Design-Driven Development Approach

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

7.6. Analysis of How People with Intellectual Disabilities Organize Information Using Computerized Guidance

Access to residential settings for people with intellectual disabilities (ID) contributes to their social participation, but presents particular challenges. Assistive technologies can help people perform activities of daily living. However, the majority of the computerized solutions offered use guidance modes with a fixed, unchanging sequencing that leaves little room for self-determination to emerge. The objective of the project was to develop a flexible guidance mode and to test it with participants, to describe their information organization methods. This research used a descriptive exploratory design and conducted a comparison between five participants with ID and five participants with no ID. The results showed a difference in the information organization methods for both categories of participants. The people with ID used more diversified organization methods (categorical, schematic, action-directed) than the neurotypical participants (visual, action-directed). These organization methods varied depending on the people, but also on the characteristics of the requested task. Furthermore, several people with ID presented difficulties when switching from virtual to real mode. These results demonstrate the importance of developing flexible guidance modes adapted to the users' cognitive strategies, to maximize their benefits. Studies using experimental designs will have to be conducted to determine the impacts of more-flexible guidance modes.

PLEIADE Team

6. New Results

6.1. Inference of metabolic networks

Participants: David Sherman [correspondant], Razanne Issa, Pascal Durrens.

We are particularly interested in incremental modeling of metabolic networks, where the target organism to be modeled is demonstrably similar to other organisms for which whole or partial models are available. The other organisms are typically strains of the same species as the target, or species with a close phylogenetic relation to the target species. The similarity is measured genomically at different scales: sequence polymorphisms, expansions and contractions in conserved protein families, and genome rearrangements. We have defined and refined two complementary methods for inferring metabolic models for target species.

In the same way that comparative analysis of genomes and proteomes makes it possible to define protein families that summarize protein-coding genes into phyletic patterns [24], comparative analysis of related metabolic models makes it possible to define network generalizations [26] that factor families of reactions and metabolites into summary graphs that preserve stoichiometry. These summaries can be used for expert curation and visualization [5]. An online demonstration tool is made available at <http://mimoza.bordeaux.inria.fr/>.

Starting from an existing reference metabolic network and measures of similarity between the reference and the target organisms' genomes, we can use knowledge-based inference to rewrite the reference network based on these differences, and thus obtain a draft network for the metabolisms of the target organism [2]. This rewriting, formalized in the Pantograph system, can be extended to an abductive logic framework as described in Razanne Issa's thesis [19]. Current work aims at extending the Pantograph and ab-Pantograph frameworks to leverage reaction classifications obtained by network generalization.

6.2. Bio-medicine and biotechnology

Participants: Pascal Durrens [correspondant], David Sherman.

6.2.1. Genome assembly for bio-medicine

We performed the assembly of the *Clavispora lusitaniae* (aka *Candida lusitaniae*) genome. Yeasts from the genus *Candida* are opportunistic human pathogens in immunocompromised patients, linked to a high mortality rate. Although *Candida albicans* is the major pathogen, related species are more and more isolated, such as *Clavispora lusitaniae* which is responsible for candidaemia in newborn babies and in onco-hematology patients.

Even though the genome of a *Clavispora lusitaniae* strain (ATCC 42720), isolated from a patient, has already been sequenced by the Broad Institute, we achieved the genome assembly of the wild type reference strain (CBS 6936) as patient isolates tend to harbor genome modifications. The assembly was computed from Illumina reads with a coverage of 30X, using the MINIA assembler from Inria GENSCALE team. We also looked for single nucleotide polymorphisms (SNPs) in the reads coming from 3 hypovirulent mutants impaired in the beta oxidation metabolic pathway. Some detected SNPs are now under experimental validation and we are going to make a Genome Announcement for the CBS 6936 genome.

6.2.2. Transcriptome assembly for bio-technology

We carried out the assembly of transcriptomes from different tissues of the African oil palm tree *Elaeis guineensis*. The goal of this project is twofold: (i) Select the most relevant genes involved in oil synthesis in order to implement heterologous expression of some of these genes in a cultivated plant recipient such as tobacco. Preliminary results on heterologous expression of 2-3 key genes/factors ended in 15% of dry weight of oil synthesis. New expression technology allows for simultaneous expression of 15-20 genes. Identifying the best candidates for co-expression will permit efficient heterologous oil synthesis. (ii) Identify the polymorphism of genes in a panel of 25 wild type isolates and of 5 production lineages of *Elaeis guineensis* in relation to the oil yield in different environmental conditions. In addition to a high variability of oil quantity (1-12 tons/ha/year), the relative amount of unsaturated fatty acids spans widely (15-55% dry weight) among the 30 *Elaeis guineensis* strains. Identification of polymorphisms will pave the way to genome-wide association genetics (GWAS) for the improvement of the oil resource.

In a first step, we produced assembled transcriptomes of ca. 300 million reads from 3 tissues (leaf, mesocarp, kernel) coming from a single tree, using state-of-the-art assembler TRINITY. Tuning of the software parameters was performed on the Inria PLAFRIM computation platform. About 20% of the assembled sequences revealed to be tissue-specific. Computation of the protein sequences deduced from the assembled transcripts gave a protein repertoire which was annotated using related sequences available in public databases. These transcript and protein sets will be used as a framework in the polymorphism studies.

6.3. Biodiversity and ecology

Participants: Alain Franc [correspondant], Jean-Marc Frigerio, Philippe Chaumeil, Razanne Issa, Leyla Mirvakhabova.

Our main activity has been code development in the framework of a research project with ONEMA, and preparing future development. Code development has been fostered with the work of Razanne Issa (CDD ONEMA) in the last three months of 2015, and preparation of further development has been fostered by welcoming Leyla Mirvakhabova (L3, National Research University Higher School of Economics (NRU HSE), Moscow, mathematics). Declic is a python library providing tools for analysing molecular data for biodiversity studies. The main object is a distance matrix, from which one can either build a point cloud in a Euclidean space with distances between points as close as possible from distances between reads (Multidimensional scaling), or to build a graph with edges between reads when their distance is smaller than a given threshold. Meanwhile, the team has developed the network around the Galaxy server where an early version of tools has been installed and made available, especially with SLU at Uppsala (Maria Khalert). Alain Franc has developed a collaboration with Olivier Coulaud and Pierre Blanchard (Hiepac) for efficient computation of eigenvectors and eigenvalues of large, dense and symmetric matrices, needed for scaling in Multidimensional scaling.

The work of Razanne Issa has made it possible to extend the declic library in the direction of machine learning, by incorporating tools from support vector machines through library sklearn. This development will be pursued in 2016. The work of Leyla Mirvakhabova has permitted a first incursion into topological data analysis as a possible approach for studying the shape of point clouds produced by multidimensional scaling. The collaboration with NRU HSE on this topic will be pursued in 2016.

POSET Team

7. New Results

7.1. Efficient interactive score

We have proposed a solution to the problem of real-time performance for interactive multimedia applications, specifically in the interpretation of interactive multimedia scores. For that, we have proposed a new parallel implementation of interactive scores on a reconfigurable hardware. We take advantage of the parallelism and reliability provided by Field Programmable Gate Arrays (FPGAs) to perform in real-time the hardware representation of scores. The results of the simulations show that our approach allows the system to react instantaneously to user interactions. Moreover, the real-time constraints of the score are satisfied [21].

7.2. Modeling with tile

In [3], [8] it has been observed that musical objects are conveniently modeled by tiles. These modeling experiments have been continued this year by showing, in particular, how both high-level music modeling and low-level signal combination can be modeled by means of tiles [23]. This has been further extended relating classical musical constructs with tile modeling features [34].

7.3. Tiles and inverse semigroups

In [10] it has already been observed that the theory of inverse semigroups⁰ is the adequate mathematical framework to define and study tiles and their languages. In this direction, we have shown that strings, trees and even many types of graphs can be unified into a notion of higher-dimensional strings [24], [35]. Using techniques of partial algebra [4], this notion recovers advanced results on formal languages of graphs of bounded tree-width⁰, which shows the robustness of the approach.

