RESEARCH CENTER
Paris - Rocquencourt

FIELD

Activity Report 2015

Section New Results

Edition: 2016-03-21

ALPAGE Project-Team

7. New Results

7.1. Playing with DyALog-based parsers

Participants: Éric Villemonte de La Clergerie, Nicholas Parslow.

Éric de la Clergerie has continued the development of two DyALog-based parsers, namely DyALog-SR, a transition-based dependency parser, and FRMG, a wide-coverage French TAG based on an underlying meta-grammar.

The coverage of FRMG has been extended to cover more (rare) syntactic phenomena. A new conversion scheme has been added for the French version of the Universal Dependency Scheme. Preliminary evaluation experiments have been conducted on the French UD corpus, with both FRMG and DyALog-SR. FRMG has also been evaluated on the French SPMRL corpus, alone and with coupling with DyALog-sr

A new notion of secondary edges has been investigated in FRMG metagrammar and parser to provide additional dependency edges, helpful for understanding parsing outputs. In particular, secondary edges are used to denote controls between a verb and its hidden subject.

FRMG's disambiguation tuning is learned from CONLL-like treebanks using supervised learning method. We have conducted preliminary experiments to use unsupervised learning methods with observed accuracy gains between 1 to 1.5 points w.r.t. the no tuning case. However, trying to mix supervised and unsupervised methods have shown no significant gain w.r.t. the supervised case.

The hybridation of FRMG and DyALog-SR have been tried on a larger spectrum of treebanks.

FRMG has also been exploited during the Master internship of Nicholas Parslow about the use of NLP tools to provide feedback information and correlations on essays written by non-native French learners. In particular, the correction mechanism of FRMG has been extended to cover more cases of frequent errors and provide more explicit messages.

7.2. Linear-time discriminant syntactico-semantic parsing

Participants: Benoit Crabbé, Maximin Coavoux, Rachel Bawden.

In this module we study efficient and accurate models of statistical phrase structure parsing. We focus on linear time lexicalized parsing algorithms (shift reduce) with approximations entailing linear time processing. The existing prototype involves a global discriminant parsing model of the large margin family (Perceptron,Mira, SVM) able to parse user defined structured input tokens [62]. Thus the model can take into account various sources of information for taking decisions such as word form, part of speech, morphology or semantic classes inter alia.

Our model has been generalized in a multilingual setting where we are among the state of the art systems and state of the art on some languages [23]. To our knowledge the parser is one of the fastest existing multilingual phrase structure parser. In order to ease model design for multilingual settings, we currently study efficient feature selection procedures for automating model adaptation to new languages.

We have also extended our model to continuous representations by means of deep learning methods. We currently have a neural network based decision procedure for parsing [22]. It involves both greedy search and beam based search techniques. Current work focuses on the design of dynamic oracles for improving greedy search procedures. This framework is currently tested in the multilingual setting too.

Further work involves to tackle the knowledge acquisition bottleneck problem by integrating either symbolic knowledge such as dictionaries or semi-supervised procedures for improving the formal representation of lexical dependencies in order to leverage data sparsity and estimation issues recurrent in lexicalized parsing.

7.3. French Deep Syntactic Dependency Parsing

Participants: Corentin Ribeyre, Djamé Seddah, Éric Villemonte de La Clergerie, Marie Candito.

At Alpage, we used two distinct but complementary approaches to parse and produce deep syntactic dependency graphs from the DeepSequoia and the DeepFTB (crossref here). The first one was developped by using OGRE [87], [86], a graph rewriting system (crossref here). We developed a set of rewriting rules to transform surfacic syntactic dependency trees into deep syntactic dependency graphs, then we applied this set of rules on previously parsed surfacic trees. Those trees were produced using up to three different surfacic syntactic parsers: FRMG [109], DyALog-SR [109] and Mate [47]. The results were convincing and on par with what we got on English.

The second approach was based on the work made last year regarding the English broad-coverage semantic dependency parsing. We reused our two graph parsers (the first one is based on a previous work on DAG parsing [89] and the second one on the FRMG surfacic syntactic parser [109]) to parse the same graphs. As we previously have shown on English, the use of a mix of syntactic features (tree fragments from a constituent syntactic parser [80], dependencies from a syntactic parser [47], elementary spinal trees using a spine grammar [102], etc.) improve our results. Our intuition is that syntax and semantic are not independent of each other and using syntax could improve semantic parsing. Finally, we extended a dual-decomposition third-order graph parser [76] to incorporate our syntactic feature set and we were able to reach the best performances to this day on the task for both English [28] and French (Ribeyre et al, to appear).

7.4. Towards a French FrameNet

Participants: Marie Candito, Marianne Djemaa, Benoît Sagot.

The ASFALDA project [0] is an ANR project coordinated by Marie Candito. 5 partners collaborate on the project, on top of Alpage : the Laboratoire d'Informatique Fondamentale de Marseille(LIF), the Laboratoire de Linguistique Formelle (LLF), the MELODI team (IRIT - Toulouse) and the CEA-List. The project started in October 2012, and will end in march 2016. Its objective is to build semantic resources (generalizations over predicates and over the semantic arguments of predicates) and a corresponding semantic analyzer for French. We chose to build on the work resulting from the FrameNet project [45], [0] which provides a structured set of prototypical situations, called *frames*, along with a semantic characterization of the participants of these situations (called *frame elements*). The resulting resources will consist of :

1. a French lexicon in which lexical units are associated to FrameNet frames,
2. a semantic annotation layer added on top of existing syntactic French treebanks
3. and a frame-based semantic analyzer, focused on joint models for syntactic and semantic analysis.

In 2015, we continued the corpus annotation phase, which started in 2014. We currently have about 90 frames and 790 lexical units with at least one annotated occurrence, totalizing about $11,000$ annotated occurrences. We also set up :

- procedures for checking the coherence of the annotations
- a procedure for extracting the "annotated lexicon", namely extract quantitative information about the annotated lexical units, and syntax/semantics interface information (in terms of the probabilistic distributions of the syntactic paths used to express a given semantic role)
- the graphical vizualization of the annotated corpus

We also just started a collaboration with the LIF laboratory for using deep syntactic representations for predicting semantic frames and roles.

7.5. Development of Verb∋net

Participants: Laurence Danlos, Quentin Pradet, Lucie Barque.

[0] https://sites.google.com/site/anrasfalda/
[0] https://framenet.icsi.berkeley.edu/

VerbNet is an English lexical resources for verbs, which is internationally known and widely used in numerous NLP applications [74]. Verb∋net is a French adaptation of this resource. It is semi-automatically developed thanks to the use of two French existing resources created in the 70's: LG, Lexique-Grammaire developed at LADL under the supervision of Maurice Gross, and LVF, Lexique des verbs du français by Dubois and Dubois-Charlier. The idea is to map English classes, which gather verbs with a common syntactic and semantic behavior, into classes of LG and LVF, then to manually adapt the syntactic frames according to French grammar while keeping the thematic roles and the semantic information, [84], [68] [14]. A first version of this work has been achieved in June 2015 in collaboration with Takuya Nakamura (Institut Gaspard Monge) [33].

The next step was to verify the coherence of the resource. A particular focus has been to check the way alternations have been encoded and to document this encoding. A journal article extracted from this documentation has been submitted to the *TAL* journal and Verb∋net will be released after getting the feedback of the editorial board.

7.6. Development of the French Discourse TreeBank (FDTB)

Participants: Laurence Danlos, Margot Colinet, Jacques Steinlin, Pierre Magistry.

FDTB1 is the first step towards the creation of the French Discourse Tree Bank (FDTB) with a discourse layer on top of the syntactic one which is available in the French Tree Bank (FTB). In this first step, we have identified all the words or phrases in the corpus that are used as "discourse connectives". The methodology was the following: first, we highlighted all the items in the corpus that are recorded in LexConn [88], a lexicon of French connectives with 350 items, next we eliminated some of these items with the following criteria:

1. first, we filtered out the LexConn items that are annotated in FTB with parts of speech incompatible with a connective use, e.g. *bref* annotated as *Adj* instead of *Adv*, *en fait* annotated as *Pro V* instead of (compound) *Adv*;
2. second, as we lay down for theoretical and pratical reasons that elementary arguments of connectives must be clauses or VPs, we filtered out e.g. LexConn prepositions that introduce NPs;
3. last, we filtered out LexConn prepositions and adverbials with a non-discursive function.

The last criterion requires a manual work contrarily to the two others. For example the preposition *pour (to)*, is ambiguous between a connective use *(Fred s'est dépeché pour être à la gare à 17h (Fred hurried to be at the station at 17h))* and a preposition introducing a complement *(Fred s'est dépeché pour aller à la gare (Fred hurried to go to the station))*, and the disambiguation between the two uses is subtle and so the topic of a long paper [58], whose results have been used to enhance Le*fff* [93].

FDTB1 identifies 9 833 discourse connectives (among 18 535 sentences). This ressource is freely available and has been released in May 2015 [36].

FDTB2 is the next step in the creation of the FDTB. It consists in annotating the arguments of the discourse connectives identified in FDTB1 as well as the senses of these connectives (senses expressed through a set of discourse relations). This resource is still worked on.

7.7. Discourse Parsing

Participants: Chloé Braud, Laurence Danlos.

Discourse parsing goal is to reflect the rhetorical structure of a document, how pieces of text are linked in order to form a coherent document. Understanding such links could benefits to several other natural language applications (summarization, language generation, information extraction...).

A discourse parser corresponds to two major subtasks: a segmentation step wherein discourse units (DUs) are extracted, and a parsing step wherein these DUs are (recursively) related through "discourse (rhetorical) relations". The most difficult task in discourse parsing is the labeling of the relations between DUs, especially when no so-called connective overtly marks the relation (we then talk about implicit relations as opposed to explicit ones).

In her PhD, defended in December 2015, Chloé Braud develops a discourse relation classifier, carrying experiments on French and English. Focusing on the problem on implicit relation identification, this work explores ways of using raw data in combination with the available manually annotated data: this work led to systems based on domain adaptation methods exploiting automatically annotated explicit relations – demonstrating improvements on the French corpus Annodis and on the English corpus PDTB –, and to systems using word embeddings built from raw text to efficiently transform a word based representation of the data – leading to state-of-the art performance or above on the English corpus PDTB without the need of hand-crafted resources [21].

7.8. Towards a morpho-semantic resource for French designed for Word Sense Disambiguation

Participant: Lucie Barque.

The most promising WSD methods are those relying on external knowledge resources [78] but semantic resources for French are scarse. Moreover, existing resources offer fine grained sense dictinctions that do not fit to WSD. Our aim is to provide the NLP community with a broad-coverage morpho-semantic lexicon for French that relies on coarse-grained sense distinctions for polysemic units. Preliminary results concern nouns, on which we have first focused because their semantic description, compared to verbs, crucially lacks (for information retrieval, for instance) and because the regular polysemy phenomenon (recurring cases of polysemy within semantic classes) mainly occurs in nominal semantic classes:

- We proposed a linguistically motivated description of general semantic labels for nouns, that will allow for coarse-grained sense distinctions [107]
- Regular polysemy of nouns that can denote an event or a participant of this event has also been described for a large number of French nouns in [46]
- From a morphological point of view, nouns denoting events in French are mostly deverbal nouns (eg. *conversation* 'conversation', *promenade* 'stroll'), but there are also underived event nouns (eg. *guerre* 'war', *séisme* 'earthquake'). We compared their semantic properties in [35].
- Some lexical meanings are not easily captured by ontological semantic classes and a closer look has to be taken at them. Relational meanings in relational nouns are one of them [15].

7.9. Development of the Corpus de Référence du Français

Participants: Stéphane Riou, Benoît Sagot.

The 'Initiative Corpus de Référence du Français' (ICRF) is a project of Institut de Linguistique Française (ILF-FR2393 CNRS), coordinated by its director Franck Neveu and by Benoît Sagot.

The purpose of the ICRF is the development of a first prototype of the future French Reference Corpus, so as to assess the feasibility of this project and evaluate its potential impact. ICRF reuses existing freely-available French corpora, supplemented by additional data in an opportunistic fashion (e.g. a French media critic corpus and the corpus of talks given at an workshop on ethics and neurodegenerative diseases). ICRF preserves copyright and authorship of all corpora used. These corpora have been or will be part-of-speech tagged with MElt, converted to XML (TEI-P5-compliant) and made accessible via a web interface. The aim of ICRF is not to replace individual corpora and the interface will therefore allow, whenever possible, to easily recover access to each individual corpus. ICRF adds 5 metadata tags to categorize each individual corpus: spoken/written, text type and genre, linguistic competence level, date and linguistic area.

In 2015, the normalisation, tagging and conversion to XML of individual corpora has started, following the design of format specifications. The development of the web interface has already started, and a prototype is now available. Users can perform queries (search by tokens and/or POS) and use basic linguistic tools on the corpora (e.g. a concordancer). It is therefore more than a simple search interface or a download site: it improves research and selection of corpus.

7.10. Word order variation in Old French

Participants: Benoit Crabbé, Alexandra Simonenko.

As participant of the strand *Experimental Grammar* of the Labex EFL project *Empirical Foundations of Linguistics*[0] we study word order issues on Old French and more specifically the relative ordering of complements of ditransitive verbs. The inquiry seeks to identify several factors influencing the ordering of Old French complementation in different texts (varying in dates and genres) by carrying quantitative and statistical work from annotated Old French data. [0]

The first quantitative results [29] will be compared with what is known from corpus studies on the relative ordering of subject and complement in Old French [75]. It will also be compared to the quantitative results obtained on the relative ordering of complements of ditransitive verbs in Modern French [8] and modern English [53]. This comparative perspective is expected to provide new insights on French language evolution.

7.11. Cross linguistic factors governing word order

Participant: Benoit Crabbé.

In many languages, flexible word order often has a pragmatic role and marks the introduction of new information, a focus or a topic shift. Other cases of language- internal word order variation are alternations between two options such as *Mary gave John a book* and *Mary gave a book to John*, which are conditioned on syntactic and semantic factors such as the complexity of the constituents (as in *Mary gave John a book she had read ten times*), their animacy or the meaning of the verb [52].

One of the goals of this module is to investigate the connection between the quantitative aspects of word order variation across languages and the quantitative aspects of word order variation within a language. We study the corresponding patterns in language-internal variation by looking at the syntactically annotated corpora of various languages. Focusing on the variation of the internal word order of the noun-phrase as a case study [25], we explore, in collaboration with Kristina Gulordava (PhD at the University of Geneva, former international visitor at Alpage), to which extent a computational corpus-based analysis can provide new evidence not only for empirical, but also for theoretical linguistic research.

[0] www.labex-efl.org
[0] SRCMF corpus: http://srcmf.org/; MCVF: http://www.voies.uottawa.ca

ALPINES Project-Team

7. New Results

7.1. Communication avoiding algorithms for dense linear algebra

Our group continues to work on algorithms for dense linear algebra operations that minimize communication. During this year we focused on improving the performance of communication avoiding QR factorization as well as designing algorithms for computing rank revealing and low rank approximations of dense and sparse matrices.

In [2] we discuss the communication avoiding QR factorization of a dense matrix. The standard algorithm for computing the QR decomposition of a tall and skinny matrix (one with many more rows than columns) is often bottlenecked by communication costs. The algorithm which is implemented in LAPACK, ScaLAPACK, and Elemental is known as Householder QR. For tall and skinny matrices, the algorithm works column-by-column, computing a Householder vector and applying the corresponding transformation for each column in the matrix. When the matrix is distributed across a parallel machine, this requires one parallel reduction per column. The TSQR algorithm, on the other hand, performs only one reduction during the entire computation. Therefore, TSQR requires asymptotically less inter-processor synchronization than Householder QR on parallel machines (TSQR also achieves asymptotically higher cache reuse on sequential machines). However, TSQR produces a different representation of the orthogonal factor and therefore requires more software development to support the new representation. Further, implicitly applying the orthogonal factor to the trailing matrix in the context of factoring a square matrix is more complicated and costly than with the Householder representation.

We show how to perform TSQR and then reconstruct the Householder vector representation with the same asymptotic communication efficiency and little extra computational cost. We demonstrate the high performance and numerical stability of this algorithm both theoretically and empirically. The new Householder reconstruction algorithm allows us to design more efficient parallel QR algorithms, with significantly lower latency cost compared to Householder QR and lower bandwidth and latency costs compared with Communication-Avoiding QR (CAQR) algorithm. Experiments on supercomputers demonstrate the benefits of the communication cost improvements: in particular, our experiments show substantial improvements over tuned library implementations for tall-and-skinny matrices. We also provide algorithmic improvements to the Householder QR and CAQR algorithms, and we investigate several alternatives to the Householder reconstruction algorithm that sacrifice guarantees on numerical stability in some cases in order to obtain higher performance.

In [4] we introduce CARRQR, a communication avoiding rank revealing QR factorization with tournament pivoting. Revealing the rank of a matrix is an operation that appears in many important problems as least squares problems, low rank approximations, regularization, nonsymmetric eigenproblems. In practice the QR factorization with column pivoting often works well, and it is widely used even if it is known to fail, for example on the so-called Kahan matrix. However in terms of communication, the QR factorization with column pivoting is sub-optimal with respect to lower bounds on communication. If the algorithm is performed in parallel, then typically the matrix is distributed over P processors by using a two-dimensional block cyclic partitionning. This is indeed the approach used in the `psgeqpf` routine from ScaLAPACK. At each step of the decomposition, the QR factorization with column pivoting finds the column of maximum norm and permutes it to the leading position, and this requires exchanging $O(n)$ messages, where n is the number of columns of the input matrix. For square matrices, when the memory per processor used is on the order of $O(n^2/P)$, the lower bound on the number of messages to be exchanged is $\Omega(\sqrt{P})$. The number of messages exchanged during the QR factorization with column pivoting is larger by at least a factor of $n/\sqrt{P}$ than the lower bound.

In this paper we introduce CARRQR, a communication optimal (modulo polylogarithmic factors) rank revealing QR factorization based on tournament pivoting. The factorization is based on an algorithm that computes the decomposition by blocks of b columns (panels). For each panel, tournament pivoting proceeds in two steps. The first step aims at identifying a set of b candidate pivot columns that are as well-conditioned as possible. These columns are permuted to the leading positions, and they are used as pivots for the next b steps of the QR factorization. To identify the set of b candidate pivot columns, a tournament is performed based on a reduction operation, where at each node of the reduction tree b candidate columns are selected by using the strong rank revealing QR factorization. The idea of tournament pivoting has been first used to reduce communication in Gaussian elimination, an algorithm referred to as CALU.

We show that CARRQR reveals the numerical rank of a matrix in an analogous way to QR factorization with column pivoting (QRCP). Although the upper bound of a quantity involved in the characterization of a rank revealing factorization is worse for CARRQR than for QRCP, our numerical experiments on a set of challenging matrices show that this upper bound is very pessimistic, and CARRQR is an effective tool in revealing the rank in practical problems.

Our main motivation for introducing CARRQR is that it minimizes data transfer, modulo polylogarithmic factors, on both sequential and parallel machines, while previous factorizations as QRCP are communication sub-optimal and require asymptotically more communication than CARRQR. Hence CARRQR is expected to have a better performance on current and future computers, where commmunication is a major bottleneck that highly impacts the performance of an algorithm.

7.2. Algebraic preconditioners

Our work focused on the design of robust algebraic preconditioners and domain decomposition methods to accelerate the convergence of iterative methods.

In [5] we present a communication avoiding ILU0 preconditioner for solving large linear systems of equations by using iterative Krylov subspace methods. Recent research has focused on communication avoiding Krylov subspace methods based on so called s-step methods. However there is no communication avoiding preconditioner yet, and this represents a serious limitation of these methods. Our preconditioner allows to perform s iterations of the iterative method with no communication, through ghosting some of the input data and performing redundant computation. It thus reduces data movement by a factor of $3s$ between different levels of the memory hierarchy in a serial computation and between different processors in a parallel computation. To avoid communication, an alternating reordering algorithm is introduced for structured and unstructured matrices, that requires the input matrix to be ordered by using a graph partitioning technique such as kway or nested dissection. We show that the reordering does not affect the convergence rate of the ILU0 preconditioned system as compared to kway or nested dissection ordering, while it reduces data movement and should improve the expected time needed for convergence. In addition to communication avoiding Krylov subspace methods, our preconditioner can be used with classical methods such as GMRES or s-step methods to reduce communication.

7.3. A robust coarse space for Optimized Schwarz methods SORAS-GenEO-2

Optimized Schwarz methods (OSM) are very popular methods which were introduced by P.L. Lions for elliptic problems and Després for propagative wave phenomena. In [18], we have built a coarse space for which the convergence rate of the two-level method is guaranteed regardless of the regularity of the coefficients. We do this by introducing a symmetrized variant of the ORAS (Optimized Restricted Additive Schwarz) algorithm and by identifying the problematic modes using two different generalized eigenvalue problems instead of only one as for the ASM (Additive Schwarz method), BDD (balancing domain decomposition) or FETI (finite element tearing and interconnection) methods.

7.4. Time-dependent wave splitting and source separation

Starting from classical absorbing boundary conditions, we propose, in [17], a method for the separation of time-dependent scattered wave fields due to multiple sources or obstacles. In contrast to previous techniques, our method is local in space and time, deterministic, and also avoids a priori assumptions on the frequency spectrum of the signal. Numerical examples in two space dimensions illustrate the usefulness of wave splitting for time-dependent scattering problems.

7.5. Boundary integral formulations of wave scattering

We have continued to develop and further analyze new boundary integral formulation for wave scattering by complex objects.

In [13] we considered acoustic scattering of time-harmonic waves at objects composed of several homogeneous parts. Some of those may be impenetrable, giving rise to Dirichlet boundary conditions on their surfaces. We started from the second-kind boundary integral approach of [X. Claeys, and R. Hiptmair, and E. Spindler. A second-kind Galerkin boundary element method for scattering atcomposite objects. BIT Numerical Mathematics, 55(1):33-57, 2015] and extended it to this new setting. Based on so-called global multi-potentials, we derived variational second-kind boundary integral equations posed in $L^2(\Sigma)$, where Σ denotes the union of material interfaces. To suppress spurious resonances, we introduced a combined-field version (CFIE) of our new method. We conducted thorough numerical tests that highlighted the low and mesh-independent condition numbers of Galerkin matrices obtained with discontinuous piecewise polynomial boundary element spaces. They also confirmed competitive accuracy of the numerical solution in comparison with the widely used first-kind single-trace approach.

We spent much effort investigating the potentialities of multi-trace formulations in terms of domain decomposition. We considered multi-trace formulations in this perspective. Indeed Multi-Trace Formulations are based on a decomposition of the problem domain into subdomains, and thus domain decomposition solvers are of interest. The fully rigorous mathematical MTF can however be daunting for the non-specialist. In [12], we introduced MTFs on simple model problems using concepts familiar to researchers in domain decomposition. This allowed us to get a new understanding of MTFs and a natural block Jacobi iteration, for which we determined optimal relaxation parameters. We then showed how iterative multitrace formulation solvers are related to a well known domain decomposition method called optimal Schwarz method: a method which used Dirichlet to Neumann maps in the transmission condition. We finally showed that the insight gained from the simple model problem leads to remarkable identities for Calderón projectors and related operators, and the convergence results and optimal choice of the relaxation parameter we obtained is independent of the geometry, the space dimension of the problem, and the precise form of the spatial elliptic operator, like for optimal Schwarz methods. We confirmed this analysis with numerical experiments.

This work was extended in [10]. Considering pure transmission scattering problems in piecewise constant media, we derived an exact analytic formula for the spectrum of the corresponding local multi-trace boundary integral operators in the case where the geometrical configuration does not involve any junction point and all wave numbers equal. We deduced from this the essential spectrum in the case where wave numbers vary. Numerical evidences of these theoretical results were obtained in 2D.

Finally, in connection with boundary integral formulations, we extended the past work of [X. Claeys and R. Hiptmair, *Integral equations on multi-screens.* Integral Equations and Operator Theory, 77(2):167–197, 2013] where we had developed a framework for the analysis of boundary integral equations for acoustic scattering at so-called multi-screens, which are arbitrary arrangements of thin panels made of impenetrable material. In [3] we extended these considerations to boundary integral equations for electromagnetic scattering.

Viewing tangential multi-traces of vector fields from the perspective of quotient spaces we introduced the notion of single-traces and spaces of jumps. We also derived representation formulas and established key properties of the involved potentials and related boundary operators. Their coercivity were proved using a splitting of jump fields. Another new aspect emerged in the form of surface differential operators linking various trace spaces.

7.6. Asymptotic models for time harmonic wave propagation

Asymptotic models oriented toward more efficient numerical simulation methods have been investigated in three different directions.

In [8] we considered the Poisson equation in a domain with a small hole of size δ, and presented a simple numerical method, based on an asymptotic analysis, which allows to approximate robustly the far field of the solution as δ goes to zero without meshing the small hole. We proved the stability of the scheme and provide error estimates. This was confirmed with numerous numerical experiments illustrating the efficiency of the technique.

In [11] we considered a Laplace problem with Dirichlet boundary condition in a three dimensional domain containing an inclusion taking the form of a thin tube with small thickness. We proved convergence in operator norm of the resolvent of this problem as the thickness goes to 0, establishing that the perturbation on the resolvent induced by the inclusion is not greater than some (negative) power of the logarithm of the thickness. From this we deduced convergence of the eigenvalues of the perturbed operator toward the limit operator.

In [9] we investigated the eigenvalue problem $-\mathrm{div}(\sigma\nabla u) = \lambda u$ $(\mathcal{P})$ in a 2D domain Ω divided into two regions $\Omega_\pm$. We were interested in situations where σ takes positive values on Ω_+ and negative ones on Ω_-. Such problems appear in time harmonic electromagnetics in the modeling of plasmonic technologies. In a recent work [L. Chesnel, X. Claeys, and S.A. Nazarov. *A curious instability phenomenon for a rounded corner in presence of a negative material*. Asymp. Anal., 88(1):43–74, 2014], we had highlighted an unusual instability phenomenon for the source term problem associated with $(\mathcal{P})$: for certain configurations, when the interface between the subdomains $\Omega_\pm$ presents a rounded corner, the solution may depend critically on the value of the rounding parameter. In [9] we explained this property studying the eigenvalue problem $(\mathcal{P})$. We provided an asymptotic expansion of the eigenvalues and prove error estimates. We established an oscillatory behaviour of the eigenvalues as the rounding parameter of the corner tends to zero. This work was ended with numerical illustrations.

7.7. New results related to FreeFem++

In [6], we consider a model of soil water and nutrient transport with plant root uptake. The geometry of the plant root system is explicitly taken into account in the soil model. We first describe our modeling approach. Then, we introduce an adaptive mesh refinement procedure enabling us to accurately capture the geometry of the root system and small-scale phenomena in the rhizosphere. Finally, we present a domain decomposition technique for solving the problems arising from the soil model as well as some numerical results.

In [15], we study an interface transport scheme of a two-phase flow of an incompressible viscous immiscible fluid. The problem is discretized by the characteristics method in time and finite elements in space. The interface is captured by the Level-Set function. Appropriate boundary conditions for the problem of mould filling are investigated, a new natural boundary condition under pressure effect for the transport equation is proposed and an algorithm for computing the solution is presented. Finally, numerical experiments show and validate the effectiveness of the proposed scheme.

ANGE Project-Team

7. New Results

7.1. Modelling of complex flows

7.1.1. Non-hydrostatic models

Participant: Martin Parisot.

A new shallow water type model involving non-hydrostatic effects is derived in [37]. Under the assumption that the horizontal velocity is close to its vertical mean value, the model enables to recover the energy from the Euler system before integration. Link with the non-hydrostatic published in [18] is identified. Compared to the aforementioned models, the new system consists of more equations (6). However, the numerical strategy presented in the paper does not induce extra computational time.

7.1.2. Seismic activities: energy radiated by elastic waves

Participants: Anne Mangeney, Jacques Sainte-Marie.

Estimating the energy loss in elastic waves during an impact is an important problem in seismology and in industry. Three complementary methods to estimate the elastic energy radiated by bead impacts on thin plates and thick blocks from the generated vibration are proposed in [30]. The first two methods are based on the direct wave front and are shown to be equivalent. The third method makes use of the diffuse regime. These methods are shown to be relevant to establish the energy budget of an impact. The radiated elastic energy estimated with the presented methods is quantitatively validated by Hertz's model of elastic impact.

7.1.3. Layer-averaged Euler and Navier-Stokes systems

Participants: Marie-Odile Bristeau, Bernard Di Martino, Cindy Guichard, Jacques Sainte-Marie.

In [25] we propose a strategy to approximate incompressible free surface Euler and Navier-Stokes models. The main advantage of the proposed models is that the water depth is a dynamical variable of the system and hence the model is formulated over a fixed domain.

The proposed strategy extends previous works approximating the Euler and Navier-Stokes systems using a multilayer description. Here, the needed closure relations are obtained using an energy-based optimality criterion instead of an asymptotic expansion. Moreover, the layer-averaged description is successfully applied to the Navier-Stokes system with a general form of the Cauchy stress tensor.

7.2. Applications to marine energies

7.2.1. Partially free surface flow

Participants: Martin Parisot, Fabien Wahl.

In view of taking into account interactions with buoys, a new formulation of the shallow water model is derived with a constraint corresponding to a static roof. A relaxation approach is considered to adapt the standard numerical schemes. A particular attention is paid to the energy law whether it be for the original model with constraint or the relaxed version.

7.2.2. Swell energy

Participants: Sebastian Reyes-Riffo, Julien Salomon.

The internship consisted in designing an optimisation algorithm to determine advantageous topographies in view of producing energy from swell. This approach corresponds to the coupling between a shallow water type model with iterative updates of the topography. Stability of the numerical scheme is a critical point and requires the tuning of parameters.

7.3. Analysis of models in Fluid Mechanics

7.3.1. Weak solutions of multilayer models

Participants: Bernard Di Martino, Ethem Nayir, Yohan Penel.

Proving the existence of global weak solutions is a difficult problem for Navier-Stokes type equations, particularly in case of a degenerate viscosity (viscosity term can vanish if density or thickness goes to zero). In some recent works, Vasseur and Yu [46], have proved this existence for 2D shallow water equations. For the multilayer model, a collaboration with Boris Haspot (Univ. Paris-Dauphine) lead to stability results for the system with a focus on the difficulty to construct a sequence of approximate solutions that conserve all a priori estimates.

7.3.2. Strong solutions of multilayer models

Participants: Emmanuel Audusse, Ethem Nayir, Yohan Penel.

The existence and uniqueness of strong solutions of the multilayer model proposed in [41] was previously proven in the case of boundary conditions. We extended this result to an unbounded domain for short times, overcoming the issue of integrability often barely evoked in similar investigations. Current works deal with the long time existence by a continuation process which requires a particular care of the short time solution at the end of its existence interval.

7.3.3. Hyperbolic problems under constraints

Participant: Nicolas Seguin.

In [21], we study a family of linear hyperbolic systems whose solution must satisfy a constraint (e.g. a simplified model of river flows taking risk of flooding into account). We analyse relaxed models based on a penalisation. This theoretical approach could be used to derive numerical methods.

7.3.4. Entropy-satisfying finite volume schemes

Participant: Nicolas Seguin.

In [44], we carry out an analysis of 1st-order entropy-satisfying finite volume schemes for hyperbolic systems. More precisely, we investigate the numerical dissipation on unstructured meshes under relevant stability conditions. This results in a minimal convergence order towards smooth solutions.

7.3.5. Global existence for Green-Naghdi type equations

Participant: Dena Kazerani.

In [31], we consider the Cauchy problem for the Green-Naghdi equations with viscosity, for small initial data. It is well-known that adding a second order diffusion term to a hyperbolic system leads to the existence of global smooth solutions, as soon as the hyperbolic system is symmetrizable and the so-called Kawashima-Shizuta condition is satisfied. In a previous work, we have proved that the Green-Naghdi equations can be written in a symmetric form, using the associated Hamiltonian. This system being dispersive, in the sense that it involves third order derivatives, the symmetric form is based on symmetric differential operators. We use this structure for an appropriate change of variable to prove that adding viscosity effects through a second order term leads to global existence of smooth solutions, for small data. We also deduce that constant solutions are asymptotically stable.

7.4. Numerical methods for free-surface flows

7.4.1. Godunov schemes for the low Froude regime

Participants: Emmanuel Audusse, Do Minh Hieu, Yohan Penel.

We investigated in [29] the behaviour of collocated Godunov type finite volume schemes when applied to the 1d linear wave equation with Coriolis force in collaboration with S. Dellacherie and P. Omnes (CEA). Accuracy for short time and stability were proven for different versions of the classical Godunov schemes, including some schemes already proposed in the literature (Bouchut *et al*, [42]). Next step will be to include linear advection and then to study the fully non linear shallow water model. Then results will be extended to 2d problems for which geometrical constraints should be taken into account.

7.4.2. Numerical method for non-hydrostatic models

Participants: Nora Aïssiouene, Marie-Odile Bristeau, Edwige Godlewski, Jacques Sainte-Marie.

In [1], a numerical method based on a prediction-correction scheme in one dimension has been developed and compared to experimental data and analytical solutions. The issue is then to extend the method in higher dimensions. We propose a variational framework for the resolution of a non-hydrostatic Saint-Venant type model with bottom topography. This model is a shallow water type approximation of the free surface incompressible Euler system and slightly differs from the Green-Naghdi model. The resolution of the incompressibility constraint leads to an elliptic problem involving the non-hydrostatic part of the pressure. This step uses a variational formulation of a shallow water version of the incompressibility condition. Several numerical experiments are performed to confirm the relevance of our approach. This work is exposed in [18].

7.4.3. Uncertainties with the topography

Participants: Emmanuel Audusse, Nicole Goutal, Philippe Ung.

We propose to study the uncertainty related to the Saint-Venant system. A perturbation is introduced in the bottom topography such that the topography deviation is characterized by two parameters: its amplitude and its smoothness. In particular, we extend the work previously done with periodic boundary conditions and suggest a treatment of the physical ones. In doing so, we are interested in the influence of the topography deviation on the hydraulic quantities, and in particular, we numerically exhibit a relationship between the spatial correlations of the topography and the water height. Furthermore, we complete the study by a comparison of the outputs between the two flow regimes – fluvial and torrential.

7.4.4. Coupled Stokes-Exner model

Participant: Nora Aïssiouene.

In the framework of the 2015 CEMRACS session (Coupling Multi-Physics Models involving Fluids), we explored an approach to model the sediment transport. In [17], we consider a coupling between the Exner equation and the Stokes system to model sediments in geophysical flow phenomena. We focus on a model without free surface and used some numerical tests to evaluate the relevance of the method. The fluid structure interaction theory and methods have been applied on the coupled system and the objective is to test the proposed method which can be extend to a free surface model. The library Feel++ and the high computing performance embedded have been used to test the solution method. Therefore, the goal of this project is to understand the impact of the sediment transport on the flow using Navier-Stokes with a free surface system coupled with the standard Exner equation. This work has been done in collaboration with Tarik Amtout, Matthieu Brachet, Emmanuel Frenod, Romain Hild, Christophe Prud'homme, Antoine Rousseau and Stéphanie Salmon.

7.5. Software developments and assessments

7.5.1. Improvements in the FRESHKISS3D code

Participants: Marie-Odile Bristeau, David Froger, Raouf Hamouda, Jacques Sainte-Marie.

Several tasks have been achieved in the FRESHKISS3D software:

- The parallelisation of FRESHKISS3D with MPI is achieved for the Eulerian description and the explicit time scheme.
- The paddle wheel vertical effect is now taken into account.
- Vertical and time dependent flow rates can be customised.
- Unit tests have been improved and functional tests have been added.
- Software dependencies are packaged in SED-Paris repository.
- Online documentation is being written.
- A prototype of the software implemented in Cython is under discussion.
- Code executing time's loop is being refactored into multiple classes.
- Various improvements (build system, continuous integrations, coding rules) have been provided.

ANTIQUE Project-Team

6. New Results

6.1. Memory Abstraction

6.1.1. Abstraction of arrays based on non contiguous partitions

Participants: Jiangchao Liu, Xavier Rival [correspondant].

Abstract interpretation, Memory abstraction, Array abstract domains. In [19], we studied array abstractions.

Array partitioning analyses split arrays into contiguous partitions to infer properties of cell sets. Such analyses cannot group together non contiguous cells, even when they have similar properties. We proposed an abstract domain which utilizes semantic properties to split array cells into groups. Cells with similar properties will be packed into groups and abstracted together. Additionally, groups are not necessarily contiguous. This abstract domain allows to infer complex array invariants in a fully automatic way. Experiments on examples from the Minix 1.1 memory management demonstrated its effectiveness.

6.1.2. Static analysis for unstructured sharing

Participants: Huisong Li, Bor-Yuh Evan Chang [University of Colorado, Boulder, USA], Xavier Rival [correspondant].

Abstract interpretation, Memory abstraction, Separation logic. In [18], we studied the abstraction of shared data-structures.

Shape analysis aims to infer precise structural properties of imperative memory states and has been applied heavily to verify safety properties on imperative code over pointer-based data structures. Recent advances in shape analysis based on separation logic has leveraged summarization predicates that describe unbounded heap regions like lists or trees using inductive definitions. Unfortunately, data structures with *unstructured sharing*, such as graphs, are challenging to describe and reason about in such frameworks. In particular, when the sharing is unstructured, it cannot be described inductively in a local manner. In this work, we proposed a global abstraction of sharing based on set-valued variables that when integrated with inductive definitions enables the specification and shape analysis of structures with unstructured sharing.

6.1.3. Synthesizing short-circuiting validation of data structure invariants

Participants: Yi-Fan Tsai, Devin Coughlin, Bor-Yuh Evan Chang [University of Colorado, Boulder, USA], Xavier Rival [correspondant].