7.4. Reactive programming with tile

The first step towards programming music with tiles is proposed as a Domain-Specific Language : the T-calculus [9]. Further collaboration with Paul Hudak [7] led us to various implementation experiments on top of Haskell [30], [33], [29]. Within the ADT “Tuilage” and S. Archipoff’s PhD thesis in progress, we eventually managed to integrate tile modeling into reactive programming as illustrated, in December 2015, by the first concert of the Idex Arts & Science project “Sound of Algorithm” in collaboration with the musician Edwin Buger.

7.5. Behavioral properties of higher-order programs

In a series of results [28], [27], we have been able to cast to traditional denotational semantics the behavioral properties captured by Monadic Second Order Logic (MSOL) and weak MSOL. The main difficulty was to represent infinitary properties in finitary models. From a foundational point of view, these results exhibit once more the robustness of approaches based on recognizability to capture complex properties of programs. They also make salient the problem of program evaluation in finite models as a milestone towards effective model-checking of higher-order programs.

⁰See [43] for general presentation of inverse semigroup theory, and [45], [44] for graph-based representation of inverse semigroup elements.

⁰See [38] for an up-to-date presentation of the formal language theory of graphs.

7.6. Art & Science project

This year has seen the members of PoSET involved in a number of Art & Science projection, especially some granted by Idex Bordeaux, including but not limited to : *Illumination* created Aurelio Edler-Copes, in partnership with compagnie Eclats, performed in November 2015 at Molière Theater, *Mobiles* and *Le Chant du filament #2* respectively created by Renaud Rubiano and Nicolas Villenave, performed or displayed in November 2015 during FACTS festival, the Art and Science Festival of Bordeaux University, and, *Le son des algorithmes* with Edwin Bugar that led to a first musical performance in December 2015.

POTIOC Project-Team

7. New Results

7.1. Pointing in Spatial Augmented Reality from 2D Pointing Devices

Participants: Renaud Gervais, Jérémy Frey, Martin Hachet.



Figure 3. A user reaches a target displayed on a spatially augmented object with an indirect input device

Spatial Augmented Reality (SAR) opens interesting perspectives for new generations of mixed reality applications. Compared to traditional human-computer interaction contexts, there is little work that studies user performance in SAR. In this project, we present an experiment that compares pointing in SAR versus pointing in front of a screen, from standard pointing devices (mouse and graphics tablet). The results showed that the participants tend to interact in SAR in a way that is similar to the screen condition, without a big loss of performance [30] (See Figure 3).

7.2. Tangible Viewports

Participants: Renaud Gervais, Joan Sol Roo, Martin Hachet.

Spatial augmented reality and tangible interaction enrich the standard computer I/O space. Systems based on such modalities offer new user experiences and open up interesting perspectives in various fields. On the other hand, such systems tend to live outside the standard desktop paradigm and, as a consequence, they do not benefit from the richness and versatility of desktop environments. In this work, we propose to join together physical visualization and tangible interaction within a standard desktop environment. We introduce the concept of Tangible Viewport, an on-screen window that creates a dynamic link between augmented objects and computer screens, allowing a screen-based cursor to move onto the object in a seamless manner (Figure 4). We describe an implementation of this concept and explore the interaction space around it. A preliminary evaluation shows that the metaphor is transparent to the users while providing the benefits of tangibility [31].

7.3. Tobe

Participants: Renaud Gervais, Jérémy Frey, Alexis Gay, Fabien Lotte, Martin Hachet.




Figure 4. A user interacts with an object located in front of the screen as if the object was rendered on screen




Figure 5. Two users are relaxing together using Tobe as a biofeedback for heartrate and breathing

We propose Tobe, a toolkit for creating Tangible Out-of-Body Experiences: exposing the inner states of users using physiological signals such as heart rate or brain activity. Tobe can take the form of a tangible avatar displaying live physiological readings to reflect on ourselves and others. Such a toolkit could be used by researchers and designers to create a multitude of potential tangible applications, including (but not limited to) educational tools about Science Technologies Engineering and Mathematics (STEM) and cognitive science, medical applications or entertainment and social experiences with one or several users or Tobes involved. Through a co-design approach, we investigated how everyday people picture their physiology and we validated the acceptability of Tobe in a scientific museum. We also give a practical example where two users relax together, with insights on how Tobe helped them to synchronize their signals and share a moment, as illustrated in Figure 5 [29].

7.4. Inner Garden

Participants: Joan Sol Roo, Renaud Gervais, Martin Hachet.




Figure 6. Inner garden is an augmented sandbox which depicts an evolving world reflecting the inner state of the user

We present a prototype of an augmented sandbox where the sand is used to create a miniature living world, designed as an ambient display for contemplation and self-reflection. The landscape can be reshaped at any time. Once the sand is left still for a moment, the world starts evolving – vegetation grows, water flows and creatures move around – according to the user’s internal state. We use a consumer-grade EEG and breathing sensors to reflect on frustration and meditative states of users, which they can monitor by looking at the sandbox (Figure 6) [49].

7.5. Augmented geographic maps

Participants: Julia Chatain, Marie Demangeat, Anke Brock, Martin Hachet.

Interactive geographic maps are today widely available, but remain mostly limited to standard interaction contexts. We introduce SyMAPse [48], a spatial augmented reality map, based on the **PapARt framework**. In our prototype, we use augmented reality to display a virtual map on a physical piece of paper, thus keeping features of both media. Thanks to the digital map base, users can pan, zoom and even change the basemap. At the same time, the paper base allows users to manipulate the map physically and so to interact in a more “natural” way, as well as to draw on the paper using regular pens. In a preliminary study with visitors of the "Cap Sciences" science center, we compared interaction techniques based on touch, tangible and spatial modalities for these three common map functions: zooming, panning, and changing the basemap. Our results suggest that object-based and spatial interaction may be advantageous over touch in our augmented reality setup.



Figure 7. Interacting with an augmented geographic map using tangible, spatial and multitouch interaction




Figure 8. HOBIT mixes physical and virtual elements to teach optics.

7.6. HOBIT - Hybrid Optical Bench for Innovative Teaching

Participants: David Furio, Benoit Coulais, Martin Hachet.

Experiments in optics are essential for learning and understanding physical phenomena. The problem with these experiments is that they are generally time consuming for both their construction and their maintenance, potentially dangerous through the use of laser sources, and often expensive due to high technology optical components. We propose to simulate such experiments by using hybrid systems that exploit both spatial augmented reality and tangible interaction (See Figure 8). In particular, we focus on one of the most popular optical experiments: Michelson interferometer. In our approach, we target a highly interactive system where students are able to interact in real time with the Augmented Michelson Interferometer (AMI) to observe, test hypotheses and then to enhance their comprehension. Compared to a fully digital simulation, we are investigating an approach that benefits from both physical and virtual elements, and where the students experiment by manipulating 3D-printed physical replicas of optical components (e.g. lenses and mirrors). Our objective is twofold. First, we want to ensure that the students will learn with our simulator the same concepts and skills that they learn with traditional methods. Second, we hypothesis that such a system opens new opportunities to teach optics in a way that was not possible before, by manipulating concepts beyond the limits of observable physical phenomena. To reach this goal, we have built a complementary team composed of experts in the field of optics, human-computer interaction, computer graphics, sensors and actuators, and education science. HOBIT is a joint project between Inria and Université de Bordeaux (I dex CPU – LAPHIA), in collaboration with Université de Lorraine (team PErSEUs). [28]

7.7. Mixed Reality to improve children's interaction with astronomical concepts

Participant: Martin Hachet.

This project stands on a collaboration with Stéphanie Fleck from Université de Lorraine. To make astronomical learning more efficient for young pupils, we have designed an Augmented Inquiry-Based Learning Environment (AIBLE): HELIOS. Because manipulations in astronomy are intrinsically not possible, we propose to manipulate the underlying model. With HELIOS, virtual replicas of the sun, moon and earth are directly manipulated from tangible manipulations. This digital support combines the possibilities of Augmented Reality (AR) while maintaining intuitive interactions following the principles of didactic of sciences. Light properties are taken into account and shadows of Earth and Moon are directly produced by an omnidirectional light source associated to the virtual Sun. This AR environment provides users with experiences they would otherwise not be able to experiment in the physical world. Our main goal is that students can take active control of their learning, express and support their ideas, make predictions and hypotheses, and test them by conducting investigations. [24][23]

7.8. Combining and Revealing Spaces for Musical Performances

Participant: Martin Hachet.

In collaboration with University of Bristol (Florent Berthaut, Diego Martinez, and Sriram Subramanian) we have designed a mixed-reality environment for musical performances that allows for freely displaying virtual content on stage, such as 3D virtual musical interfaces or visual augmentations of instruments and performers. This environment, called Reflets, relies on spectators and performers revealing virtual objects by slicing through them with body parts or objects, and on planar slightly reflective transparent panels that combine the stage and audience spaces. It allows for placing virtual content anywhere on large stages, even overlapping with physical elements and provides a consistent rendering of this content for large numbers of spectators. It also preserves non-verbal communication between the audience and the performers, and is inherently engaging for the spectators. Reflets opens musical performance opportunities such as augmented interaction between musicians and novel techniques for 3D sound shapes manipulation [20].

7.9. Improving User-Training for Brain-Computer Interfaces

Participants: Martin Hachet, Emilie Jahanpour, Camille Jeunet, Fabien Lotte, Boris Masencal, Julia Schumacher.

While Mental Imagery based Brain-Computer Interfaces (MI-BCIs) are promising for many applications, ranging from assistive technologies for motor disabled patients to video games, their usability “out-of-the-lab” has been questioned due to their lack of reliability: literature reports that 15% to 30% of users cannot control such a technology, while most of the remaining users obtain only modest performances. As controlling an MI-BCI requires the acquisition of specific skills (i.e., producing stable and distinct brain-activity patterns), an adapted training is necessary. Thus, the main objective of our project is to improve the user training to facilitate the acquisition of MI-BCI related skills. In order to do so, we focused on two axes [18]: (1) the impact of the user-profile and (2) the impact of the protocol on MI-BCI performance.

Concerning the impact of the user-profile, our results ([40], [14]) suggested an important impact of some aspects of the personality (such as the tension and autonomy levels) as the spatial abilities (i.e., the ability to produce, interpret and transform mental imageries). On the one hand, we are working on learning companions, whose goal would be to provide the learners with a specific emotional support, based on their profile and on their cognitive state. On the other hand, we are currently implementing and testing a spatial ability training in order to test the hypothesis of a causal effect of the spatial abilities on MI-BCI performance [39]. In other words, we would like to know if increasing spatial abilities would result in better MI-BCI performance. One application of such a research is stroke rehabilitation. Indeed, motor after-effects are usual following a stroke. MI-BCI have been shown very useful to facilitate the rehabilitation process, which consists in enhancing brain plasticity through motor-imagery, as they enable to visualise the BCI activity while the patients perform MI-tasks. However, MI-tasks tend to increase the depressive state of the patients as they remind them they lost the ability to move their limb. Thus, as spatial ability exercises (e.g., mental rotation) activate the motor cortex, they could be used as more transparent rehabilitation exercises to trigger brain plasticity.

Second, concerning the impact of the protocol, we completed a study (see activity report 2014) in which we asked the participants to use the standard MI-BCI training protocol to learn to perform simple motor tasks: drawing circles and triangles on a graphic tablet. As it would have been the case for an MI-BCI experiment, they had to find the right strategy so that the system recognises the task they were performing. Seventeen percent of the participants (N=54) showed difficulties in performing these tasks. Also, when we selected the 10 best and 10 worst performers of this experiment and asked them to use an MI-BCI (by imagining left and right-hand movements), it appeared that the ones who had difficulties in performing the simple motor tasks improved in terms of performance during the MI-BCI experiment, while the participants who performed well during the motor experiment did not progress during the second. Furthermore, we have shown that tactile feedback was more efficient than an equivalent visual feedback in a multitasking context [32]. Based on a literature review, this could be due to an increased sense of agency (i.e., the feeling to be in control). We are thus currently exploring the impact of the sense of agency on MI-BCI performance. Finally, still regarding the feedback, we explored what kind of information could help the user to perform better mental imagery tasks. As such, we look for physiological features that could predict whether a mental task will be correctly recognized by the BCI, and that could be understood by the user. Among the different features we explored, it appears that the user’s relaxation (from a muscular point of view), as measured in EMG activity collected by EEG channels, is one of such features. We are currently building and exploring new BCI training protocols that provide additional information about the user’s muscular relaxation as complementary feedback [34].

7.10. EEG Signal Processing

Participant: Fabien Lotte.

To make BCI practical and useful, we need to make them reliable, i.e., able to recognize the users’ mental commands, despite noise and non-stationarities [42]. We also need to reduce their calibration time, as current systems need many examples from each user to calibrate the system for this specific user. This year we addressed these two issues with two different studies.

In order to reduce BCI calibration times, we first surveyed existing approaches, these approaches being notably based on regularization, user-to-user transfer, semi-supervised learning and a-priori physiological information. We then proposed new tools to reduce BCI calibration time. In particular, we proposed to generate artificial EEG trials from the few EEG trials initially available, in order to augment the training set size. These artificial EEG trials are obtained by relevant combinations and distortions of the original trials available. We proposed 3 different methods to do so. We also proposed a new, fast and simple approach to perform user-to-user transfer for BCI. Finally, we studied and compared offline different approaches, both old and new ones, on the data of 50 users from 3 different BCI data sets. This enabled us to identify guidelines about how to reduce or suppress calibration time for BCI [16].

In order to increase BCI robustness, we performed an empirical comparison of covariance matrix averaging methods for EEG signal classification. Indeed, averaging EEG signal covariance matrices is a key step in designing brain-computer interfaces (BCI) based on the popular common spatial pattern (CSP) algorithm. BCI paradigms are typically structured into trials and we argue that this structure should be taken into account. Moreover, the non-Euclidean structure of covariance matrices should be taken into consideration as well. We reviewed several approaches from the literature for averaging covariance matrices in CSP and compared them empirically on three publicly available data sets. Our results showed that using Riemannian geometry for averaging covariance matrices improves performances for small dimensional problems, but also the limits of this approach when the dimensionality increases [36].

7.11. ECoG-based analysis of Speech processes

Participant: Fabien Lotte.

Acoustic speech output results from coordinated articulation of dozens of muscles, bones and cartilages of the vocal mechanism. While we commonly take the fluency and speed of our speech productions for granted, the neural mechanisms facilitating the requisite muscular control are not completely understood. Previous neuroimaging and electrophysiology studies of speech sensorimotor control has typically concentrated on speech sounds (i.e. phonemes, syllables and words) in isolation; sentence-length investigations have largely been used to inform coincident linguistic processing. In this study, we examined the neural representations of segmental features in the context of fluent, continuous speech production. We used recordings from the cortical surface (electrocorticography (ECoG)) to simultaneously evaluate the spatial topography and temporal dynamics of the neural correlates of speech articulation that may mediate the generation of hypothesized gestural or articulatory scores. We found some aspects of speech production (place of articulation) involved broad networks of brain regions during all phases of speech production: preparation, execution and monitoring. Other aspects (manner of articulation and voicing status) were dominated by auditory cortical responses after speech had been initiated. These results provide a new insight into the articulatory and auditory processes underlying speech production in terms of their motor requirements and acoustic correlates (see Figure 9 , [15]).

7.12. Toward a portable tangible EEG interface

Participants: Maxime Duluc, Thibault Laine, Jérémy Frey, Renaud Gervais, Fabien Lotte, Martin Hachet.

Last year we presented **Teegi**, the first interface that combines electroencephalographic (EEG) recordings and tangible interaction in order to let novices learn about how their brain works. By displaying EEG activity in real time on a support that is easy to manipulate and to comprehend, Teegi is a good tool for scientific outreach, that raises public interest.

Yet, the gap between research projects and the field is not often bridged. While our past prototype used an external projector and a supplementary tracking device to display information onto the head of the puppet, over the course of the year we developed a semi-spherical display based on LEDs (see Figure 10). By embedding all the electronics into the puppet, Teegi will be easier to bring outside the laboratory. Thanks to these technological advances, real-life applications of the system are finally within reach.




Figure 9. Exemple of the ECoG signature of vowel phonemes.




Figure 10. Development version of a portable version of Teegi, a Tangible EEG Interface. An array of LEDs and a diffuser replace the use of spatial augmented reality.

7.13. Electroencephalography-based evaluation of user experience

Participants: J r my Frey, Maxime Daniel, Dennis Wobrock, Julien Castet, Martin Hachet, Fabien Lotte.