In [28], we studied the synthesis of short-circuiting validators for data-structure invariants.

This work introduces *incremental verification-validation*, a novel approach for checking rich data structure invariants expressed as separation logic assertions. Incremental verification-validation combines static verification of separation properties with efficient, *short-circuiting* dynamic validation of arbitrarily rich data constraints. A data structure invariant checker is an inductive predicate in separation logic with an executable interpretation; a short-circuiting checker is an invariant checker that stops checking whenever it detects at *run time* that an assertion for some sub-structure has been fully proven *statically*. At a high level, our approach does two things: it statically proves the separation properties of data structure invariants using a static shape analysis in a standard way but then leverages this proof in a novel manner to synthesize short-circuiting dynamic validation of the data properties. As a consequence, this approach enables dynamic validation to make up for imprecision in sound static analysis while simultaneously leveraging the static verification to make the remaining dynamic validation efficient. This work has shown empirically that short-circuiting can yield asymptotic improvements in dynamic validation, with low overhead over no validation, even in cases where static verification is incomplete.

6.2. Abstract domains

6.2.1. Abstract domains and solvers for set reasoning

Participants: Arlen Cox, Bor-Yuh Evan Chang [University of Colorado, Boulder, USA], Huisong Li, Xavier Rival [correspondant].

In [15], we studied the abstraction and inferrence of set properties.

When constructing complex program analyses, it is often useful to reason about not just individual values, but collections of values. Symbolic set abstractions provide building blocks that can be used to partition elements, relate partitions to other partitions, and determine the provenance of multiple values, all without knowing any concrete values. To address the simultaneous challenges of scalability and precision, we formalized and implemented an interface for symbolic set abstractions and constructed multiple abstract domains relying on both specialized data structures and off-the-shelf theorem provers. We developed techniques for lifting existing domains to improve performance and precision. We evaluated these domains on real-world data structure analysis problems.

6.2.2. Abstraction of optional numerical values

Participants: Jiangchao Liu, Xavier Rival [correspondant].

In [20], we designed a functor to lift a numerical abstract domain into an abstract domain that accounts for *optional* numerical values.

We proposed a technique to describe properties of numerical stores with optional values, that is, where some variables may have no value. Properties of interest include numerical equalities and inequalities. Our approach lifts common linear inequality based numerical abstract domains into abstract domains describing stores with optional values. This abstraction can be used in order to analyze languages with some form of *option* scalar type. It can also be applied to the construction of abstract domains to describe complex memory properties that introduce symbolic variables, e.g., in order to summarize unbounded sets of program variables, and where these symbolic variables may be undefined, as in some array or shape analyses. We described the general form of abstract states, and propose sound and automatic static analysis algorithms. We evaluated our construction in the case of an array abstract domain.

6.3. Static analysis of JavaScript applications

6.3.1. Desynchronized multi-state abstractions for open programs in dynamic languages

Participants: Arlen Cox [correspondant], Bor-Yuh Evan Chang [University of Colorado, Boulder, USA], Xavier Rival.

Abstract interpretation, Dynamically typed languages, Verification In [16], we have studied desynchronized multi-state abstractions for open programs in dynamic languages (libraries).

Dynamic language library developers face a challenging problem: ensuring that their libraries will behave correctly for a wide variety of client programs without having access to those client programs. This problem stems from the common use of two defining features for dynamic languages: callbacks into client code and complex manipulation of attribute names within objects. To remedy this problem, we introduced two state-spanning abstractions. To analyze callbacks, the first abstraction desynchronizes a heap, allowing partitions of the heap that may be affected by a callback to an unknown function to be frozen in the state prior to the call. To analyze object attribute manipulation, building upon an abstraction for dynamic language heaps, the second abstraction tracks attribute name/value pairs across the execution of a library. We implemented these abstractions and use them to verify modular specifications of class-, trait-, and mixin-implementing libraries.

6.4. Static analysis of spreadsheet applications

Participants: Tie Cheng [correspondant], Xavier Rival.

Abstract interpretation, Spreadsheet applications, Verification In [14], we have proposed a static analysis to detect type unsafe operations in spreadsheet applications including formulas and macros.

Spreadsheets are widely used, yet are error-prone: they use a weak type system, allowing certain operations that will silently return unexpected results, like comparisons of integer values with string values. However, discovering these issues is hard, since data and formulas can be dynamically set, read or modified. We defined a static analysis that detects all run-time type-unsafe operations in spreadsheets. It is based on an abstract interpretation of spreadsheet applications, including spreadsheet tables, global re-evaluation and associated programs. Our implementation supports the features commonly found in real-world spreadsheets. We ran our analyzer on the EUSES Spreadsheet Corpus. This evaluation shows that our tool is able to automatically verify a large number of real spreadsheets, runs in a reasonable time and discovers complex bugs that are difficult to detect by code review or by testing.

6.5. Distributed systems verification and programming language

Participants: Cezara Drăgoi [correspondant], Thomas Henzinger [IST Austria, Austria], Damien Zufferey [MIT, CSAIL, USA].

Fault-tolerant distributed systems, Programming languages, Verification Fault-tolerant distributed algorithms play an important role in many critical/high-availability applications. These algorithms are notoriously difficult to implement correctly, due to asynchronous communication and the occurrence of faults, such as the network dropping messages or computers crashing. Noteworthy is the lack of automated verification techniques for distributed systems, highly contrasting the mass distribution and development of distributed software. Therefore, our main motivation is to increase the confidence we have in distributed systems using formal verification methods. However, due to the complexity distributed systems have reached, we believe it is no longer realistic nor efficient to assume that high level specifications can be proved when development and verification are two disconnected steps in the software production process. We think that the difficulty does not only come from the algorithms but from the way we think about distributed systems. Therefore, we are interested in finding an appropriate programming model for fault-tolerant distributed algorithms, that increases the confidence we have distributed software. We introduced PSYNC, a domain specific language based on the Heard-Of model, which views asynchronous faulty systems as synchronous ones with an adversarial environment that simulates asynchrony and faults by dropping messages. We defined a runtime system for PSYNC that efficiently executes on asynchronous networks. We formalize the relation between the runtime system and PSYNC in terms of observational refinement. PSYNC introduces a high-level lockstep abstraction (on top of the standard asynchronous semantics), which simplifies the design and implementation of fault-tolerant distributed algorithms and enables automated formal verification. We have implemented an embedding of PSYNC in the SCALA programming language with a runtime system for asynchronous networks. We showed the applicability of PSYNC by implementing several important fault-tolerant distributed algorithms and we compared the implementation of consensus algorithms in PSYNC against implementations in other languages in terms of code size, runtime efficiency, and verification.

6.6. Derivation of Qualitative Dynamical Models from Biochemical Networks

Participants: Wassim Abou-Jaoudé [IBENS], Jérôme Feret [correspondant], Denis Thieffry [IBENS].

Systems biology, Logical models, Automatic derivation As technological advances allow a better identification of cellular networks, more and more molecular data are produced allowing the construction of detailed molecular interaction maps. One strategy to get insights into the dynamical properties of such systems is to derive compact dynamical models from these maps, in order to ease the analysis of their dynamics.

Starting from a case study, we present in [13] a methodology for the derivation of qualitative dynamical models from biochemical networks. Properties are formalized using abstract interpretation. We first abstract states and traces by quotienting the number of instances of chemical species by intervals. Since this abstraction is too coarse to reproduce the properties of interest, we refine it by introducing additional constraints. The resulting abstraction is able to identify the dynamical properties of interest in our case study.

6.7. Annotation of rule-based models with formal semantics to enable creation, analysis, reuse and visualization

Participants: G. Misirli, M. Cavaliere, W. Waites, M. Pocock, C. Madsen, O. Gifellon, R. Honorato-Zimmer, P. Zuliani, V. Danos [correspondant], A. Wipat.

In [35] We present an annotation framework and guidelines for annotating rule-based models, encoded in the commonly used Kappa and BioNetGen languages. Biological systems are complex and challenging to model and therefore model reuse is highly desirable. To promote model reuse, models should include both information about the specifics of simulations and the underlying biology in the form of metadata. The availability of computationally tractable metadata is especially important for the effective automated interpretation and processing of models. Metadata are typically represented as machine-readable annotations which enhance programmatic access to information about models. Rule-based languages have emerged as a modelling framework to represent the complexity of biological systems. Annotation approaches have been widely used for reaction-based formalisms such as SBML. However, rule-based languages still lack a rich annotation framework to add semantic information, such as machine-readable descriptions, to the components of a model. We introduced an annotation framework and guidelines for annotating rule-based models, encoded in the commonly used Kappa and BioNetGen languages. We adapted widely adopted annotation approaches to rule-based models. We initially proposed a syntax to store machine-readable annotations and describe a mapping between rule-based modelling entities, such as agents and rules, and their annotations. We then described an ontology to both annotate these models and capture the information contained therein, and demonstrate annotating these models using examples. Finally, we presented a proof of concept tool for extracting annotations from a model that can be queried and analyzed in a uniform way. The uniform representation of the annotations can be used to facilitate the creation, analysis, reuse and visualization of rule-based models. Although examples are given, using specific implementations the proposed techniques can be applied to rule-based models in general.

6.8. Quantitative genomic analysis of RecA protein binding during DNA double-strand break repair reveals RecBCD action in vivo

Participants: Charlotte Cockram, Milana Filatenkova, Vincent Danos [correspondant], Meriem Karoui, Leach David.

Understanding molecular mechanisms in the context of living cells requires the development of new methods of in vivo biochemical analysis to complement established in vitro biochemistry. A critically important molecular mechanism is genetic recombination, required for the beneficial reassortment of genetic information and for DNA double-strand break repair (DSBR). Central to recombination is the RecA (Rad51) protein that assembles into a spiral filament on DNA and mediates genetic exchange. Here we developed a method that combines chromatin immunoprecipitation with next-generation sequencing (ChIP-Seq) and mathematical modeling to quantify RecA protein binding during the active repair of a single DSB in the chromosome of Escherichia coli. In [29] we have used quantitative genomic analysis to infer the key in vivo molecular parameters governing RecA loading by the helicase/ nuclease RecBCD at recombination hot-spots, known as Chi. Our genomic analysis has also revealed that DSBR at the lacZ locus causes a second RecBCD-mediated DSBR event to occur in the ter- minus region of the chromosome, over 1 Mb away.

6.9. Moment Semantics for Reversible Rule-Based Systems

Participants: Vincent Danos [correspondant], Tobias Hinder, Ricardo Honorato-Zimmer, Sandro Stuck.

In [34] we developed a notion of stochastic rewriting over marked graphs – i.e. directed multigraphs with degree constraints. The approach is based on double-pushout (DPO) graph rewriting. Marked graphs are expressive enough to internalize the 'no-dangling-edge' condition inherent in DPO rewriting. Our main result is that the linear span of marked graph occurrence-counting functions – or motif functions – form an algebra which is closed under the infinitesimal generator of (the Markov chain associated with) any such rewriting

system. This gives a general procedure to derive the moment semantics of any such rewriting system, as a countable (and recursively enumerable) system of differential equations indexed by motif functions. The differential system describes the time evolution of moments (of any order) of these motif functions under the rewriting system. We illustrate the semantics using the example of preferential attachment networks; a well-studied complex system, which meshes well with our notion of marked graph rewriting. We show how in this case our procedure obtains a finite description of all moments of degree counts for a fixed degree.

6.10. Dirichlet is Natural

Participants: Vincent Danos [correspondant], Ilias Garnier.

In [32] the authors reconstruct a family of higher-order probabilities known as the Dirichlet process.

Giry and Lawvere's categorical treatment of probabilities, based on the probabilistic monad G, offer an elegant and hitherto unexploited treatment of higher-order probabilities. The goal of this paper is to follow this formulation to reconstruct a family of higher-order probabilities known as the Dirichlet process. This family is widely used in non-parametric Bayesian learning.

Given a Polish space X, we build a family of higher-order probabilities in G(G(X)) indexed by M(X), the set of non-zero finite measures over X. The construction relies on two ingredients. First, we develop a method to map a zero-dimensional Polish space X to a projective system of finite approximations, the limit of which is a zero-dimensional compactification of X. Second, we use a functorial version of Bochner's probability extension theorem adapted to Polish spaces, where consistent systems of probabilities over a projective system give rise to an actual probability on the limit. These ingredients are combined with known combinatorial properties of Dirichlet processes on finite spaces to obtain the Dirichlet family on X. We prove that the Dirichlet family is a natural transformation from the monad M to GG over Polish spaces, which in particular is continuous in its parameters. This is an improvement on extant constructions of Dirichlet.

6.11. Mechanistic links between cellular trade-offs, gene expression, and growth

Participants: Andrea Weisse, Diego Oyarzun, Vincent Danos [correspondant], Peter Swain.

Intracellular processes rarely work in isolation but continually interact with the rest of the cell. In microbes, for example, we now know that gene expression across the whole genome typically changes with growth rate. The mechanisms driving such global regulation, however, are not well understood. In [36] we considered three trade-offs that, because of limitations in levels of cellular energy, free ribosomes, and proteins, are faced by all living cells and we construct a mechanistic model that comprises these trade-offs. Our model couples gene expression with growth rate and growth rate with a growing population of cells. We show that the model recovers Monod's law for the growth of microbes and two other empirical relationships connecting growth rate to the mass fraction of ribosomes. Further, we can explain growth-related effects in dosage compensation by paralogs and predict host–circuit interactions in synthetic biology. Simulating competitions between strains, we find that the regulation of metabolic pathways may have evolved not to match expression of enzymes to levels of extracellular substrates in changing environments but rather to balance a trade-off between exploiting one type of nutrient over another. Although coarse-grained, the trade-offs that the model embodies are fundamental, and, as such, our modeling framework has potentially wide application, including in both biotechnology and medicine.

6.12. Thermodynamic graph-rewriting

Participants: Vincent Danos [correspondant], Russell Harmer, Ricardo Honorato-Zimmer.

In [33] we developped a new thermodynamic approach to stochastic graph-rewriting. The ingredients are a finite set of reversible graph-rewriting rules called generating rules, a finite set of connected graphs P called energy patterns and an energy cost function. The idea is that the generators define the qualitative dynamics, by showing which transformations are possible, while the energy patterns and cost function specify the long-term probability π of any reachable graph. Given the generators and energy patterns, we construct a finite set of rules which (i) has the same qualitative transition system as the generators; and (ii) when equipped with suitable rates, defines a continuous-time Markov chain of which π is the unique fixed point. The construction relies on the use of site graphs and a technique of 'growth policy' for quantitative rule refinement which is of independent interest. This division of labour between the qualitative and long-term quantitative aspects of the dynamics leads to intuitive and concise descriptions for realistic models (see the examples in S4 and S5). It also guarantees thermodynamical consistency (AKA detailed balance), otherwise known to be undecidable, which is important for some applications. Finally, it leads to parsimonious parameterizations of models, again an important point in some applications.

6.13. Kappa Rule-Based Modelling in Synthetic Biology

Participants: John Wilson-Kanamori, Vincent Danos [correspondant], Ty Thomson, Ricardo Honorato-Zimmer.

This [37] is a chapter of a book that provides complete coverage of the computational approaches currently used in Synthetic Biology. Rule-based modeling, an alternative to traditional reaction-based modeling, allows us to intuitively specify biological interactions while abstracting from the underlying combinatorial complexity. One such rule-based modeling formalism is Kappa, which we introduce to readers in this chapter. We discuss the application of Kappa to three modeling scenarios in synthetic biology: a unidirectional switch based on nitrosylase induction in Saccharomyces cerevisiae, the repressilator in Escherichia coli formed from BioBrick parts, and a light-mediated extension to said repressilator developed by the University of Edinburgh team during iGEM 2010. The second and third scenarios in particular form a case-based introduction to the Kappa BioBrick Framework, allowing us to systematically address the modeling of devices and circuits based on BioBrick parts in Kappa. Through the use of these examples, we highlight the ease with which Kappa can model biological interactions both at the genetic and the protein–protein interaction level, resulting in detailed stochastic models accounting naturally for transcriptional and translational resource usage. We also hope to impart the intuitively modular nature of the modeling processes involved, supported by the introduction of visual representations of Kappa models. Concluding, we explore future endeavors aimed at making modeling of synthetic biology more user-friendly and accessible, taking advantage of the strengths of rule-based modeling in Kappa.

This Chapters focus on computational methods and algorithms for the design of bio-components, insight on CAD programs, analysis techniques, and distributed systems. Written in the highly successful Methods in Molecular Biology series format, the chapters include the kind of detailed description and implementation advice that is crucial for getting optimal results in the laboratory.

Authoritative and practical, Computational Methods in Synthetic Biology serves as a guide to plan in silico the in vivo or in vitro construction of a variety of synthetic bio-circuits.

AOSTE Project-Team

7. New Results

7.1. CCSL as a Logical Clock Calculus Algebra: expressiveness and decidability results

Participants: Robert de Simone, Julien Deantoni, Frédéric Mallet, Qingguo Xu.

CCSL is a language dedicated to the expression of time constraints, based on so-called logical clocks. Its declarative nature is akin to the Lustre or (even closer to) the Signal language, but without values (to clock/event occurrences) and with both synchronous and asynchronous constraints. Solving a set of CCSL constraints amounts to the production of a feasible schedule of the system. While the TimeSquare tool may attempt to generate such a schedule trace by insightful simulation, it is not guaranteed to be complete in its search.So the issue of expressiveness and decidability was left open to this day.

Still, in previous years, we had established the CCSL constraints could be translated into parallel products (extended, transition-labelled) Büchi machines, but some of these machines had to contain integer shift counters, and were thus not fully FSMs. Our (misled) conjecture that CCSL had semilinear, Presburger-arithmetic power was defeated by a new translation expressing (unitary then general) Petri Nets and Vector Addition Systems into CCSL by encoding. The new conjecture that CCSL was then as powerful as Petri Nets was again defeated by a construction interpreting the features of *inhibitor arcs* in CCSL. As such inhibitor arcs extend the expressive power of Petri Nets to become universal (Turing-complete), CCSL enjoys the same universal property (which makes it unfortunately impossible to solve automatically in general).

Despite this negative result we could show that, under natural restrictions such as the assumption that "input" clocks have bounded jitter around a mean rate, and even if those bounds are not exactly known (but may be used as a parameter), then expressiveness remains in the semi-linear, Presburger-arithmetic range.

As a side-effect of this work we provided the translation of CCSL constraints into Büchi components by using a well-defined fragment of the Esterel syntax to express the Buchi automata.

Preliminary results are exposed in a research report. A much more ambitious article is in preparation.

As part of Professor Xu sabbatical in Aoste, we also considered the topic of machine-assisted proof of schedulability using theorem-provers (in our case PVS) [54]).

7.2. Industrial design flow for Embedded System Engineering

Participants: Julien Deantoni, Frédéric Mallet, Marie Agnes Peraldi Frati, Robert de Simone, Ales Mishchenko.

As part of the PIA LEOC Clarity collaborative project we attempt to instill some of our theoretical and methodological ideas into the framework of the (open-source, Polarsys Eclipse) Capella environment. This environment was initially developed inside Thales, under the name ARCADIA/Melody, as a modeling tool flow for System-Level Design in-the-large. As such, several aspects were not fully considered, specially those regarding safe sound simulation semantics at this level, or the role of states and modes in variability regarding both the software applicative and hardware architectural platform models. This research is in part motivated by concrete needs as expressed by end-users such as Airbus, Areva/EDF and Thales.
Results on methodological enhancements are described

7.3. Coordination of heterogeneous Models of Computation as Domain-Specific Languages

Participants: Matias Vara Larsen, Julien Deantoni, Frédéric Mallet.

In the context of the collaborative ANR GEMOC project (9.2.1.2 , we investigated the way the multiview approach generally promoted in Aoste could deal with analysis and simulation of systems specified using multiple heterogeneous languages. Coordinated use of heterogeneous domain specific languages (DSL) led to so-called globalization of modeling language. We wrote a chapter related to these concerns [50], as part of a book dedicated to the challenges of the field, gathering industrial and academic contributors.

This goal was achieved in two steps. First step consisted in specifying a language able to support appropriate information (*i.e.,* the one required for the coordination) in a *Language Behavioral Interface (LBI)*. Second step consisted in using the LBI to define coordination patterns from which the coordination of models can be automatically inferred. Design is supported by an heterogeneous simulation engine that has been developed and integrated in the Gemoc studio environment. Gemoc Studio, enhanced with our new research ideas, won the 9^{th} execution tool contest at ...

We also developed MoCCML (Model of Concurrency and Communication Modeling Language), an imperative extension of the CCSL language in the form of constraint automata [28]. MoCCML defines the concurrent and communication part of the semantics of a language, and is used by the LBI to exhibit internal causalities and synchronizations. Finally, we defined a protocol combining the concurrency aspects and the execution functions (*i.e.,* the rewriting rules) so as to be able to develop, in a modular way, the whole behavioral semantics of a language [30], [31].

Our work this on coordination of heterogeneous languages produced two major results. The first one is the development of BCOoL (Behavioral Coordination Operator Language [33]). BCOoL is a language dedicated to the specification of coordination patterns between heterogeneous languages. It comes with a tool chain allowing the generation of the coordination given a BCOoL operator and specific models. Our second result is the development of an heterogeneous execution engine, integrated to Gemoc studio, to run conjointly different models [44]. Both works were mainly realized by Matias Vara Larsen, as part of his upcoming PhD.

7.4. SoC multiview (meta)modeling for performance, power, and thermal aspects

Participants: Amani Khecharem, Robert de Simone, Emilien Kofman, Julien Deantoni.

In the framework of the ANR HOPE project we progressed the definition of multiview metamodels for the design of Systems-on-Chip) (SoC systems integrating performance, power and thermal aspects. The main concern was to stress regularity and commonality between those views, each developed on "domains" defined as partitions of the original block diagram (clock domains, voltage domains, floorplans,...), and with finite state machine controlers setting the levels of these domains; links between distinct views are originally provided by laws of physics, but then usually identified on discrete allowable values by engineers. The application view, meant to provide typical use-cases to help dimension the SoC platform by abstract simulation, also fits in this framework. This methodological work was presented in the local forum SAME (Sophia-Antipolis MicroElectronices) [53]. It is supposed to work in two ways, both by allowing the appplication of analytic methods to compute an optimized mapping of application tasks onto platform resources, and then to translate these results towards sophisticated simulation environments (such as MCO Platform Architect by Synopsys or ACEplorer by Docea Power/Intel, both partners in the HOPE consortium) which consider non-functional aspects of power and thermal modeling in their simulation environments. The various approaches considered in Aoste to define mapping constraints and solve them algorithmically are presented elsewhere. All this should sonn be rported in Ameni Khacharem PhD document.

7.5. Networks-on-Board: between NoCs and rack connector buses

Participants: Amine Oueslati, Robert de Simone, Albert Savary, Emilien Kofman.

The recent paradigm of Massively Parallel Processor Arrays (MPPA), or more generally manycore Systems-on-Chip, rely on the existence of a high-throughput on-Chip Network (NoC) to interconnect the various cores and processing clusters. Despite its benefits, it requires that all components are put on the same dye, and thus designed monolithically. On the other end, supercomputers are built by assembling racks or blades of processors, connected by fast buses (fast ethernet or infinyband usually), with low predictivity of throughput. A third, intermediate path is explored in the context of the FUI Clistine project, based on a notion of Network-on-Board (or Network-in-Package), aiming at the benefits of NoCs brought to the level of a single PCB board, where the various components can be assembled in a modular fashion. We consider the applicationof our previous expertise on modeling and analysis of NoC-based architecture, with their implications on the optimized mapping of dataflow models of applications onto such interconnects, to adapt them in this new context. The objective is to consider alternative network topologies, and to transate optimal mappings into the concrete network operations on a prototype implementation realized by SynergieCAD, the company heading the project. This topic reflects the PhD thesis of Amine Oueslati, and the engineering work of Albert Savary.

7.6. Solving AAA constraints analytically

Participants: Emilien Kofman, Dumitru Potop Butucaru, Thomas Carle, Raul Gorcitz, Robert de Simone, Mohamed Bergach, Amine Oueslati.

Given two abstract modeling descriptions, one of a dataflow process network for the application, one of a block diagram structure for the computing platform and its interconnects, together with cost functions for the elementary compuations and communications, one is bound to seek optimal mappings pairing the two. Amongst all the possible techniques, an obvious one consists in solving constraint using general solvers (real, integer, or boolean constraint programming, SMT solvers, etc). Given the NP-hard nature of the problem, the issue here is to scale to the dimensions of realistic problems. We conducted extensive experiments on several case studies, with as extra objective the concern of studying how the formulation of constraints, or the exploiation of additional information (in concurency or exclusion of tasks, structural symmetries,...) could impact favorably or negatively the process. Results were compiled in a publication [57].

In the framework of the PhD thesis of Mohamed Bergach, under CIFRE funding with Kontron Toulon, we studied how to adjust a radar application, that typically computes extensively FFT convolutions, on an hybrid CPU/GPU architecture such as IntelCore IvyBridge and Haswell processors. This approach works in two stages: first we considered how to implement a FFT redex as large as possible in exactly one core (either a CPU core or a GPU Execution Unit), so as to make full use of the local register memories and SIMD/vectorial instructions. Not by accident certainly FFT blocks of size exactly 8 and 16 respectively can so be fitted on a GPU (resp. CPU) block. This provides a new "compound" instruction, on which to build modularly and optimization the allocation of larger aplications based on such basic block. This is fully described in Mohamed Bergach PhD document [16].

7.7. Stochastic extension of MARTE/CCSL for CPS modeling

Participant: Frédéric Mallet.

This work was conducted during the sabbatical period of Frédéric Mallet at ECNU Shanghai, in the context of the associated team FM4CPS (9.4.1.1).

As a declarative language, CCSL allows the specification of causal and temporal properties of systems expressed as constraints in a specific syntax. While each constraint reduces the set of possible behaviors, there may still be multiple (schedule) solutions, or none at all. When several solutions remain feasible, our TimeSquare tool allows to set up a resolving policy, to choose whether we want to attempt exploring exhaustively all these solutions, or else narrow the solution space according to an auxiliary criterion.

The extension of CCSL with stochastic features and probabilistic information is meant to help provide such an additional criterion, while modeling temporal constraints on the environment which are not necessarily well-known or controllable, specially in the domain of Cyber-Physical Systems. Then, such features should help reducing the set of possible behaviors, narrowing for instance to the most likely ones (in a formal quantitative meaning).

We are currently relying on UPPAAL SMC (Stochastic Model-Checking) toolset as prototype analyzer for the resulting specifications.

7.8. Coupling SystemC and FMI for co-simulation of Cyber-Physical Systems

Participants: Stefano Centomo, Julien Deantoni, Robert de Simone.

In the context of Stefano Centoma master internship, and in collaboration with his global supervisor Professor Davide Quaglia, from the University of Verona, we considered the possibility to build heterogenous, multi-physics co-simulation schemes for hybrid continuous-discrete Cyber-Physical systems. The first step consisted in extracting relevant interface information from IP component described in the SystemC language; it was naturally inspired from some of our former work. But currently IP-XACT is meant to address easy component assembly at the *structural* (static) level, and is not concerned with dynamical aspects of behavior simulation. This extension, and the proper combination with the FMI standard for its purpose, allowing hybrid and multiform co-simulation of SystemC components (and also others describing the continuous physical environment) are the next-step objective being currently tacked.

7.9. Code generation for time-triggered platforms based on Real-Time Scheduling

Participants: Dumitru Potop Butucaru, Raul Gorcitz, Yves Sorel.

We have continued this year the work on real-time scheduling and code generation for time-triggered platforms. Much of this work was carried out as part of a trilateral collaboration with Airbus DS and the CNES, which have funded an (onerous) TTEthernet-based test platform and partly funded the post-doctorate of Raul Gorcitz. The remainder of Raul Gorcitz' post-doc has been funded by the ITEA3 Assume project.

This year, the objective has been to allow code generation on an industry-grade platform comprising ARINC 653-based computers connected through a TTEthernet network. The novelty with respect to previous years comes from the time-triggered TTEthernet network, whose scheduling properties raise new problems. Unlike in classical field buses, resource reservation in a TTEthernet network is done at the level of directed links (physical wires that connect routers and end stations). Each of these links is controlled by an arbiter that determines the scheduling of both time-triggered data transfers and control messages needed to ensure the global time synchronization. This year we have built a model of the TTEthernet network allowing precise real-time scheduling, and worked on code generation aspects. We expect to have a fully running prototype in the next 2 months, and to demonstrate it to our funders. Relevant publications are [18], [38].

For teaching purposes and to achieve a finer understanding of ARINC 653-based operating systems, we have also developed an implementation of the standard on inexpensive RaspberryPi platforms, and published a scientific vulgarization paper [55].

7.10. Real-time systems compilation

Participants: Dumitru Potop Butucaru, Keryan Didier, Mihail Asavoae.

This research line develops over various results of the team over the years, its aim being to develop fully automatic implementation flows going fully automatically from functional and non-functional specification to correct and efficient running implementation. We advocate for a real-time systems compilation approach that combines aspects of both real-time scheduling and compilation of both classical and synchronous languages. Like a classical compiler such as GCC, a real-time systems compiler should use fast and efficient scheduling and code generation heuristics, to ensure scalability. Similarly, it should provide traceability support under the form of informative error messages enabling an incremental trial-and-error design style, much like that of classical application software. This is more difficult than in a classical compiler, given the complexity of the transformation flow (creation of tasks, allocation, scheduling, synthesis of communication and synchronization code, etc.), and requires a full formal integration along the whole flow, including the crucial issue of correct hardware abstraction. A real-time systems compiler should perform precise, conservative timing

accounting along the whole scheduling and code generation flow, allowing it to produce safe and tight real-time guarantees. More generally, and unlike in classical compilers, the allocation and scheduling algorithms must take into account a variety of non-functional requirements, such as real-time constraints, criticality, partitioning, preemptability, allocation constraints, etc. As the accent is put on the respect of requirements (as opposed to optimization of a metric, like in classical compilation), resulting scheduling problems are quite different.

We are currently building such a real-time systems compiler, called Lopht. The construction of the Lopht tool, which takes into account complex functional and non-functional specifications is discussed in the corresponding section and in [17].

This year, we have initiated work on two fundamental topics. The first one is sound architecture abstraction – ensuring that the platform models used for real-time scheduling and code generation are conservative abstractions of the real hardware and basic software, allowing the generation of implementations that are functionally and non-functionally correct. This work is performed in the framework of the LEOC Capacites project, which funds the post-doc of Mihail Asavoae. The second line of work aims at formally proving that the output of Lopht is correct with respect to its input models (including functional specification and platform model). This work is performed in the ITEA3 Assume project, which funds the PhD thesis of Keryan Didier. Together with the Parkas team-project we have also considered the implementation of mixed-criticality systems [26].

7.11. Uniprocessor Real-Time Scheduling

Participants: Mamadou Diallo, Yves Sorel, Walid Talaboulma, Robert Davis.

In the context of the master internship of Mamadou Diallo we implemented the offline time trigger scheduler proposed in his PhD thesis by Falou Ndoye on a development board based on an ARM Cortex M4. We used this ARM version since it is better suited to embedded systems, since more predictable, than the ARM 7 we used last year. Especially, it allows to determine more accurately the cost of the scheduler and of the preemptions we use in our offline schedulability analysis. We remind that the schedulability analysis provides a scheduling table which is exploited by the scheduler during the real-time execution of the tasks. This approach allows a low and fixed cost for the scheduler and the preemptions whereas these costs are variable in the case of classical online schedulers. For several task sets we compared the timing diagrams predicted by the schedulability analysis with the real-time timing diagrams measured on the ARM Cortex M4. It turns out that those timings are very close, as expected.

A new direction opened with the arrival of Rob Davis was to consider by studying the impact of the non-preemptivity constraints on the optimality of the schedulers [37], or by considering fixed priorities while scheduling messages in the context of Control Area Networks [36].

7.12. Multiprocessor Real-Time Scheduling

Participants: Aderraouf Benyahia, Laurent George, Salah Eddine Saidi, Yves Sorel, Robert Davis, Liliana Cucu.

In the context of the PhD thesis of Salah Eddine Saidi we considered the co-simulation of several process models specified in continuous time and several controllers models specified in discrete time according to a real-time hardware in the loop approach. These models specified with different tools such as Simulink, AMEsim, Modelica, etc., cooperate according to the FMI standard. They are translated in a dataflow graph that is compliant with the conditioned repetitive dataflow model of our AAA methodology for functional specification. Each model considers the feed-through function as well as the functions which depend of the state, and the state computation itself. In order to meet the real-time constraints of such complex co-simulation we need to execute them on multicore platforms. We studied the limitations of greedy and local search distributed real-time scheduling heuristics we developped in the past for control applications. The first limitation is related to the FMI standard which requires that the functions belonging to a model are allocated to the same core. We first try to introduce additional semaphores in the real-time code generated automatically

to avoid these situations. Unfortunately, this solution decreases significantly the acceleration brought by the multicore. Therefore, we started to investigate graph based techniques that add non directed edges to specify the FMI relation and search solutions where some non oriended edges can be oriented to minimize locally the makespan.

In the context of the master internship of Mamadou Diallo we studied the possibilities to extend the offline time trigger scheduler implemented on a uniprocessor to the multiprocessor case. Since the embedded board based on the ARM Cortex M4 we utilize features an ethernet interface, we conducted several experimentations on ethernet switches to measure the end-to-end communication time between several real-time tasks running on such boards with such schedulers.

We completed the work on the gateway with modeling languages for certified code generation carried out in the P FUI project 9.2.2 which ended in June 2015. Mainly, we tested the P modeling language to SynDEx gateway on four industrial use cases provided by AdaCore, Continental and Aboard Engeneering. We specified these applications with the P language and translated them in the SynDEx format. With SynDEx we analysed the schedulability and automatically generated the corresponding code for an Intel 8 cores Xeon ES-1620v2 3.70Ghz. For these applications ranging from 103 to 1403 bloks we obtained an acceleration factor equal to the number of cores.

Thanks in part to the arrival of Rob Davis, our team has participated to the proposition of a new framework in the context of multicore platforms: *Multicore Response Time Analysis framework* [34]. This proposal was made in close collaboration with academic partners such as the University of Luxembourg, Verimag and ISEP Porto. The framework is extensible to different multicore architectures, with various types and arrangements of local memory, and different arbitration policies for the common interconnects. The MRTA framework provides a general approach to timing verification for multicore systems, parametric in the hardware configuration, and so can be used architectural design stage to compare the guaranteed levels of performance that can be obtained with different hardware configurations. The MRTA framework decouples response time analysis from a reliance on context independent WCET values. Instead, the analysis specifies response times directly according to requirements on different hardware resources.

7.13. Probabilistic and statistical temporal analysis

Participants: Liliana Cucu, Robert Davis, Adriana Gogonel, Walid Talaboulma, Dorin Maxim, Cristian Maxim.

Real-time constraint guarantees require worst-case reasoning to provide sound solutions. We have proposed to define and use worst-case reasoning in different contexts: optimal scheduling algorithms, response time analysis, estimation of worst-case execution times. These results have laid the foundations for certifiable probabilistic solutions to real-time systems.

In particular, we have studied the probabilistic response time analysis for systems with multiple probabilistic parameters, either by using bounds based on real-time calculus, extreme value theory, direct calculation or in a context of component-based systems. Generally, probabilistic methods have high complexity cost; using upper-bounds for the input probability distributions we provide conservative(safe) results faster. Worst-case reasoning is also provided for the statical estimation of a task probabilistic worst-case execution time.

Results were published in [22], [24], [58], [56], [42], [46], [23], [42], [43], [40]

7.14. Parametric and Non-Parametric Statistics for Program Performance Analysis and Comparison

Participant: Sid Touati.

This research activity is a continuation of our joint research effort with Julien Worms, Assistant Professor at University of Versailles Saint-Quentin (UVSQ), dealing with statistical program performance analysis and comparison, in presence of performance variability. In the previous study (called Speedup-Test), we gave a rigorous statistical methodology for analysis of program speedups based on mean or median performance metrics: execution time, energy consumption, etc. However mean or median observed performances do not always reflect the user's feeling of performance, especially when the performances are really unstable. In the current study, we propose additional precise performance metrics, based on performance modeling using gaussian mixtures. We explore the difference between parametric and non parametric statistics applied on program performance analysis. Our additional statistical metrics for analysing and comparing program performances give to the user more precise decision tools to select best code versions, not necessarily based on mean or median numbers. Also, we provide a new metric to estimate performance variability based on gaussian mixture model. Our statistical methods are implemented in R, and distributed as open source code. A research report is under completion, before submission as article.

ARAMIS Project-Team

7. New Results

7.1. Learning spatiotemporal trajectories from manifold-valued longitudinal data

Participants: Jean-Baptiste Schiratti [Correspondant], Stéphanie Allassonniere, Olivier Colliot, Stanley Durrleman.

We propose a Bayesian mixed-effects model to learn typical scenarios of changes from longitudinal manifold-valued data, namely repeated measurements of the same objects or individuals at several points in time. The model allows the estimation of a group-average trajectory in the space of measurements. Random variations of this trajectory result from spatiotemporal transformations, which allow changes in the direction of the trajectory and in the pace at which trajectories are followed. The use of the tools of Riemannian geometry allows to derive a generic algorithm for any kind of data with smooth constraints, which lie therefore on a Riemannian manifold. Stochastic approximations of the Expectation-Maximization algorithm is used to estimate the model parameters in this highly non-linear setting.

The method is used to estimate a data-driven model of the progressive impairments of cognitive functions during the onset of Alzheimer's disease. Experimental results show that the model correctly put into correspondence the age at which each individual was diagnosed with the disease, thus validating the fact that it effectively estimated a normative scenario of disease progression. Random effects provide unique insights into the variations in the ordering and timing of the succession of cognitive impairments across different individuals.

More details in [30] and [31].

7.2. Joint Morphometry of Fiber Tracts and Gray Matter structures using Double Diffeomorphisms

Participants: Pietro Gori [Correspondant], Olivier Colliot, Linda Marrakchi-Kacem, Yulia Worbe, Alexandre Routier, Cyril Poupon, Andreas Hartmann, Nicholas Ayache, Stanley Durrleman.

This work proposes an atlas construction method to jointly analyse the relative position and shape of fiber tracts and gray matter structures. It is based on a double diffeomorphism which is a cascade of two diffeomorphisms. The first deformation acts only on the white matter keeping fixed the gray matter of the atlas. The resulting white matter, together with the gray matter, are then deformed by the second diffeomorphism which puts into correspondence the homologous anatomical structures across subjects. The first diffeomorphism makes the fiber bundles slide on the fixed gray matter revealing the variability in structural connectivity within the population, namely both the changes in the connected areas and in the geometry of the pathway of the tracts. Fiber bundles are approximated with weighted prototypes using the metric of weighted currents. The algorithm is based on a Bayesian framework which allows the automatic estimation of the covariance matrix of deformation parameters and of the noise variance of each structure. This approach is applied to patients with Tourette syndrome and controls showing a variability in the structural connectivity of the left cortico-putamen circuit.

More details in [26].

7.3. Bayesian Mixed Effect Atlas Estimation with a Diffeomorphic Deformation Model

Participants: Stanley Durrleman [Correspondant], Stéphanie Allassonniere, Estelle Kuhn.

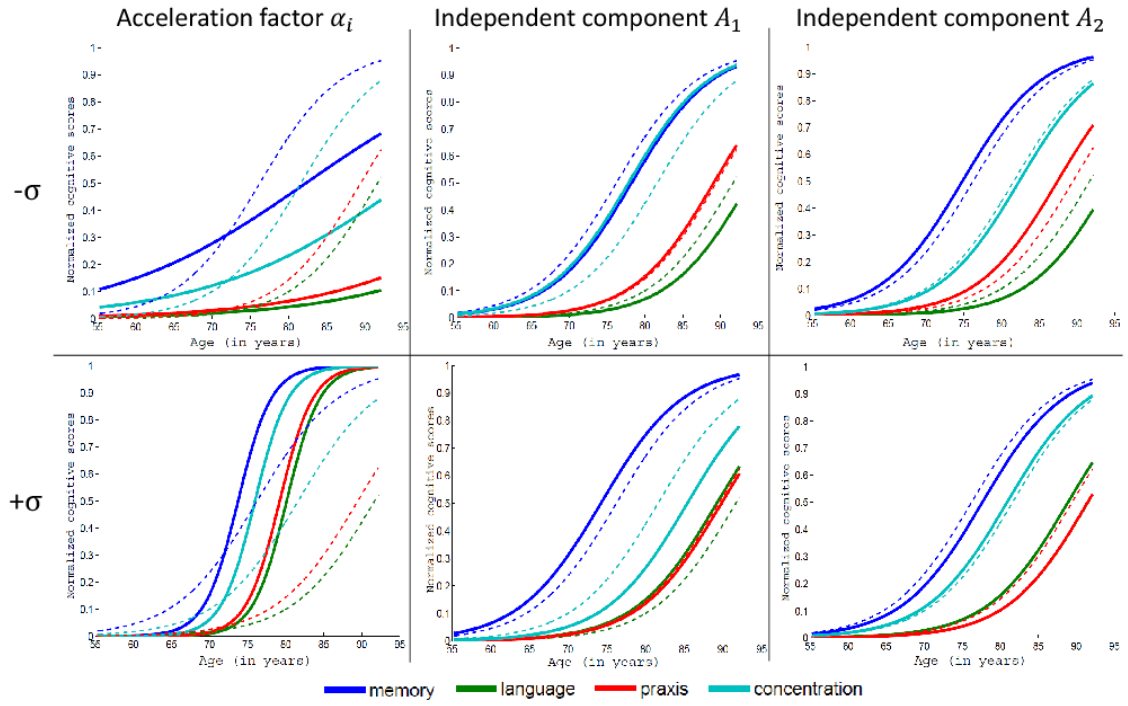

Figure 1. Disease progression model obtained from neuropsychological assessments of 248 patients observed at multiple times (from 3 to 11 times) who converted from Mild Cognitive Impairment stage to Alzheimer's disease during the observation period. Dashed lines represent the average scenario of disease progression (same in all plots). Solid lines represent the variability of this scenario within the observed population in terms of pace of disease progression (left) and relative timing and ordering of the decline of cognitive functions (middle and right).

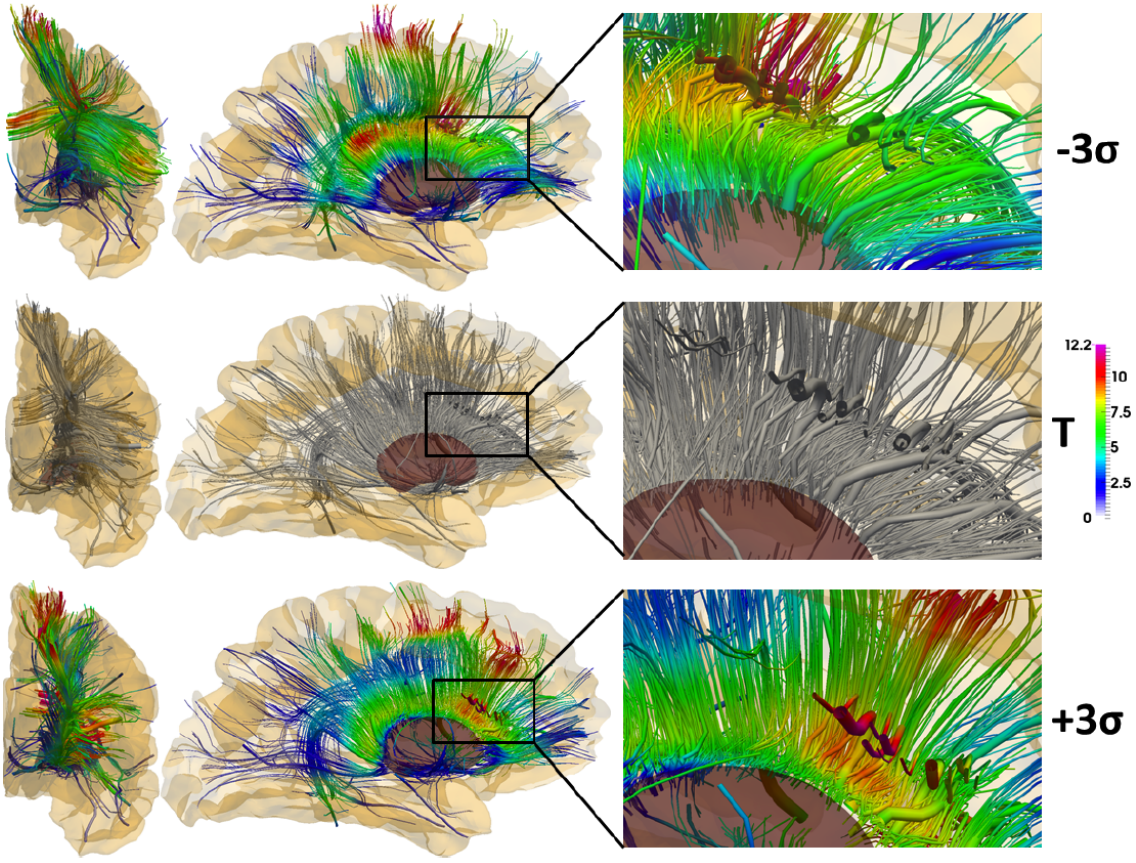

Figure 2. Estimation of a virtual representation of brain structure from anatomical and diffusion images of 3 patients with Gilles de la Tourette syndrome and 2 control subjects. Deformation of the white matter fiber bundle along the first mode of variability is shown while the estimated grey matter frame is kept fixed. Colors refer to the magnitude of displacement during deformation.

In this work, we introduced a diffeomorphic constraint on the deformations considered in the deformable Bayesian Mixed Effect (BME) Template model. Our approach is built on a generic group of diffeomorphisms, which is parameterized by an arbitrary set of control point positions and momentum vectors. This enables us to estimate the optimal positions of control points together with a template image and parameters of the deformation distribution which compose the atlas. We propose to use a stochastic version of the Expectation-Maximization (EM) algorithm where the simulation is performed using the Anisotropic Metropolis Adjusted Langevin Algorithm (AMALA). We propose also an extension of the model including a sparsity constraint to select an optimal number of control points with relevant positions. Experiments are carried out on the USPS database, on mandibles of mice, and on 3D murine dendrite spine images.

More details in [2].

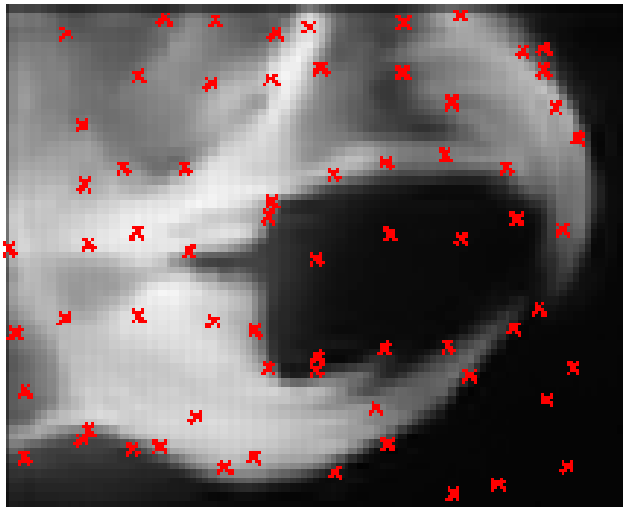

Figure 3. Template image of mouse mandible obtained from 36 X-ray image using 70 control points.

7.4. A sub-Riemannian modular approach for diffeomorphic deformations

Participants: Barbara Gris [Correspondant], Stanley Durrleman, Alain Trouvé.

We develop a generic framework to build large deformations from a combination of base modules. These modules constitute a dynamical dictionary to describe transformations. The method, built on a coherent sub-Riemannian framework, defines a metric on modular deformations and characterises optimal deformations as geodesics for this metric. We present a generic way to build local affine transformations as deformation modules, and display examples.

More details in [27].

7.5. Results of a multicenter randomized placebo-controlled clinical trial in prodromal Alzheimer's disease

Participants: Bruno Dubois, Marie Chupin, Harald Hampel, Simone Lista, Enrica Cavedo, Bernard Croisille, Guy Louis Tisserand, Jacques Touchon, Alain Bonafé, Pierre-Jean Ousset, Amir Ait Ameur, Olivier Rouaud,

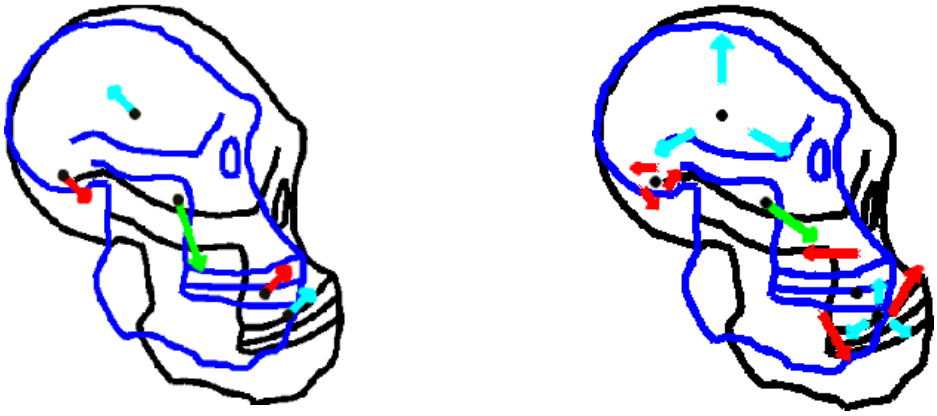

Figure 4. Initial position of deformation modules and their control parameters (left) leads to the construction of local scaling (cyan), rotation (red) and translation (green) (right), which combine together to deform the blue shape into the black one.

Frédéric Ricolfi, Alain Viguetto, Florence Pasquier, Christine Delmaire, Mathieu Ceccaldi, Nadine Girard, Carole Dufouil, Stéphane Lehéricy, Isabelle Tonelli, Françoise Duveau, Olivier Colliot, Line Garnero, Marie Sarazin, Didier Dormont [Correspondant].

Our team coordinated neuroimage acquisition and analysis of a multicenter randomized placebo-controlled clinical trial aiming to assess the efficacy of donepezil in prodromal Alzheimer's disease. Subjects underwent two brain magnetic resonance imaging scans (baseline and final visit). The primary efficacy outcome was the annualized percentage change (APC) of total hippocampal volume (left + right) measured by the software (see Section SACHA 6.3) developped by our team. Two-hundred and sixteen only subjects were randomized across 28 French expert clinical sites. In the per protocol population (placebo = 92 and donepezil = 82), the donepezil group exhibited a significant reduced rate of hippocampal atrophy (APC= -1.89%) compared with the placebo group (APC= -3.47%), P <.001. There was no significant difference in neuropsychological performance between treatment groups. A 45% reduction of rate of hippocampal atrophy was observed in prodromal AD following 1 year of treatment with donepezil compared with placebo.

This new approach opens interesting perspectives for the evaluation of treatments in neurodegenerative diseases.

More details in [12].

7.6. Sulcal morphology as a new imaging marker for the diagnosis of early onset Alzheimer's disease

Participants: Lorraine Hamelin, Bruno Dubois, Marie Chupin, Olivier Colliot [Correspondant], Marie Sarazin.

We investigated the utility of sulcal width measures in the diagnosis of Alzheimer's disease (AD). Sixty-six biologically confirmed AD patients (positive amyloid positron emission tomography [PET] and/or AD cerebrospinal fluid profile) were contrasted to 35 controls with negative amyloid PET. Patients were classified into prodromal or dementia stages as well as into late onset (LOAD, n = 31) or early onset (EOAD, n = 35) subgroups according to their age of onset. An automated method was used to calculate sulcal widths and hippocampal volumes (HV). In EOAD, the greatest ability to differentiate patients from age-matched controls, regardless of severity, was displayed by sulcal width of the temporoparietal cortex. In this region, diagnosis

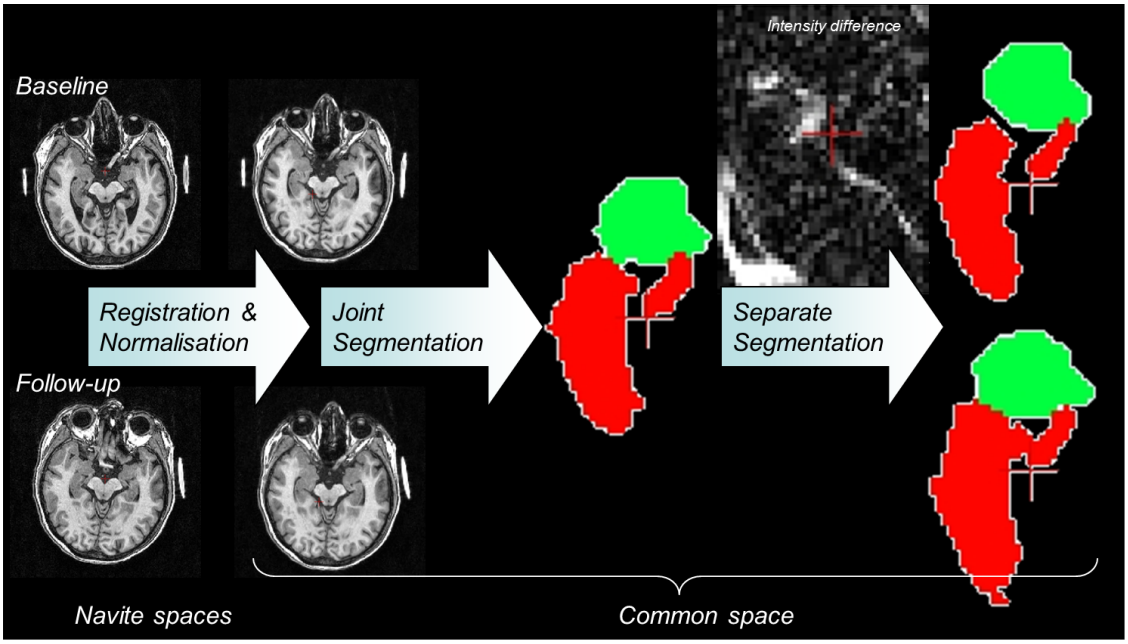

Figure 5. Hippocampus longitudinal segmentation method illustrating preliminary registration of the baseline and final visit magnetic resonance imaging (MRI) scans in a common space followed by normalization of the intensities of both scans. The baseline and final visit MRI scans were then segmented jointly. The resulting segmentation was then used as an initialization of separate segmentations while keeping the two segmentations consistent between the two time-points.

accuracy was better than the HV, especially at prodromal stage. In LOAD, HV provided the best discrimination power from age-matched controls. In conclusion, sulcal width measures are better markers than the HV for identifying prodromal AD in patients aged <65 years. In contrast, in older patients, the risk of over-diagnosis from using only sulcal enlargement is important.

More details in [14].

7.7. Imaging Markers of the Presymptomatic GRN Disease

Participants: Paola Caroppo, Stanley Durrleman, Alexandre Routier, Olivier Colliot [Correspondant], Alexis Brice, Isabelle Le Ber.

The preclinical stage of frontotemporal lobar degeneration (FTLD) is not well characterized. We conducted a brain metabolism (FDG-PET) and structural (cortical thickness) study to detect early changes in asymptomatic GRN mutation carriers (aGRN+) that were evaluated longitudinally over a 20-month period. At baseline, a left lateral temporal lobe hypometabolism was present in aGRN+ without any structural changes. Importantly, this is the first longitudinal study and, across time, the metabolism more rapidly decreased in aGRN+ in lateral temporal and frontal regions. The main structural change observed in the longitudinal study was a reduction of cortical thickness in the left lateral temporal lobe in carriers (Figure 6). A limit of this study is the relatively small sample (n=16); nevertheless, it provides important results. First, it evidences that the pathological processes develop a long time before clinical onset, and that early neuroimaging changes might be detected approximately 20 years before the clinical onset of disease. Second, it suggests that metabolic changes are detectable before structural modifications and cognitive deficits. Third, both the baseline and longitudinal studies provide converging results implicating lateral temporal lobe as early involved in GRN disease. Finally, our study demonstrates that structural and metabolic changes could represent possible biomarkers to monitor the progression of disease in the presymptomatic stage toward clinical onset.

More details in [6].

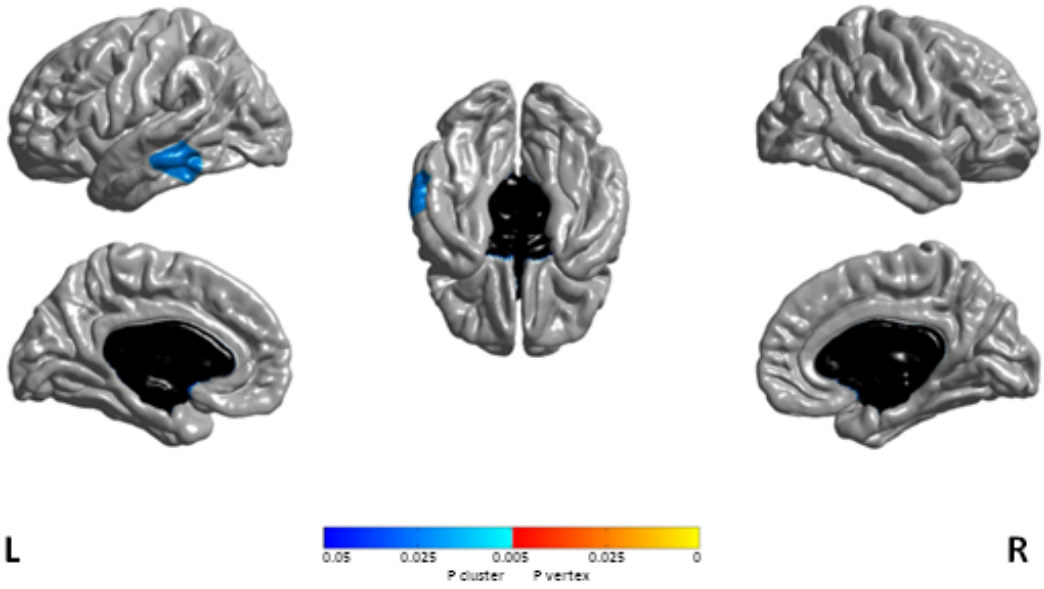

Figure 6. Cluster with significant cortical thickness changes in aGRN+ between the two time-points ($p < 0.05$ corrected). L, left; R, right.

7.8. Incomplete Hippocampal Inversions in healthy subjects: a comprehensive study of over 2000 participants

Participants: Claire Cury [Correspondant], Joan Glaunès, Dominique Hasboun, Fanny Cohen, Jorge Samper-González, Roberto Toro, Vincent Frouin, Gunter Schumann, Olivier Colliot.

The incomplete-hippocampal-inversion (IHI), also known as malrotation, is an atypical anatomical pattern of the hippocampus, which has been reported in healthy subjects in different studies. However, extensive characterization of IHI in a large sample has not yet been performed. Furthermore, it is unclear whether IHI are restricted to the medial-temporal lobe or are associated with more extensive anatomical changes. Here, we studied the characteristics of IHI in a community-based sample of 2008 subjects of the IMAGEN database and their association with extra-hippocampal anatomical variations. The presence of IHI was assessed on T1-weighted anatomical magnetic resonance imaging (MRI) using visual criteria. We assessed the association of IHI with other anatomical changes throughout the brain using automatic morphometry of cortical sulci. We found that IHI were much more frequent in the left hippocampus (left: 17%, right: 6%, χ^2-test, $p < 10^{-28}$). Compared to subjects without IHI, subjects with IHI displayed morphological changes in several sulci located mainly in the limbic lobe. Our results demonstrate that IHI are a common left-sided phenomenon in normal subjects and that they are associated with morphological changes outside the medial temporal lobe.

More details in [9].

7.9. Analysis of anatomical variability using diffeomorphic iterative centroid in patients with Alzheimer's disease

Participants: Claire Cury [Correspondant], Joan Glaunès, Marie Chupin, Olivier Colliot.

We proposed a new approach for template-based analysis of anatomical variability in populations, in the framework of Large Deformation Diffeomorphic Metric Mappings and mathematical currents. We proposef a fast approach in which the template is computed using an diffeomorphic iterative centroid method. Statistical analysis is then performed on the initial momenta that define the deformations between the centroid and each individual subject. We applied the approach to study the variability of the hippocampus in 134 patients with Alzheimer's disease (AD) and 160 elderly control subjects. We show that this approach can describe the main modes of variability of the two populations and can predict the performance to a memory test in AD patients.

More details in [8].

7.10. Innovation-based sparse estimation of functional connectivity from multivariate autoregressive models

Participants: Fabrizio de Vico Fallani [Correspondant], Stéphanie Allassonniere [Correspondant], Francois Deloche.

One of the main limitations of functional connectivity estimators of brain networks is that they can suffer from statistical reliability when the number of areas is large and the available time series are short. To estimate directed functional connectivity with multivariate autoregressive (MVAR) model on sparse connectivity assumption, we propose a modified Group Lasso procedure with an adapted penalty. Our procedure includes the innovation estimates as explaining variables. This approach is inspired by two criteria that are used to interpret the coefficients of the MVAR model, the Directed Transfer Function (DTF) and the Partial Directed Coherence (PDC). A causality measure can be deduced from the output coefficients which can be understood as a synthesis of PDC and DTF. We demonstrate the potential of our method and compare our results with the standard Group Lasso on simulated data.

More details in [25]

7.11. Lucid Dreaming in Narcolepsy

Participants: Pauline Daudet, Mario Chavez [Correspondant], Smaranda Leu-Semenescu, Jean-Louis Golmard, Isabelle Arnulf.

Lucid dreaming is the experience of being aware of dreaming while asleep and continuing to dream. Lucid dreams generally arise in REM sleep. Compared to non-lucid REM sleep, lucid REM sleep is associated with local frontal lobe EEG changes in the 40 Hz band, increased brain coherence, and increased activity on functional MRI in the bilateral precuneus, cuneus, parietal lobules, and prefrontal and occipito-temporal cortices, which may correspond to restored reflective consciousness. We decided to study the frequency and determinants of lucid dreaming in narcolepsy and to challenge patients' alleged ability to achieve lucid dreaming using sleep monitoring during nighttime and daytime sleep. Compared to 53 healthy controls, the 53 narcolepsy patients reported more frequent dream recall, nightmares and recurrent dreams. The frequency of cataplexy, hallucinations, sleep paralysis, dyssomnia, positivity, and the severity of sleepiness were similar in narcolepsy with and without lucid dreaming. The delta power in the electrode average, in delta, theta, and alpha powers in C4, and coherences between frontal electrodes were lower in lucid than non-lucid REM sleep in spectral EEG analysis. The duration of REM sleep was longer, the REM sleep onset latency tended to be shorter, and the percentage of atonia tended to be higher in lucid vs. non-lucid REM sleep; the arousal index and REM density and amplitude were unchanged. Our results suggest that narcoleptics have a high propensity for lucid dreaming without differing in REM sleep characteristics from people without narcolepsy. This also suggests that narcolepsy patients may provide useful information in future studies on the nature of lucid dreaming.

More details in [11]

7.12. An Algebraic Topological Method for Multimodal Brain Networks Comparisons

Participants: Tiago Simas, Mario Chavez [Correspondant], Pablo Rodriguez, Albert Diaz-Guilera.

Understanding brain connectivity is one of the most important issues in neuroscience. Nonetheless, connectivity data can reflect either functional relationships of brain activities or anatomical connections between brain areas. Although both representations should be related, this relationship is not straightforward. We have devised a powerful method that allows different operations between networks that share the same set of nodes, by embedding them in a common metric space, enforcing transitivity to the graph topology. Here, we apply this method to construct an aggregated network from a set of functional graphs, each one from a different subject. Once this aggregated functional network is constructed, we use again our method to compare it with the structural connectivity to identify particular brain regions that differ in both modalities (anatomical and functional). Remarkably, these brain regions include functional areas that form part of the classical resting state networks. We conclude that our method -based on the comparison of the aggregated functional network- reveals some emerging features that could not be observed when the comparison is performed with the classical averaged functional network.

More details in [23]

7.13. Steady state visual evoked potentials-based patient interface under breathing constraints

Participants: Xavier Navarro [Correspondant], Sebastien Campion, Fabrizio de Vico Fallani [Correspondant], Pierre Pouget, Thomas Similowski, Mathieu Raux, Mario Chavez.

Steady state visual evoked potentials (SSVEP) have been widely utilized in brain computer interfacing (BCI) in last years. In this paper, we present a study exploring the possibilities of SSVEP to manage the communication between patients suffering respiratory disorders and health care providers. By imposing different breathing constraints, five healthy subjects communicated their breathing sensations (breathing well/breathing bad) using a visual frequency tagging paradigm: two visual stimuli with different flickering frequencies (15 and 20 Hz) were simultaneously presented on a screen. Using electroencephalographic (EEG) signals from only three EEG electrodes, two spectral features were extracted by a spatial filter in a sliding window, then classified by an unsupervised algorithm based on k-medians. Average detection success rates were of 70% during breathing

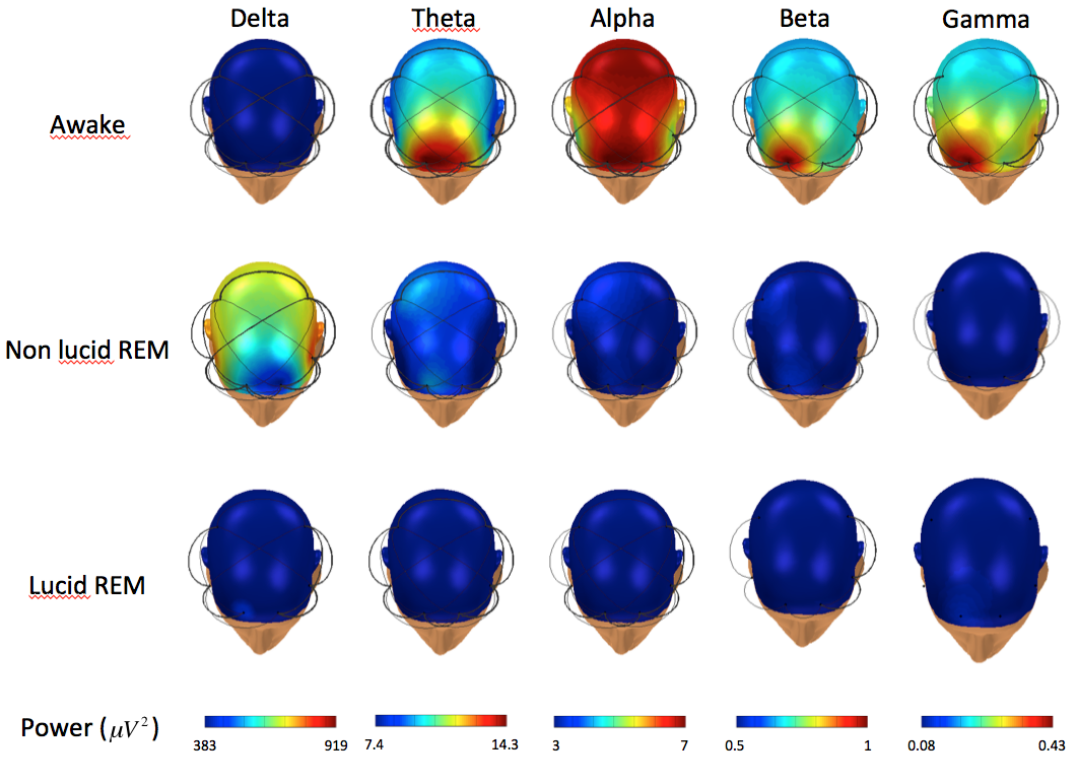

Figure 7. Topographical distribution (obtained by a spherical spline interpolation) of EEG spectral power during wakefulness (top row), non-lucid (middle row) and lucid (bottom row) REM sleep for different frequency bands. Signi cant couplings between the electrodes are indicated by the black links (the thickness is proportional to the coherence value). Colors from dark blue (lower EEG power) to dark red (higher EEG power) indicated for each EEG band in the Power line (bottom row).

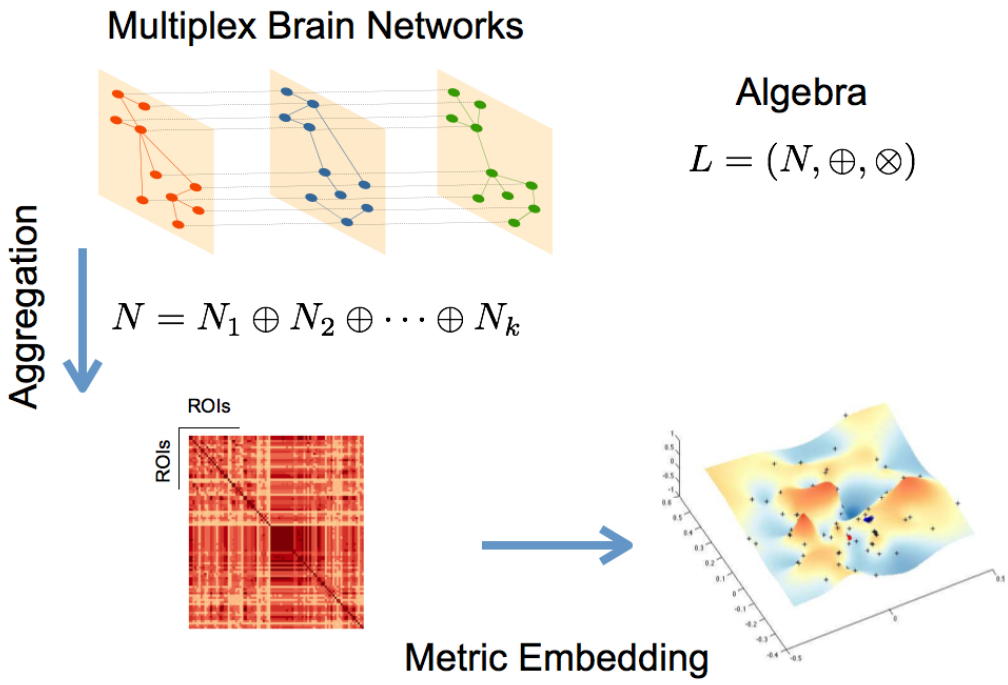

Figure 8. Schematic representation of the main steps for the described networks aggregation and metric embedding (defined here for the algebra L)

discomfort, and of 83% when subjects breathed comfortably. Results suggest that SSVEP-based BCI may be a promising choice to improve patient-caregiver communication in situations of breathing discomfort when verbal communication is difficult.

More details in [29]

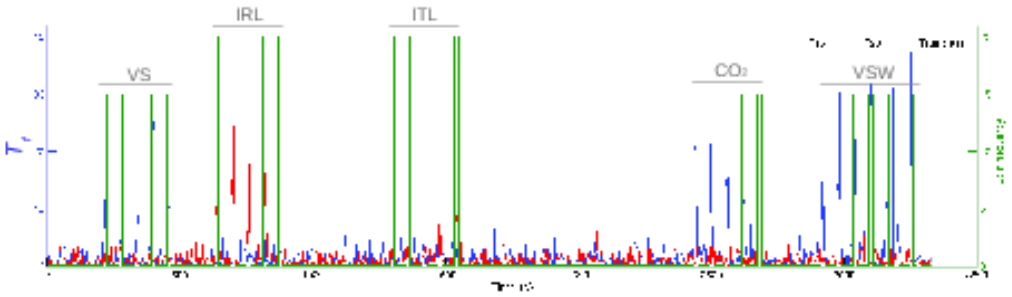

Figure 9. An example of T_f obtained after applying the spatial filters on 15 Hz (blue curve) and 20 Hz (red curve) during the experiment in subject 3. The statistics T_f reflects the signal-to-noise ratio at frequency f with respect to the no-stimulus power.

CASCADE Project-Team

6. New Results

6.1. Results

All the results of the team have been published in journals or conferences (see the list of publications). They are all related with the research program (see before) and the research projects (see after):

- New zero-knowledge proofs
- Advanced families of hash proofs
- More efficient constructions with lattices
- New e-cash constructions
- Advanced primitives for the privacy in the cloud
- Efficient functional encryption
- Various predicate encryption schemes
- Cryptanalysis of symmetric primitives
- New leakage-resilient primitives
- Stronger security with related-key security

CLIME Project-Team

6. New Results

6.1. Simulation, observation and state estimation for analysis and forecast

The objective of Clime is the merging of simulation and observations, with data assimilation methods, for state estimation in environmental applications. However, this aim previously requires, as seen in some of the next subsection, to collect the observations and carry out the simulations.

6.1.1. Assimilation of drifter data in the East Mediterranean Sea

Participants: Julien Brajard, Milad Fakhri [CNRS, Lebanon], Daniel Hayes [Oceanography Centre, Cyprus], Leila Issa [Lebanese American University, Lebanon], Laurent Mortier [LOCEAN], Pierre-Marie Poulain [Oceanography Institute of Trieste, Italy].

Surface velocity fields of the ocean in the Eastern Levantine Mediterranean are estimated by blending altimetry and surface drifters data. The method is based on a variational assimilation approach for which the velocity is corrected by matching real drifters positions with those predicted by a simple advection model, while taking into account the wind effect. The velocity correction is done in a time-continuous fashion by assimilating at once a whole trajectory of drifters using a sliding time window. A divergence-free regularization term is added to the cost function mnimized during the assimilation process in order to constrain the velocity field. First results show that with few drifters, the method improves the estimation of the surface velocity: an eddy between the Lebanese coast and Cyprus is better assessed and the values of velocities along the Lebanese coast are more accurate.

6.1.2. Traffic simulation

Participants: Vivien Mallet, Vincent Aguiléra [CEREMA], Ruiwei Chen [CEREA].

The ANR project ESTIMAIR aims at propagating uncertainties in the complete simulation chain of air quality at urban scale. A key step in the chain lies in traffic assignment and the computation of the corresponding emissions. We take part to the simulation of traffic in the streets of Clermont-Ferrand metropolitan area, with the dynamic traffic assignment model LADTA. The simulations are evaluated against observations from loop counters and also against the simulations of the reference static model VISUM.

From the traffic assignment, the emissions are computed for nitrogen dioxide and particulate matter, using COPERT IV formulae. Preliminary work shows large uncertainties in the emissions due to the fleet composition.

6.1.3. Observation of noise pollution

Participants: Vivien Mallet, Raphaël Ventura, Valérie Issarny [MiMove], Pierre-Guillaume Raverdy [Ambientic], Fadwa Rebhi [MiMove].

Exposure to noise pollution is highly variable in space. As a consequence, it is very difficult to determine individual exposure using only numerical simulations of noise levels. Together with the MiMove Inria project-team, we take part to the SoundCity project that aims at collecting noise observations from smartphones and better evaluating the individual exposure. We assist MiMove in the development of an Android application that automatically senses noise along the day and collects the data (when the user agrees) for the improvement of simulated noise maps. Clime especially contributes to the calibration of the application. Comparisons between the measurements of smartphones and a sound meter allow us to estimate the bias of the main smartphones available on the market.