Designing user interfaces requires adequate evaluation tools to ensure good usability and user experience. While many evaluation tools are already available and widely used, existing approaches generally cannot provide continuous and objective measures of usability qualities during interaction without interrupting the user. On the other hand, the measure of brain activity by the mean of electroencephalography (EEG) is mature enough to assess mental states. Combined with existing methods, such tool can be used to strengthen the understanding of user experience.




Figure 11. Schematic view of a user performing 3D manipulations tasks with the CubTile input device. His/her mental effort are monitored based on brain signals (electroencephalography).

In [35] we studied 3D object manipulation tasks. We showed how mental workload can be estimated from EEG, and then measured it on 8 participants during an actual 3D object manipulation task with an input device known as the CubTile (see figure 11). These first results suggested that we could continuously assess the 3DUI and/or interaction technique ease-of-use.




Figure 12. A keyboard (left) can be compared with a touch interface (middle) using a continuous measure of mental workload through electroencephalography (right).

We pushed further these finding in a second study [26], where we have developed a set of methods to continuously estimate the user's mental workload, attention level and recognition of interaction errors during

different interaction tasks. We validated these measures in a controlled virtual environment and showed how they can be used to compare different interaction techniques – for instance a keyboard and a touch-based interface (see Figure 12).

Thanks to such framework, EEG becomes a useful addition to the repertoire of available evaluation tools, enabling a finer grain assessment of the ergonomic qualities of computer systems.

7.14. Classifying EEG Signals during Stereoscopic Visualization to Estimate Visual Comfort

Participants: Jérémy Frey, Aurélien Appriou, Fabien Lotte, Martin Hachet.




Figure 13. Setup of the experiment, with a subject being presented with stereoscopic images while his EEG signals are being recorded.

With stereoscopic displays a sensation of depth that is too strong can impede visual comfort and may result in fatigue or pain. We used Electroencephalography (EEG) to develop a novel brain-computer interface that monitors users' states in order to reduce visual strain. We present the first system that discriminates comfortable conditions from uncomfortable ones during stereoscopic vision of still images using EEG [13], [25] – see Figure 13). In particular, we show that changes in event-related potentials' (ERPs) amplitudes following stereoscopic objects presentation can be used to estimate visual comfort. Our system reacts within 1 second to depth variations, achieving 63% accuracy on average (up to 76%) and 74% on average when 7 consecutive variations are measured (up to 93%). Performances are stable ($\approx 62.5\%$) when a simplified signal processing is used to simulate online analyses or when the number of EEG channels is lessened. This study could lead to adaptive systems that automatically suit stereoscopic displays to users and viewing conditions. For example, it could be possible to match the stereoscopic effect with users' state by modifying the overlap of left and right images according to the classifier output

REALOPT Project-Team

7. New Results

7.1. Improving Branch-and-Price Methods

We have made progress on stabilization techniques and math-heuristics that have become essential components for Branch-and-Price methods.

Smoothing and proximal methods based on penalizing the deviation from the incumbent dual solution have become standards of the domain. Interpreting column generation as cutting plane strategies in the dual problem, we analyze in [26] the mechanisms on which stabilization relies. In particular, the link is established between smoothing and in-out separation strategies to derive generic convergence properties. For penalty function methods as well as for smoothing, we describe proposals for parameter self-adjusting schemes. Such schemes make initial parameter tuning less of an issue as corrections are made dynamically. Such adjustments also allow to adapt the parameters to the phase of the algorithm. We provide extensive test reports that validate our self-adjusting parameter scheme and highlight their performances. Our results also show that using smoothing in combination with penalty function yields a cumulative effect on convergence speed-ups.

Effects of stabilization techniques can be seen in practice. Routing and logistics applications are often viewed as intractable for exact optimization tools. Although such problems are naturally suited for a decomposition approach, branch-and-price-and-cut algorithms of the literature typically do not scale to the size of real-life instances. Some recent progress in stabilization techniques amongst other advances (such as diving heuristics, strong branching, and the combination with cutting plane approaches) generate new ambitions for column generation approach in solving approximately very large scale instances. Let us for instance point to the new benchmarks for the Capacitated Vehicle Routing Problem (CVRP) in [62]. The paper [24] illustrates this trend, showing exact results for freight transportation instances of a scale never considered before. Our column generation algorithm yields dual bounds and serves as the core procedure for a primal heuristic. The overall procedure is quite competitive in great part due to the convergence speed-ups resulting from efficient stabilization schemes. It typically provides optimal solutions as primal and dual bounds tend to be equal. The very large scale freight transportation instances (with up to 1,025 stations, 5,300 demands, and 12,651 rail cars) were submitted to us by our Russian partner Freight-One.

Math-heuristics have become an essential component in mixed integer programming (MIP) solvers. Extending generic MIP heuristics, our study in [28] outlines generic procedures to build primal solutions in the context of a Branch-and-Price approach and reports on their performance. Rounding the linear relaxation solution of the Dantzig-Wolfe reformulation, which is typically tighter than that of the original compact formulation, sometimes produces better solutions than state-of-the-art specialised heuristics as revealed by our numerical experiments. We focus on the so-called diving methods and their combination with diversification-intensification paradigms such as Limited Discrepancy Search, sub-MIPing, relaxation induced neighbourhood search, local branching, and strong branching. The dynamic generation of variables inherent to a column generation approach requires specific adaptation of heuristic paradigms. Our contribution lies in proposing simple strategies to get around these technical issues. Our numerical results on Generalized Assignment, Cutting Stock, and Vertex Coloring problems sets new benchmarks, highlighting the performance of diving heuristics as generic procedures in a column generation context.

7.2. Dual feasible functions

Dual-feasible functions have proved to be very effective for generating fast lower bounds and valid inequalities for integer linear programs with knapsack constraints. However, a significant limitation is that they are defined only for positive arguments. Extending the concept of dual-feasible function to the general domain and range R is not straightforward. In [10], we propose the first construction principles to obtain general functions with domain and range R , and we show that they lead to non-dominated maximal functions.

7.3. Allocation algorithms in Cloud platforms

In the context of service hosting in large-scale datacenters, we provide [11] a deep analysis of a cluster data trace recently released by Google and we focus on a number of questions which have not been addressed in previous studies. In particular, we describe the characteristics of job resource usage in terms of dynamics (how it varies with time), of correlation between jobs (identify daily and/or weekly patterns), and correlation inside jobs between the different resources (dependence of memory usage on CPU usage). From this analysis, we derive scalable formalizations of the allocation problem which encompass most job features. In [19], [22], we study one such model, where long-running services experience demand variations with a periodic (daily) pattern. Such services account for most of the overall CPU demand. This leads to an allocation problem where the classical Bin-Packing issue is augmented with the possibility to co-locate jobs whose peaks occur at different times of the day, which is bound to be more efficient than the usual approach that consists in over-provisioning for the maximum demand. We propose mathematical formulations, column generation approaches, and analyze their performance compared to standard packing heuristics (such as Best-Fit or First-Fit Decreasing). We show that taking periodicity of demand into account allows for a substantial improvement on machine utilization in the context of large-scale, state-of-the-art production datacenters, and that column generation allows to obtain quasi-optimal solutions in reasonable time.

7.4. Scheduling and placement for HPC

With the complexification of the architecture of HPC nodes (multicores, non uniform memory access, GPU and accelerators), a recent trend in application development is to explicitly express the computations as a task graph, and rely on a specialized middleware stack to make scheduling decisions and implement them. Traditional algorithms used in this community are dynamic heuristics, to cope with the unpredictability of execution times. In [17], [18] we analyze the performance of static and hybrid strategies, obtained by adding more static (resp. dynamic) features into dynamic (resp. static) strategies. Our conclusions are somehow unexpected in the sense that we prove that static-based strategies are very efficient, even in a context where performance estimations are not very good.

Another study [13] focuses on the memory-constrained case, where tasks may produce large data. A task can only be executed if all input and output data fit into memory, and a data can only be removed from memory after the completion of the task that uses it as an input data. Trees of such tasks arise in the multifrontal method of sparse matrix factorization. Minimizing the peak memory required on a single processor is well studied, [13] extends the problem to multiple processors, where both makespan and memory need to be minimized. We study the computational complexity of this problem and provide inapproximability results even for unit weight trees. We design a series of practical heuristics achieving different trade-offs between the minimization of peak memory usage and makespan. Some of these heuristics are able to process a tree while keeping the memory usage under a given memory limit. The different heuristics are evaluated in an extensive experimental evaluation using realistic trees.

In [20], we perform another study of static, dynamic and hybrid strategies in the context of load balancing and data placement for matrix multiplication in heterogeneous machines. Through a set of extensive simulations, we analyze the behavior of static, dynamic, and hybrid strategies, and we assess the possible benefits of introducing more static knowledge and allocation decisions in runtime libraries. In [21], we consider the purely static problem, modeled as a partitioning of a square into a set of zones of prescribed areas, while minimizing the overall size of their projections onto horizontal and vertical axes. We combine two ideas from the literature (recursive partitioning, and optimal solution structure for low number of processors) to obtain a non-rectangular recursive partitioning (NRRP), whose approximation ratio is $\frac{2}{\sqrt{3}} \simeq 1.15$, improving over the previous 1.25 ratio. Moreover, we observe on a large set of realistic platforms built from CPUs and GPUs that this proposed NRRP algorithm allows to achieve very efficient partitionings on all considered cases.

7.5. Production scheduling

Together with Shunji Tanaka, from Kyoto University, we developed Lagrangian relaxation-based methods for solving min-sum shop scheduling problems. In our studies, large scale network flow formulations of the problems are suggested together with strong Lagrangian bounds based on these formulations.

In [23], we consider the flow-shop problem on two machines with sequence-independent setup times to minimize total completion time. To cope with the size of the network, filtering procedures are developed. To solve the problem to optimality, we embed the Lagrangian bounds into two branch-and-bound algorithms. The best algorithm is able to solve all 100-jobs instances of our testbed with and without setup times, thus significantly outperforming the best algorithms in the literature, which were limited to instances with 30 and 45 jobs respectively.

In [25], we propose a new dual bound for the job-shop problem with the objective of minimizing the sum of completion costs of the operations. The bound is obtained by a Lagrangian relaxation that decomposes the problem into two types of large network flow problems: one dealing with the precedence constraints among operations of a same job, and the other one satisfying the disjunctive constraints related to the machines. Numerical experiments on the just-in-time job-shop problem show that the method is able to improve the existing lower bounds significantly.

7.6. Clustering problems

Clustering problems, and in particular partitioning problems, are widespread in combinatorial optimization. The goal is to partition a set of items in subset satisfying various constraints such as knapsack constraints, cardinality constraints, connectivity constraints, and so on. Beside the PhD thesis of Jérémy Guillot that aims to develop aggregating techniques to handles large scale instances for partitioning problems, the team also study some particular versions.

In [15] we present the application of branch-and-price approaches to the automatic version of the Software Clustering Problem. To tackle this problem, we apply the Dantzig-Wolfe decomposition to a formulation from literature. Given this, we present two Column Generation (CG) approaches to solve the linear programming relaxation of the resulting reformulation: the standard CG approach, and a new approach, which we call Staged Column Generation (SCG). Also, we propose a modification to the pricing subproblem that allows to add multiple columns at each iteration of the CG. We test our algorithms in a set of 45 instances from the literature. The proposed approaches were able to improve the literature results solving all these instances to optimality. Furthermore, the SCG approach presented a considerable performance improvement regarding computational time, number of iterations and generated columns when compared with the standard CG as the size of the instances grows.

In collaboration with researchers from University Paris 6 and Paris 13, we also study the problem of partitioning a geographical area in connected parcels. A first step of this study was to cut the area in two connected parcels while minimizing the dissimilarities inside each parcels. Such partitioning is also called a bond. It happens that in series-parallel graph, a bond correspond to a circuit in the dual graph. In [12], we give a full description of the circuit polytope on series-parallel graphs. We first show the existence of a compact extended formulation. Though not being explicit, its construction process helps us to inductively provide the description in the original space. As a consequence, using the link between bonds and circuits in planar graphs, we also describe the bond polytope on series-parallel graphs.

7.7. Tour scheduling with multi-skill heterogeneous workforce

In [14], we address a multi-activity tour scheduling problem with time varying demand. The objective is to compute a team schedule for a fixed roster of employees in order to minimize the over-coverage and the under-coverage of different parallel activity demands along a planning horizon of one week. Numerous complicating constraints are present in our problem: all employees are different and can perform several different activities during the same day-shift, lunch breaks and pauses are flexible, demand is given for 15 minutes periods. Employees have feasibility and legality rules to be satisfied, but the objective function does not account for any quality measure associated with each individual's schedule. More precisely, the problem mixes simultaneously

days-off scheduling, shift scheduling, shift assignment, activity assignment, pause and lunch break assignment. To solve this problem, we developed four methods: a compact Mixed Integer Linear Programming model, a branch-and-price like approach with a nested dynamic program to solve heuristically the subproblems, a diving heuristic and a greedy heuristic based on our subproblem solver. The computational results, based on both real cases and instances derived from real cases, demonstrate that our methods are able to provide good quality solutions in a short computing time. Our algorithms are now embedded in a commercial software, which is already in use in a mini-mart company.

7.8. Traffic routing in optical networks

In [16], we consider a multi-layer network design model arising from a real-life telecommunication application where traffic routing decisions imply the installation of expensive nodal equipment. Customer requests come in the form of bandwidth reservations for a given origin destination pair. Bandwidth demands are expressed as multiples of nominal granularities. Each request must be single-path routed. Grooming several requests on the same wavelength and multiplexing wavelengths in the same optical stream allow a more efficient use of network capacity. However, each addition or withdrawal of a request from a wavelength requires optical to electrical conversion and the use of cross-connect equipment with expensive ports of high densities. The objective is to minimize the number of required ports of the cross-connect equipment. We deal with backbone optical networks, therefore with networks with a moderate number of nodes (14 to 20) but thousands of requests. Further difficulties arise from the symmetries in wavelength assignment and traffic loading. Traditional multi-commodity network flow approaches are not suited for this problem. Instead, four alternative models relying on Dantzig-Wolfe and/or Benders' decomposition are introduced and compared. The formulations are strengthened using symmetry breaking restrictions, variable domain reduction, zero-one discretization of integer variables, and cutting planes. The resulting dual bounds are compared to the values of primal solutions obtained through hierarchical optimization and rounding procedures. For realistic size instances, our best approaches provide solutions with optimality gap of approximately 5% on average in around two hours of computing time.

7.9. Dense sphere packing

In [27], we consider the sphere packing problem in arbitrary dimension: what is the maximum fraction Δ_n of the Euclidean space \mathbb{R}_n that can be covered by unit balls with pairwise disjoint interiors?

Δ_n is known for only for some small values of n , and when n grows, we only have lower bounds. A trivial lower bound states that for every n , $\Delta_n \geq 2^{-n}$. Minkowski and Hlwaka's Theorem (1905) improves this lower bound by a factor 2: $\Delta_n \geq 2 \times 2^{-n}$. Asymptotic improvements of this bound were obtained (from Rogers, 1947 up to Ball, 1992), all of them being of the form $\Delta_n \geq cn2^{-n}$ where c is a constant.

This problem has a natural reformulation in graph theoretic terms as follows: let G denote the graph whose vertices are the points of the Euclidean space and edges are pair of vertices at distance at most 2 one from the other. The independent sets of G are the sphere packings: so, finding a maximum-density sphere packing is the same as finding a maximum-density independent set in this infinite graph. By using graph theoretic arguments only, Krivelevich et al. established that $\Delta_n \geq 0.01n2^{-n}$ for sufficiently large n .

In a recent breakthrough, Venkatesh introduced the first superlinear improvement: there are infinitely many n such that $\Delta_n \geq cn \log \log n 2^{-n}$, where c is a constant. Venkatesh's result is however non-constructive.

In this joint work with C. Bachoc and P. Moustrou, we give a constructive proof of Venkatesh's lower bound.

This study has been carried out with financial support from the French State, managed by the French National Research Agency (ANR) in the frame of the " Investments for the future " Programme IdEx Bordeaux - CPU (ANR-10-IDEX-03-02).

SISTM Project-Team

7. New Results

7.1. Time-Course Gene Set Analysis for Longitudinal Gene Expression Data

The application of TcGSA methodology has revealed the commitment of inflammatory pathways and T-cell pathway in response of DC-based vaccine.