The SoundCity application was launched in July 2015 with Bernard Jomier, deputy mayor responsible for health, disability, and relations with Paris public hospital system, during a press conference organized by Paris City. The application received a positive coverage in the media, so that the application gained about 2500 users. About one million observations are collected every four days and ongoing work tries to process these data to correct Paris noise maps.

6.1.4. Evaluation of fire models

Participants: Jérémy Lefort, Vivien Mallet, Jean-Baptiste Filippi [CNRS].

In the field of forest fires risk management, important challenges exist in terms of people and goods preservation. Answering to strong needs from different actors (firefighters, foresters), researchers focus their efforts to develop operational decision support system tools that may forecast wildfire behavior. This requires the evaluation of model performance.

We carry out the evaluation of several fire propagation models based on over 500 real fires. We use the data as they would be available in operational conditions, so as to avoid any tuning that would be incompatible with real-time forecasting. The study shows significant performance difference between the models, despite the poor data quality.

6.2. Image assimilation

Sequences of images, such as satellite acquisitions, display structures evolving in time. This information is recognized of major interest by forecasters (meteorologists, oceanographers, etc.) in order to improve the information provided by numerical models. However, the satellite images are mostly assimilated in geophysical models on a point-wise basis, discarding the space-time coherence visualized by the evolution of structures such as clouds. Assimilating in an optimal way image data is of major interest and this issue should be considered in two ways:

- from the model's viewpoint, the location of structures on the observations is used to control the state vector.
- from the image's viewpoint, a model of the dynamics and structures is built from the observations.

6.2.1. Model error and motion estimation

Participants: Isabelle Herlin, Dominique Béréziat [UPMC].

Data assimilation technics are used to retrieve motion from image sequences. These methods require a model of the underlying dynamics, displayed by the evolution of image data. In order to quantify the approximation linked to the chosen dynamic model, an error term is included in the evolution equation of motion and a weak formulation of 4D-Var data assimilation is designed. The cost function to be minimized depends simultaneously on the initial motion field, at the beginning of the studied temporal window, and on the error value at each time step. The result allows to assess the model error and analyze its impact on motion estimation. The approach is used to estimate geophysical forces (gravity, Coriolis, diffusion) from images in order to better assess the surface dynamics and forecast the displacement of structures like oilspill.

6.2.2. Tracking of structures from an image sequence

Participants: Isabelle Herlin, Yann Lepoittevin, Dominique Béréziat [UPMC].

The research concerns an approach to estimate velocity on an image sequence and simultaneously segment and track a given structure. It relies on the underlying dynamics' equations of the studied physical system. A data assimilation method is designed to solve evolution equations of image brightness, those of motion's dynamics. The method is for instance applied on meteorological satellite data, in order to track tropical clouds on image sequences and estimate their motion, as seen on Fig. 2 .

Data assimilation is performed either with a 4D-Var variational approach or with a Kalman ensemble method. In the last case, the initial ensemble is obtained from a set of optical flow methods of the literature with various parameters values.

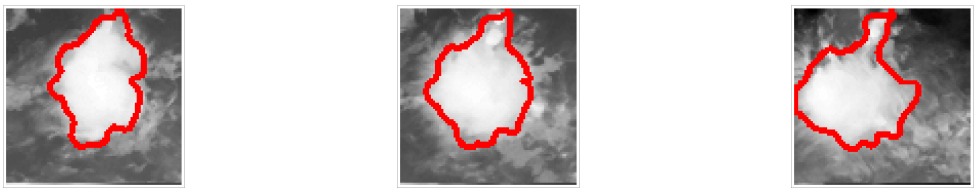

Figure 2. Tracking a tropical cloud. Frames 3, 9, 18 of the sequence.

Various ways for representing the structures are studied and compared.

- For the variational approach, we consider: 1- a distance map modeling the tracked structures, which is added to the state vector, 2- anisotropic regularization terms, which are added to the cost function minimized during the assimilation process, 3- covariances between pixels, which are included in the background error covariance matrix.
- For the filtering approach, we focus either on domain decomposition or on explicit localization, which are both related to the displayed structures.

6.2.3. Applying POD on a model output dabase for defining a reduced motion model

Participants: Isabelle Herlin, Etienne Huot.

Dimension reduction may be obtained by determining a small size reduced basis computed by Proper Orthogonal Decomposition (POD) of a motion fields database and applying the Galerkin projection. This database is constructed for characterizing accurately the surface circulation of the studied area, so that linear combinations of the basis elements obtained by POD accurately describe the motion function observed on satellite image sequences. The database includes the geostrophic motion fields obtained from Sea Level Anomaly reanalysis maps that are available from the MyOcean European project website (http://marine.copernicus.eu/). Fig. 3 displays such SLA maps and the associated motion fields.

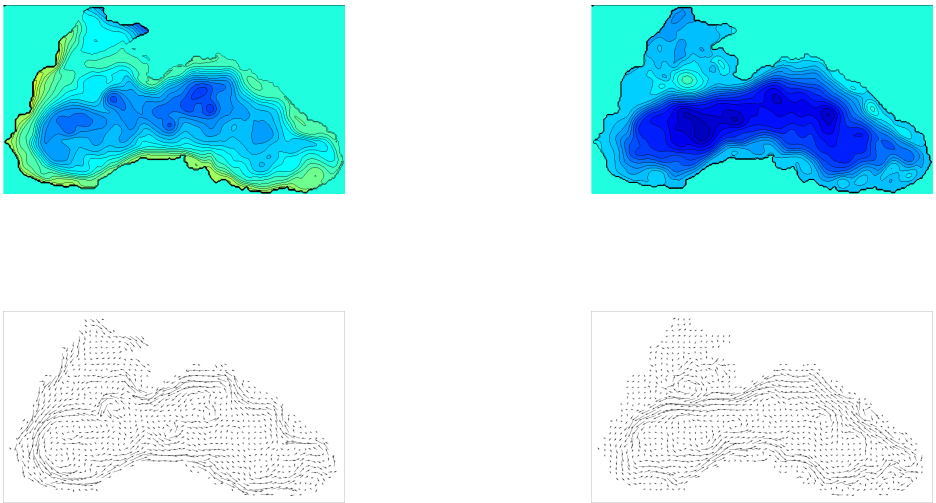

Figure 3. Top: reanalysis of SLA. Bottom: geostrophic motion.

Image assimilation with the POD reduced model allows estimating motion as displayed on Fig. 4 .

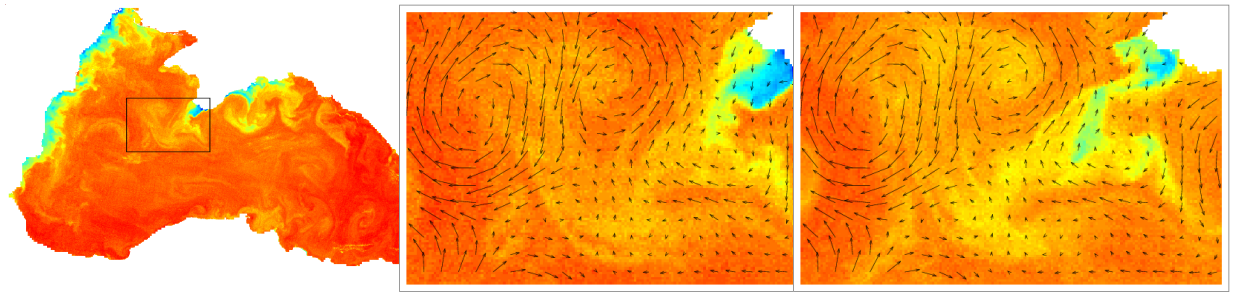

Figure 4. Zoom on a region of interest and motion estimation superposed on two consecutive images.

6.2.4. Rain nowcasting from radar image acquisitions

Participants: Isabelle Herlin, Yann Lepoittevin.

This research concerns the design of an operational method for rainfall nowcasting that aims at mitigating flash floods. The nowcasting method is composed of two main components:

- a data assimilation method, based on radar images, estimates the state of the atmosphere: this is the estimation phase.
- a forecast method uses this estimation to extrapolate the state of the atmosphere in the future: this is the forecast phase.

The method is transfered to the industrial company Weather Measures.

Current research concerns the use of object components in the state vector in order to get an improved motion estimation and a better localization of endangered regions. Assimilation of pluviometers measures in the nowcasting system is also investigated.

6.3. Uncertainty quantification and risk assessment

The uncertainty quantification of environmental models raises a number of problems due to:

- the dimension of the inputs, which can easily be 10^5-10^8 at every time step;
- the dimension of the state vector, which is usually 10^5-10^7;
- the high computational cost required when integrating the model in time.

While uncertainty quantification is a very active field in general, its implementation and development for geosciences requires specific approaches that are investigated in Clime. The project-team tries to determine the best strategies for the generation of ensembles of simulations. In particular, this requires addressing the generation of large multimodel ensembles and the issue of dimension reduction and cost reduction. The dimension reduction consists in projecting the inputs and the state vector to low-dimensional subspaces. The cost reduction is carried out by emulation, i.e., the replacement of costly components with fast surrogates.

6.3.1. Application of sequential aggregation to meteorology

Participants: Paul Baudin, Vivien Mallet, Gilles Stoltz [CNRS].

Nowadays, it is standard procedure to generate an ensemble of simulations for a meteorological forecast. Usually, meteorological centers produce a single forecast, out of the ensemble forecasts, computing the ensemble mean (where every model receives an equal weight). It is however possible to apply aggregation methods. When new observations are available, the meteorological centers also compute analyses. Therefore, we can apply the ensemble forecast of analyses, which consists in weighting the ensemble of forecasts to better forecast the forthcoming analyses. Before any forecast, the weights are updated with past observations and past forecasts. The performance of the aggregated forecast is guaranteed, in the long run, to perform at least as well as any linear combination of the forecasts with constant weights.

Ensembles of forecasts for mean sea level pressure, from the THORPEX Interactive Grand Global Ensemble, are aggregated with a forecast error decreased by 18% compared to the best individual forecast. The approach is also proved to be efficient for wind speed. The contribution of the ensembles (from different meteorological centers) to the performance increase are evaluated.

6.3.2. Sequential aggregation with uncertainty quantification and application to photovoltaics production

Participants: Paul Baudin, Vivien Mallet, Jean Thorey, Christophe Chaussin [EDF R&D], Gilles Stoltz [CNRS].

We study the aggregation of ensembles of solar radiations and photovoltaic productions. The aggregated forecasts show a 20% error decrease compared to the individual forecasts. They are also able to retrieve finer spatial patterns than the ones found in the individual forecasts (see Figure 5).

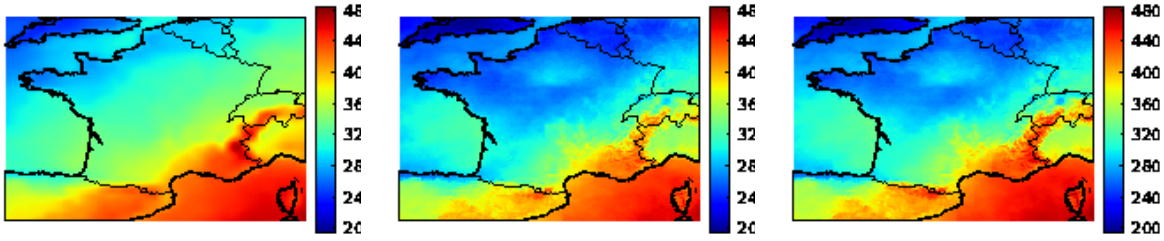

(a) (b) (c)

Figure 5. Yearly average of the map of downward shortwave solar radiation in $\mathrm{W m^{-2}}$*, for an ensemble mean (a), for our aggregated forecasts (b) and observed (c).*

An important issue is the estimation of the uncertainties associated with the aggregated forecasts. We devise a new approach to predict a probability density function or a cumulative distribution function instead of a single aggregated forecast. In practice, the aggregation procedure aims at forecasting the cumulative distribution function of the observations which is simply a Heaviside function centered at the observed value. Our forecast is the weighted empirical cumulative distribution function based on the ensemble of forecasts. The method guarantees that, in the long run, the forecast cumulative distribution function has a Continuous Ranked Probability Score (CRPS) at least as good as the best weighted empirical cumulative function with weights constant in time.

The CRPS is a classical score to evaluate the probabilistic forecasts. However, applying the CRPS on weighted empirical distribution functions (derived from the weighted ensemble) introduces a bias because of which minimizing the CRPS does not produce the optimal weights. Thus, we propose an unbiased version of the CRPS which relies on clusters of members and is strictly proper.

6.3.3. Sensitivity analysis in the dispersion of radionuclides

Participants: Sylvain Girard, Vivien Mallet, Irène Korsakissok [IRSN].

We carry out a sensitivity analysis of the dispersion of radionuclides during Fukushima disaster. We considered the dispersion at regional scale, with the Eulerian transport model Polair3D from Polyphemus. Simulations of the atmospheric dispersion of radionuclides involve large uncertainties originating from the limited knowledge of meteorological input data, composition, amount and timing of emissions and some model parameters. We studied the relative influence of each uncertain input on several outputs. In practice, we used the variance-based sensitivity analysis method of Sobol. This method requires a large number of model evaluations which are not achievable directly due to the high computational cost of the model. To circumvent this issue, we built a mathematical approximation of the model using Gaussian process emulation.

In previous studies, the uncertainties in the meteorological forecasts were crudely modeled by homogeneous and constant perturbations on the fields. Hence, we started investigating the use of ensembles of meteorological forecasts instead of just one base meterological forecast. Including such ensembles allows to better represent the directions along which meteorological uncertainties should lie.

6.3.4. Fire risk assessment

Participants: Jérémy Lefort, Vivien Mallet, Jean-Baptiste Filippi [CNRS].

During days with extreme weather conditions, every wildland fire must be fought within minutes of its occurrence. This means that sufficient firefighting force is available at the right place and at the right time. In practice, firefighters wait at different critical locations, so that they can act quickly. For efficient preventive positioning of the firefighters, forecasting the risks of ignition of large fires is essential. This requires to predict where a fire may start, to estimate its potential size, to evaluate fighting scenarios and to anticipate which urban or protected areas may be under threat.

We designed a surrogate propagation model based on Gaussian process emulation of the model ForeFire. This surrogate model is fast enough to be run all over a region with high fire risk, e.g., Corsica. It can even be used for Monte Carlo simulations, with perturbations in the meteorological conditions and vegetation state, over Corsica. It is then possible to generate a risk map that identifies all the locations where a new fire can grow large.

6.3.5. Ensemble variational data assimilation

Participants: Julien Brajard, Isabelle Herlin, Marc Bocquet [CEREA], Jérôme Sirven [LOCEAN], Olivier Talagrand [LMD, ENS], Sylvie Thiria [LOCEAN].

The general objective of ensemble data assimilation is to produce an ensemble of analysis from observations and a numerical model which is representative of the uncertainty of the system. In a bayesian framework, the ensemble represents a sampling of the state vector probability distribution conditioned to the available knowledge of the system, denoted the a-posteriori probability distribution.

Ensemble variational data assimilation (EnsVar) consists in producing such an ensemble, by perturbating N times the observations according to their error law, and run a standard variationnal assimilation for each perturbation. An ensemble of N members is then produced. In the case of linear models, there is a theoretical guarantee that this ensemble is a sampling of the a-posteriori probability. But there is no theoretical result in the non-linear case.

Numerical experiments using non-linear numerical models suggest that the conclusion reached for linear models still stands for non-linear toy models.

The objective of this work is to study the ability of EnsVar to produce "good" ensemble (i.e. that sampled the a posteriori probablility) on a more realistic model: a shallow-water model. Some statistical properties of the ensemble are presented, and the sensitivity to the main features of the assimilation system (number, distribution of observations, size of the assimilation window, ...) are also studied.

CRYPT Team (section vide)

DEDUCTEAM Team

7. New Results

7.1. Termination

In [15], Frédéric Blanqui showed how to extend the notion of reducibility introduced by Girard for proving the termination of β-reduction in the polymorphic λ-calculus, to prove the termination of various kinds of rewrite relations on λ-terms, including rewriting modulo some equational theory and rewriting with matching modulo $\beta\eta$, by using the notion of computability closure. This provides a powerful termination criterion for various higher-order rewriting frameworks, including Klop's Combinatory Reductions Systems with simple types and Nipkow's Higher-order Rewrite Systems.

In [16], Frédéric Blanqui, together with Jean-Pierre Jouannaud and Albert Rubio, introduced the computability path ordering (CPO), a recursive relation on terms obtained by lifting a precedence on function symbols. A first version, core CPO, is essentially obtained from the higher-order recursive path ordering (HORPO) by eliminating type checks from some recursive calls and by incorporating the treatment of bound variables as in the so-called computability closure. The well-foundedness proof shows that core CPO captures the essence of computability arguments à la Tait and Girard, therefore explaining its name. We further show that no more type check can be eliminated from its recursive calls without loosing well-foundedness, but one for which we found no counterexample yet. Two extensions of core CPO are then introduced which allow one to consider: the first, higher-order inductive types; the second, a precedence in which some function symbols are smaller than application and abstraction.

Another extension of CPO, to dependently typed terms, has been developed by Jean-Pierre Jouannaud and Jianqi Li in [50].

Jean-Pierre Jouannaud and Albert Rubio showed in [51] how to modify recursive path orders for higher-order terms which, like CPO, include $\beta\eta$-reductions, into orders that are compatible with $\beta\eta$-conversion. The result is a powerful order for proving termination of higher-order rewrite rules based on higher-order pattern matching.

Gaëtan Gilbert and Olivier Hermant have introduced a constructive way to perform proof normalization through completeness proofs [23].

Frédéric Blanqui formalized Ramsey's proof of the (infinite) Ramsey's theorem [54] (see http://color.inria.fr/).

7.2. Confluence

Jean-Pierre Jouannaud, in collaboration with Jiaxiang Liu, has started a program in order to enable confluence proofs in $\lambda\Pi$ modulo, investigating several open confluence problems for non-terminating relations. In [27], together with Mizuhito Ogawa, they introduced the new class of layered rewrite systems, and showed that their confluence can be reduced to that of their critical pairs computed by using unification over infinite rational terms when they do not increase the layer-depth of terms. This shows why an old example of non-terminating, left non-linear, critical pair free rewrite system due to Klop was non-confluent: it indeed had a critical pair in infinite rational trees. In the same paper, they also give an example of a non-confluent, layer-depth increasing system which has no critical pairs, hence showing that layer-depth plays a key role.

7.3. Automated theorem proving

In [25], Guillaume Bury, Raphaël Cauderlier and Pierre Halmagrand presented the extension of the automated theorem prover Zenon to ML-style polymorphism.

In [20], Guillaume Bury, David Delahaye, Damien Doligez, Pierre Hamalgrand and Olivier Hermant introduced an encoding of the set theory of the B method using polymorphic types and deduction modulo, used for the automated verification of proof obligations in the framework of the BWare project.

In [24], Kailiang Ji designed a strategy to translate model-checking problems into proving the satisfiability of a set of first-order formulas. The focus is to give an encoding of temporal properties expressed in CTL as first-order formulas, by translating the logical equivalence between temporal operators into rewrite rules. In this way, proof-search algorithms designed for Deduction Modulo, such as Resolution Modulo or Tableaux Modulo, can be used to verify temporal properties of finite transition systems. This strategy is implemented in iProver Modulo, and the testing results show that Resolution Modulo can be considered as a new way to quickly determine whether a temporal property is violated or not in transition system models.

7.4. $\lambda\Pi$ modulo and Dedukti

Gaëtan Gilbert, supervised by Arnaud Spiwack, wrote a prototype of a principle unification and type inference mechanism for Dedukti, based on a monadic API. This prototype separates with an abstraction barrier a unifier kernel which implements correct unification primitives from the unification algorithm and heuristics. The unification algorithm is written in a style which closely mirrors a pen-and-paper deduction rule presentation.

Éric Uzena, supervised by David Delahaye and Arnaud Spiwack, wrote a prototype of an extension of Dedukti with associative and commutative symbols and rewriting modulo associativity and commutativity of these symbols.

7.5. Encodings into Dedukti and interoperability

Ali Assaf, Guillaume Burel, Raphaël Cauderlier, David Delahaye, Gilles Dowek, Catherine Dubois, Frédéric Gilbert, Pierre Hamalgrand, Olivier Hermant, and Ronan Saillard have written a synthetic paper on the Dedukti system and on the expression of theories in this system. This paper is submitted to publication.

Ali Assaf [32] proved that Cousineau and Dowek's embedding of functional pure type systems [41] is conservative with respect to the original systems, using a new notion of reducibility called relative normalization. Together with Cousineau and Dowek's original result on the preservation of typing, this result justifies the use of the $\lambda\Pi$-calculus modulo as a logical framework.

Ali Assaf's translation of the calculus of inductive constructions to the $\lambda\Pi$-calculus modulo, which was presented at the TYPES conference in 2014, has been published in the postproceedings of TYPES 2014 [39]. This translation, which is based on the translation of pure type systems by Cousineau and Dowek [41], is implemented in the automated translation tool Coqine.

Ali Assaf and Guillaume Burel presented their translation of HOL to Dedukti at the PxTP 2015 workshop [18]. This translation, which is based on the translation of pure type systems by Cousineau and Dowek [41], is implemented in the automated translation tool Holide.

Raphaël Cauderlier and Catherine Dubois' translation of object calculus and subtyping to Dedukti, which was presented at the TYPES conference in 2014, has been published in the post-proceedings of TYPES 2014 [34].

In [26], Raphaël Cauderlier and Pierre Halmagrand presented a shallow embedding into Dedukti of proofs produced by ZenonModulo, an extension of the tableau-based first-order theorem prover Zenon to deduction modulo and typing.

In [33], Ali Assaf and Raphaël Cauderlier have combined simple developments written in Coq and HOL using Dedukti and the existing translation tools Coqine and Holide. This work is a first step towards using Dedukti as a framework for proof interoperability.

7.6. Proof theory

Guillaume Burel, Gilles Dowek and Ying Jiang have introduced a general framework to prove the decidability of reachability and provability problems. This framework uses an analogy between the objects recognized by an automaton and cut-free proofs. Various aspects of this work have been published at FroCoS [19], LPAR [21], and another paper is in preparation.

Gilles Dowek's paper on the definition of the classical connectives and quantifiers has been published [30].

Arnaud Spiwack gave a predicative shallow embedding of a weak version of system U^- in dependent type theory, for Hurkens's paradox to hold. He also showed that a variety of incarnations of Hurkens's paradox are straightforward instantiations of this encoding, greatly simplifying existing proofs.

Arnaud Spiwack developped a topos-theoretic methodology to reason equationally on circuit languages. Results that hold for combinational circuits are lifted to sequential circuits thanks to a transfer principle. This approach allows, in particular, to simplify reasoning about more complex temporal gates than the unit delay. These results aim at enriching the compiler of the Faust audio signal processing programming language, which features such complex temporal gates.

For the sake of reliability, the kernels of Interactive Theorem Provers (ITPs) are kept relatively small in general. On top of the kernel, additional symbols and inference rules are defined. Some dependency analysis of symbols of HOL Light indicates that the depth of dependency could be reduced by introducing a few more symbols to the kernel. Shuai Wang showed that extending the kernel of HOL Light is a successful attempt to reduce proof size and speed up proof-checking. More specifically, symbols and inference rules of universal quantification and implication were added to the kernel. This approach has been proved to give equivalent proof-checking results with the size of the proof files reduced to 24% on average and a speedup of 38% for proof-checking overall.

7.7. Computation models

Pablo Arrighi and Gilles Dowek have studied the expression of mecanic motions in cellular automata. Part of this work has been published in TPNC [17] and another paper is in preparation.

Arnaud Spiwack developped a variant of Turing machine where the tape is replaced by an unlabeled tree. The additional structure makes combining machines much easier, making it tractable to give explicit descriptions of rather complex machines. The cost model of these machines models that of purely functional programming languages, making it possible to compare mathematically the complexity of imperative algorithms and of purely functional algorithms.

DYOGENE Project-Team

7. New Results

7.1. Evaluation and optimization of the quality of service perceived by mobile users for new services in cellular networks

The goal of this thesis[1] defended in 2015 is to develop tools and methods for the evaluation of the QoS (Quality of Service) perceived by users, as a function of the traffic demand, in modern wireless cellular networks. This complex problem, directly related to network dimensioning, involves modeling dynamic processes at several time-scales, which due to their randomness are amenable to probabilistic formalization. Firstly, on the ground of information theory, we capture the performance of a single link between a base station and a user in the context of a cellular network with orthogonal channels and MIMO technology. We prove and use some lower bounds of the information-theoretic ergodic capacity of such a link, which account also for the fast channel variability caused by multi-path propagation. These bounds give robust basis for further user QoS evaluation. Next, one considers several (possibly mobile) users, arriving in the network and requesting some service from it. We consider variable (elastic) bit-rate services, in which transmissions of some amounts of data are realized in a best-effort manner, or constant bit-rate services, in which a certain transmission rate needs to be maintained during requested times. On the ground of queuing theory, one captures this traffic demand and service process using appropriate (multi-class) processor sharing (PS) or loss models. In this thesis, we adapt existing PS models and develop a new loss model for wireless streaming traffic, in which the aforementioned information-theoretic capacities of single links describe the instantaneous user service rates. The multi-class models are used to capture the spatial heterogeneity of user channels, which depends on the user geographic locations and propagation shadowing phenomenon. Finally, on top of the queueing-theoretic processes, one needs to consider a multi-cellular network, whose base stations are not necessarily regularly placed, and whose geometry is further perturbed by the shadowing phenomenon. We address this randomness aspect by using some models from stochastic geometry, notably Poisson point processes and Palm formalism applied to the typical cell of the network. Applying the above three-fold approach, supposed to represent all crucial mechanisms and engineering parameters of cellular networks (such as LTE), we establish some macroscopic relations between the traffic demand and the user QoS metrics for some elastic and constant bit-rate services. These relations are mostly obtained in a semi-analytic way, i.e., they only involve static simulations of a Poisson point process (modeling the locations of base stations) in order to evaluate its characteristics which are not amenable to analytic expressions. More precisely, regarding the data traffic (the elastic bit-rate service), we capture the inter-cell interference, making the PS queue models of individual cells dependent, via some system of cell-load equations. These equations allow one to determine the mean user throughput, the mean number of users and the mean cell load in a large network, as a function of the traffic demand. The spatial distribution of these QoS metrics in the network is also studied. We validate our approach by comparing the obtained results with those measured from live-network traces. We observe a remarkably good agreement between the model predictions and the statistical data collected in several deployment scenarios. Regarding constant bit-rate services, we propose a new stochastic model to evaluate the frequency and the number of interruptions during real-time streaming calls in function of user radio conditions. Despite some fundamental similarities with the classical Erlang loss model, a more adequate model was required for in this case, where the denial of service is not definitive for a given call: it takes the form of, hopefully short, interruptions or outage periods. Our model allows one to take into account realistic implementations of the considered streaming service. We use it to study the quality of service metrics in function of user radio conditions in LTE networks. All established results contribute to the development of network dimensioning methods and are currently used in Orange internal tools for network capacity calculations.

7.2. Interference and SINR coverage in spatial non-slotted Aloha networks

In [8] we propose two analytically tractable stochastic-geometric models of interference in ad-hoc networks using pure (non-slotted) Aloha as the medium access. In contrast the slotted model, the interference in pure Aloha may vary during the transmission of a tagged packet. We develop closed form expressions for the Laplace transform of the empirical average of the interference experienced during the transmission of a typical packet. Both models assume a power-law path-loss function with arbitrarily distributed fading and feature configurations of transmitters randomly located in the Euclidean plane according to a Poisson point process. Depending on the model, these configurations vary over time or are static. We apply our analysis of the interference to study the Signal-to-Interference-and-Noise Ratio (SINR) outage probability for a typical transmission in pure Aloha. The results are used to compare the performance of non-slotted Aloha to the slotted one, which has almost exclusively been previously studied in the same context of mobile ad-hoc networks.

7.3. Random linear multihop relaying in a general field of interferers using spatial Aloha

In [9] we study, as a basic model, a stationary Poisson pattern of nodes on a line embedded in an independent planar Poisson field of interfering nodes. Assuming slotted Aloha and the signal-to-interference-and-noise ratio capture condition, with the usual power-law path loss model and Rayleigh fading, we explicitly evaluate several local and end-to-end performance characteristics related to the nearest-neighbor packet relaying on this line, and study their dependence on the model parameters (the density of relaying and interfering nodes, Aloha tuning and the external noise power). Our model can be applied in two cases: the first use is for vehicular ad-hoc networks, where vehicles are randomly located on a straight road. The second use is to study a typical route traced in a (general) planar ad-hoc network by some routing mechanism. The approach we have chosen allows us to quantify the non-efficiency of long-distance routing in pure ad-hoc networks and evaluate a possible remedy for it in the form of additional fixed relaying nodes, called road-side units in a vehicular network. It also allows us to consider a more general field of interfering nodes and study the impact of the clustering of its nodes the routing performance. As a special case of a field with more clustering than the Poison field, we consider a Poisson-line field of interfering nodes, in which all the nodes are randomly located on random straight lines. The comparison to our basic model reveals a paradox: clustering of interfering nodes decreases the outage probability of a single (typical) transmission on the route, but increases the mean end-to-end delay.

7.4. Studying the SINR process of the typical user in Poisson networks by using its factorial moment measures

Based on a stationary Poisson point process, a wireless network model with random propagation effects (shadowing and/or fading) is considered in [7] in order to examine the process formed by the signal-to-interference-plus-noise ratio (SINR) values experienced by a typical user with respect to all base stations in the down-link channel. This SINR process is completely characterized by deriving its factorial moment measures, which involve numerically tractable, explicit integral expressions. This novel framework naturally leads to expressions for the k-coverage probability, including the case of random SINR threshold values considered in multi-tier network models. While the k-coverage probabilities correspond to the marginal distributions of the order statistics of the SINR process, a more general relation is presented connecting the factorial moment measures of the SINR process to the joint densities of these order statistics. This gives a way for calculating exact values of the coverage probabilities arising in a general scenario of signal combination and interference cancellation between base stations. The presented framework consisting of mathematical representations of SINR characteristics with respect to the factorial moment measures holds for the whole domain of SINR and is amenable to considerable model extension.

7.5. Performance laws of large heterogeneous cellular networks

In [24] we propose a model for heterogeneous cellular networks assuming a space-time Poisson process of call arrivals, independently marked by data volumes, and served by different types of base stations (having different

transmission powers) represented by the superposition of independent Poisson processes on the plane. Each station applies a processor sharing policy to serve users arriving in its vicinity, modeled by the Voronoi cell perturbed by some random signal propagation effects (shadowing). Users' peak service rates depend on their signal-to-interference-and-noise ratios (SINR) with respect to the serving station. The mutual-dependence of the cells (due to the extra-cell interference) is captured via some system of cell-load equations impacting the spatial distribution of the SINR. We use this model to study in a semi-analytic way (involving only static simulations, with the temporal evolution handled by the queuing theoretic results) network performance metrics (cell loads, mean number of users) and the quality of service perceived by the users (mean throughput) served by different types of base stations. Our goal is to identify macroscopic laws regarding these performance metrics, involving averaging both over time and the network geometry. The reveled laws are validated against real field measurement in an operational network.

7.6. Wireless networks appear Poissonian due to strong shadowing

Geographic locations of cellular base stations sometimes can be well fitted with spatial homogeneous Poisson point processes. In [6] we make a complementary observation: In the presence of the log-normal shadowing of sufficiently high variance, the statistics of the propagation loss of a single user with respect to different network stations are invariant with respect to their geographic positioning, whether regular or not, for a wide class of empirically homogeneous networks. Even in perfectly hexagonal case they appear as though they were realized in a Poisson network model, i.e., form an inhomogeneous Poisson point process on the positive half-line with a power-law density characterized by the path-loss exponent. At the same time, the conditional distances to the corresponding base stations, given their observed propagation losses, become independent and log-normally distributed, which can be seen as a decoupling between the real and model geometry. The result applies also to Suzuki (Rayleigh-log-normal) propagation model. We use Kolmogorov-Smirnov test to empirically study the quality of the Poisson approximation and use it to build a linear-regression method for the statistical estimation of the value of the path-loss exponent.

7.7. What frequency bandwidth to run cellular network in a given country? - a downlink dimensioning problem

In [25] we propose an analytic approach to the frequency bandwidth dimensioning problem, faced by cellular network operators who deploy/upgrade their networks in various geographical regions (countries) with an inhomogeneous urbanization. We present a model allowing one to capture fundamental relations between users' quality of service parameters (mean downlink throughput), traffic demand, the density of base station deployment, and the available frequency bandwidth. These relations depend on the applied cellular technology (3G or 4G impacting user peak bit-rate) and on the path-loss characteristics observed in different (urban, sub-urban and rural) areas. We observe that if the distance between base stations is kept inversely proportional to the distance coefficient of the path-loss function, then the performance of the typical cells of these different areas is similar when serving the same (per-cell) traffic demand. In this case, the frequency bandwidth dimensioning problem can be solved uniformly across the country applying the mean cell approach proposed in [Blaszczyszyn et al. WiOpt2014]. We validate our approach by comparing the analytical results to measurements in operational networks in various geographical zones of different countries.

7.8. Optimal Geographic Caching In Cellular Networks

In [23] we consider the problem of an optimal geographic placement of content in wireless cellular networks modelled by Poisson point processes. Specifically, for the typical user requesting some particular content and whose popularity follows a given law (e.g. Zipf), we calculate the probability of finding the content cached in one of the base stations. Wireless coverage follows the usual signal-to-interference-and noise ratio (SINR) model, or some variants of it. We formulate and solve the problem of an optimal randomized content placement policy, to maximize the user's hit probability. The result dictates that it is not always optimal to follow the standard policy "cache the most popular content, everywhere". In fact, our numerical results regarding three different coverage scenarios, show that the optimal policy significantly increases the chances of hit under high-coverage regime, i.e., when the probabilities of coverage by more than just one station are high enough.

7.9. Spatial distribution of the SINR in Poisson cellular networks with sector antennas

In [5] we consider a model of cellular networks where the base station locations constitute a Poisson point process and each base station is equipped with three sectorial antennas is proposed. This model permits to study the spatial distribution of the SINR in the downlink. In particular, this distribution is shown to be insensitive to the distribution of antenna azimuths. Moreover, the effect of horizontal sectorisation is shown to be equivalent to that of shadowing. Assuming ideal vertical antenna pattern, an explicit expression of the Laplace transform of the inverse of SINR is given. The model is validated by comparing its results to measurements in an operational network. It is observed numerically that, in the case of dense urban regions where interference is preponderant, one may neglect the effect of the vertical sectorization when calculating the distribution of the SINR, which provides considerable tractability. Combined with queuing theory results, the SINR's distribution permits to express the user's quality of service as function of the traffic demand. This permits in particular to operators to predict the required investments to face the continual increase of traffic demand.

7.10. Theoretical expression of link performance in OFDM cellular networks with MIMO compared to simulation and measurements

The objective of [18] is to establish a theoretical expression of the link performance in the downlink of a multiple input multiple output (MIMO) cellular network and compare it to the real Long-Term Evolution (LTE) performance. In order to account for the interference, we prove that the worst additive noise process in the MIMO context is the white Gaussian one. Based on this theoretical result, we build an analytic expression of the link performance in LTE cellular networks with MIMO. We study also the minimum mean square error (MMSE) scheme currently implemented in the field, as well as its improvement MMSE-SIC (successive interference cancellation) known to achieve the MIMO capacity. Comparison to simulation results as well as to measurements in the field shows that the theoretical expression predicts well practical link performance of LTE cellular networks. This theoretical expression of link performance is the basis of a global analytic approach to the evaluation of the quality of service perceived by the users in the long run of their arrivals and departures.

7.11. Information Theory: Boolean model in the Shannon Regime

In a paper accepted for publication in the Journal of Applied Probability, F. Baccelli and V. Anantharam consider a family of Boolean models, indexed by integers $n \geq 1$. The n-th model features a Poisson point process in $\mathbb{R}^n$ of intensity $e^{n\rho_n}$ and balls of independent and identically distributed radii distributed like $\overline{X}_n\sqrt{n}$. Assume that $\rho_n \to \rho$ as $n \to \infty$, and that $\overline{X}_n$ satisfies a large deviations principle. It is shown that there then exist three deterministic thresholds: τ_d the degree threshold; τ_p the percolation probability threshold; and τ_v the volume fraction threshold, such that asymptotically as n tends to infinity, we have the following features. (i) For $\rho < \tau_d$, almost every point is isolated, namely its ball intersects no other ball; (ii) for $\tau_d < \rho < \tau_p$, the mean number of balls intersected by a typical ball converges to infinity and nevertheless there is no percolation; (iii) for $\tau_p < \rho < \tau_v$, the volume fraction is 0 and nevertheless percolation occurs; (iv) for $\tau_d < \rho < \tau_v$, the mean number of balls intersected by a typical ball converges to infinity and nevertheless the volume fraction is 0; (v) for $\rho > \tau_v$, the whole space covered. The analysis of this asymptotic regime is motivated by problems in information theory, but it could be of independent interest in stochastic geometry. The relations between these three thresholds and the Shannon–Poltyrev threshold are discussed.

7.12. Stochastic Geometry: Wireless Modeling

In an Infocom'15 paper, F. Baccelli and X. Zhang (Qualcomm) have introduced an analytically tractable stochastic geometry model for urban wireless networks, where the locations of the nodes and the shadowing are highly correlated and different path loss functions can be applied to line-of-sight (LOS) and non-line-of-sight (NLOS) links.

Using a distance-based LOS path loss model and a blockage (shadowing)-based NLOS path loss model, one can derive the distribution of the interference observed at a typical location and the joint distribution at different locations of the network. When applied to cellular networks, this model leads to tractable coverage probabilities (SINR distribution) expressions. This model captures important features of urban wireless networks, which were difficult to analyze using existing models.

This model was lately extended in a joint work by the same authors and Robert Heath (UT Austin) in a paper presented at IEEE Globecom'15 where it received the best paper award.

7.13. Information Theory: SIMO

In a paper to be published in IEEE Transactions of Information Theory, F. Baccelli, N. Lee and Robert Heath consider large random wireless networks where transmit-and-receive node pairs communicate within a certain range while sharing a common spectrum. By modeling the spatial locations of nodes as Poisson point processes, analytical expressions for the ergodic spectral efficiency of a typical node pair are derived as a function of the channel state information available at a receiver (CSIR) in terms of relevant system parameters: the density of communication links, the number of receive antennas, the path loss exponent, and the operating signal-to-noise ratio. One key finding is that when the receiver only exploits CSIR for the direct link, the sum spectral efficiency increases linearly with the density, provided the number of receive antennas increases as a certain super-linear function of the density. When each receiver exploits CSIR for a set of dominant interfering links in addition to that of the direct link, the sum spectral efficiency in creases linearly with both the density and the path loss exponent if the number of antennas is a linear function of the density. This observation demonstrates that having CSIR for dominant interfering links provides an order gain in the scaling law. It is also shown that this linear scaling holds for direct CSIR when incorporating the effect of the receive antenna correlation, provided that the rank of the spatial correlation matrix scales super-linearly with the density. These scaling laws are derived from integral representations of the distribution of the Signal to Interference and Noise Ratio, which are of independent interest and which in turn derived from stochastic geometry and more precisely from the theory of Shot Noise fields.

7.14. Theory of point processes

In a joint work with Mir-Omid Haji-Mirsadeghi, Sharif University, Department of Mathematics, F. Baccelli studied a class of non-measure preserving dynamical systems on counting measures called point-maps. This research introduced two objects associated with a point map f acting on a stationary point process Φ:

- The f-probabilities of Φ, which can be interpreted as the stationary regimes of the action of f on Φ. These probabilities are defined from the compactification of the action of the semigroup of point-map translations on the space of Palm probabilities. The f-probabilities of Φ are not always Palm distributions.
- The f-foliation of Φ, a partition of the support of Φ which is the discrete analogue of the stable manifold of f, i.e., the leaves of the foliation are the points of Φ with the same asymptotic fate for f. These leaves are not always stationary point processes. There always exists a point-map allowing one to navigate the leaves in a measure-preserving way.

Two papers on the matter available. The first one is under revision for Annals of Probability.

7.15. Cross-Technology Interference Mitigation in Body Area Networks: An Optimization Approach

In recent years, wearable devices and wireless body area networks have gained momentum as a means to monitor people's behavior and simplify their interaction with the surrounding environment, thus representing a key element of the body-to-body networking (BBN) paradigm. Within this paradigm, several transmission technologies, such as 802.11 and 802.15.4, that share the same unlicensed band (namely, the industrial, scientific, and medical band) coexist, dramatically increasing the level of interference and, in turn, negatively

affecting network performance. In this paper, we analyze the cross-technology interference (CTI) caused by the utilization of different transmission technologies that share the same radio spectrum. We formulate an optimization model that considers internal interference, as well as CTI to mitigate the overall level of interference within the system, explicitly taking into account node mobility. We further develop three heuristic approaches to efficiently solve the interference mitigation problem in large-scale network scenarios. Finally, we propose a protocol to compute the solution that minimizes CTI in a distributed fashion. Numerical results show that the proposed heuristics represent efficient and practical alternatives to the optimal solution for solving the CTI mitigation (CTIM) problem in large-scale BBN scenarios.

7.16. Body-to-Body Area Networks

The ongoing evolution of wireless technologies has fostered the development of innovative network paradigms like the Internet of Things (IoT). Wireless Body Area Networks, and more specifically Body-to-Body Area Networks (BBNs), are emerging solutions for the monitoring of people's behavior and their interaction with the surrounding environment. These networks represent a key building block of the IoT paradigm. In BBNs several transmission technologies like 802.11 and 802.15.4 that share the same unlicensed band (namely the industrial, scientific and medical (ISM) radio band) coexist, increasing dramatically the level of interference and, in turn, negatively affecting network's performance. In [14], we investigate the Cross-Technology Interference Mitigation (CTIM) problem caused by the utilization of different transmission technologies that share the same radio spectrum, from a centralized and distributed point of view, respectively.

7.17. Exact Worst-Case Delay in FIFO-Multiplexing Feed-Forward Networks

In [11], we compute the actual worst-case end-to-end delay for a flow in a feed-forward network of FIFO-multiplexing service curve nodes, where flows are shaped by piecewise-affine concave arrival curves, and service curves are piecewise affine and convex. We show that the worst-case delay problem can be formulated as a mixed integer-linear programming problem, whose size grows exponentially with the number of nodes involved. Furthermore, we present approximate solution schemes to find upper and lower delay bounds on the worst-case delay. Both only require to solve just one linear programming problem, and yield bounds which are generally more accurate than those found in the previous work, which are computed under more restrictive assumptions.

7.18. Fast symbolic computation of the worst-case delay in tandem networks and applications

Computing deterministic performance guarantees is a defining issue for systems with hard real-time constraints, like reactive embedded systems. In [10], we use burst-rate constrained arrivals and rate-latency servers to deduce tight worst-case delay bounds in tandem networks under arbitrary multiplexing. We present a constructive method for computing the exact worst-case delay, which we prove to be a linear function of the burstiness and latencies; our bounds are hence symbolic in these parameters. Our algorithm runs in quadratic time in the number of servers. We also present an application of our algorithm to the case of stochastic arrivals and server capacities. For a generalization of the exponentially bounded burstiness (EBB) model, we deduce a polynomial-time algorithm for stochastic delay bounds that strictly improve the state-of-the-art separated flow analysis (SFA) type bounds.

7.19. Ancillary Service to the Grid Using Intelligent Deferrable Loads

Renewable energy sources such as wind and solar power have a high degree of unpredictability and time-variation, which makes balancing demand and supply challenging. One possible way to address this challenge is to harness the inherent flexibility in demand of many types of loads. Introduced in [19] is a technique for decentralized control for automated demand response that can be used by grid operators as ancillary service for maintaining demand-supply balance. A randomized control architecture is proposed, motivated by the need for decentralized decision making, and the need to avoid synchronization that can lead to large and detrimental

spikes in demand. An aggregate model for a large number of loads is then developed by examining the mean field limit. A key innovation is a linear time-invariant (LTI) system approximation of the aggregate nonlinear model, with a scalar signal as the input and a measure of the aggregate demand as the output. This makes the approximation particularly convenient for control design at the grid level.

7.20. Spectral Decomposition of Demand-Side Flexibility for Reliable Ancillary Services in a Smart Grid

[22] describes a new way of thinking about demand-side resources to provide ancillary services to control the grid. It is shown that loads can be classified based on the frequency bandwidth of ancillary service that they can offer. If demand response from loads respects these frequency limitations, it is possible to obtain highly reliable ancillary service to the grid, while maintaining strict bounds on the quality of service (QoS) delivered by each load. It is argued that automated demand response is required for reliable control. Moreover, some intelligence is needed at demand response loads so that the aggregate will be reliable and controllable.

7.21. State Estimation for the Individual and the Population in Mean Field Control with Application to Demand Dispatch

[29] concerns state estimation problems in a mean field control setting. In a finite population model, the goal is to estimate the joint distribution of the population state and the state of a typical individual. The observation equations are a noisy measurement of the population. The general results are applied to demand dispatch for regulation of the power grid, based on randomized local control algorithms. In prior work by the authors it has been shown that local control can be carefully designed so that the aggregate of loads behaves as a controllable resource with accuracy matching or exceeding traditional sources of frequency regulation. The operational cost is nearly zero in many cases. The information exchange between grid and load is minimal, but it is assumed in the overall control architecture that the aggregate power consumption of loads is available to the grid operator. It is shown that the Kalman filter can be constructed to reduce these communication requirements, and to provide the grid operator with accurate estimates of the mean and variance of quality of service (QoS) for an individual load.

7.22. Perfect sampling of Jackson queueing networks

In [12], we consider open Jackson networks with losses with mixed finite and infinite queues and analyze the efficiency of sampling from their exact stationary distribution. We show that perfect sampling is possible, although the underlying Markov chain may have an infinite state space. The main idea is to use a Jackson network with infinite buffers (that has a product form stationary distribution) to bound the number of initial conditions to be considered in the coupling from the past scheme. We also provide bounds on the sampling time of this new perfect sampling algorithm for acyclic or hyper-stable networks. These bounds show that the new algorithm is considerably more efficient than existing perfect samplers even in the case where all queues are finite. We illustrate this efficiency through numerical experiments. We also extend our approach to variable service times and non-monotone networks such as queueing networks with negative customers.

7.23. Speeding up Glauber Dynamics for Random Generation of Independent Sets

The maximum independent set (MIS) problem is a well-studied combinatorial optimization problem that naturally arises in many applications, such as wireless communication, information theory and statistical mechanics. MIS problem is NP-hard, thus many results in the literature focus on fast generation of maximal independent sets of high cardinality. One possibility is to combine Gibbs sampling with coupling from the past arguments to detect convergence to the stationary regime. This results in a sampling procedure with time complexity that depends on the mixing time of the Glauber dynamics Markov chain. We propose in [37] an adaptive method for random event generation in the Glauber dynamics that considers only the events that are effective in the coupling from the past scheme, accelerating the convergence time of the Gibbs sampling algorithm.

7.24. Approximate optimality with bounded regret in dynamic matching models

In [28], we consider a dynamic matching model with random arrivals. In prior work, authors have proposed policies that are stabilizing, and also policies that are approximately finite-horizon optimal. This paper considers the infinite-horizon average-cost optimal control problem. A relaxation of the stochastic control problem is proposed, which is found to be a special case of an inventory model, as treated in the classical theory of Clark and Scarf. The optimal policy for the relaxation admits a closed-form expression. Based on the policy for this relaxation, a new matching policy is proposed. For a parameterized family of models in which the network load approaches capacity, this policy is shown to be approximately optimal, with bounded regret, even though the average cost grows without bound.

7.25. Perfect sampling for multiclass closed queueing networks

In [27] we present an exact sampling method for multiclass closed queuing networks. We consider networks for which stationary distribution does not necessarily have a product form. The proposed method uses a compact representation of sets of states, that is used to derive a bounding chain with significantly lower complexity of one-step transition in the coupling from the past scheme. The coupling time of this bounding chain can be larger than the coupling time of the exact chain, but it is finite in expectation. Numerical experiments show that coupling time is close to that of the exact chain. Moreover, the running time of the proposed algorithm outperforms the classical algorithm.

7.26. Fast and Memory Optimal Low-Rank Matrix Approximation

In this paper, we revisit the problem of constructing a near-optimal rank k approximation of a matrix $M \in [0,1]^{m\times n}$ under the streaming data model where the columns of M are revealed sequentially. We present SLA (Streaming Low-rank Approximation), an algorithm that is asymptotically accurate, when $ks_{k+1}(M) = o(\sqrt{mn})$ where $s_{k+1}(M)$ is the $(k+1)$-th largest singular value of M. This means that its average mean-square error converges to 0 as m and n grow large (i.e., $\|\widehat{M}^{(k)} - M^{(k)}\|_F^2 = o(mn)$ with high probability, where $\widehat{M}^{(k)}$ and $M^{(k)}$ denote the output of SLA and the optimal rank k approximation of M, respectively). Our algorithm makes one pass on the data if the columns of M are revealed in a random order, and two passes if the columns of M arrive in an arbitrary order. To reduce its memory footprint and complexity, SLA uses random sparsification, and samples each entry of M with a small probability δ. In turn, SLA is memory optimal as its required memory space scales as $k(m+n)$, the dimension of its output. Furthermore, SLA is computationally efficient as it runs in $O(\delta kmn)$ time (a constant number of operations is made for each observed entry of M), which can be as small as $O(k\log(m)^4 n)$ for an appropriate choice of δ and if $n \geq m$.

7.27. Combinatorial Bandits Revisited

[42] investigates stochastic and adversarial combinatorial multi-armed bandit problems. In the stochastic setting under semi-bandit feedback, we derive a problem-specific regret lower bound, and discuss its scaling with the dimension of the decision space. We propose ESCB, an algorithm that efficiently exploits the structure of the problem and provide a finite-time analysis of its regret. ESCB has better performance guarantees than existing algorithms, and significantly outperforms these algorithms in practice. In the adversarial setting under bandit feedback, we propose COMBEXP, an algorithm with the same regret scaling as state-of-the-art algorithms, but with lower computational complexity for some combinatorial problems.

7.28. Non-backtracking spectrum of random graphs: community detection and non-regular Ramanujan graphs

A non-backtracking walk on a graph is a directed path such that no edge is the inverse of its preceding edge. The non-backtracking matrix of a graph is indexed by its directed edges and can be used to count non-backtracking walks of a given length. It has been used recently in the context of community detection and has appeared previously in connection with the Ihara zeta function and in some generalizations of Ramanujan graphs. In [26], we study the largest eigenvalues of the non-backtracking matrix of the Erdos-Renyi random graph and of the Stochastic Block Model in the regime where the number of edges is proportional to the number of vertices. Our results confirm the "spectral redemption" conjecture that community detection can be made on the basis of the leading eigenvectors above the feasibility threshold.

7.29. Designing Adaptive Replication Schemes in Distributed Content Delivery Networks

In [32], we address the problem of content replication in large distributed content delivery networks, composed of a data center assisted by many small servers with limited capabilities and located at the edge of the network. The objective is to optimize the placement of contents on the servers to offload as much as possible the data center. We model the system constituted by the small servers as a loss network, each loss corresponding to a request to the data center. Based on large system / storage behavior, we obtain an asymptotic formula for the optimal replication of contents and propose adaptive schemes related to those encountered in cache networks but reacting here to loss events, and faster algorithms generating virtual events at higher rate while keeping the same target replication. We show through simulations that our adaptive schemes outperform significantly standard replication strategies both in terms of loss rates and adaptation speed.

7.30. Spectral Detection in the Censored Block Model

In [36], we consider the problem of partially recovering hidden binary variables from the observation of (few) censored edge weights, a problem with applications in community detection, correlation clustering and synchronization. We describe two spectral algorithms for this task based on the non-backtracking and the Bethe Hessian operators. These algorithms are shown to be asymptotically optimal for the partial recovery problem, in that they detect the hidden assignment as soon as it is information theoretically possible to do so.

7.31. A spectral method for community detection in moderately-sparse degree-corrected stochastic block models

In the ordinary stochastic block model, all degrees in a cluster have the same expected degree. The Degree-Corrected Stochastic Block Models (DC-SBM) is a generalization of the former where the expected degrees of individual nodes follow a prescribed degree-sequence. We consider community detection in the DC-SBM in a paper currently in preparation [43]. We perform spectral clustering on a suitably normalized adjacency matrix. This leads to consistent recovery of the block-membership of all but a vanishing fraction of nodes, in the regime where the lowest degree is of order $\log(n)$ or higher. The main contributions of this paper are (i) the fact that recovery succeeds for very heterogeneous degree-distributions and (ii) a clean analysis for the DC-SBM, which is a messy model.

7.32. An Impossibility Result for Reconstruction in a Degree-Corrected Planted-Partition Model

In a paper currently in preparation [44], we consider a degree-corrected planted-partition model: a random graph on n nodes with two equal-sized clusters. The model parameters are two constants $a, b > 0$ and an i.i.d. sequence $(\phi_i)_{i=1}^n$, with finite second moment Φ^2. Vertices i and j are joined by an edge with probability $\frac{\phi_i\phi_j}{n}a$ whenever they are in the same class and with probability $\frac{\phi_i\phi_j}{n}b$ otherwise. We prove that the underlying community structure cannot be accurately recovered from observations of the graph when $(a-b)^2\Phi^2 \leq 2(a+b)$.

7.33. Universality in polytope phase transitions and message passing algorithms

In [], we consider a class of nonlinear mappings $\mathsf{F}_{A,N}$ in $\mathbb{R}^N$ indexed by symmetric random matrices $A \in \mathbb{R}^{N\times N}$ with independent entries. Within spin glass theory, special cases of these mappings correspond to iterating the TAP equations and were studied by Bolthausen [Comm. Math. Phys. 325 (2014) 333-366]. Within information theory, they are known as "approximate message passing" algorithms. We study the high-dimensional (large N) behavior of the iterates of F for polynomial functions F, and prove that it is universal; that is, it depends only on the first two moments of the entries of A, under a sub-Gaussian tail condition. As an application, we prove the universality of a certain phase transition arising in polytope geometry and compressed sensing. This solves, for a broad class of random projections, a conjecture by David Donoho and Jared Tanner.

7.34. Contagions in Random Networks with Overlapping Communities

In [13], we consider a threshold epidemic model on a clustered random graph with overlapping communities. In other words, our epidemic model is such that an individual becomes infected as soon as the proportion of her infected neighbors exceeds the threshold q of the epidemic. In our random graph model, each individual can belong to several communities. The distributions for the community sizes and the number of communities an individual belongs to are arbitrary. We consider the case where the epidemic starts from a single individual, and we prove a phase transition (when the parameter q of the model varies) for the appearance of a cascade, i.e. when the epidemic can be propagated to an infinite part of the population. More precisely, we show that our epidemic is entirely described by a multi-type (and alternating) branching process, and then we apply Sevastyanov's theorem about the phase transition of multi-type Galton-Watson branching processes. In addition, we compute the entries of the matrix whose largest eigenvalue gives the phase transition.

7.35. The Diameter of Weighted Random Graphs.

In [3], we study the impact of random exponential edge weights on the distances in a random graph and, in particular, on its diameter. Our main result consists of a precise asymptotic expression for the maximal weight of the shortest weight paths between all vertices (the weighted diameter) of sparse random graphs, when the edge weights are i.i.d. exponential random variables.

EVA Team

7. New Results

7.1. Wireless Sensor Networks

7.1.1. Time slot and channel assignment in multichannel Wireless Sensor Networks

Participants: Pascale Minet, Ridha Soua, Erwan Livolant.

Wireless sensor networks (WSNs) play a major role in industrial environments for data gathering (convergecast). Among the industrial requirements, we can name a few like 1) determinism and bounded convergecast latencies, 2) throughput and 3) robustness against interferences. The classical IEEE 802.15.4 that has been designed for low power lossy networks (LLNs) partially meets these requirements. That is why the IEEE 802.15.4e MAC amendment has been proposed recently. This amendment combines a slotted medium access with a channel hopping (i.e. Time Slotted Channel Hopping TSCH). The MAC layer orchestrates the medium accesses of nodes according to a given schedule. Nevertheless, this amendment does not specify how this schedule is computed. We propose a distributed joint time slot and channel assignment, called $Wave$ for data gathering in LLNs. This schedule targets minimized data convergecast delays by reducing the number of slots assigned to nodes. Moreover, $Wave$ ensures the absence of conflicting transmissions in the schedule provided. In such a schedule, a node is awake only during its slots and the slots of its children in the convergecast routing graph. Thus, energy efficiency is ensured. We describe in details the functioning of $Wave$, highlighting its features (e.g. support of heterogeneous traffic, support of a sink equipped with multiple interfaces) and properties in terms of worst case delays and buffer size. We discuss its features with regard to a centralized scheduling algorithm like $TMCP$ and a distributed one like $DeTAS$. Simulation results show the good performance of $Wave$ compared to $TMCP$. Since in an industrial environment, several routing graphs can coexist, we study how $Wave$ supports this coexistence.

7.1.2. Centralized Scheduling in TSCH-based Wireless Sensor Networks

Participants: Erwan Livolant, Pascale Minet, Thomas Watteyne.

Scheduling in an IEEE802.15.4e TSCH(Time Slotted Channel Hopping 6TiSCH) low-power wireless network can be done in a centralized or distributed way. When using centralized scheduling, a scheduler installs a communication schedule into the network. This can be done in a standards-based way using CoAP. In this study, we compute the number of packets and the latency this takes, on real-world examples. The result is that the cost is very high using today's standards, much higher than when using an ad-hoc solution such as OCARI. We conclude by making recommendations to drastically reduce the number of messages and improve the efficiency of the standardized approach.

7.1.3. Distributed and Optimized Deployment of WSNs

Participants: Ines Khoufi, Pascale Minet.

This is a joint work with Telecom SudParis: Anis Laouiti.
We are witnessing the deployment of many wireless sensor networks in various application domains such as pollution detection in the environment, intruder detection at home, preventive maintenance in industrial process, monitoring of temporary industrial worksites, damage assessment after a disaster.... Wireless sensor networks are deployed to monitor physical phenomena. The accuracy of the information collected depends on the position of sensor nodes. These positions must meet the application requirements in terms of coverage and connectivity, which therefore requires the use of deployment algorithms. We distinguish two cases: firstly when the nodes are autonomous, and secondly when they are static and the deployment is assisted by mobile robots. In both cases, this deployment must not only meet the application's coverage and connectivity requirements, but must also minimize the number of sensors needed while satisfying various constraints (e.g. obstacles, energy, fault-tolerant connectivity). We distinguished two cases: autonomous and mobile wireless sensor nodes on the one hand, and static wireless sensor nodes on the other hand.

We propose a distributed and optimized deployment of mobile and autonomous sensor nodes to ensure full coverage of the 2D-area considered, as well as network connectivity. With the full coverage of the area, any event occurring in this area is detected by at least one sensor node. In addition, the connectivity ensures that this event is reported to the sink in charge of analyzing the data gathered from the sensors and acting according to these data. This distributed algorithm, called OA-DVFA, can run in an unknown area with obstacles dicovered dynamically. We distinguish two types of obstacles: the transparent ones like ponds in outdoor environment, or tables in an indoor site that only prevent the location of sensor nodes inside them; whereas the opaque obstacles like walls or trees prevent the sensing by causing the existence of hidden zones behind them: such zones may remain uncovered. Opaque obstacles are much more complex to handle than transparent ones and require the deployment of additional sensors to eliminate coverage holes. OA-DVFA is based on virtual forces to obtain a fast spreading of sensor nodes and uses a virtual grid to stop node oscillations and save energy by making sleep redundant nodes. It automatically detects when the maximum area coverage is reached.
We also considered 3D volumes and proposed an algorithm, called 3D-DVFA, also based on virtual forces, to ensure full coverage of 3D volumes and ensure network connectivity. This is a joint work with Nadya Boufares from ENSI, Tunisia. Since applications of such 3D deployments may be limited, we focus on 3D surface covering, where the objective is to deploy wireless sensor nodes on a 3D-surface (e.g. a mountain) to ensure full area coverage and network connectivity. To reach this goal we propose 3D-DVFA-SC, a distributed deployment algorithm based on virtual forces strategy to move sensor nodes.

7.1.4. WSN deployment assisted by mobile robots

Participants: Ines Khoufi, Pascale Minet.

This is a joint work with Telecom SudParis: Anis Laouiti.
Autonomous deployment may be expensive when the number of mobile sensor nodes is very high. In this case, an assisted deployment may be necessary: the nodes' positions being pre-computed and given to mobile robots that place a static sensor at each position. In order to reduce both the energy consumed by the robots, their exposure time to a hostile environment, as well as the time at which the wireless network becomes operational, the optimal tour of robots is this minimizing the delay. This delay must take into account not only the time needed by the robot to travel the tour distance but also the time spent in the rotations performed by the robot each time it changes its direction. This problem is called the Multiple Robot Deploying Sensor nodes problem, in short MRDS. We first show how this problem differs from the well-known traveling salesman problem. We adopt two approaches to optimize the deployment duration. The first one is based on game theory to optimize the length of the tours of two robots (TRDS), and the second is based on a multi-objective optimization, for multiple robots (MRDS). The objectives to be met are: optimizing the duration of the longest tour, balancing the durations of the robot tours and minimizing the number of robots used, while bypassing obstacles.

The TRDS problem is modeled as a non-cooperative game with two players representing the mobile robots, these robots compete for the selection of the sensor nodes to deploy. Each robots tends to maximize its utility function.

We then propose an integer linear program formulation of the MRDS problem. We propose various algorithms relevant to iterative improvement by exchanging tour edges, genetic approach and hybridization. The solutions provided by these algorithms are compared and their closeness to the optimal is evaluated in various configurations.

7.1.5. Sinks Deployment and Packet Scheduling for Wireless Sensor Networks

Participants: Nadjib Achir, Paul Muhlethaler.

The objective of this work is to propose an optimal deployment and distributed packet scheduling of multi-sink Wireless Sensors networks (WNSs). We start by computing the optimal deployment of sinks for a given maximum number of hops between nodes and sinks. We also propose an optimal distributed packet scheduling in order to estimate the minimum energy consumption. We consider the energy consumed due to reporting, forwarding and overhearing. In contrast to reporting and forwarding, the energy used in overhearing is difficult to estimate because it is dependent on the packet scheduling. In this case, we determine the lower-bound of overhearing, based on an optimal distributed packet scheduling formulation. We also propose another estimation of the lower-bound in order to simulate non interfering parallel transmissions which is more tractable in large networks. We note that overhearing largely predominates in energy consumption. A large part of the optimizations and computations carried out in this work are obtained using ILP (Integer Linear Programming) formalization.

7.1.6. Security in wireless sensor networks

Participants: Selma Boumerdassi, Paul Muhlethaler.

Sensor networks are often used to collect data from the environment where they are located. These data can then be transmitted regularly to a special node called a *sink*, which can be fixed or mobile. For critical data (like military or medical data), it is important that sinks and simple sensors can mutually authenticate so as to avoid data to be collected and/or accessed by fake nodes. For some applications, the collection frequency can be very high. As a result, the authentication mechanism used between a node and a sink must be fast and efficient both in terms of calculation time and energy consumption. This is especially important for nodes which computing capabilities and battery lifetime are very low. Moreover, an extra effort has been done to develop alternative solutions to secure, authenticate, and ensure the confidentiality of sensors, and the distribution of keys in the sensor network. Specific researches have also been conducted for large-scale sensors. At present, we work on an exchange protocol between sensors and sinks based on low-cost shifts and xor operations.

7.1.7. Massive MIMO Cooperative Communications for Wireless Sensor Networks

Participants: Nadjib Achir, Paul Muhlethaler.

This work is a collaboration with Mérouane Debbah (Supelec, France).

The objective of this work is to propose a framework for massive MIMO cooperative communications for Wireless Sensor Networks. Our main objective is to analyze the performances of the deployment of a large number of sensors. This deployment should cope with a high demand for real time monitoring and should also take into account energy consumption. We have assumed a communication protocol with two phases: an initial training period followed by a second transmit period. The first period allows the sensors to estimate the channel state and the objective of the second period is to transmit the data sensed. We start analyzing the impact of the time devoted to each period. We study the throughput obtained with respect to the number of sensors when there is one sink. We also compute the optimal number of sinks with respect to the energy spent for different values of sensors. This work is a first step to establish a complete framework to study energy efficient Wireless Sensor Networks where the sensors collaborate to send information to a sink. Currently, we are exploring the multi-hop case.

7.2. Cognitive Radio Networks

7.2.1. Multichannel time slot assignment in Cognitive Radio Sensor Networks

Participants: Ons Mabrouk, Pascale Minet.

This is a joint work with Hanen Idoudi and Leila Saidane from ENSI, Tunisia.
The unlicensed spectrum bands become overcrowded causing an increased level of interference for current wireless sensor nodes. Cognitive Radio Sensor Networks (CRSNs) overcome this problem by allowing sensor nodes to access opportunistically the underutilized licensed spectrum bands. The sink assigns the spectrum holes to the secondary users (SUs). Therefore, it must rely on reliable information about the spectrum holes to protect the primary users (PUs). We focused on the MultiChannel Time Slot Assignment problem in CRSN and tackled this problem: first at the level of a cluster (i.e. Intra-cluster multichannel scheduling), second at the level of several clusters coexisting in the same CRSN (i.e. inter-cluster multichannel scheduling).

In 2013, we proposed an Opportunistic centralized TIme slot assignment in COgnitive Radio sensor networks (OTICOR) for the Intra-cluster multichannel scheduling. OTICOR differs from the existing schemes in its ability to allow non-interfering cognitive sensors to access the same channel and time slot pair. OTICOR takes advantages of spatial reuse, multichannel communication and multiple radio interfaces of the sink. We proved through simulations that a smaller schedule length improves the throughput. Applying OTICOR, we showed that, even in the presence of several PUs, the average throughput granted to SUs remains important. We also showed how to get the best performances of OTICOR when the channel occupancy by PUs is known.

In 2014, we extended this Intra-cluster multichannel scheduling algorithm by proposing two ways for the sink to determine the available channels and alert the SUs if an unexpected activity of PU occurs. Our objective is to design an algorithm able to detect the unexpected presence of PUs in the multi-hop network while maximizing the throughput. If the estimation of PU presence is accurate, a channel sensing at the beginning of the slotframe is sufficient. This algorithm, called Frame-ICMS (Frame Intra-Cluster Multichannel Scheduling), takes advantage of the slots dedicated to the control period by allowing noninterfering cognitive sensors to access the control/data channel and time slot pair. We showed through simulations that using the control period also for data transmission minimizes the schedule length and maximizes the throughput. However, if the estimation of PU presence is not accurate, channel sensing should be done before each slot. We proposed the Slot-ICMS algorithm.

In 2015, we focused on inter-cluster multichannel scheduling algorithm. We considered the coexistence of different clusters in a same CSRN, each cluster having an intra-cluster multichannel scheduling algorithm. Our goal is to obtain a better scalability without loosing the properties provided by OTICOR:

- collision-free schedule,
- minimized data gathering delays,
- sleeping periods per node to save node's energy.

However, the co-existence of several clusters in the same environment may lead to conflicts in the allocation of time slots and channels between these clusters. To avoid inter-cluster collisions, we donot require that different clusters use different channels, we assign distinct channels only to nodes having one-hop neighbors out of their cluster. Once the problem of inter-cluster collision is avoided, each cluster head schedules the transmissions of its members independently. This whole solution exhibits good performances as shown by the simulation results.

7.3. Learning for an efficient and dynamic management of network resources and services

7.3.1. Learning in networks

Participants: Dana Marinca, Nesrine Ben Hassine, Pascale Minet, Selma Boumerdassi.

This work is a joint work with Dominique Barth (University of Versailles-Saint-Quentin). To guarantee an efficient and dynamic management of network resources and services we intend to use a powerful mathematical tool: prediction and learning from prediction. Prediction will be concerned with guessing the evolution of network or network components state, based on knowledge about the past elements and/or other available information. Basically, the prediction problem could be formulated as follows: a forecaster observes the values of one or several metrics giving indications about the network state (generally speaking the network represents the environment). At each time t, before the environment reveals the new metric values, the forecaster predicts the new values based on previous observations. Contrary to classical methods where the environment evolution is characterized by stochastic process, we suppose that the environment evolution follows an unspecified mechanism, which could be deterministic, stochastic, or even adaptive to a given behavior. The prediction process should adapt to unpredictable network state changes due to its non-stationary nature. To properly address the adaptivity challenge, a special type of forecasters is used: the experts. These experts analyse the previous environment values, apply their own computation and make their own prediction. The experts predictions are given to the forecaster before the next environment values are revealed.

The forecaster can then make its own prediction depending on the experts' "advice". The risk of a prediction may be defined as the value of a loss function measuring the discrepancy between the predicted value and the real environment value. The principal notion to optimize the behavior of the forecasters is the regret, seen as a difference between the forecaster's accumulated loss and that of each expert. To optimize the prediction process means to construct a forecasting strategy that guarantees a small loss with respect to defined experts. Adaptability of the forecaster is reflected in the manner in which it is able to follow the better expert according to the context.

Our purpose is to apply on-line learning strategies to:

- Wireless Sensor Networks (WSNs) to predict the quality of a wireless link in a WSN, based on the LQI metric for instance and take advantage of wireless links with the best possible quality to improve the packet delivery rate. We model this problem as a forecaster prediction game based on the advice of several experts. The forecaster learns on-line how to adjust its prediction to better fit the environment metric values. A forecaster estimates the LQI value using the advice of experts.
- Content Delivery Networks (CDNs) to predict the number of solicitations of video contents to cache the contents with the highest popularity.
- Data centers require a huge amount of energy. As an example, in 2014, the electric consumption of all date centers will be larger than 42 TWh, and after 2020 the CO2 production will be larger then 1.27 GTons, ie. more than the aeronautic industry (GeSI SMARTer 2020 report). These "frightening" figures led the research community to work on the management of energy consumption. Several tracks have been explored, among which the optimization of computation and load balancing of servers. At present, we work on tools dedicated to traffic prediction, thus allowing a better management of servers. Our work consists in modeling the traffic specific to data centers and apply different statistical prediction methods.

7.3.2. Tools for learning and prediction

Participants: Dana Marinca, Nesrine Ben Hassine, Pascale Minet.

In 2015, Nesrine Ben Hassine developped an extraction tool to provide real traces from YouTube. these real traces are used as a learning sample by the different prediction algorithms used.

Nesrine Ben Hassine and Dana Marinca extended their simulation tool developed in Python to integrate:

- various prediction strategies SES (Single Explonential Smoothing), DES (Double Exponential Smoothing), Basic and enhanced basic, strategies based on averages (e.g. Average on a Moving Window), regressions (e.g. polynomial or Savitzky Golay), as well as prediction strategies adapting dynamically their parameters according to the loss obtained.
- various loss functions (e.g. absolute value, square). The prediction accuracy is evaluated by a loss function as the discrepancy between the prediction value and the real number obtained.
- different forecaster strategies: Best expert, Exponential Weighted Average, K Best-Experts, etc.

With these tools, we can now tune parameters of prediction strategies and evaluate them.

7.3.3. Popularity prediction in CDNs

Participants: Dana Marinca, Nesrine Ben Hassine, Pascale Minet.

To predict the popularity of video contents, expressed as the number of solicitations, we compared three prediction strategies: Single Exponential Smooting (SES), Double Exponential Smoothing (DES) and Basic. The best tuning of each strategy is determined, depending on the considered phase of the solicitation curve. For DES, values of the smoothing factor close to 1 probide the best results. We study the behavior of each strategy within a phase and around a phase change, where a phase is defined as an interval of time during which a measured metric remains relatively stable.

Basic expert makes large errors at the phase change, but it quickly corrects its prediction and it is the expert having the closest prediction to the real value within a phase. DES expert provides also good quality predictions within a phase. Since DES and Basic experts outperform the SES expert, we recommend the use of on the one hand, the best DES expert per phase within a phase and on the other hand, the Basic expert to automatically detect phase changes, because of its better reactivity. This self-learning and prediction method can be applied to optimize resources allocation in service oriented architectures and self-adaptive networks, more precisely for cache management in CDNs.

7.3.4. Automatic phase detection in popularity evolution of video contents

Participants: Dana Marinca, Nesrine Ben Hassine, Pascale Minet.

In Content Delivery Networks (CDNs) where experts predict the number of solicitations of video contents, simulations based on real YouTube traces show that the accuracy of prediction is improved by splitting the video content profile in contiguous phases. A phase is an interval of time during which a measured metric remains relatively stable. The best expert per phase outperforms the best expert on the whole video content profile. Different prediction methods are compared and also different phase change-points detection methods are evaluated:

- the R tool using Bayesian inference,
- the Basic expert (an important loss may indicate a phase change),
- a fixed time interval (e.g. each week).

The goal is to identify the method (or method parameters) minimizing the cumulated discrepancy compared to real solicitations of video contents. The use of this machine learning method allows the Content Delivery Network to self-adapt to users solicitations by caching the most popular contents near the end users. More generally, such method can be applied to decide which contents should be replicated to improve the performance of audio and video applications and maximize the satisfaction degree of users.

7.4. VANETs

7.4.1. Protocols for VANETs

Participants: Nadjib Achir, Younes Bouchaala, Mohamed Elhadad Or Hadded, Paul Muhlethaler, Oyunchimeg Shagdar.

7.4.1.1. Synthetic study of TDMA protocols for VANETs

Recently several Time Division Multiple Access (TDMA)-based medium access control protocols have been proposed for VANETs in an attempt to ensure that all the vehicles have enough time to send safety messages without collisions and to reduce the end-to-end delay and the packet loss ratio. In this paper, we identify the reasons for using the collision-free medium access control paradigm in VANETs. We then present a novel topology-based classification and we provide an overview of TDMA-based MAC protocols that have been proposed for VANETs. We focus on the characteristics of these protocols, as well as on their benefits and limitations. Finally, we give a qualitative comparison, and we discuss some open issues that need to be tackled in future studies in order to improve the performance of TDMA-based MAC protocols for vehicle to vehicle (V2V) communications.

7.4.1.2. A stable clustering protocol for VANETs

VANETs have a highly dynamic and portioned network topology due to the constant and rapid movement of vehicles. Currently, clustering algorithms are widely used as the control schemes to make VANET topology less dynamic for Medium Access Control (MAC), routing and security protocols. An efficient clustering algorithm must take into account all the necessary information related to node mobility. In this paper, we propose an Adaptive Weighted Clustering Protocol (AWCP), specially designed for vehicular networks, which takes the highway ID, direction of vehicles, position, speed and the number of neighboring vehicles into account in order to enhance the stability of the network topology. However, the multiple control parameters of our AWCP, make parameter tuning a nontrivial problem. In order to optimize the protocol, we define a

multi-objective problem whose inputs are the AWCP's parameters and whose objectives are: providing stable cluster structures, maximizing data delivery rate, and reducing the clustering overhead. We address this multi-objective problem with the Nondominated Sorted Genetic Algorithm version 2 (NSGA-II). We evaluate and compare its performance with other multi-objective optimization techniques: Multi-objective Particle Swarm Optimization (MOPSO) and Multi-objective Differential Evolution (MODE). The experiments reveal that NSGA-II improves the results of MOPSO and MODE in terms of spacing, spread, ratio of non-dominated solutions, and inverse generational distance, which are the performance metrics used for comparison.