STORM Team

7. New Results

7.1. Memory Management and Distributed Computing with StarPU

Task-based programming models manage to abstract away much of the architecture complexity while efficiently meeting the performance challenge, even at large scale. While computation scheduling has been well studied, the dynamic management of memory resource subscription inside such run-times has however been little explored, despite the fact that the lookahead, anticipative capabilities offered by the decoupled task submission/task execution steps may sometime come with a high memory cost, especially in distributed context where buffers for receiving incoming contributions have to be accounted for. We therefore studied the cooperation between a task-based application code and a run-time system engine to control the memory subscription levels throughout the execution. We showed that the task paradigm allows to control the memory footprint of the application by throttling the task submission flow rate, striking a compromise between the performance benefits of anticipative task submission and the resulting memory consumption.

7.2. Simulation Execution with StarPU and SimGrid

The combination of StarPU and SimGrid allows to fast, accurate, and reproducible simulation the execution of task-based HPC applications.

This has proved to be very useful for theoretical research on scheduling heuristics [10]. It notably allowed to modify the simulated platform in order to easily observe and understand which parts of the platform (bandwidth, computation power) cause a bottleneck. It also allowed to remove some parts of the problem, such as the cost of data transfers, to simplify the problem and be able to deeply study scheduling solutions and compare them with optimum solutions in a simple environment before tackling the complete platform.

We have also extended the modelization of computation nodes, to take into account the PCI hierarchy of the system. This allows to get a more accurate simulation of systems which have dedicated channels between GPUs.

Last but not least, we have started to extend the StarPU+Simgrid combination to StarPU+MPI+Simgrid, to simulate the execution of HPC applications on *clusters* of heterogeneous systems. The preliminary results seem to show good accuracy. This will allow to easily study how applications scale, and study for instance how network performance have impact on it.

7.3. Scheduling heuristics for dense linear algebra

In the context of Suraj KUMAR's PhD thesis, we are studying the scheduling of the Cholesky factorization on heterogeneous systems.

We have started to introduce communication costs into the constraint programming. Since this increases resolution time a lot, we had to optimize the expression of the data transfers to simplify the resolution. We modified the StarPU runtime system to be able to inject not only a static schedule of tasks, but also a static schedule of data transfers. This allows to inject the schedule optimized by constraint programming into real executions.

We have also shown how static schedules and dynamic scheduling strategies compare on heterogeneous platforms, and notably in the context of varying task execution time can typically be a problem for static scheduling. Static schedules have actually proven to be robust against variation in execution time. We have also studied injecting static information into dynamic schedulers, which improves the resulting performance with little offline analysis.

7.4. Out-of-core support for task graphs

In the context of the Hi-BOX project, Airbus factorizes very large compressed matrices, which can not fit in the main memory, and most of the data thus have to be temporarily transferred to the disk, and loaded on-demand. We have thus extended the StarPU out-of-core support to the case of compressed matrices, and improved the eviction heuristics, so as to transfer data to the disk in advance.

7.5. Parallel Tasks within StarPU

One of the biggest challenge raised by the design of high performance task-based applications on top of heterogeneous accelerator-based machines lies in choosing the adequate granularity of tasks. Indeed, GPUs generally exhibit better performance when executing kernels featuring numerous threads whereas regular CPU cores typically reach their peak performance with fine grain tasks working on a reduced memory footprint. As a consequence, task-based applications running on such heterogeneous platforms have to adopt an intermediate granularity, chosen as a trade-off between coarse-grain and fine-grain tasks. We have tackle this granularity problem via resource aggregation : our approach consists in reducing the performance gap between accelerators and single cores by scheduling parallel tasks over cluster of CPUs. For this purpose, we have extended the concept of scheduling context, which we introduced in a previous work, so that the main runtime system only sees virtual computing resources on which it can schedule parallel tasks (e.g. BLAS kernels). The execution of tasks inside such clusters can virtually rely on any thread-based runtime system, and does not interfere with the outer scheduler. As a proof of concept we allow the interoperability of StarPU and OpenMP to co-manage task parallelism. We showed that our approach is able to outperform the magma, dplasma and chameleon state-of-the-art dense linear algebra libraries when dealing with matrices of small and medium size.

7.6. Running Compliant OpenMP Applications on top of StarPU with the Klang-Omp Compiler

Several robust runtime systems proposed recently have shown the benefits of task-based parallelism models. However, the common weakness of these supports is to tie applications to specific APIs. The OpenMP specification aims at providing a common parallel programming means for shared-memory platforms. It appears a good candidate to address this issue. We assessed the effectiveness and limits of this approach on the ScalFMM library developed by Inria HiePACS team, implementing fast multipole methods (FMM) algorithms. We showed that OpenMP dependent tasks allow for significant performance improvements over OpenMP loops and independent tasks for this application. We also demonstrated that suitable, targeted language extensions can further improve performances by a important margin in some cases.

7.7. Task-based Seismology Simulation on top of StarPU

Understanding three-dimensional seismic wave propagation in complex media is still one of the main challenges of quantitative seismology. Because of its simplicity and numerical efficiency, the finite-differences method is one of the standard techniques implemented to consider the elastodynamics equation. Additionally, this class of modelling heavily relies on parallel architectures in order to tackle large scale geometries including a detailed description of the physics. Last decade, significant efforts have been devoted towards efficient implementation of the finite-differences methods on emerging architectures. These contributions have demonstrated their efficiency leading to robust industrial applications. The growing representation of heterogeneous architectures combining general purpose multicore platforms and accelerators leads to re-design current parallel application. We thus considered the StarPU task-based runtime system in order to harness the power of heterogeneous CPU+GPU computing nodes. Preliminary results demonstrate significant speedups in comparison with the best implementation suitable for homogeneous cores.

7.8. Interfacing the MPC Parallel Framework with StarPU

CEA has developed a framework named MPC that transforms MPI+OpenMP applications into a lightweight thread-based program which can flexibly and efficiently exploit multicore architectures. StarPU, on its side, is mainly dedicated to schedule coarse grain tasks over accelerators, and is less suited to fine grain task scheduling. We have thus initiated a software interoperability effort between StarPU and MPC. The first step was to implement a new StarPU task scheduling strategy based on a NUMA-aware adaptative task granularity according to the target device (GPU or CPU). We observed significant performance gains for a Cholesky application in comparison to an eager strategy, thanks to the NUMA-aware aspect. However more work is still needed with respect to task decomposition as it implies data partitioning during the execution. We are also working on a variable granularity task programming interface in order to simplify the developer's coding effort. Finally, we develop a mechanism in StarPU to isolate some parts of the computing platform for another runtime. We used nested *scheduling contexts* to redirect some tasks to a scheduling component that StarPU may or may not control. The idea is to associate a scheduling subcontext to a runtime, for instance MPC, that would access to a dedicated set of computing resources for executing parallel kernels.

7.9. A Domain Specific Framework for Executing Stencil Kernels on Accelerated Platforms with SYCL

Stencil kernels arise in many scientific codes as the result from discretizing natural, continuous phenomena. Many research works have designed stencil frameworks to help programmer optimize stencil kernels for performance, and to target CPUs or accelerators. However, existing stencil kernels, either library-based or language-based necessitate to write distinct source codes for accelerated kernels and for the core application, or to resort to specific keywords, pragmas or language extensions. SYCL is a C++ based approach designed by the Khronos Group to program the core application as well as the application kernels with a single unified, C++ compliant source code. A SYCL application can then be linked with a CPU-only runtime library or processed by a SYCL-enabled compiler to automatically build an OpenCL accelerated application. We designed a stencil-dedicated domain specific embedded language (DSEL) which leverage SYCL together with expression template techniques to implement statically optimized stencil applications able to run on platforms equipped with OpenCL devices, while preserving the single source benefits from SYCL. Our stencil DSL has been tested using the SYCL compiler ComputeCpp from the CodePlay company on an accelerated platform, as well as with the TriSYCL library designed as a compilerless approach for CPU-only prototyping.