7.4.1.3. Using Road IDs to Enhance Clustering in Vehicular Ad hoc Networks

Vehicular ad hoc networks (VANETs) where vehicles act as mobile nodes is an instance of Mobile Ad hoc NETworks (MANETs), which are essentially developed for intelligent transportation systems. A challenging problem when designing communication protocols in VANETs is coping with high vehicle mobility, which causes frequent changes in the network topology and leads to frequent breaks in communication. The clustering technique is being developed to reduce the impact of mobility between neighboring vehicles. In this paper, we propose an Adaptive Weighted Cluster Protocol for VANETs, which is a road map dependent and uses road IDs and movement direction in order to make the clusters structure as stable as possible. The experimental results reveal that AWCP outperforms four other most commonly used clustering protocols in terms of control packet overhead, the packet delivery ratio, and the average cluster lifetime, which are the most usual metrics used for comparing performance.

7.4.2. Models for VANETs

Participants: Nadjib Achir, Younes Bouchaala, Guy Fayolle, Paul Muhlethaler, Oyunchimeg Shagdar.

7.4.2.1. Model of IEEE 802.11 broadcast scheme with infinite queue

We have analyzed the so-called back-off technique of the IEEE 802.11 protocol in broadcast mode with waiting queues. In contrast to existing models, packets arriving when a station (or node) is in back-off state are not discarded, but are stored in a buffer of infinite capacity. As in previous studies, the key point of our analysis hinges on the assumption that the time on the channel is viewed as a random succession of transmission slots (whose duration corresponds to the length of a packet) and mini-slots during which the back-off of the station is decremented. These events occur independently, with given probabilities. The state of a node is represented by a two-dimensional Markov chain in discrete-time, formed by the back-off counter and the number of packets at the station. Two models are proposed both of which are shown to cope reasonably well with the physical principles of the protocol. The stabillity (ergodicity) conditions are obtained and interpreted in terms of maximum throughput. Several approximations related to these models are also discussed.

7.4.2.2. Model and optimization of CSMA

We have studied the maximum throughput of CSMA in scenarios with spatial reuse. The nodes of our network will be a Poisson Point Process (PPP) of a one or two dimensional space. The one dimensional well fits VANETs. To model the effect of Carrier Sense Multiple Access (CSMA), we give random marks to our nodes and to elect transmitting nodes in the PPP we choose the nodes with the smallest marks in their neighborhood, this is the Matern hardcore selection process. To describe the signal propagation, we use a signal with power-law decay and we add a random Rayleigh fading. To decide whether or not a transmission is successful, we adopt the Signal-over-Interference Ratio (SIR) model in which a packet is correctly received if its transmission power divided by the interference power is above a capture threshold. We assume that each node in our PPP has a random receiver at a typical distance. We choose the average distance to its closest neighbor. We also assume that all the network nodes always have a pending packet. With all these assumptions, we analytically study the density of throughput of successful transmission and we show that it can be optimized with the carrier-sense threshold.

7.4.2.3. Performance analysis of IEEE 802.11 broadcast schemes

We have analyzed different broadcast strategies in IEEE 802.11p Vehicular Ad-hoc NETworks (VANETs). The first strategy is the default IEEE 802.11p strategy. Using a model derived from the Bianchi model, we provide the network performance in terms of throughput and success rate. The second strategy is to use an

acknowledgment technique similar to the acknowledgment with point-to-point traffic. A node will send its broadcast packet as in the default case,but it requires an acknowledgment from a neighbor node. This node may be a random neighbor or may be selected according to precise rules. We analyze this second strategy in terms of throughput and success rate. Somewhat surprisingly, we show that this second strategy improves the delivery ratio of the transmitted packets but reduces the overall throughput. This means that if the CAM messages (Cooperative Awareness Messages) are broadcasted, the total number of packets actually delivered will be greater with the default strategy than with the improved strategy. We propose a third strategy which consists in using the default strategy for normal packets, but we add random redundant transmissions to ensure greater reliability for very important packets. We show that with this simple technique, not only do we obtain suitable reliability, but we also achieve larger global throughput than with the acknowledgment-oriented technique. We have also computed network performance in terms of throughput and success rate with respect to the network parameters and to analyze their impact on performances.

7.5. Models for wireless networks

7.5.1. Interference and SINR coverage in spatial non-slotted Aloha networks

Participants: Bartek Blaszczyszyn, Paul Muhlethaler.

We propose two analytically tractable stochastic-geometric models of interference in adhoc networks using pure (non-slotted) Aloha as the medium access. In contrast the slotted model, the interference in pure Aloha may vary during the transmission of a tagged packet. We develop closed form expressions for the Laplace transform of the empirical average of the interference experienced during the transmission of a typical packet. Both models assume a power-law path-loss function with arbitrarily distributed fading and feature configurations of transmitters randomly located in the Euclidean plane according to a Poisson point process. Depending on the model, these configurations vary over time or are static. We apply our analysis of the interference to study the Signal-to- Interference-and-Noise Ratio (SINR) outage probability for a typical transmission in pure Aloha. The results are used to compare the performance of non-slotted Aloha to the slotted one, which has almost exclusively been previously studied in context of wired ad-hoc networks.

7.5.2. Random linear multihop relaying in a general field of interferers using spatial Aloha

Participants: Bartek Blaszczyszyn, Paul Muhlethaler.

We study a stationary Poisson pattern of nodes on a line embedded in an independent planar Poisson field of interfering nodes. Assuming slotted Aloha and the signal-to-interference-and-noise ratio capture condition, with the usual power-law path loss model and Rayleigh fading, we explicitly evaluate several local and end-to-end performance characteristics related to the nearest-neighbor packet relaying on this line, and study their dependence on the model parameters (the density of relaying and interfering nodes, Aloha tuning and the external noise power). Our model can be applied in two cases: the first use is for vehicular ad-hoc networks, where vehicles are randomly located on a straight road. The second use is to study a “typical” route traced in a (general) planar ad-hoc network by some routing mechanism. The approach we have chosen allows us to quantify the non-efficiency of long-distance routing in “pure ad-hoc” networks and evaluate a possible remedy for it in the form of additional “fixed” relaying nodes, called road-side units in a vehicular network. It also allows us to consider a more general field of interfering nodes and study the impact of the clustering of its nodes on the routing performance. As a special case of a field with more clustering than the Poison field, we consider a Poisson-line field of interfering nodes, in which all the nodes are randomly located on random straight lines. In this case our analysis rigorously (in the sense of Palm theory) corresponds to the typical route of this network. The comparison to our basic model reveals a paradox: clustering of interfering nodes decreases the outage probability of a single (typical) transmission on the route, but increases the mean end-to-end delay

GALLIUM Project-Team

7. New Results

7.1. Formal verification of compilers and static analyzers

7.1.1. The CompCert formally-verified compiler

Participants: Xavier Leroy, Jacques-Henri Jourdan, François Pottier, Bernhard Schommer [AbsInt GmbH].

In the context of our work on compiler verification (see section 3.3.1), since 2005 we have been developing and formally verifying a moderately-optimizing compiler for a large subset of the C programming language, generating assembly code for the PowerPC, ARM, and x86 architectures [6]. This compiler comprises a back-end, which translates the Cminor intermediate language to PowerPC assembly and is reusable for source languages other than C [5], and a front-end, which translates the CompCert C subset of C to Cminor. The compiler is mostly written within the specification language of the Coq proof assistant, from which Coq's extraction facility generates executable OCaml code. The compiler comes with a 50000-line, machine-checked Coq proof of semantic preservation establishing that the generated assembly code executes exactly as prescribed by the semantics of the source C program.

This year, we improved the CompCert C compiler in several directions:

- The generation of debugging information in DWARF format was implemented by Bernhard Schommer at AbsInt. Consequently, CompCert-compiled programs can now be debugged using standard debuggers. Xavier Leroy extended the back-end compilation passes and their proofs to propagate debugging information throughout the compilation pipeline.
- The CompCert formal semantics was made more precise in order to increase confidence. We tightened the semantics of pointer comparisons against the null pointer. We formalized the distinction between public and private (`static`) global definitions, and used it to prove the correctness of the "Unusedglob" pass that removes unreferenced private definitions.
- The calling conventions used to pass function arguments and results of `struct` and `union` types were revised in order to comply with the Application Binary Interfaces of the target platforms.
- We added partial support for extended inline assembly, an extension of the C language popularized by the GCC compiler and often used in low-level code.
- Detailed explanations of syntax errors are now produced. This usability feature builds on François Pottier's work on error reporting in LR parsers (see section 7.4.4).
- The PowerPC back-end was extended to support the PowerPC 64-bit extensions and the Freescale E5500 variant.

We released two versions of CompCert, integrating these enhancements: version 2.5 in June and version 2.6 in December. This is the public version of CompCert, available for evaluation and research purposes. In parallel, our industrial partner, AbsInt Angewandte Informatik GmbH, sells a commercial version of CompCert with long-term maintenance.

7.1.2. Formal verification of static analyzers based on abstract interpretation

Participants: Jacques-Henri Jourdan, Xavier Leroy, Sandrine Blazy [team Celtique], Vincent Laporte [team Celtique], David Pichardie [team Celtique], Sylvain Boulmé [Grenoble INP, VERIMAG], Alexis Fouilhé [Université Joseph Fourier de Grenoble, VERIMAG], Michaël Périn [Université Joseph Fourier de Grenoble, VERIMAG].

In the context of the ANR Verasco project, we are investigating the formal specification and verification in Coq of a realistic static analyzer based on abstract interpretation. This static analyzer handles a large subset of the C language (the same subset as the CompCert compiler, minus recursion and dynamic allocation); supports a combination of abstract domains, including relational domains; and should produce usable alarms. The long-term goal is to obtain a static analyzer that can be used to prove safety properties of real-world embedded C code. The overall architecture and specification of Verasco is described in a paper that was presented at POPL 2015 [19].

This year, Jacques-Henri Jourdan continued the development of this static analyzer, with two goals. First, Jacques-Henri Jourdan improved the precision and analysis time of the existing abstract domains. The existing communication system between domains was instantiated to the cooperation between the abstract domain of intervals and the abstract domain of congruences. Second, Jacques-Henri Jourdan implemented and formalized in our static analyzer the Octagon abstract domain of Miné [46]. This led to new results in the theory behind this abstract domain, allowing Jourdan to use sparse data structures for representing octagons.

7.1.3. A SPARK Front-end for CompCert

Participants: Pierre Courtieu, Zhi Zang [Kansas University].

SPARK is a language, and a platform, dedicated to developing and verifying critical software. It is a subset of the Ada language. It shares with Ada a strict typing discipline and gives strict guarantees in terms of safety. SPARK goes one step further by disallowing certain "dangerous" features, that is, those that are too difficult to statically analyze (aliasing, references, etc). Given its dedication to safety critical software, we think that the SPARK platform can benefit from a certified compiler. We are working on adding a SPARK front-end to the CompCert verified compiler.

Defining a semantics for SPARK in Coq is previous joint work with Zhi Zang from Kansas University. The current front-end is based on this semantics. The compiler has been written and tested, and the proofs of correctness are currently under way.

7.1.4. Verified JIT compilation of Coq

Participants: Maxime Dénès, Xavier Leroy.

Last year, we started the Coqonut project, whose objective is to develop and formally verify an efficient, compiled implementation of Coq's reduction. This year, we made progress on this verification effort:

- We ported our OCaml prototype to Coq and started its verification, notably of the first phase of the compiler which involves uncurrying, using untyped step-indexed logical relations.
- We adapted (part of) the Coq x86 macro assembler by Andrew Kennedy, Nick Benton, Jonas B. Jensen and Pierre-Evariste Dagand to x86-64. This macro assembler framework is used in Coqonut's backend to generate assembly or machine code.

7.2. Language design and type systems

7.2.1. Full reduction in the presence of inconsistent assumptions

Participants: Didier Rémy, Gabriel Scherer.

Gabriel Scherer and Didier Rémy continued their work on assumption hiding and presented it at ESOP 2015 [22]. This work aims at restoring confluence when mixing full and weak reduction and providing a continuum between consistent and inconsistent abstraction. Assumption hiding supports fine-grained control of dependencies between computations and the logical hypotheses they depend on. Although studied for a language of coercions, the solution is more general and should be applicable to any language with abstraction over propositions that are left implicit, either for the user's convenience in a surface language or because they have been erased prior to computation in an internal language.

7.2.2. Equivalence and normalization of lambda-terms with sums

Participants: Gabriel Scherer, Guillaume Munch-Maccagnoni [Université Paris-Diderot, laboratoire PPS].

Gabriel Scherer presented at TLCA 2015 his work on understanding equivalence of sum types using the proof-theoretical technique of focusing [24]. Independently, his collaboration with Guillaume Munch-Maccagnoni resulted in a presentation of sum equivalence using an abstract machine calculus [33]. This approach allows for a more concise and cleaner definition of the equivalence relation, and a finer-grained understanding of the role of purity assumptions in the program equivalence relation.

7.2.3. Types with unique inhabitants for code inference

Participants: Gabriel Scherer, Didier Rémy.

Gabriel Scherer and Didier Rémy presented at ICFP 2015 [23] an algorithm to decide whether a type has a unique inhabitant in the simply-typed lambda-calculus with sum types. This algorithm comes along with a prototype implementation. This minimal setting is not representative of the expressiveness of realistic programming languages, but already covers a first few interesting code inference scenarios for polymorphic libraries in functional languages with prenex polymorphism: for instance, we can infer the "bind" function of the exception monad.

7.2.4. Refactoring with ornaments in ML

Participants: Thomas Williams, Didier Rémy.

Thomas Williams and Didier Rémy continued working on ornaments for program refactoring and program transformation in ML. Ornaments have been introduced as a way to describe some changes in data type definitions that preserve their recursive structure, reorganizing, adding, or dropping some pieces of data. After a new data structure has been described as an ornament of an older one, some functions operating on the bare structure can be partially or sometimes totally lifted into functions operating on the ornamented structure.

We have previously described an algorithm to perform this lifting in ML. This description was informal. This year, we improved this algorithm by decomposing it in several steps and we formalized it. Using ornament inference, we first elaborate an ML program into a generic program, which can be seen as a template for all possible liftings of the original program. The generic program is defined in a superset of ML. It can then be instantiated with specific ornaments, and simplified back into an ML program. We also studied the properties of lifting, particularly the preservation of complexity and effects, with the aim of characterizing more precisely the syntactic liftings that can be produced by our algorithm.

On the practical side, our prototype ornamentation tool has been improved with an implementation of ornament inference. The generalized program gives a description of all possible extension points that must be filled by providing patches. In practice, a few heuristics are enough to automate most of the patching work. The rest can be filled interactively by the programmer. In the case of refactoring (the representation of a data type is modified without adding any data), the transformation is fully automatic.

7.2.5. The Mezzo programming language

Participants: Thibaut Balabonski [Université Paris Sud], François Pottier, Jonathan Protzenko.

Mezzo is a programming language proposal whose untyped foundation is very much like OCaml (i.e., it is equipped with higher-order functions, algebraic data structures, mutable state, and shared-memory concurrency) and whose type system offers flexible means of describing ownership policies and controlling side effects.

A comprehensive paper, which contains both a tutorial introduction to Mezzo and a description of its formal definition and proof, was submitted to TOPLAS in 2014. This year, after a round of reviewing, it was revised and accepted for publication [11]. A reflection on the design of Mezzo was presented at SNAPL 2015 [21].

7.3. Shared-memory parallelism

7.3.1. Weak memory models

Participants: Luc Maranget, Jade Alglave [Microsoft Research, Cambridge], Patrick Cousot [New York University], Keryan Didier.

Modern multi-core and multi-processor computers do not follow the intuitive "Sequential Consistency" model that would define a concurrent execution as the interleaving of the executions of its constituent threads and that would command instantaneous writes to the shared memory. This situation is due both to in-core optimisations such as speculative and out-of-order execution of instructions, and to the presence of sophisticated (and cooperating) caching devices between processors and memory. Luc Maranget took part in an international research effort to define the semantics of the computers of the multi-core era, and more generally of shared-memory parallel devices or languages, with a clear focus on devices.

More precisely, in 2015, Luc Maranget collaborated with Jade Alglave and Patrick Cousot to extend "Cats", a domain-specific language for defining and executing weak memory models. A precise semantics for "Cats" is the core of a submitted journal article that also includes a study and formalisation of the HSA memory model — the Heterogeneous System Architecture foundation is an industry standards body targeting heterogeneous computing devices (see http://www.hsafoundation.com/). The new extensions of the Cats language have been integrated in the released version of the **diy** tool suite (see section 6.2).

Luc Maranget also co-authored a paper that will be presented at POPL 2016 [18]. This work describes an operational semantics for the new generation ARM processors. It is joint work with many researchers, including S. Flur and other members of P. Sewell's team (University of Cambridge) and W. Deacon (ARM Ltd).

During his M2 internship, supervised by Luc Maranget, Keryan Didier significantly improved the **diy** tool suite, in particular by writing front-ends for ARMv8 and for a subset of the C language. Keryan Didier also wrote a new (as yet unreleased) tool to translate between various input languages, in particular from machine assemblers to generic assembler and back.

7.3.2. Algorithms and data structures for parallel computing

Participants: Umut Acar, Vitalii Aksenov, Arthur Charguéraud, Mike Rainey, Filip Sieczkowski.

The ERC Deepsea project, with principal investigator Umut Acar, started in June 2013 and is hosted by the Gallium team. This project aims at developing techniques for parallel and self-adjusting computation in the context of shared-memory multiprocessors (i.e., multicore platforms). The project is continuing work that began at Max Planck Institute for Software Systems between 2010 and 2013. As part of this project, we are developing a C++ library, called PASL, for programming parallel computations at a high level of abstraction. We use this library to evaluate new algorithms and data structures. We obtained three major results this year.

Our result on the development of fast and robust parallel graph traversal algorithms based on depth-first-search has been presented at the ACM/IEEE Conference on High Performance Computing [15]. This algorithm leverages a new sequence data structure for representing the set of edges remaining to be visited. In particular, it uses a balanced split operation for partitioning the edges of a graph among the processors involved in the computation. Compared with prior work, the new algorithm is designed to be efficient not just for particular classes of graphs, but for all input graphs.

Our second result is a calculus for parallel computing on hardware shared memory computers such as modern multicores. Many languages for writing parallel programs have been developed. These languages offer several distinct abstractions for parallelism, such as fork-join, async-finish, futures, etc. While they may seem similar, these abstractions lead to different semantics, language design and implementation decisions. In this project, we consider the question of whether it would be possible to unify these approaches to parallelism. To this end, we propose a calculus, called the *DAG-calculus*, which can encode existing approaches to parallelism based on fork-join, async-finish, and futures, and possibly others. We have shown that the approach is realistic by presenting an implementation in C++ and by performing an empirical evaluation. This work has been submitted for publication.

Our third result concerns the development of parallel dynamic algorithms. This year, we started developing a parallel dynamic algorithm for tree computations. The algorithm is dynamic in the sense that it admits changes to the underlying tree in the form of insertions and deletions of edges and vertices and updates the computation by doing total work that is linear in the size of the changes, but only logarithmic in the size of the tree. The

algorithm is parallel in the sense that the updates take place in parallel. Parallel algorithms have been studied extensively in the past, but few of these are dynamic. Similarly, dynamic algorithms have also been studied extensively in the past, but few of these are parallel. Our work thus explores what in retrospect seems like an obvious gap in the literature. A paper describing this work is in preparation.

7.4. The OCaml language and system

7.4.1. The OCaml system

Participants: Damien Doligez, Alain Frisch [Lexifi SAS], Jacques Garrigue [University of Nagoya], Fabrice Le Fessant, Xavier Leroy, Luc Maranget, Gabriel Scherer, Mark Shinwell [Jane Street], Leo White [Jane Street], Jeremy Yallop [OCaml Labs, Cambridge University].

This year, we released versions 4.02.2 and 4.02.3 of the OCaml system. These are minor releases that fix about 100 bugs and implement 12 minor new features, including support for nonrecursive type definitions and a higher-level interface with documentation generation tools.

Most of our activity was devoted to preparing the next major release of OCaml, version 4.03.0, which is expected in the first quarter of 2016. The novelties we worked on include:

- Inline record types as arguments to constructors of sum types, combining the clarity and extensibility brought by named record fields with the compact in-memory representation of unnamed constructor arguments.
- Improved redudancy and exhaustiveness checks for pattern-matching over generalized algebraic data types (GADTs) [41].
- Improved unboxing optimizations for numbers, including the ability to mark arguments and results of external C functions as unboxed.
- The garbage collector was made more incremental, so as to reduce the worst-case GC pause times.
- The native-code compiler was ported to two new architectures: PowerPC 64 bits (including IBM's new little-endian variant) and IBM zSystems.

On the organization side, we switched to Github as the central repository for the OCaml development sources. Github facilitates collaborative work among the growing community of contributors to the OCaml code base. In 2015, more than 100 contributors proposed small or large improvements to the OCaml compiler distribution.

7.4.2. Memory profiling OCaml applications

Participants: Fabrice Le Fessant, Çagdas Bozman [OCamlPro], Albin Coquereau [OCamlPro].

Most modern languages make use of automatic memory management to discharge the programmer from the burden of explicitly allocating and releasing chunks of memory. As a consequence, when an application exhibits an unexpected usage of memory, programmers need new tools to understand what is happening and how to solve such an issue. In OCaml, the compact representation of values, with almost no runtime type information, makes the design of such tools more complex.

In the past, we have experimented with different tools to profile the memory usage of real OCaml applications, in particular one that saves snapshots of the heap after every garbage collection. Snapshots can then be analysed to display the evolution of memory usage, with detailed information on the types of values, where they were allocated and from where they are still reachable.

This year, we experimented in three new directions, mostly driven by the size of the snapshots to be analysed:

- We studied several ways of displaying snapshots. Because of the large amount of information contained in a snapshot, it is hard for a typical user to find what he or she is looking for. We tried multiple filtering methods, based on graph algorithms, to remove the least significant information from the reports given to the user.

- We experimented with new algorithms to compress and analyse *huge* memory snapshots, i.e., snapshots that are too big to fit in the computer's memory. Indeed, standard analyses on snapshots bigger than the available memory are too long to run in practice because of random disk accesses. Thus, we tried several compression methods for snapshots and graph-reduced them to fit in memory, without losing any information, reaching a 50x speedup in complete analysis time.
- We implemented a new graph algorithm to merge sets of blocks in memory by the sets of roots they are reachable from. Such a computation was heretofore supposed to be untractable in practice, but could actually be computed in our case on huge compressed snapshots in reasonable time.

7.4.3. Advanced development tools for OCaml

Participants: Fabrice Le Fessant, Pierre Chambart [OCamlPro], Michael Laporte [OCamlPro].

In order to promote the use of OCaml in industrial contexts, we have worked on improving the tools that accompany OCaml:

- We developed the first prototype of a native debugger for OCaml, based on the LLDB debugging framework on top of LLVM. For that, we first generated a full OCaml binding for the LLDB library, by parsing the C++ headers of the libraries and automatically generating OCaml and C++ stubs. We were then able to use the OCaml binding to develop several tools, ranging from a simple tool that displays the internal GC information of a finished OCaml application, to an almost complete debugger, which displays OCaml values using runtime type information added for memory profiling.
- We also developed a new profiling framework for OCaml, called *operf*. The framework is composed of two tools: *operf-micro* can be used to run micro-benchmarks directly from inside modified OCaml compiler sources, while the *operf-macro* tool can be used to evaluate the impact of a new compiler optimization on a large set of OPAM packages.
- Finally, we came up with new ideas for *ocp-build*, a generic building tool with OCaml-specific support, to improve the expressiveness of its package description language and to easily describe cross-compilation of OCaml packages.

7.4.4. Error diagnosis in Menhir parsers

Participant: François Pottier.

LR parsers are powerful and efficient, but traditionally have done a poor job of explaining syntax errors. Although it is easy to report where an error was detected, it seems difficult to explain what has been understood so far and what is expected next. The OCaml and CompCert compilers, until now, have offered little information to the user beyond the traditional "syntax error" message.

In 2003, Jeffery proposed associating a fixed diagnostic message with every state of the LR automaton (therefore ignoring the automaton's stack). This simple approach may seem tempting. However, a typical automaton has hundreds or thousands of states. Not all of them can trigger an error, but it is difficult to tell which can, and which cannot. Furthermore, for certain states, it is difficult (or even impossible) to write an accurate diagnostic message, because some vital contextual information resides in the stack, which Jeffery's method cannot access.

In 2015, François Pottier proposed a reachability algorithm for LR automata, which he implemented in the Menhir parser generator (see section 6.3). This algorithm allows finding out which states can trigger an error and (therefore) require writing a diagnostic message. Furthermore, Pottier proposed two mechanisms for influencing where errors are detected. If used appropriately, these mechanisms make it easier (or possible) to write an accurate diagnostic message.

Pottier applied this approach to the C grammar in the front-end of the CompCert compiler, therefore allowing CompCert to produce better diagnostic messages when a C program is syntactically incorrect.

A short paper describing this work will be presented at JFLA 2016 [29]. A longer paper is in submission.

7.4.5. Improvements to Menhir

Participants: Frédéric Bour [independent consultant], Jacques-Henri Jourdan, François Pottier, Yann Régis-Gianas [team πr^2], Gabriel Scherer.

In 2015, The Menhir parser generator (see section 6.3) was extended with many new features, several of which originated in the Merlin IDE for OCaml and were ported back into Menhir.

- The parsers generated by Menhir are now incremental: they can be stopped and resumed at any point, at essentially no cost. This is exploited in Merlin, where the text is re-parsed after every keystroke.
- The state of the parser can be inspected by the user. This allows building custom libraries, outside Menhir, for error diagnosis, error recovery, etc. This is exploited in Merlin, where a valid abstract syntax tree is built (and passed to the OCaml type-checker) even if the text contains syntax errors.
- A reachability algorithm has been implemented (see section 7.4.4). It allows finding out which states can trigger an error and (therefore) require a diagnostic message to be written. It is accompanied with several tools that help maintain the database of diagnostic messages as the grammar evolves.
- Compatibility with `ocamlyacc` has been improved, in particular insofar as the computation of locations is concerned. This should help port the OCaml parser from `ocamlyacc` to Menhir, a transition that we envision making in the near future. This should help improve the quality of OCaml's syntax error messages.

7.5. Software specification and verification

7.5.1. Machine-checked proofs of programs, including time complexity

Participants: Arthur Charguéraud, Armaël Guéneau, François Pottier.

In a security-critical setting, it is important to prove that a program is correct, and to do so formally, that is, via a machine-checked proof. It is also important, one may argue, to prove that the program does not require more resources than expected (where a "resource" may be time, memory space, disk space, network bandwidth, etc.). Otherwise, even though the program is "correct" in theory, it may turn out to be unusable in practice.

Separation Logic, extended with the notion of a "time credit", a permission to perform one step of computation, allows reasoning about the correctness and the (amortized) time complexity of a program. Using this approach, which Charguéraud implemented in the CFML tool, Charguéraud and Pottier produced a machine-checked proof of the correctness and time complexity of a Union-Find data structure, implemented as an OCaml module. This demonstrates that this approach scales up to difficult complexity analyses and down to the level of actual executable code (as opposed to pseudo-code). This work was presented at ITP 2015 [17].

During his M2 internship, Armaël Guéneau extended this approach so as to allow working conveniently with the big-O notation. He extended the CFML library and verified the time complexity of a binary random access list data structure due to Okasaki. This work has not been published yet.

7.5.2. Verified property-based random testing

Participants: Zoe Paraskevopoulou [ENS Cachan, team Prosecco], Cătălin Hriţcu [team Prosecco], Maxime Dénès, Leonidas Lampropoulos [U. of Pennsylvania], Benjamin C. Pierce [U. of Pennsylvania].

Property-based random testing has been popularized in the functional programming community by tools like QuickCheck. Its integration with a proof assistant creates an interesting opportunity: reusable or tricky testing code can be formally verified using the proof assistant itself.

We introduced a novel methodology for formally verified property-based testing and implemented it as a foundational verification framework for QuickChick, a port of QuickCheck to Coq. Our framework enables one to verify that the executable testing code is testing the right Coq property. To make verification tractable, we provided a systematic way for reasoning about the set of outcomes a random data generator can produce with non-zero probability, while abstracting away from the actual probabilities.

We also applied this methodology to a complex case study on testing an information-flow control abstract machine, demonstrating that our verification methodology is modular and scalable and that it requires minimal changes to existing code.

Maxime Dénès more specifically contributed to the development of the QuickChick Coq plug-in, to the development of Coq libraries for reasoning on the set of outcomes of random generators and to the verification of QuickChick's combinator library.

This work was presented at ITP 2015 [20].

7.5.3. Tools for TLA+

Participants: Damien Doligez, Leslie Lamport [Microsoft Research], Martin Riener [team VeriDis], Stephan Merz [team VeriDis].

Damien Doligez is head of the "Tools for Proofs" team in the Microsoft-Inria Joint Centre. The aim of this project is to extend the TLA+ language with a formal language for hierarchical proofs, formalizing Lamport's ideas [43], and to build tools for writing TLA+ specifications and mechanically checking the proofs.

This year, we released version 1.4.3 of the TLA+ Proof System (TLAPS) [40], the part of the TLA+ tools that handles mechanical checking of TLA+ proofs.

This was the last year of the ADN4SE project, which develops tools for rapid development of real-time software based on the PharOS real-time kernel developed by CEA. Within this project we built, in collaboration with CEA, a formal proof of determinacy of the message-passing subsystem of PharOS. We used this experience to improve our TLA+ tools and libraries.

We have started a rewrite of TLAPS from scratch, which will make it possible to handle all aspects of the TLA+ language, including temporal formulas and their proofs.

7.5.4. Certified distributed algorithms for autonomous mobile robots

Participants: Pierre Courtieu, Xavier Urbain [ENSIIE], Sébastien Tixeuil [U. Pierre et Marie Curie], Lionel Rieg [Collège de France].

The variety and complexity of the tasks that can be performed by autonomous robots are increasing. Many applications envision groups of mobile robots that self-organise and cooperate toward the resolution of common objectives, in the absence of any central coordinating authority.

We are developing a Coq-based verification platform for distributed algorithms for autonomous robots. This year, we mechanically proved and slightly generalized a non-trivial proof of impossibility of such an algorithm under certain hypotheses [14]. We also proved several algorithms in the literature, demonstrating the viability of the platform [13].

7.5.5. Contributions to ProofGeneral, an IDE for Coq

Participant: Pierre Courtieu.

User interface is a crucial issue for theorem provers like Coq. ProofGeneral [38], an emacs-based prover interface, is widely used among Coq users. In addition to synchronizing with the evolutions of Coq itself, we contributed many improvements to ProofGeneral during the past year, among which: a better debugging mode and message printing, user assistance for naming hypotheses and indenting proof scripts, and more.

GAMMA3 Project-Team

5. New Results

5.1. Serendipity and reduced elements

Participants: Paul Louis George [correspondant], Houman Borouchaki, Nicolas Barral.

We give a method to constructing Serendipity elements for quads and hexes with full symmetry properties and indicate the reading of their shape functions. We show that, since the degree 5, the Serendipity elements are no longer symmetric but we propose a method resulting in a Lagrange element of degree 5 with full symmetry properties after adding an adequate number of additional nodes.

On the other hand, we show how to guarantee the geometric validity of a given curved element (seen as a patch) of a mesh. This is achieved after writing the patch in a Bézier setting (Bernstein polynomials and control points). In addition, we discuss the case of patch derived from a transfinite interpolation and it is proved that only some of them are Serendipity elements indeed, we return to the same elements as above

We also give a method to constructing Lagrange Serendipity (or reduced) simplices with a detailed description of the triangles of degree 3 and 4. We indicate that higher order triangles are not candidate apart if we impose a restricted polynomial space. We show that a tetrahedron of degree 3 is a candidate while high order elements are not candidate even if a restriction in the polynomial space is considered. In addition, we propose a method for the validation of such elements, in a given mesh, where the validation means the positiveness of the jacobian.

5.2. Validity of transfinite and Bézier-Serendipity patches

Participants: Paul Louis George [correspondant], Houman Borouchaki, Nicolas Barral.

We define generalized transfinite patches for quads and hexes with full symmetry properties. We give a way of constructing those patches by considering the Bézier setting using linear combinations of tensor-product patches of various degree. Those patches are exactly the Bézier-Serendipity patches recently introduced

ASsfor reduced quadrilateral patches, we introduce the so called "Bézier-Serendip" patches. After some recalls about standard Bézier patches, we propose a method to constructing those reduced patches. The corresponding Bernstein polynomials are written by means of linear combinations of the standard Bernstein polynomials. We give a full description of the patches of degree 2, 3, 4 and 5. Since degree 5, the location of the control points is no longer symmetric and to remedy this problem, we propose adding a number of control points which results in *extended* Bézier-Serendip patches. Those reduced patches are in the Bézier framework what the Serendipity elements are in the finite element framework.

A technical report and a paper have been published [16].

5.3. Meshing Strategies and the Impact of Finite Element Quality on the Velocity Field in Fractured Media

Participants: Patrick Laug [correspondant], Géraldine Pichot.

For calculating flow in a fracture network, the mixed hybrid finite element (MHFE) method is a method of choice as it yields a symmetric, positive definite linear system. However, a drawback to this method is its sensitivity to bad aspect ratio elements. For poor-quality triangles, elementary matrices are ill-conditioned, and inconsistent velocity vectors are obtained by inverting these local matrices. In this work, different strategies have been proposed for better reconstruction of the velocity field.

5.4. Automatic Mesh Generation of Multiface Models on Multicore Processors

Participant: Patrick Laug [correspondant].

This work started in September 2014, as part of a sabbatical year at Polytechnique Montréal. In a previous study, a parallel version of an indirect approach for meshing composite surfaces – also called multiface models – was developed. However, this methodology could be inefficient in practice, as the memory management of most existing CAD (computer aided design) systems use static global caches to save information. In a first approach, CAD queries are fully parallelized, using the Pirate library from Polytechnique Montréal (this library provides a set of C++ classes that implement STEP-compliant B-Rep geometric and topological entities, as well as classes to represent meshes and solutions). In a second approach, the CAD system is completely disconnected from the mesh generator, using a discrete geometric support.

5.5. Applications du maillage et développements de méthodes avancées pour la cryptographie

Participants: Thomas Grosges [correspondant], Dominique Barchiesi, Michael François.

L'utilisation des nombres (pseudo)-aléatoires a pris une dimension importante ces dernières décennies. De nombreuses applications dans le domaine des télécommunications, de la cryptographie, des simulations numériques ou encore des jeux de hasard, ont contribué au développement et à l'usage de ces nombres. Les méthodes utilisées pour la génération de tels nombres (pseudo)-aléatoires proviennent de deux types de processus : physique et algorithmique. Ce projet de recherche a donc pour objectif principal le développement de nouveaux procédés de génération de clés de chiffrement, dits "exotiques", basés sur des processus physiques, multi-échelles, multi-domaines assurant un niveau élevé de sécurité. Deux classes de générateurs basés sur des principes de mesures physiques et des processus mathématiques ont été développé.

La première classe de générateurs exploite la réponse d'un système physique servant de source pour la génération des séquences aléatoires. Cette classe utilise aussi bien des résultats de simulation que des résultats de mesures interférométriques pour produire des séquences de nombres aléatoires. L'application du maillage adaptatif sert au contrôle de l'erreur sur la solution des champs physiques (simulés ou mesurés). A partir de ces cartes physiques, un maillage avec estimateur d'erreur sur l'entropie du système est appliqué. Celui-ci permet de redistribuer les positions spatiales des noeuds. L'étude (locale) de la réduction d'entropie des clés tout au long de la chaîne de création et l'étude (globale) de l'entropie de l'espace des clés générées sont réalisées à partir de tests statistiques.

La seconde classe de générateurs porte sur le développement de méthodes avancées et est basée sur l'exploitation de fonctions chaotiques en utilisant les sorties de ces fonctions comme indice de permutation sur un vecteur initial. Ce projet s'intéresse également aux systèmes de chiffrement pour la protection des données et deux algorithmes de chiffrement d'images utilisant des fonctions chaotiques sont développés et analysés. Ces Algorithmes utilisent un processus de permutation-substitution sur les bits de l'image originale. Une analyse statistique approfondie confirme la pertinence des cryptosystèmes développés.

5.6. Développement de méthodes avancées et maillages appliqués à l'étude de la nanomorphologie des nanotubes-fils en suspension liquide

Participants: Thomas Grosges [correspondant], Dominique Barchiesi, Abel Cherouat, Houman Borouchaki, Laurence Giraud-Moreau, Anis Chaari.

Ce projet de recherche (NANOMORPH) a pour objet principal le développement et la mise au point d'une instrumentation optique pour déterminer la distribution en tailles et le coefficient de forme de nanofils (NF) ou de nanotubes (NT) en suspension dans un écoulement. Au cours de ce projet, deux types de techniques optiques complémentaires sont développées. La première, basée sur la diffusion statique de la lumière, nécessite d'étudier au préalable la physico-chimie de la dispersion, la stabilisation et l'orientation des nanofils dans les milieux d'étude. La seconde méthode, basée sur une méthode opto-photothermique pulsée, nécessite en sus,

la modélisation de l'interaction laser/nanofils, ainsi que l'étude des phénomènes multiphysiques induits par ce processus. L'implication de l'équipe-projet GAMMA3 concerne principalement la simulation multiphysique de l'interaction laser-nanofils et l'évolution temporelle des bulles et leurs formations. L'une des principales difficultés de ces problématiques est que la géométrie du domaine est variable (à la fois au sens géométrique et topologique). Ces simulations ne peuvent donc être réalisées que dans un schéma adaptatif de calcul nécessitant le remaillage tridimensionnel mobile, déformable avec topologie variable du domaine (formation et évolution des bulles au cours du temps et de l'espace).

5.7. Applications du maillage à des problèmes multi-physiques, développement de méthodes de résolutions avancées et modélisation électromagnetisme-thermique-mécanique á l'échelle mesoscopique

Participants: Dominique Barchiesi [correspondant], Abel Cherouat, Thomas Grosges, Houman Borouchaki, Laurence Giraud-Moreau, Sameh Kessentini, Anis Chaari, Fadhil Mezghani.

Le contrôle et l'adaptation du maillage lors de la résolution de problèmes couplés ou/et non linéaires reste un problème ouvert et fortement dépendant du type de couplage physique entre les EDP à résoudre. Notre objectif est de développer des modèles stables afin de calculer les dilatations induites par l'absorption d'énergie électromagnétique, par des structures matérielles inférieures au micron. Les structures étudiées sont en particulier des nanoparticules métalliques en condition de résonance plasmon. Dans ce cas, un maximum d'énergie absorbée est attendu, accompagné d'un maximum d'élévation de température et de dilatation. Il faut en particulier développer des modèles permettant de simuler le comportement multiphysique de particules de formes quelconques, pour une gamme de fréquences du laser d'éclairage assez étendue afin d'obtenir une étude spectroscopique de la température et de la dilatation. L'objectif intermédiaire est de pouvoir quantifier la dilatation en fonction de la puissance laser incidente. Le calcul doit donc être dimensionné et permettre finalement des applications dans les domaines des capteurs et de l'ingénierie biomédicale. En effet, ces nanoparticules métalliques sont utilisées á la fois pour le traitement des cancers superficiels par nécrose de tumeur sous éclairage adéquat, dans la fenêtres de transparence cellulaire. Déposées sur un substrat de verre, ces nanoparticules permettent de construire des capteurs utilisant la résonance plasmon pour être plus sensibles (voir projet européen *Nanoantenna* et l'activité génération de nombres aléatoires). Cependant, dans les deux cas, il est nécessaire, en environnement complexe de déterminer la température locale, voire la dilatation de ces nanoparticules, pouvant conduire à un désaccord du capteur, la résonance plasmon étant très sensible aux paramètres géométriques et matériels des nanostructures. Dans ce sens, l'étude permet d'aller plus loin que la "simple" interaction électromagnétique avec la matière du projet européen *Nanoantenna*.

Le travail de l'année 2014 a constitué en la poursuite de l'étude des spécificités de ce type de problème multiphysique pour des structures de forme simple et la mise en place de fonctions test, de référence, pour les développements de maillage adaptatifs pour les modèles multiphysiques éléments finis. Nous espérons pouvoir proposer un projet ANR couplant les points de vue microscopiques et macroscopiques dans les deux années qui viennent.

5.8. Visualization and modification of high-order curved meshes

Participants: Alexis Loyer, Adrien Loseille [correspondant].

During the partnership between Inria and Distene, a new visualization software has been designed. It address the typical operations that are required to quickly assess the newly algorithm developed in the team. In particular, interactive modifications of high-order curved mesh and hybrid meshes has been addressed. The software VIZIR is freely available at https://www.rocq.inria.fr/gamma/gamma/vizir/.

5.9. Mesh adaptive ALE numerical simulation

Participants: Frédéric Alauzet [correspondant], Nicolas Barral, Adrien Loseille.

Running highly accurate numerical simulations with moving geometries is still a challenge today due to their prohibitive cost in CPU time. Using anisotropic mesh adaptation is one way to drastically reduce the size of the problem and to reach the desired accuracy. Previously, we have developed an ALE formulation using mesh connectivity change in order to achieve any complex displacement. Then, this method has been coupled with the unsteady anisotropic mesh adaptation using the fixed-point algorithm. The key point of this work is the use of an ALE metric that takes into account the mesh motion in the metric field definition.

5.10. Mesh adaptation for Navier-Stokes Equations

Participants: Frédéric Alauzet, Victorien Menier, Adrien Loseille [correspondant].

Adaptive simulations for Navier-Stokes equations require to propose accurate error estimates and design robust mesh adaptation algorithms (for boundary layers).

For error estimates, we design new estimates suited to accurately capture the speed profile in the boundary layers. For mesh adaptation, we design a new method to generate structured boundary layer meshes which are mandatory to accurately compute compressible flows a high Reynolds number (several millions). It couple the specification of the optimal boundary layer from the geometry boundary and moving mesh techniques to extrude the boundary layer in an already existing mesh. The main advantage of this approach is its robustness, *i.e.,* at each step of the algorithm we have always a valid mesh [23].

5.11. Adaptive multigrid strategies

Participants: Frédéric Alauzet [correspondant], Victorien Menier, Adrien Loseille.

Multigrid is a well known technique used to accelerate the convergence of linear system solutions. Using a multigrid strategy to solve non-linear problems improves the robustness and the convergence of each Newton step, the accelerating overall the whole process. In particular, larger time step can be considered. This of main importance when solving turbulent Navier-Stokes equations on complex geometries. First, we developed the classical multigrid method on non-nested meshes. Then, we have pointed out the similarity between the Full MultiGrid (FMG) algorithm and the mesh adaptation algorithm. We have proposed a new Adaptive Full MultiGrid algorithm which improve the overall robustness of the adaptive process and its overall efficiency [23].

5.12. Metric-orthogonal and metric-aligned mesh adaptation

Participants: Frédéric Alauzet, Victorien Menier, Adrien Loseille [correspondant].

A new algorithm to derive adaptive meshes has been introduced through new cavity-based algorithms. It allows to generate anisotropic surface and volume mesh that are aligned along the eigenvector directions. This allows us to improv the quality of the meshes and to deal naturally with boundary layer mesh generation.

5.13. Parallel mesh adaptation

Participants: Frédéric Alauzet, Victorien Menier, Adrien Loseille [correspondant].

We devise a strategy in order to generate large-size adapted anisotropic meshes $O(10^8 - 10^9)$ as required in many fields of application in scientific computing. We target moderate scale parallel computational resources as typically found in R&D units where the number of cores ranges in $O(10^2 - 10^3)$. Both distributed and shared memory architectures are handled. Our strategy is based on typical domain splitting algorithm to remesh the partitions in parallel. Both the volume and the surface mesh are adapted simultaneously and the efficiency of the method is independent of the complexity of the geometry. The originality of the method relies on (i) a metric-based static load-balancing, (ii) dedicated mesh partitioning techniques to (re)split the (complex) interfaces meshes, (iii) anisotropic Delaunay cavity to define the interface meshes, (iv) a fast, robust and generic sequential cavity-based mesh modification kernel, and (v) out-of-core storing of completing parts to reduce the memory footprint. We are able to generate (uniform, isotropic and anisotropic) meshes with more than 1 billion tetrahedra in less than 20 minutes on 120 cores.

5.14. Unsteady adjoint computation on dynamic meshes

Participants: Eléonore Gauci, Frédéric Alauzet [correspondant].

Adjoint formulations for unsteady problems are less common in unsteady methodologies due to the extra complexity inherent in the numerical solution and storage but these methods are a great option in engineering because it takes more into account the cost function we want to minimize. Moreover the engineering applications involve moving elements and this motion must be taken into account by the governing flow equations. We develop a model of unsteady adjoint solver on moving mesh problems. The derivation of the adjoint formulation based on the ALE form of the equations requires consideration of the dynamic meshes. Our model takes into account the DGCL.

5.15. Line solver for efficient stiff parse system resolution

Participants: Loïc Frazza, Frédéric Alauzet [correspondant].

Afin d'accélérer la résolution des problèmes raides, un line-solver a ´été développé. Cette méthode extrait tout d'abord des lignes dans le maillage du problème selon des critères géométriques ou physiques. Le problème peut alors Ítre résolu exactement le long des ces lignes ‡ moindre cout. Cette méthode est particulièrement bien adaptée aux cas o˘ l'information se propage selon une direction privilégiée tels que les chocs, les couches limites ou les sillages. Ces cas sont généralement associés ‡ des maillages très étirés ce qui conduit ‡ des problèmes raides mais quasi-unidimensionnels. Ils peuvent donc être résolu efficacement par un line-solver, réduisant ainsi les temps de calculs tout en gagnant en robustesse.

5.16. Error estimate for high-order solution field

Participants: Olivier Coulaud, Adrien Loseille [correspondant].

Afin de produire des solveurs d'ordre élevé, et ainsi répondre aux exigences inhérentes à la résolution de problèmes physiques complexes, nous développons une méthode d'adaptation de maillage d'ordre élevé. Celle-ci est basée sur le contrôle par une métrique de l'erreur d'interpolation induite par le maillage du domaine. Plus précisément, pour une solution donnée, l'erreur d'interpolation d'ordre k est paramétrée par la différentielle kè de cette solution, et le problème se réduit à trouver la plus grande ellipse incluse dans une ligne de niveau de cette différentielle. S'il reste encore quelques difficultés techniques à résoudre avant l'exploitation numérique de notre méthode, les résultats sont très encourageants, tant en terme d'optimalité de la métrique obtenue que de temps de calcul. Il n'y a que peu de doutes sur le fait que ce projet aboutisse prochainement.

GANG Project-Team

7. New Results

7.1. Graph and Combinatorial Algorithms

7.1.1. Rainbow matchings in hypergraphs

A rainbow matching for (not necessarily distinct) sets $F_1, ..., F_k$ of hypergraph edges is a matching consisting of k edges, one from each F_i. In [8], we give some order to the multitude of conjectures that relate to this concept, as well as introduce some new conjectures. We also present some partial results on one of these conjectures, that seems central among them – the so-called Ryser-Brualdi-Stein conjecture.

7.1.2. A graph formulation of the union-closed sets conjecture

In 1979, Frankl conjectured that in a finite non-trivial union-closed collection of sets there has to be an element that belongs to at least half the sets. In [7], we show that this is equivalent to the conjecture that in a finite non-trivial graph there are two adjacent vertices, each belonging to at most half of the maximal stable sets. In this graph formulation other special cases become natural. The conjecture is trivially true for non-bipartite graphs and we show that it also holds for the classes of chordal bipartite graphs, subcubic bipartite graphs, bipartite series-parallel graphs and bipartitioned circular interval graphs.

7.1.3. Cops-and-robber games on k-chordal graphs

The cops-and-robber games, introduced by Winkler and Nowakowski (in Discrete Math. 43, 1983) and independently defined by Quilliot (in J. Comb. Theory, Ser. B 38, 1985), concern a team of cops that must capture a robber moving in a graph. In [20], we consider the class of k-chordal graphs, i.e., graphs with no induced (chordless) cycle of length greater than k, $k \geq 3$. We prove that $k-1$ cops are always sufficient to capture a robber in k-chordal graphs. This leads us to our main result, a new structural decomposition for a graph class including k-chordal graphs.

We present a polynomial-time algorithm that, given a graph G and $k \geq 3$, either returns an induced cycle larger than k in G, or computes a tree-decomposition of G, each bag of which contains a dominating path with at most $k-1$ vertices. This allows us to prove that any k-chordal graph with maximum degree Δ has treewidth at most $(k-1)(\Delta-1)+2$, improving the $O(\Delta(\Delta-1)k-3)$ bound of Bodlaender and Thilikos (Discrete Appl. Math. 79, 1997). Moreover, any graph admitting such a tree-decomposition has small hyperbolicity). As an application, for any n-vertex graph admitting such a tree-decomposition, we propose a compact routing scheme using routing tables, addresses and headers of size $O(k \log \Delta + \log n)$ bits and achieving an additive stretch of $O(k \log \Delta)$. As far as we know, this is the first routing scheme with $O(k \log \Delta + \log n)$-routing tables and small additive stretch for k-chordal graphs.

7.1.4. Distinguishing views in symmetric networks

The view of a node in a port-labeled network is an infinite tree encoding all walks in the network originating from this node. In [16], we prove that for any integers $n \geq D \geq 1$, there exists a port-labeled network with at most n nodes and diameter at most D, which contains a pair of nodes whose (infinite) views are different, but whose views truncated to depth $\Omega(D \log (n/D))$ are identical.

7.1.5. Vertex elimination orderings for hereditary graph classes

In [3], we provide a general method to prove the existence and compute efficiently elimination orderings in graphs. This method relies on several tools that were known before, but that were not put together so far: the algorithm LexBFS due to Rose, Tarjan and Lueker, its additional properties discovered by Berry and Bordat, and a local decomposition property of graphs discovered by Maffray, Trotignon and Vušković. We use this method to prove the existence of elimination orderings in several classes of graphs, and to compute them in linear time. Some of the classes have already been studied, namely even-hole-free graphs, square-theta-free Berge graphs, universally signable graphs and wheel-free graphs. Some other classes are new. It turns out that all the classes that we consider can be defined by excluding some of the so-called Truemper configurations. For several classes of graphs, we obtain directly bounds on the chromatic number, or fast algorithms for the maximum clique problem or the coloring problem.

7.1.6. Fast collaborative graph exploration

In [14], we study the following scenario of online graph exploration. A team of k agents is initially located at a distinguished vertex r of an undirected graph. We ask how many time steps are required to complete exploration, i.e., to make sure that every vertex has been visited by some agent. As our main result, we provide the first strategy which performs exploration of a graph with n vertices at a distance of at most D from r in time $O(D)$, using a team of agents of polynomial size $k = Dn^{1+\epsilon} < n^{2+\epsilon}$, for any $\epsilon > 0$. Our strategy works in the local communication model, in which agents can only exchange information when located at a vertex, without knowledge of global parameters such as n or D.

We also obtain almost-tight bounds on the asymptotic relation between exploration time and team size, for large k, in both the local and the global communication model.

7.1.7. Position discovery for a system of bouncing robots

In [11], we consider a scenario in which a collection of n anonymous mobile robots is deployed on a unit-perimeter ring or a unit-length line segment. Every robot starts moving at constant speed, and bounces each time it meets any other robot or segment endpoint, changing its walk direction. We study the problem of position discovery, in which the task of each robot is to detect the presence and the initial positions of all other robots. The robots cannot communicate or perceive information about the environment in any way other than by bouncing nor they have control over their walks which are determined by their initial positions and their starting directions. Each robot has a clock allowing it to observe the times of its bounces. We give complete characterizations of all initial configurations for both the ring and the segment in which no position detection algorithm exists and we design optimal position detection algorithms for all feasible configurations.

7.1.8. Rendezvous of mobile agents in edge-weighted networks

In [15], we introduce a variant of the deterministic rendezvous problem for a pair of heterogeneous agents operating in an undirected graph, which differ in the time they require to traverse particular edges of the graph. Each agent knows the complete topology of the graph and the initial positions of both agents. The agent also knows its own traversal times for all of the edges of the graph, but is unaware of the corresponding traversal times for the other agent. The goal of the agents is to meet on an edge or a node of the graph. In this scenario, we study the time required by the agents to meet, compared to the meeting time T_{OPT} in the offline scenario in which the agents have complete knowledge about each others' speed characteristics. When no additional assumptions are made, we show that rendezvous in our model can be achieved after time $O(nT_{OPT})$ in a n-node graph, and that such time is essentially in some cases the best possible. However, we prove that the rendezvous time can be reduced to $\Theta(T_{OPT})$ when the agents are allowed to exchange $\Theta(n)$ bits of information at the start of the rendezvous process. We then show that under some natural assumption about the traversal times of edges, the hardness of the heterogeneous rendezvous problem can be substantially decreased, both in terms of time required for rendezvous without communication, and the communication complexity of achieving rendezvous in time $\Theta(T_{OPT})$.

7.1.9. Monitoring a graph using faulty mobile robots

In the scenario studied in [27], a team of k mobile robots is deployed on a weighted graph whose edge weights represent distances. The robots perpetually move along the domain, represented by all points belonging to the graph edges, not exceeding their maximal speed. The robots need to patrol the graph by regularly visiting all points of the domain. Here, we consider a team of robots (patrolmen), at most f of which may be unreliable, i.e. they fail to comply with their patrolling duties.

What algorithm should be followed so as to minimize the maximum time between successive visits of every edge point by a reliable patrolmen? The corresponding measure of efficiency of patrolling called idleness has been widely accepted in the robotics literature. We extend it to the case of untrusted patrolmen; we denote by $I_k^f(G)$ the maximum time that a point of the domain may remain unvisited by reliable patrolmen. The objective is to find patrolling strategies minimizing $I_k^f(G)$.

We investigate this problem for various classes of graphs. We design optimal algorithms for line segments, which turn out to be surprisingly different from strategies for related patrolling problems proposed in the literature. We then use these results to study the case of general graphs. For Eulerian graphs G, we give an optimal patrolling strategy with idleness $I_k^f(G) = (f+1)|E|/k$, where $|E|$ is the sum of the lengths of the edges of G. Further, we show the hardness of the problem of computing the idle time for three robots, at most one of which is faulty, by reduction from 3-edge-coloring of cubic graphs — a known NP-hard problem. A byproduct of our proof is the investigation of classes of graphs minimizing idle time (with respect to the total length of edges); an example of such a class is known in the literature under the name of Kotzig graphs.

7.1.10. Limit behavior of the rotor-router system

The rotor-router model, also called the Propp machine, was introduced as a deterministic alternative to the random walk. In this model, a group of identical tokens are initially placed at nodes of the graph. Each node maintains a cyclic ordering of the outgoing arcs, and during consecutive turns the tokens are propagated along arcs chosen according to this ordering in round-robin fashion. The behavior of the model is fully deterministic. Yanovski et al. (Algorithmica, 2003) proved that a single rotor-router walk on any graph with m edges and diameter D stabilizes to a traversal of an Eulerian circuit on the set of all $2m$ directed arcs on the edge set of the graph, and that such periodic behaviour of the system is achieved after an initial transient phase of at most $2mD$ steps.

The case of multiple parallel rotor-routers was studied experimentally, leading Yanovski et al. to the experimental observation that a system of $k > 1$ parallel walks also stabilizes with a period of length at most $2m$ steps. In our work [26] we disprove this observation, showing that the period of parallel rotor-router walks can in fact, be superpolynomial in the size of graph. On the positive side, we provide a characterization of the periodic behavior of parallel router walks, in terms of a structural property of stable states called a subcycle decomposition. This property provides us the tools to efficiently detect whether a given system configuration corresponds to the transient or to the limit behavior of the system. Moreover, we provide polynomial upper bounds of $O(m^4D^2 + mD\log k)$ and $O(m^5k^2)$ on the number of steps it takes for the system to stabilize. Thus, we are able to predict any future behavior of the system using an algorithm that takes polynomial time and space. In addition, we show that there exists a separation between the stabilization time of the single-walk and multiple-walk rotor-router systems, and that for some graphs the latter can be asymptotically larger even for the case of $k = 2$ walks.

7.2. Distributed Computing

7.2.1. Self-stabilizing verification and computation of an MST

In the work [19], we demonstrate the usefulness of distributed local verification of proofs, as a tool for the design of self-stabilizing algorithms. In particular, it introduces a somewhat generalized notion of distributed local proofs, and utilizes it for improving the time complexity significantly, while maintaining space optimality. As a result, we show that optimizing the memory size carries at most a small cost in terms of time, in the context of Minimum Spanning Tree (MST). That is, we present algorithms that are both time and space efficient for

both constructing an MST and for verifying it. This involves several parts that may be considered contributions in themselves.

First, we generalize the notion of local proofs, trading off the time complexity for memory efficiency. This adds a dimension to the study of distributed local proofs, which has been gaining attention recently. Specifically, we design a (self-stabilizing) proof labeling scheme which is memory optimal (i.e., $O(\log n)$ bits per node), and whose time complexity is $O(\log^2 n)$ in synchronous networks, or $O(\Delta \log^3 n)$ time in asynchronous ones, where Δ is the maximum degree of nodes. This answers an open problem posed by Awerbuch and Varghese (FOCS 1991). We also show that $\Omega(\log n)$ time is necessary, even in synchronous networks. Another property is that if f faults occurred, then, within the required detection time above, they are detected by some node in the O(f log n) locality of each of the faults. Second, we show how to enhance a known transformer that makes input/output algorithms self-stabilizing. It now takes as input an efficient construction algorithm and an efficient self-stabilizing proof labeling scheme, and produces an efficient self-stabilizing algorithm. When used for MST, the transformer produces a memory optimal self-stabilizing algorithm, whose time complexity, namely, $O(n)$, is significantly better even than that of previous algorithms. (The time complexity of previous MST algorithms that used $O(\log^2 n)$ memory bits per node was $O(n^2)$, and the time for optimal space algorithms was $O(n|E|)$.) Inherited from our proof labelling scheme, our self-stabilising MST construction algorithm also has the following two properties: (1) if faults occur after the construction ended, then they are detected by some nodes within $O(\log^2 n)$ time in synchronous networks, or within $O(\Delta \log^3 n)$ time in asynchronous ones, and (2) if f faults occurred, then, within the required detection time above, they are detected within the $O(f \log n)$ locality of each of the faults. We also show how to improve the above two properties, at the expense of some increase in the memory.

7.2.2. Clock synchronization and distributed estimation in highly dynamic networks

In [21], we consider the External Clock Synchronization problem in dynamic sensor networks. Initially, sensors obtain inaccurate estimations of an external time reference and subsequently collaborate in order to synchronize their internal clocks with the external time. For simplicity, we adopt the drift-free assumption, where internal clocks are assumed to tick at the same pace. Hence, the problem is reduced to an estimation problem, in which the sensors need to estimate the initial external time. In this context of distributed estimation, this work is further relevant to the problem of collective approximation of environmental values by biological groups.

Unlike most works on clock synchronization that assume static networks, the setting considered here is an extreme case of highly dynamic networks. We do however impose a restriction on the dynamicity of the network. Specifically, we assume a non-adaptive scheduler adversary that dictates an arbitrary, yet independent, meeting pattern. Such meeting patterns fit, for example, with short-time scenarios in highly dynamic settings, where each sensor interacts with only few other arbitrary sensors.

We propose an extremely simple clock synchronization (or an estimation) algorithm that is based on weighted averages, and prove that its performance on any given independent meeting pattern is highly competitive with that of the best possible algorithm, which operates without any resource or computational restrictions, and further knows the whole meeting pattern in advance. In particular, when all distributions involved are Gaussian, the performances of our scheme coincide with the optimal performances. Our proofs rely on an extensive use of the concept of Fisher information. We use the Cramér-Rao bound and our definition of a Fisher Channel Capacity to quantify information flows and to obtain lower bounds on collective performance. This opens the door for further rigorous quantifications of information flows within collaborative sensors.

7.2.3. Wait-freedom with advice

In [13], we motivate and propose a new way of thinking about failure detectors which allows us to define what it means to solve a distributed task wait-free using a failure detector. In our model, the system is composed of computation processes that obtain inputs and are supposed to produce outputs and synchronization processes that are subject to failures and can query a failure detector. Under the condition that correct (never failing) synchronization processes take sufficiently many steps, they provide the computation processes with enough advice to solve the given task wait-free: every computation process outputs in a finite number of its own

steps, regardless of the behavior of other computation processes. Every task can thus be characterized by the weakest failure detector that allows for solving it, and we show that every such failure detector captures a form of set agreement. We then obtain a complete classification of tasks, including ones that evaded comprehensible characterization so far, such as renaming or weak symmetry breaking.

7.2.4. Linear-space bootstrap communication schemes

In [12], we consider a system of n processes with ids that are drawn from a large space. How can these n processes communicate to solve a problem? It is shown that linear number of Multi-Writer Multi-Reader (MWMR) registers are sufficient to solve any read-write wait-free solvable problem and needed to solve some read-write wait-free solvable problem. This contrasts with the existing possible solution borrowed from adaptive algorithms that require $\Theta(n^{3/2})$ MWMR registers.

To obtain the sufficiency result, we show how the processes can non-blockingly emulate a system of n Single-Writer Multi-Reader (SWMR) registers on top of n Multi-Writer Multi-Reader (MWMR) registers. For the necessity result, we show it is impossible to do such an emulation with $n-1$ MWMR registers.

We also presents a wait-free emulation, using $2n-1$ rather than just n registers. The emulation can be used to solve an infinite sequence of tasks that are sequentially dependent (processes need the previous task's outputs in order to proceed to the next task). A non-blocking emulation cannot be used in this case, because it might starve a process forever.

7.2.5. Space complexity of set agreement

The k-set agreement problem is a generalization of the classical consensus problem in which processes are permitted to output up to k different input values. In a system of n processes, an m-obstruction-free solution to the problem requires termination only in executions where the number of processes taking steps is eventually bounded by m. This family of progress conditions generalizes wait-freedom ($m = n$) and obstruction-freedom ($m = 1$). In [29], we prove upper and lower bounds on the number of registers required to solve m-obstruction-free k-set agreement, considering both one-shot and repeated formulations. In particular, we show that repeated k set agreement can be solved using $n + 2m - k$ registers and establish a nearly matching lower bound of $n + m - k$.

7.2.6. Consensus capability of distributed systems

A fundamental research theme in distributed computing is the comparison of systems in terms of their ability to solve basic problems such as consensus that cannot be solved in completely asynchronous systems. In particular, in a seminal work (ACM Trans. Program. Lang. Syst. 13, 1991), Herlihy compares shared-memory systems in terms of the shared objects that they have: he proved that there are shared objects that are powerful enough to solve consensus for n processes, but are too weak to solve consensus for $n + 1$ processes; such objects are placed at level n of a wait-free hierarchy.

Similarly as in that work, in [30] we compare shared-memory systems with respect to their ability to solve consensus for n processes. But instead of comparing systems defined by the shared objects that they have, we compare read-write systems defined by the set of process schedules that can occur in these systems. Defining systems this way can capture many types of systems, e.g., systems whose synchrony ranges from fully synchronous to completely asynchronous, several systems with failure detectors, and "obstruction-free" systems. Here, we consider read-write systems defined in terms of sets of process schedules, and investigate the following fundamental question: Is there a system of $n + 1$ processes such that consensus can be solved for every subset of n processes in the system, but consensus cannot be solved for the $n + 1$ processes of the system? We show that the answer to the above question is "yes", and so these systems can be classified into a hierarchy akin to Herlihy's hierarchy.

7.2.7. Shared whiteboard models of distributed systems

In [4], we study distributed algorithms on massive graphs where links represent a particular relationship between nodes (for instance, nodes may represent phone numbers and links may indicate telephone calls).

Since such graphs are massive they need to be processed in a distributed way. When computing graph-theoretic properties, nodes become natural units for distributed computation. Links do not necessarily represent communication channels between the computing units and therefore do not restrict the communication flow. Our goal is to model and analyze the computational power of such distributed systems where one computing unit is assigned to each node. Communication takes place on a whiteboard where each node is allowed to write at most one message. Every node can read the contents of the whiteboard and, when activated, can write one small message based on its local knowledge. When the protocol terminates its output is computed from the final contents of the whiteboard. We describe four synchronization models for accessing the whiteboard. We show that message size and synchronization power constitute two orthogonal hierarchies for these systems. We exhibit problems that separate these models, i.e., that can be solved in one model but not in a weaker one, even with increased message size. These problems are related to maximal independent set and connectivity. We also exhibit problems that require a given message size independently of the synchronization model.

7.2.8. Discrete Lotka-Volterra population protocols

In [28], we focus on a natural class of population protocols whose dynamics are modeled by the discrete version of Lotka-Volterra equations with no linear term. In such protocols, when an agent a of type (species) i interacts with an agent b of type (species) j with a as the initiator, then b's type becomes i with probability P_{ij}. In such an interaction, we think of a as the predator, b as the prey, and the type of the prey is either converted to that of the predator or stays as is. Such protocols capture the dynamics of some opinion spreading models and generalize the well-known Rock-Paper-Scissors discrete dynamics. We consider the pairwise interactions among agents that are scheduled uniformly at random.

We start by considering the convergence time and show that any Lotka-Volterra-type protocol on a n-agent population converges to some absorbing state in time polynomial in n, w.h.p., when any pair of agents is allowed to interact. By contrast, when the interaction graph is a star, there exist protocols of the considered type, such as Rock-Paper-Scissors, which require exponential time to converge. We then study threshold effects exhibited by Lotka-Volterra-type protocols with 3 and more species under interactions between any pair of agents. We present a simple 4-type protocol in which the probability difference of reaching the two possible absorbing states is strongly amplified by the ratio of the initial populations of the two other types, which are transient, but "control" convergence. We then prove that the Rock-Paper-Scissors protocol reaches each of its three possible absorbing states with almost equal probability, starting from any configuration satisfying some sub-linear lower bound on the initial size of each species. That is, Rock-Paper-Scissors is a realization of a "coin-flip consensus" in a distributed system. Some of our techniques may be of independent value.

7.2.9. Deterministic load-balancing

In [23], we consider the problem of deterministic load balancing of tokens in the discrete model. A set of n processors is connected into a d-regular undirected network. In every time step, each processor exchanges some of its tokens with each of its neighbors in the network. The goal is to minimize the discrepancy between the number of tokens on the most-loaded and the least-loaded processor as quickly as possible. Rabani et al. (FOCS 1998) present a general technique for the analysis of a wide class of discrete load balancing algorithms. Their approach is to characterize the deviation between the actual loads of a discrete balancing algorithm with the distribution generated by a related Markov chain. The Markov chain can also be regarded as the underlying model of a continuous diffusion algorithm. Rabani et al. showed that after time $T = O(\log(Kn)/\mu)$, any algorithm of their class achieves a discrepancy of $O(d \log n/\mu)$, where μ is the spectral gap of the transition matrix of the graph, and K is the initial load discrepancy in the system.

In this work we identify some natural additional conditions on deterministic balancing algorithms, resulting in a class of algorithms reaching a smaller discrepancy. This class contains well-known algorithms, e.g., the rotor-router. Specifically, we introduce the notion of cumulatively fair load-balancing algorithms where in any interval of consecutive time steps, the total number of tokens sent out over an edge by a node is the same (up to constants) for all adjacent edges. We prove that algorithms which are cumulatively fair and where every node retains a sufficient part of its load in each step, achieve a discrepancy of $O(d\sqrt{\log n/\mu}, d\sqrt{n})$ in time $O(T)$. We also show that in general neither of these assumptions may be omitted without increasing discrepancy. We

then show by a combinatorial potential reduction argument that any cumulatively fair scheme satisfying some additional assumptions achieves a discrepancy of $O(d)$ almost as quickly as the continuous diffusion process. This positive result applies to some of the simplest and most natural discrete load balancing schemes.

7.2.10. Randomized local network computing

In [32], we have carried on investigating the line of research questioning the power of randomization for the design of distributed algorithms. In their seminal paper, Naor and Stockmeyer [STOC 1993] established that, in the context of network computing, in which all nodes execute the same algorithm in parallel, any construction task that can be solved locally by a randomized Monte-Carlo algorithm can also be solved locally by a deterministic algorithm. This result however holds in a specific context. In particular, it holds only for distributed tasks whose solutions that can be locally checked by a deterministic algorithm. We have extended the result of Naor and Stockmeyer to a wider class of tasks. Specifically, we proved that the same derandomization result holds for every task whose solutions can be locally checked using a 2-sided error randomized Monte-Carlo algorithm. This extension finds applications to, e.g., the design of lower bounds for construction tasks which tolerate that some nodes compute incorrect values. In a nutshell, we have showed that randomization does not help for solving such resilient tasks.

7.2.11. Proof-labeling schemes: randomization and self-stabilization

We have also carried on investigating the power of randomization for the design of proof-labeling schemes. Recall that a proof-labeling scheme, introduced by Korman, Kutten and Peleg [PODC 2005], is a mechanism enabling to certify the legality of a network configuration with respect to a boolean predicate. Such a mechanism finds applications in many frameworks, including the design of fault-tolerant distributed algorithms. In a proof-labeling scheme, the verification phase consists of exchanging labels between neighbors. The size of these labels depends on the network predicate to be checked. There are predicates requiring large labels, of poly-logarithmic size (e.g., MST), or even polynomial size (e.g., Symmetry). In [22], we introduce the notion of randomized proof-labeling schemes. By reduction from deterministic schemes, we show that randomization enables the amount of communication to be exponentially reduced. As a consequence, we show that checking any network predicate can be done with probability of correctness as close to one as desired by exchanging just a logarithmic number of bits between neighbors. Moreover, we design a novel space lower bound technique that applies to both deterministic and randomized proof-labeling schemes. Using this technique, we establish several tight bounds on the verification complexity of classical distributed computing problems, such as MST construction, and of classical predicates such as acyclicity, connectivity, and cycle length.

Next, we have established the formal connections between self-stabilization and proof-labeling scheme. Recall that self-stabilizing algorithms are distributed algorithms supporting transient failures. Starting from any configuration, they allow the system to detect whether the actual configuration is legal, and, if not, they allow the system to eventually reach a legal configuration. In the context of network computing, it is known that, for every task, there is a self-stabilizing algorithm solving that task, with optimal space-complexity, but converging in an exponential number of rounds. On the other hand, it is also known that, for every task, there is a self-stabilizing algorithm solving that task in a linear number of rounds, but with large space-complexity. It is however not known whether for every task there exists a self-stabilizing algorithm that is simultaneously space-efficient and time-efficient. In [24], we make a first attempt for answering the question of whether such an efficient algorithm exists for every task, by focussing on constrained spanning tree construction tasks. We present a general roadmap for the design of silent space-optimal self-stabilizing algorithms solving such tasks, converging in polynomially many rounds under the unfair scheduler. By applying our roadmap to the task of constructing minimum-weight spanning tree (MST), and to the task of constructing minimum-degree spanning tree (MDST), we provide algorithms that outperform previously known algorithms designed and optimized specifically for solving each of these two tasks.

7.2.12. Role of node identifiers in local decision

We have also investigated the role of IDs in network computing. This role is well understood as far as symmetry breaking is concerned. However, the unique identifiers also leak information about the computing environment

— in particular, they provide some nodes with information related to the size of the network. It was recently proved that in the context of local decision, there are some decision problems such that (1) they cannot be solved without unique identifiers, and (2) unique node identifiers leak a sufficient amount of information such that the problem becomes solvable. In [33] we study what is the minimal amount of information that we need to leak from the environment to the nodes in order to solve local decision problems. Our key results are related to scalar oracles f that, for any given n, provide a multi-set $f(n)$ of n labels; then the adversary assigns the labels to the n nodes in the network. This is a direct generalization of the usual assumption of unique node identifiers. We give a complete characterization of the weakest oracle that leaks at least as much information as the unique identifiers. Our main result is the following dichotomy: we classify scalar oracles as large and small, depending on their asymptotic behavior, and show that (1) any large oracle is at least as powerful as the unique identifiers in the context of local decision problems, while (2) for any small oracle there are local decision problems that still benefit from unique identifiers.

7.2.13. Geometry on the utility space

In [31], we study the geometrical properties of the utility space (the space of expected utilities over a finite set of options), which is commonly used to model the preferences of an agent in a situation of uncertainty. We focus on the case where the model is neutral with respect to the available options, i.e. treats them, a priori, as being symmetrical from one another. Specifically, we prove that the only Riemannian metric that respects the geometrical properties and the natural symmetries of the utility space is the round metric. This canonical metric allows to define a uniform probability over the utility space and to naturally generalize the Impartial Culture to a model with expected utilities.

7.3. Network Algorithms and Analysis

7.3.1. Information dissemination on social networks

In [17], we model an online social network as a network formation game. We study convergence of selfish dynamics and show that somewhat natural metric assumption enable fast convergence towards an equilibrium with efficient collaborative filtering of content.

7.3.2. Verification of network forwarding tables

In [25], we investigate the problem of verifying forwarding network tables. We show that it is sufficient to test few representative headers when the set of rules applied by routers is complete under intersection.

7.3.3. Refreshing old datasets in a network: LiveRank

In [18], we consider the problem of refreshing a dataset. More precisely, given a collection of nodes gathered at some time (Web pages, users from an online social network) along with some structure (hyperlinks, social relationships), we want to identify a significant fraction of the nodes that still exist at present time. The liveness of an old node can be tested through an online query at present time. We call LiveRank a ranking of the old pages so that active nodes are more likely to appear first. The quality of a LiveRank is measured by the number of queries necessary to identify a given fraction of the active nodes when using the LiveRank order. We study different scenarios from a static setting where the LiveRank is computed before any query is made, to dynamic settings where the LiveRank can be updated as queries are processed. Our results show that building on the PageRank can lead to efficient LiveRanks, for Web graphs as well as for online social networks.

7.3.4. Exploiting user movement for position detection

The major issue of indoor localization system is the trade-off between implementation cost and accuracy. A low-cost system which demands only few hardware devices could save the cost but often it turns out to be less reliable. Aiming at improving classical triangulation method that requires several reference points, we propose in [34] a new method, called Two-Step Movement (2SM), which requires only one reference point (RP) by exploiting useful information given by the position change of a mobile terminal (MT), or the user movement. This method can minimize the number of reference points required in a localization system or

navigation service and reduce system implementation cost. Analytical result shows that the user position can be thus derived and given in simple closed-form expression. Finally, simulation is conducted to demonstrate its effectiveness under noisy environment.

Then, in [35], we build on 2SM. We first improve the positioning performance through multi-sampling technique to combat measurement noise. Secondly, we propose the Generalized Two-Step Movement (G2SM) method for device-to-device (D2D) systems in which both the mobile terminal (MT) and RP can be mobile device. The mobile user's position can be derived analytically and given in simple closed-form expression. Its effectiveness in the presence of noise is shown in simulation results.