7.10. Combining Code Generation and Template Specialization Techniques in the P-EDGE Generic Polar Error Correction Code Framework

Error Correction Code decoding algorithms for consumer products such as *Internet of Things* (IoT) devices are usually implemented as dedicated hardware circuits. As processors are becoming increasingly powerful and energy efficient, there is now a strong desire to perform this processing in software to reduce production costs and time to market. The recently introduced family of Successive Cancellation decoders for Polar codes has been shown in several research works to efficiently leverage the ubiquitous SIMD units in modern CPUs, while offering strong potentials for a wide range of optimizations. Together with the IMS Laboratory, we designed the P-EDGE framework which combines a specialized skeleton generator and a building blocks library routines to provide a generic, extensible Polar code exploration workbench. It enables ECC code designers to easily experiments with combinations of existing and new optimizations, while delivering performance close to state-of-art decoders.

7.11. Binary Kernel Analysis, Hinting and Transformation for SIMDization

SIMD processor units have become ubiquitous. Using SIMD instructions is the key for performance for many applications. Modern compilers have made immense progress in generating efficient SIMD code. However, they still may fail or SIMDize poorly, due to conservativeness, source complexity or missing

capabilities. When SIMDization fails, programmers are left with little clues about the root causes and actions to be taken. Our proposed guided SIMDization framework builds on the assembly-code quality assessment toolkit MAQAO to analyze binaries for possible SIMDization hindrances. It proposes improvement strategies and readily quantifies their impact, using *in vivo* evaluations of suggested transformation. Thanks to our framework, the programmer gets clear directions and quantified expectations on how to improve his/her code SIMDizability. We show results of our technique on TSVC benchmark.

7.12. Dynamic Granularity Adaptation of OpenCL Kernels on Heterogeneous Multi-device Systems

On-going work as part of the PhD of P. Huchant aims to transparently execute an OpenCL kernel, and further a complete task graph, on an heterogeneous multi-device system. We propose methods to split an OpenCL kernel at compile time and adapt its granularity dynamically to ensure load balance. If the kernel is executed multiple times, we propose to determine its granularity by using a linear program whose constraints are built from performance measurements collected during the first invocations of the kernel with predefined granularities. Splitting the execution of one kernel into different executions does not require additional information from the user, therefore increasing the level of portability of OpenCL codes. First experiments show the interest of our approach.

TADAAM Team

6. New Results

6.1. TreeMatch Development

This year we have modified the TreeMatch API in order to enable better integration inside application with higher-level abstractions more precise semantic. We also introduced the “over subscribing” features that allow to map more than one process on a given processing unit. We also added new metrics to measure the performance of the proposed placement. We now have three metrics: The sum of the communication cost, the maximum of the communication cost and the hope-byte.

6.2. Affinity Abstraction

This year we worked on the affinity abstraction. Often, the affinity between two processes or threads is measured by the a matrix where a high entry represent a high affinity. Such example of matrices are the number of messages and the number of bytes exchanged between processing units. However, such matrix hide many characteristics of the application such as computation/communication overlap, network contention, etc. First, we have developed a new OpenMPI PML module to gather communication matrix of a running application. Then, we have conducted an extensive study of the minighost application to understand how such communication matrix actually measure the affinity between processes. On this application it appears that the size metric better matches the performances and that the performance of process placement is highly correlated to the proportion of communication in the application.

6.3. Locality for Application Using Locks on Clusters of Multicore Platforms

The aim of this post-doc work is to study the locality for applications based on read-write locks on clusters of multi-core platforms. We focused on the implementation of the video tracking application [25] using the Ordered Read Write Locks (ORWL) [20] model of programming on multi-cores architecture. For several uses, such as, human-computer interaction, security or traffic control, the tracking application aim to locate multiple moving objects over time using a camera. Its processing can be a time consuming process due to the amount of data that is contained in high definition video which leads to decrease the throughput. To overcome this problem it is possible to speed up the processing by exploiting task parallelism of ORWL model. Indeed, the model proposes abstractions of the decomposition in parallel parts (tasks), the synchronization of and the communication between threads. However, we noted some problems which decrease the parallelism scaling thus we introduced different optimizations: stream multiplexing, multiple buffering, etc. We are now working on parallelizing long-running tasks.

6.4. Topology Aware Malleability of MPI programs

Current parallel environments aggregate large numbers of computational resources with a high rate of change in their availability and load conditions. In order to obtain the best performance in this type of infrastructures, parallel applications must be able to adapt to these changing conditions.

In collaboration with Universidade da Coruña, Spain, we have worked on automatically and transparently adapt MPI applications to available resources is proposed. The solution relies on application-level migration approach, incorporating a new scheduling algorithm, based on TreeMatch and Hwloc, to obtain well balanced nodes while preserving performance as much as possible.

The experimental evaluation shows successful and efficient operation, with an overhead of less than 1 second for the proposed scheduling algorithm, and of only a few seconds for the complete reconfiguration, which will be negligible in large applications with a realistic reconfiguration frequency.

6.5. Topology Aware Load Balancing

Charm++ is a message-passing based programming environment that uses an object-oriented approach. However, where MPI considers processes in its model, Charm++ uses finer-grain migratable objects called chares. Brought together with an adaptive runtime system, Charm++ allows to perform dynamic load balancing considering the CPU load of each chare. Our work on data locality and process placement lead us to add the benefits of our TreeMatch algorithm in a load balancing solution. Thus we developed few months ago a topology-aware load-balancer in Charm++ using TreeMatch to reduce the communication costs. During the last months, we significantly improved this load-balancer and its scalability. Particularly, our load balancing algorithm is now hierarchical and distributed. To validate this approach, we have begun to carry out experiments with a cosmological application on the Blue Waters supercomputer. The results will be published soon.

6.6. Topology Aware Resource Management

SLURM [24] is a Resource and Job Management System, a middleware in charge of delivering computing power to applications in HPC systems. Our goal is to take in account in SLURM placement process hardware topology as well as application communication pattern. We have a new selection option for the `cons_res` plugin in SLURM. In this case the usually BestFit algorithm used to choose nodes is replaced by TreeMatch to find the best placement among the free nodes list in light of a given application communication matrix.

We updated this plugin based on SLURM 2.6.5 for last version SLURM 15.08. To decrease the overhead due to our algorithm we also implemented an alternative to use a subtree of the global topology. We ran experiments to compare these different solutions using our plugin with or without subtree and the current algorithm topology-aware in SLURM.

6.7. Topology Aware Performance Monitoring

While system's scale is growing exponentially, memory hierarchy is getting larger, at various levels. Hence optimizing applications to reach an optimal usage of a machine may involve a large spectrum of performance metrics interacting at different level of the system's hierarchy. Memory bound applications showing irregular patterns lead to locality issues. Addressing those issues and getting a good schedule on complex systems is a NP hard problem and can therefore only be solved with heuristics. Although powerful algorithms using the most intuitive heuristics such as communications path reduction and/or cache contention reduction may show good results on some cases, there are still room for improvements in this direction so much the configuration of applications, systems, software stack vary and impact the execution time.

In order to step in this direction we developed a highly extendible tool to gather asynchronously performance data from different sources. This information is then aggregated into different topology objects (cache, node, processing unit, ...) in order to give a synthetic and topology aware information to drive optimization.

In brief the tool works this way: The user provide a description file with arithmetic expression of performance counters(defined into performance data plugins), and topology objects where to map the expression. A pair (expression,object) defines a monitor which will sample performance data and stored them into an history. Then others monitors can be defined as a combination of the previous. For instance we can attach a process and record on each core its L3 cache miss counter, and then add each of those monitor into an upper monitor located on the L3 cache. Several aggregation functions are already available but we aim to provide several statistical function to extend the possibility of data interpretation. Such functions allow to aggregate results in a meaningful way. Then we add a locality insight using `lstopo` tool from `hwloc` to draw the results on a topology. This has been published in [12]

6.8. Memory Hierarchy Aware Roofline Model

The increasing complexity of computer architectures, makes challenging to fully exploit computer systems' capabilities. The cost of tuning applications on such machines can raise quickly. Therefore, linking the information about a machine performance bounds and applications performance results respectively to those bounds can help finding the bottlenecks and motivate code optimization.

In 2009 the Roofline model [23] throws those bases by plotting on a 2 dimensional diagram, application performance (GFlop/s) and arithmetic intensity (Flop/Byte) with respect to the main memory bandwidth (GByte/s) and peak floating point performance (GFlop/s). In 2014 the model extended by Alexandar Illic, take into account the data movement inside the cache hierarchy to provide a finer analyse by showing application's performance results with respect to the different cache bandwidths.

With the cooperation of the Cache Aware Roofline Model authors, we have worked on extending this model to the whole memory hierarchy at NUMA scale in order to drive optimisations on next generation processors embedding different memory technologies and different memory configurations like Intel's KNL does.

While we are designing a tool based on hwloc and micro-kernels to empirically extract and validate machines bottlenecks, we also want to show with real NUMA applications that the model may be extended to such hierarchy levels, still providing insightful representation.

6.9. Topology Management and Standardization

We continued to work on the diffusion of our software and ideas in existing programming interfaces and standards tailored for HPC and parallel computing. In particular, we did integrate our TreeMatch algorithm in the Open MPI implementation of the Message Passing Interface, so as to provide enhanced Virtual Topology routines in MPI allowing the user to effectively create parallel applications taking into account both their behaviour and the characteristics of the underlying hardware. Our code is available in the master repository and should be available in an Open MPI distribution at some point in the next year (2016).

We also drafted and submitted a proposal to modify the MPI interface so that information regarding the underlying physical topology could be made available at the MPI application level. We plan to push our ideas during the next year so that our proposal can eventually make its way into the MPI standard.

6.10. Modeling Next-Generation Memory Architectures

We initiated a research topic on modeling next generation memory architectures that will mix different kinds of memories. Indeed the arrival of high-bandwidth and non-volatile memories cause computing cores to have different local memory banks with different characteristics.

The hwloc software 5.1 is being extended in collaboration with CPU vendors such as Intel and AMD to better represent these new memory technologies. We are working with Bull in the context of Nicolas Denoyelle's PhD on developing abstractions for deciding where to allocate the application buffers.

6.11. Modeling Affinity of Multithreaded Applications

With the increasing complexity and scale of multi-core processors, optimizing thread placement becomes more and more challenging. Our goal is to better understand which characteristics of a multi-threaded application can have an impact on a placement decision for a given architecture. To this end, we analyze the performance of a set of applications under different placement strategies and we try to relate the obtained results to characteristics of the applications such as the data footprint of each thread, the amount of data shared between threads, or the reuse distance.

To collect information about the characteristics of multi-threaded applications, we developed a set of tools based on the PIN dynamic binary instrumentation tools. PIN allows us to get information about all instructions executed and memory location accessed by each thread of an application during its execution, and this without modifying the source code of the application.

We used our PIN-based tools to study a representative set of applications taken from two well-known benchmark suites, namely the Mantevo benchmark suites (HPC applications based on OpenMP) and the Parsec benchmark suites (general-purpose applications based on pthreads). Analyzing the results of all our tests is an ongoing work.

6.12. Thread placement and threads policy on a multicore machine with NUMA effects.

Threads placement on multicore machine with NUMA effects is inevitable to have better performances. Threads must bind on cores to avoid thread migration and to have better cache locality. MPI non-blocking collectives can generate progress threads to complete communications. These additional threads can disturb computational threads. That is why we have implemented several thread placement algorithms into the MPC framework [22]. These algorithms allow to dedicate resources only for progress threads. Thus computational threads are not disturb. We test them with our own benchmarks which test all the MPI non-blocking collectives to compare the performances with different thread placement. We observe an improvement when resources are dedicated to progress threads and take NUMA effects into account.

We want to include a mechanism into MPC to specify thread kinds (MPI, OpenMP,...). These mechanism will allow the MPC scheduler to take threads specificity into account to improve the scheduling policy. Our goal is to increase runtime performances considering each type specific needs. We have begun to implement this mechanism.

Several MPC framework bugs have been corrected, thus we contribute to its stability.

6.13. Multithreaded Communications

To program clusters of multicores, hybrid models mixing MPI+threads, and in particular MPI+OpenMP are gaining popularity. This imposes new requirements on communication libraries, such as the need for MPI_THREAD_MULTIPLE level of multi-threading support. Moreover, the high number of cores brings new opportunities to parallelize communication libraries, so as to have proper background progression of communication and communication/computation overlap.

We have proposed PIOMan [11], a generic framework to be used by MPI implementations, that brings seamless asynchronous progression of communication by opportunistically using available cores. It uses system threads and thus is composable with any runtime system used for multithreading. Through various benchmarks, we demonstrated that our pioman-based MPI implementation exhibits very good properties regarding overlap, progression, and multithreading, and outperforms state-of-art MPI implementations.

6.14. RDMA-based Communications

High-performance network hardware is nowadays dominated by RDMA-oriented technologies. The software stack is moving too towards Remote Memory Access. However, most communication libraries stil use send/receive paradigm as a common denominator. We have proposed to study a software stack for networks the is based on remote memory access from the hardware up to the enduser API, where RDMA is first class citizen and not a compatibility layer. It is expected to obtain better performance, better scalability with regard to number of communication flows or threads, and better asynchronous progression, while optimization strategies on the packet flows such as aggregation as proposed in NewMadeleine are still possible. Work has begun as a Masters thesis [18] and continues as Romain Prou Ph.D. thesis.

6.15. Network Modeling

Netloc is a tool for hwloc [1] to find the topology of a supercomputer. For that, it discovers all the networks by exploring them by using tools specifying to the network type. The exploration gives all the machines and all the switches, with all the links between them. We improved netloc by adding the visualization of the topologies discovered. The visualization is dynamic and the user can interact with it, to get some information about the machines, the switches or the link such as the physical address, the hostname or the speed of the link. In order to be able to do optimizations that can be helping process placement, we started to class the different topologies. For now, we only handle Clos networks [21] and we are able to transform them into fat trees. The categorization in classes permits to have a clean graph and then interact with graph partitioners.

To have a complete tool, we need to handle all major classes of topologies such as meshes, torus or hypercubes. When the graph partitioning will be integrated with tools such as SCOTCH, we will be able to find a good mapping for the processes of a job. It could also help the resource scheduling to optimize the resource sharing between jobs. The visualization can be improved by showing the architecture information retrieved by `hwloc` for each machine. We can complete the visualization by giving more information especially when the original graph was transformed to simplify it, as we did to Clos networks to obtain fat trees.

6.16. Scalable mapping onto (disconnected) parts of regular target architectures

Since its inception, SCOTCH allows one to map graphs onto so-called “algorithmically-defined” target architectures. They are regular architectures such as hypercube, multi-dimensional grids and tori, butterfly networks, etc., whose characteristics are defined by subroutines which are part of the SCOTCH library. However, on today’s large-scale computer systems, software jobs do not usually run on all of the machine, but on a set of nodes assigned by the batch scheduler. Consequently, one should be able to map a process graph onto (possibly disconnected) parts of an algorithmically-defined target architecture, which was not an available feature. Only “decomposition-defined” architectures (another way to represent target architectures in SCOTCH) supported this feature, but are not scalable above a few hundred processing elements.

In order to allow SCOTCH to provide mappings onto parts of an algorithmically-defined target architecture, a new meta target architecture, called “sub”, has been created. The sub architecture allows one to restrict a regular algorithmically-defined target architecture to a subset of its vertices. Instead of using a top-down approach to build a description of the target architecture, through a recursive bipartitioning algorithm, our new algorithm uses a bottom-up approach, based on recursive matching and coarsening of neighboring vertices, much like for graph coarsening. The clustering tree is pruned of branches that lead to parts of the machine that are not allowed mapping targets. Distance between subdomains is computed using the distance function of the underlying algorithmically-defined target architecture. Preliminary results have been presented at a SIAM CS&E conference workshop [14], and a beta-version of the upcoming release 6.0.5 of SCOTCH has been shipped to early testers at Lawrence Livermore National Laboratory.

6.17. Multi-Level Parallelism in a CFD code

Code_Saturne [19] is an industrial and open source Computational Fluid Dynamics software. Developed at EDF R&D, it solves the Navier-Stokes equations for 2D, 2D-axisymmetric and 3D flows, steady or unsteady, laminar or turbulent, incompressible or weakly dilatable, isothermal or not, with scalars transport if required.

Our goal is to evaluate different ways of improving and preparing this application for the future HPC architectures. We strengthened our application knowledge by using various instrumentation tools and provided a small topology instrumentation library. As instrumentation of a full code can be a tedious thing, we provided a mini application on which to perform our future experiments. We have run experiments to determine the potential gain of topology awareness on our code by using the graph mapping solutions of PT-SCOTCH. We have also run experiments on ghost cells numbering to see the impact of their locations on cache misses.