7.3.5. Fast diameter and radius computation in real-world graphs

In [5], we propose a new algorithm that computes the radius and the diameter of a weakly connected digraph $G = (V, E)$, by finding bounds through heuristics and improving them until they are validated. Although the worst-case running time is $O(|V||E|)$, we will experimentally show that it performs much better in the case of real-world networks, finding the radius and diameter values after 10–100 runs of Breadth First Search instead of $|V|$ BFS-s (independently of the value of $|V|$), and thus having running time $O(|E|)$ in practice. As far as we know, this is the first algorithm able to compute the diameter of weakly connected digraphs, apart from the naive algorithm, which runs in time $\Omega(|V||E|)$ performing a BFS from each node. In the particular cases of strongly connected directed or connected undirected graphs, we have compared our algorithm with known approaches by performing experiments on a dataset composed by several real-world networks of different kinds. These experiments show that, despite its generality, the new algorithm outperforms all previous methods, both in the radius and in the diameter computation, both in the directed and in the undirected case, both in average running time and in robustness. Finally, as an application example, we have used the new algorithm to determine the solvability over time of the "Six Degrees of Kevin Bacon" game, and of the "Six Degrees of Wikipedia" game. As a consequence, we have computed for the first time the exact value of the radius and the diameter of the whole Wikipedia digraph.

LIFEWARE Project-Team

7. New Results

7.1. Hybrid Simulation of Heterogeneous Biochemical Models in SBML

Participants: Katherine Chiang, François Fages, Sylvain Soliman.

Models of biochemical systems presented as a set of formal reaction rules can be interpreted in different for- malisms, most notably as either deterministic Ordinary Differential Equations, stochastic continuous-time Markov Chains, Petri nets or Boolean transition systems. While the formal composition of reaction systems can be syntactically defined as the (multiset) union of the reactions, the composition and simulation of models in different formalisms remains a largely open issue. In [5], we show that the combination of reaction rules and events, as already present in SBML, can be used in a non-standard way to define stochas- tic and boolean simulators and give meaning to the hybrid composition and simulation of heterogeneous models of biochemical processes. In particular, we show how two SBML reaction models can be composed into one hybrid continuous-stochastic SBML model through a high-level interface for composing reaction models and specifying their interpretation. Furthermore, we describe dynamic strategies for automatically partitioning reactions with stochastic or continuous interpretations according to dynamic criteria. The per- formances are then compared to static partitioning. The proposed approach is illustrated and evaluated on several examples, including the reconstructions of the hybrid model of the mammalian cell cycle regulation of Singhania et al. as the composition of a Boolean model of cell cycle phase transitions with a continuous model of cyclin activation, the hybrid stochastic-continuous models of bacteriophage T7 infection of Alfonsi et al., and the bacteriophage λ model of Goutsias, showing the gain in both accuracy and simulation time of the dynamic partitioning strategy.

7.2. Theoretical and Practical Complexities of Enumerating Minimal Siphons in Petri Nets

Participants: François Fages, Thierry Martinez, Sylvain Soliman.

Petri nets are a simple formalism for modeling concurrent computa- tion. They are also an interesting tool for modeling and analysing biochemical reaction systems, bridging the gap between purely qualitative and quantita- tive models. Biological networks can indeed be complex, large, and with many unknown kinetic parameters, which makes the development of quantitative models difficult. In [9], we focus on the Petri net representation of biochemical reactions and on two structural properties of Petri nets, siphons and traps, that bring us information about the persistence of some molecular species, independently of the kinetics. We first study the theoretical time com- plexity of minimal siphon decision problems in general Petri nets, and present three new complexity results: first, we show that the existence of a siphon of a given cardinality is NP-complete; second, we prove that deciding the Siphon- Trap property is co-NP-complete; third, we prove that deciding the existence of a minimal siphon containing a given set of places, deciding the existence of a siphon of a given cardinality and deciding the Siphon-Trap property can be done in linear time in Petri nets of bounded tree-width. Then, we present a Boolean model of siphons and traps, and two methods for enumerating all minimal siphons and traps of a Petri net, by using a SAT solver and a Con- straint Logic Program (CLP) respectively. On a benchmark of 345 Petri nets of hundreds of places and transitions, extracted from biological models from the BioModels repository, as well as on a benchmark composed of 80 Petri nets from the Petriweb database of industrial processes, we show that both the SAT and CLP methods are overall faster by one or two orders of magnitude com- pared to the state-of-the-art algorithm from the Petri net community, and are in fact able to solve all the enumeration problems of our practical benchmarks. We investigate why these programs perform so well in practice, and provide some elements of explanation related to our theoretical complexity results.

7.3. Abstraction-based Parameter Synthesis for Multiaffine Systems

Participant: Grégory Batt.

Multiaffine hybrid automata (MHA) represent a powerful formalism to model complex dynamical systems. This formalism is particularly suited for the representation of biological systems which often exhibit highly non-linear behavior. In [10], we consider the problem of parameter identification for MHA. We present an abstraction of MHA based on linear hybrid automata, which can be analyzed by the SpaceEx model checker. This abstraction enables a precise handling of time-dependent properties. We demonstrate the potential of our approach on a model of a genetic regulatory network and a myocyte model.

7.4. Tropical Algebra Methods for Model Reduction

Participants: François Fages, Jonas Sénizergues, Sylvain Soliman.

Jonas Sénizergues has just started a PhD Thesis on the design of model reduction techniques for systems biology based on tropical algebra. The idea is to reason on the orders of magnitude of both kinetic parameters and molecular concentrations in order to determine particular regimes exhibiting fast-slow decomposition and amenable to model reductions. Such model reductions generalize the quasi steady-state (QSSA) and quasi-equilibrium (QE) criteria, and lead to hybrid automata for chaining the reduced dynamics. The solving of tropical equilibration equations rely on previous work using constraint programming techniques [0] with collaboration with Ovidier Radulescu (Univ. Montpellier) and Andreas Weber (University of Bonn, Germany).

7.5. Modeling the Effect of the Cell Cycle on the Circadian Clock in Mouse Embryonic Fibroblasts

Participants: François Fages, Jonas Sénizergues, Denis Thieffry, Pauline Traynard, Sylvain Soliman.

Experimental observations have put in evidence autonomous self-sustained circadian oscillators in most mammalian cells, and proved the existence of molecular links between the circadian clock and the cell cycle. Several models have been elaborated to assess conditions of control of the cell cycle by the circadian clock, in particular through the regulation by clock genes of Wee1, an inhibitor of the mitosis promoting factor, responsible for a circadian gating of mitosis and cell division period doubling phenomena. However, recent studies in individual NIH3T3 fibroblasts have shown an unexpected acceleration of the circadian clock together with the cell cycle when the milieu is enriched in FBS, the absence of such acceleration in confluent cells, and the absence of any period doubling phenomena. In [14], we try to explain these observations by a possible entrainment of the circadian clock by the cell cycle through the inhibition of transcription during mitosis. We develop a differential model of that reverse coupling of the cell cycle and the circadian clock and investigate the conditions in which both cycles are mutually entrained. We use the mammalian circadian clock model of Relogio et al. and a simple model of the cell cycle by Qu et al. which focuses on the mitosis phase. We show that our coupled model is able to reproduce the main observations reported by Feillet et al. in individual fibroblast experiments and use it for making some predictions. In [17], those hypothesis are revised in order to reproduce the phase data in addition to the period data and make new predictions.

7.6. Effects of repeated osmotic stress on gene expression and growth: from cell-to-cell variability to cellular individuality in the budding yeast Saccharomyces cerevisiae

Participants: Grégory Batt, Ewen Corre, Pascal Hersen, Artémis Llamosi.

[0] Sylvain Soliman, François Fages, Ovidiu Radulescu. A constraint solving approach to model reduction by tropical equilibration. Algorithms for Molecular Biology, 9(24), 2014.

When shifted to a stressful environment, cells are capable of complex response and adaptations. Although the cellular response to a single stress has been studied in great detail, very little is known when it comes to dynamically fluctuating stressful environments. In addition, in the context of stress response, the role of cell-to-cell variability in cellular processes and more specifically in gene expression is still unclear.

In his PhD thesis [3], Artémis Llamosi uses a systems and synthetic biology approach to investigate osmotic stress in *S. cerevisiae* at the single cell level. Combining microfluidics, fluorescent microscopy and advanced image analysis, we are able to subject cells to precise fluctuating osmolarity and monitor single-cell temporal response.

While much previous research in gene expression heterogeneity focused on its stochastic aspect, we consider here long-lasting differences between cells regarding expression kinetics. Using population models and state-of-the-art statistical analysis, we manage to represent both population and single-cell dynamics in a single concise modelling framework. This quantitative approach capturing stable individuality in gene expression dynamics can define a form of non-genetic cellular identity.

To improve our understanding of the biological interpretation of such identity, we investigate the relation between single-cell specificities in their gene expression with their phenotype and micro-environment. We then take a lineage based perspective and find this form of identity to be partially inherited.

Understanding the evolutionary consequences of inheritable non-genetic cellular identity requires a better knowledge of the impact of fluctuating stress on cell proliferation. Dissecting quantitatively the consequences of repeated stress on cell-cycle and growth gives us an overview of the energetic and temporal consequences of repeated stress. At last, technical and theoretical developments needed to carry this investigation further are presented.

7.7. Resistance to anti-cancer drugs by non-mutational mechanisms: insights from a cell-based multi-scale model of TRAIL-induced apoptosis

Participants: Virgile Andréani, Grégory Batt, François Bertaux.

The fact that tumors can acquire drug resistance by non-mutational mechanisms is increasingly gaining attention (Sharma et al., 2010; Pisco et al., 2013; Flusberg et al., 2013). Stochastic fluctuations in cellular states of different resistance and proliferative potential could play an important role in such resistance acquisition. Thus, to enable a quantitative, molecular-level understanding of those phenomena, modeling approaches that go beyond traditional, deterministic kinetic models of biological pathways are required.

An interesting and well-studied example of non-mutational resistance acquisition concerns the response of cancer cells to the agent TRAIL, a selective inducer of apoptotic cell death. In a previous work (Bertaux et al., 2014), we have developped a single-cell model of TRAIL-induced apoptosis that accounts for (1) protein-protein signaling reactions linking TRAIL exposure to commitment to apoptosis, (2) stochastic gene expression for the proteins involved in this signaling and (3) protein degradation. Under parsimonious and realistic assumptions for parameter values, fractional killing and transient resistance acquisition readily emerged from model simulations. Those two properties relating to TRAIL resistance are observed in-vitro for many different cancer cell lines.

Here, again in collaboration with Dirk Drasdo and Szymon Stoma, we investigate the long-term response of proliferating cancer cell populations repeatedly treated by TRAIL by integrating our single-cell model of TRAIL-induced apoptosis into a multi-cellular simulation framework. We predict that the long-term killing efficiency of repeated treatments is strongly reduced compared to the first treatment. A detailed analysis showed that resistance acquisition is caused mainly by the targeted degradation of activated pro-apoptotic proteins and an imbalance between the turnover of pro- and anti- apoptotic proteins. In addition, simulations of the treatment of multi-cellular spheroids suggested that limited TRAIL penetration is unlikely to be a driving cause of resistance, but that it can exacerbate the impact of cell-intrinsic resistance acquisition.

7.8. Controlling a genetic inverted pendulum

Participants: Grégory Batt, Catherine Eisenhauer, Pascal Hersen, Jean-Baptiste Lugagne.

The ability to routinely control complex genetic circuits in vivo and in real-time promises quantitative understanding of cellular processes of unprecedented precision, quality, and richness. With combined efforts in microfluidic design, microscope automation, image segmentation and analysis, and control theory, we propose a platform for real-time, single-cell, externalized in silico control and monitoring of genetic networks in E. coli. Computational framework and hardware are optimized for parallelizing the experiments and we use the platform to test and control an entire library of synthetic genetic circuits. The circuits we are trying to control are based on the genetic toggle switch, a foundational circuit in synthetic biology, which consists of two genes that repress each other. This genetic system features two stable equilibrium points where one of the genes has taken over. Our objective is to dynamically balance the circuit in single cells around a third, unstable equilibrium point at which no gene dominates and their mutual repression strengths are balanced. This is similar to the landmark problem in control theory of stabilizing an inverted pendulum. Although our work indicates that this real-time control approach can drive convoluted genetic networks towards states that are inaccessible to traditional genetic perturbations such as knock-outs and promoter induction, the a priori quantitative knowledge of the system required for achieving this control is minimal. We show that even a simple Proportional-Integral controller can stabilize the unstable point of the toggle switch in single cells. Finally, we demonstrate that manipulation, or even inversion, of the stability map of the network is possible, though counter intuitive, via the simultaneous stabilization of an entire population of toggle switch cells around their unstable point with a common dynamic input.

7.9. Synthesizing Configurable Biochemical Implementation of Linear Systems from Their Transfer Function Specifications

Participants: Katherine Chiang, François Fages, Sylvain Soliman.

The ability to engineer synthetic systems in the biochemical context is constantly being improved and has a profound societal impact. Linear system design is one of the most pervasive methods applied in control tasks, and its biochemical realization has been proposed by Oishi and Klavins and advanced further in recent years. However, several technical issues remain unsolved. Specifically, the design process is not fully automated from specification at the transfer function level, systems once designed often lack dynamic adaptivity to environmental changes, matching rate constants of reactions is not always possible, and implementation may be approximative and greatly deviate from the specifications. In [6], building upon the work of Oishi and Klavins, we overcome these issues by introducing a design flow that transforms a transfer-function specification of a linear system into a set of chemical reactions, whose input-output response precisely conforms to the specification. This system is implementable using the DNA strand displacement technique. The underlying configurability is embedded into primitive components and template modules, and thus the entire system is adaptive. Simulation of DNA strand displacement implementation confirmed the feasibility and superiority of the proposed synthesis flow.

7.10. Reconfigurable Neuromorphic Computation in Biochemical Systems

Participants: Katherine Chiang, François Fages.

Implementing application-specific computation and con- trol tasks within a biochemical system has been an important pursuit in synthetic biology. Most synthetic designs to date have focused on realiz- ing systems of fixed functions using specifically engineered components, thus lacking flexibility to adapt to uncertain and dynamically-changing environments. To remedy this limitation, an analog and modularized approach to realize reconfigurable neuromorphic computation with biochemical reactions is presented in [11]. We propose a biochemical neural network consisting of neuronal modules and interconnects that are both reconfigurable through external or internal control over the concentrations of certain molecular species. Case studies on classification and machine learning applications using the DNA strain displacement technology demonstrate the effectiveness of our design in both reconfiguration and autonomous adaptation.

7.11. Search by Constraint Propagation

Participants: François Fages, Thierry Martinez, Sylvain Soliman.

Constraint programming is traditionally presented as the combination of two components: a constraint model and a search procedure. In [13] we show that tree search procedures can be fully internalized in the constraint model with a fixed enumeration strategy. This approach has several advantages: 1) it makes search strategies declarative, and modeled as constraint satisfaction problems; 2) it makes it possible to express search strategies in existing front-end modeling languages supporting reified constraints without any extension; 3) it opens up constraint propagation algorithms to search constraints and to the implementation of novel search procedures based on constraint propagation. We illustrate this approach with a Horn clause extension of the MiniZinc modeling language and the modeling in this language of a variety of search procedures, including dynamic symmetry breaking procedures and limited discrepancy search, as constraint satisfaction problems. We show that this generality does not come with a significant overhead, and can in fact exhibit exponential speedups over procedural implementations, thanks to the propagation of the search constraints.

7.12. Execution models for Constraint Programming and Semantics Equivalence

Participants: François Fages, Thierry Martinez, Sylvain Soliman.

Logic programming and constraint programming are two declarative programming paradigms which rely on the identification of programs to theories, and programming to modeling. Execution models result from the operational interpretation of logical provability in logic programming, and of constraint propagation in constraint programming. However, the control of execution is crucial for the practicability of these schemes and extra-logical traits are thus added in those programming systems, with the classical slogans "logic program = logical theory + control", "constraint program = constraint model + search".

In his thesis [4], Thierry Martinez investigates execution models in which control and search can be shifted into the logic or the constraint model, while preserving the semantics. The three parts of the thesis correspond to the three semantics equivalence that are showed: the first between two committed-choice forward-chaining logic languages, the second between constraint logic programs and constraint models, and the third between guard semantics in angelic settings. Each of these equivalence is constructive in the sense that there exists an encoding that enables the compilation from one of the paradigm to the other.

7.13. On Translating MiniZinc Constraint Model into Fitness Functions: Application to Continuous Placement Problems.

Participants: François Fages, Thierry Martinez, Bao Duy Tran.

MiniZinc is a solver-independent constraint modeling lan- guage which is increasingly used in the constraint programming com- munity. It can be used to compare different solvers which are currently based on either constraint programming, Boolean satisfiability or mixed integer linear programming. In [12], we show how MiniZinc models can be compiled into fitness functions for evolutionary algorithms. More specifically, we describe the translation of FlatZinc models into fitness functions over the reals and their use in the Covariance Matrix Adapta- tion Evolution Strategy (CMA-ES) solver. We illustrate this approach, and evaluate it, on the modeling and solving of complex shape continuous placement problems.

MAMBA Project-Team

7. New Results

7.1. Cancer

Participants: Luís Lopes Neves de Almeida, Rebecca Chisholm, Jean Clairambault, François Delhommeau [Haematology department, St Antoine Hospital, Paris], Dirk Drasdo, Ján Eliaš, Alexandre Escargueil [Cancer biology and therapeutics lab, St Antoine Hospital, Paris], Ghassen Haddad [ENIT, Tunis], Shalla Hanson [Department of mathematics, Duke University, Durham, NC], Pierre Hirsch [Haematology department, St Antoine Hospital, Paris], Groups Invade, Lungsysii, Tim Johann, Group Klingmueller [German Cancer Center, Heidelberg], Michal Kowalczyk [Univ. Santiago de Chile], Annette Larsen [Cancer biology and therapeutics lab, St Antoine Hospital, Paris], Tommaso Lorenzi, Alexander Lorz, Benoît Perthame, Andrada Quillas Maran, Fernando Quirós [Univ. Autónoma de Madrid], Michèle Sabbah [Cancer biology and therapeutics lab, St Antoine Hospital, Paris], Min Tang [Jiaotong University, Shanghai], Emmanuel Trélat [LJLL, UPMC], Paul Van Liedekerke, Nicolas Vauchelet, Irène Vignon-Clementel [REO], Yi Yin.

7.1.1. Drug resistance

We have continued to develop our phenotypically based models of drug-induced drug resistance in cancer cell populations, representing their Darwinian or Lamarckian evolution under drug pressure by integro-differential equations. In one of them [23], a 1D space variable has been added to the phenotypic structure variable to account for drug diffusion in tumour spheroids. In another one, focusing on both Darwinian selection and Lamarckian-like (non-genetic) instruction, published in *Cancer Research* [41], where deterministic and agent-based modelling are processed in parallel, we have added advection and diffusion terms to the initial integro-differential model and considered a physiologically based 2-dimensional phenotypic structure variable. This model, designed to take account of previously published biological observations on (reversible) drug tolerance persistence in a cultured population of non-small cell lung cancer (NSCLC) cells [90], reproduces the observations and we propose to assess the model by testing biologically based hypotheses. This work, also presented in various conferences ([34], [35], [31]) is conducted in close collaboration with the INSERM-UPMC team "Cancer biology and therapeutics" (A. Larsen, A. Escargueil, M. Sabbah) at St Antoine Hospital. It has also led our postdoctoral fellows Rebecca Chisholm and Tommaso Lorenzi to prolong their work on the *Cancer Research* paper by publishing two more articles [21], [48], one of which is a joint work with Alexander Lorz. This work is currently continued from the point of view of optimal control in Camille Pouchol's PhD thesis.

7.1.2. Evolution and cancer, therapy optimisation

Guided by our goal to understand and overcome drug resistance in cancer cell populations[41], we are considering cancer as an evolutionary phenomenon at two time scales: a large time scale (billions of years) of evolution of the genomes, from unicellular organisms to organised multicellularity (viewing cancer as more an archeoplasm than a neoplasm, an evolution backwards, following Davies and Lineweaver, Phys Biol 2011, and others [78], [66], [92], [79]) with shortcomings due to malfunctions in the processes of control of cell differentiation, and a short time scale (duration of a human life) of evolution in the "epigenetic landscape" of a given genome (as advocated by Sui Huang and Angela Pisco, e.g. recently in Nature, Br J Cancer and elsewhere [76], [77], [85], [86], [94]). It leads us to propose theoretical frameworks for innovative cancer therapeutics from this evolutionary biology viewpoint, taking into account the major clinical issue of drug resistance in cancer cell populations, as presented in [31] and exposed to a medical audience at the symposium "Réseau Cancer des Points Cardinaux" (http://www.frog-oncogeriatrie.com/fichiers/evnmt_41.pdf).

7.1.3. Interactions between tumour cell populations and their cellular micro-environments

A phenotype-structured model of the interactions between a brest cancer cell population (MCF7 cultured cells, collaboration with M. Sabbah, St Antoine Hospital) and its adipocyte stroma support cell population has been developed (T. Lorenzi, C. Pouchol, J. Clairambault) in the framework of Camille Pouchol's Inria internship ([56]). It has led to hiring C. Pouchol as a PhD student at UPMC (on a university grant "Interfaces pour le Vivant") on the same subject with perspectives in optimal therapeutic control, under the supervision of J. Clairambault, M. Sabbah and E. Trélat, see below "Supervision".

7.1.4. Combining chemo- and immunotherapies

Both from the point of view of interactions with the tumour micro-environment and of innovative anticancer therapies, it is necessary to take into account the immune response in cancer. This recently developed activity, (illustrated by presentations in session 70 in ICNAAM 2016 [32]) has led to the involvement in 2015 of Shalla Hanson as a PhD student in co-tutela between Duke University, NC and UPMC, see below "Supervision".

7.1.5. Hele-Shaw model of tumour growth

The mathematical analysis of macroscopic models of tumor growth with one type of cancer cells has been continued. On the one hand, in [47], the concept of viscosity solutions has been implemented for the case with active motion. On the other hand, the regularity of the free boundary is proved in [51] using methods developed for the standard Hele-Shaw equation and a new formulation.

7.1.6. The p53 protein spatio-temporal dynamics

Our previously developed spatio-temporal models for an intracellular dynamical response of the p53 protein to DNA damage, have been exploited further, and several testable biological hypotheses have been proposed in [33]. Among them, we suggest ideas that link spatio-temporal location of the p53 protein with a specific cell fate of a single cell in [33], [2] and, based on our new oscillator relying on both positive and negative regulation of p53 by Mdm2 (in tight cooperation with MdmX), we provide molecular insights into an excitability of the p53 network, i.e., we propose a molecular explanation for a full pulsatile response of p53 independently of input (ATM) signalling, challenging thus different fates of ATM downstream targets in the regulation of p53 in response to different stimuli, such as γ- and UV-radiation.

Our mathematical models, all included in J. Eliaš's PhD thesis [2] (defended on the 1st September 2015), contribute to understanding the variability of p53 in response to single and double strand breaks and reveal some new aspects of the core p53-Mdm2 protein feedback.

7.1.7. Lung and breast cancer

We developed an image analysis software and designed image analysis pipelines which we used to quantify the invasion pattern of non-small cell lung cancer (NSCLC) cells in multicellular spheroid *in vitro* experiments [24]. Based on the analyses, we demonstrated that the concomitant over-expression of FIR (far upstream element binding protein interacting repressor) and its splice variants drives NSCLC migration and dissemination.

We developed an agent-based, centre-based model of cell migration in cancer invasion based upon experimental observations of cell shape and cell behaviour in multicellular spheroid experiments of breast and lung cancer cells. In these experiments, cells deform from a sphere into an oblong shape upon migration, and adopt a spherical shape again whenever they turn back to such spheroids. This was implemented. Moreover, we developed a 3D model for the extracellular matrix (ECM) in which the matrix is modelled by an irregular network of springs with nodes represented as elastic objects. Migrating cells anchor in the network to move, leading to network deformation. We implemented a number of different biological mechanisms of cell migration and cell-ECM interaction. We find that a relatively simple model is sufficient to explain all phenomena of a single invading cell (Palm et. al., in preparation).

The combination of image analysis and of the abovementioned refined invasion model should allow a quantitative model of multicellular invasion following the same line of research as for SK-MES-1 cells, where we inferred a multicellular spheroid growth model from image data within a pipeline of experiment, imaging, image analysis and modelling [17]. In that paper, we used spatial-temporal image data of cell nucleus distribution, cell proliferation, death, and ECM distribution for two growth conditions (oxygen and glucose) to calibrate a model which was then able to quantitatively correctly predict the growth kinetics of the tumor spheroids for two other growth conditions, one strongly glucose limited, another strongly oxygen-limited.

Finally, we developed an image analysis pipeline to estimate the number of cancer cells in a patient with non-small cell lung cancer (NSCLC) from non-invasive image modalities. The estimate bases upon cell counts from histological serial sections of the tumor which have been related to the D-value inferred from Diffusion Weighted (DW) MRI (Yi et. al., paper in preparation).

7.2. Aggregation Kinetics

Participants: Aurora Armiento, Tom Banks [CRSC, NCSU, Raleigh, USA], Thibault Bourgeron, José Antonio Carrillo [Imperial College, London, United Kingdom], Marie Doumic, Miguel Escobedo [Universidad del País Vasco, Bilbao, Spain], Sarah Eugène, Marc Hoffmann [Ceremade, Université Paris-Dauphine], François James [MAPMO, Université d'Orléans], Nathalie Krell [Université de Rennes 1], Carola Kruse, Frédéric Lagoutière [Département de mathématiques d'Orsay], Philippe Moireau [Inria Paris Saclay, M3DISIM project-team], Benoît Perthame, Stéphanie Prigent, Human Rezaei [VIM, INRA Jouy-en-Josas], Lydia Robert [Laboratoire Jean Perrin, UPMC], Philippe Robert [Inria Paris, RAP project-team], Maria Teresa Teixeira [IBCP, Paris], Nicolas Vauchelet, Min Tang [Jiaotong University, Shanghai], Zhou Xu [IBCP, Paris], Wei-Feng Xue [University of Kent, United Kingdom].

7.2.1. Heterogeneity as an intrinsic feature in biological dynamics

Combining deterministic and probabilistic approaches, we investigated in two different applications - namely senescence and protein aggregation - the impact of heterogeneity on dynamical features of the considered populations.

Yeast Senescence and Telomere replication In eukaryotes, the absence of telomerase results in telomere shortening, eventually leading to replicative senescence, an arrested state that prevents further cell divisions. While replicative senescence is mainly controlled by telomere length, the heterogeneity of its onset is not well understood. Insights on this key question may have consequences both for cancer and aging issues.

In collaboration with T. Teixeira and Z. Xue from IBCP, we proposed a mathematical model based on the molecular mechanisms of telomere replication and shortening to decipher the causes of this heterogeneity [7]. Using simulations fitted on experimental data obtained from individual lineages of senescent Saccharomyces cerevisiae cells, we decompose the sources of senescence heterogeneity into interclonal and intraclonal components, and show that the latter is based on the asymmetry of the telomere replication mechanism. We also evidence telomere rank-switching events with distinct frequencies in short-lived versus long-lived lineages, revealing that telomere shortening dynamics display important variations. Thus, the intrinsic heterogeneity of replicative senescence and its consequences find their roots in the asymmetric structure of telomeres.

These promising first results lead us to an ongoing collaboration, and hopefully will allow still more insight on complex mechanisms not yet modelled mathematically.

Variability in nucleated polymerisation

The kinetics of amyloid assembly show an exponential growth phase preceded by a lag phase, variable in duration as seen in bulk experiments and experiments that mimic the small volumes of cells. To investigate the origins and the properties of the observed variability in the lag phase of amyloid assembly currently not accounted for by deterministic nucleation dependent mechanisms, we formulated a new stochastic minimal model that is capable of describing the characteristics of amyloid growth curves despite its simplicity [44]. We then solved the stochastic differential equations of our model and gave a mathematical proof of a central limit theorem for the sample growth trajectories of the nucleated aggregation process. These results

give an asymptotic description for our simple model, from which closed-form analytical results capable of describing and predicting the variability of nucleated amyloid assembly were derived. We also demonstrated the application of our results to inform experiments in a convenient and clear way. Our model offers a new perspective and paves the way for a new and efficient approach on extracting vital information regarding the key initial events of amyloid formation.

7.2.2. Inverse Problems and Data Assimilation Applied to Protein Aggregation and other settings

As mathematical models become more complex with multiple states and many parameters to be estimated using experimental data, there is a need for critical analysis in model validation related to the reliability of parameter estimates obtained in model fitting. This leads to a fundamental question: how much information with respect to model validation can be expected in a given data set or collection of data sets?

In the biological context of amyloid formation, the question is to quantify to which extent a given model may be appropriately fitted and selected for, given relatively sparse data. Estimating reaction rates and size distributions of protein polymers is an important step towards understanding the mechanisms of protein misfolding and aggregation, a key feature for amyloid diseases. Specifically, experimental measurements often consist in the time-dynamics of a moment of the population (*i.e.*, for instance the total polymerised mass, as in Thioflavine T measurements, or the second moment measured by Static Light Scattering).

In a first study [4], in collaboration with H.T. Banks and H. Rezaei, we illustrated the use of tools (asymptotic theories of standard error quantification using appropriate statistical models, bootstrapping, model comparison techniques) in addition to sensitivity that may be employed to determine the information content in data sets. We do this in the context of recent models [87] for nucleated polymerisation in proteins, about which very little is known regarding the underlying mechanisms; thus the methodology we developed may be of great help to experimentalists.

In another study [39], related to a different biological setting (the frog olfactive tract), we use a method based on the Mellin transform, as in [64], to solve a spectral inverse problem arising from the modeling of the transduction of an odor into an electrical signal. The problem is to find the spatial distribution of CNG ion channels along the cilium of a frog, which allow a depolarizing influx of sodium ions, which initiate the electrical signal. This problem comes down to solving a Fredholm integral equation. We prove observability and continuity inequalities by estimating the Mellin transform of the kernel of this integral equation. We perform numerical computations using experimental data.

To get more insight into the estimation of reaction rates and size distributions of protein polymers, we are now developing an approach based on a data assimilation strategy. In this purpose, A. Armiento's Ph.D is focused on setting this framework problem when the experimental measurements consist in the time-dynamics of a moment of the population (*i.e.* for instance the total polymerised mass, as in Thioflavine T measurements, or the second moment measured by Static Light Scattering). In [37] we proposed a general methodology, and we solved the problem theoretically and numerically in the case of a depolymerising system. We then applied our method to experimental data of degrading oligomers, and conclude that smaller aggregates of ovPrP protein should be more stable than larger ones. This has an important biological implication, since it is commonly admitted that small oligomers constitute the most cytotoxic species during prion misfolding process.

7.2.3. Time asymptotics for growth-fragmentation equations

The long-term dynamics of fragmentation and growth-fragmentation equations has been for long an important research field for BANG then MAMBA research team. Thanks to these common efforts, these equations are now well understood. However, there remain some interesting open questions. In particular, if the generic long-time behaviour for the linear equation is known - given by a (generally exponential) trend towards a steady exponential growth described by the positive eigenvector linked to the dominant eigenvalue, see [84] for most recent results - critical cases are not yet fully understood.

With Miguel Escobedo, we focused on an important critical case, when the fragmentation is constant and the growth rate is either null or linear [43]. Using the Mellin transform of the equation, we determine the long time behaviour of the solutions and the speed of convergence, which may be either exponential or at most polynomial according to the subdomain of $(t, x) \in \mathbb{R}_+^2$ which is considered. Our results show in particular the strong dependence of this asymptotic behaviour with respect to the initial data, in contrast to the generic results. Following our study, J. Bertoin and A. Watson proposed a complementary probabilistic analysis of related models [60]. These results exemplify the continuing need for further analysis of these interesting equations.

7.2.4. Cell aggregation by chemotaxis

We follow our investigation on the kinetic model describing the chemotactic motion of bacteria. When taxis dominates the unbiased movements, the kinetic system is approximated by the aggregation equation. The study of such equation is challenging since blow-up in finite-time of solutions occurs. We have defined the notion of measure-valued solution [8] and we have proposed and studied a numerical scheme to simulate these solutions[18].

In another approach, more accuracy can be obtained with the kinetic model by adding an internal variable describing the methylation level of the internal receptors of bacteria. In [55] we have investigated the link between these kinetic models with an internal variable and the one without internal variable.

7.3. Liver modeling

Participants: Noémie Boissier, Dirk Drasdo, Géraldine Cellière, Adrian Friebel, Group Heinzle [Univ. Saarbruecken, Germany], Group Hengstler [IfADo, Germany], Stefan Hoehme, Tim Johann, Irène Reo [Vignon-Clementel], Paul Van Liedekerke, Eric Vibert [Hopital Paul Brousse], Group Zerial [Max-Planck Inst. for Molecular Genetics, Dresden, Germany], Groups Iflow, Notox, Vln.

7.3.1. Ammonia detoxification after drug-induced damage

Overdosing acetaminophen (APAP) is the main reason for acute liver failure in the US and UK. Overdose of APAP destroys the hepatocytes located in the center of each liver lobule (pericentral damage), the repetitive functional and anatomical tissue units of liver. Human has about a million of such lobules. As a consequence, the blood is not sufficiently detoxified from ammonia, which is toxic to the body and can lead to encephalopathy. In France about 1000 cases of ammonia intoxication each year. In recent papers we demonstrated by an integrated model that the widely accepted key reactions scheme of ammonia detoxification is insufficient to explain ammonia detoxification after pericentral lobule damage and predicted a missing ammonia sink [73]]. This finding has triggered new experiments leading to the identification of a widely ignored but fundamentally important ammonia sink mechanism We could show by a testing a number of different mechanisms within novel models that this sink mechanism was the only one able to explain the data [15]. The reaction turned out to have the potential to be therapeutically used by injection of a molecular cocktail triggering it. In the animal model death could be prevented using this cocktail hence providing a possible therapy approach for patients suffering from hyperammonemia. [15]. In a follow-up work, further models have been studied and classified by statistical methods to quantify model selection (Cellière et al., in preparation).

7.3.2. Concepts of modeling of liver across all scales: multiscale liver modeling

Based upon developped multiscale concepts [12], we developped a multi-level spatial temporal multiscale models of APAP (paracetamol, acetaminophen) toxicity and ammonia metabolism. In on of these models we integrated molecular pathways of APAP drug toxicity (PD); in another one, we represented the ammonia detoxification pathway into each individual hepatocyte of an agent-based model that describes the precise liver lobule architecture (compare with [73]). This allows us to study the impact of space and architecture on the drug toxicity and drug detoxification. We find in certain cases important differences between models that do represent architecture and those that do not (Cellière et al., in preparation).

7.3.3. Predicting in vivo drug toxicity from in vitro data

APAP (paracetamol, acetaminophen) *in vitro* experiments have been used to calibrate a model of APAP drug toxicity with *in vitro* data, and modify this model to predict *in vivo* toxicity. This procedure is aimed at as a general pathway among cosmetic and pharmaceutical companies to eliminate or at least reduce animal experiments and it should allow a better prediction of drug toxicity in human. Three critical differences between *in vitro* and *in vivo* settings were stepwise integrated in the model calibrated with *in vitro* toxicity data to study their impact on *in vivo* toxicity predictions: (1) The temporal drug exposure profile, (2) the temporal concentration profile of a class of key enzymes, CYP enzymes. Only in hepatocytes in which CYP enzymes are present, APAP is metabolised and downstream apoptosis can occur. (3) The liver architecture, that is responsible for critical differences in the spatial distribution of the drug. The results are in preparation for publication (Cellière et. al., in preparation).

7.3.4. Miscellaneous

In addition, regenerating lobules after partial hepatectomy were analysed by image analysis, and first simulations of blood and bile flow and molecular transport in those lobules simulated.

7.4. Miscellaneous

Participants: Noémie Boissier, Maria José Cáceres [Universidad de Granada], Julien Chevallier [Université de Nice], Géraldine Cellière, Marie Doumic, Dirk Drasdo, Adrian Friebel, Group Heinzle [Univ. Saarbruecken, Germany], Group Hengstler [IfADo, Germany], Stefan Hoehme, Tim Johann, Group Klingmueller [German Cancer Center, Heidelberg], Johannes Neitsch, Benoît Perthame, Patricia Reynaud [Université de Nice], Group Reo [Inria Paris - Rocquencourt], Paul Van Liedekerke, Eric Vibert [Hopital Paul Brousse], Yi Yin, Group Zerial [Max-Planck Inst. for Molecular Genetics, Dresden, Germany], Groups Iflow, Notox, Vln.

7.4.1. Network formation and neuroscience

Motivated by neurodevelopment and differentiation in developing tissues, a new explanation for sharp boundary formation is analysed in [25]; interestingly, this phenomenon relies on a limited diffusion of homeoproteins (collaboration with the Mycenae team).

Models for neural networks have been proposed which describe the probability to find a neuron for which a time s has elapsed since the last discharge. These are written under the form of a nonlinear age-structured equation where the total network activity modulates the firing rate. An inhomogeneous network of networks with variability on the refractory period is studied in [19].

We have also continued the analysis and numerical simulation of models for natural transportation networks formation based on an elliptic-parabolic system of partial differential equations. The model describes the pressure field using a Darcy's type equation and the dynamics of the conductance network under pressure force effects. Randomness in the material structure is represented by a linear diffusion term and conductance relaxation by an algebraic decay term [16]. Figure 1 below gives a numerical simulation of a network formed by such a model.

7.4.2. Microscopic approach of a time elapsed neural model

The spike trains are the main components of the information processing in the brain. To model spike trains several point processes have been investigated in the literature. More macroscopic approaches have also been studied, using partial differential equation models. With J. Chevallier, M. Cáceres and P. Reynaud-Bouret, we wanted to build a bridge between several point processes models (Poisson, Wold, Hawkes) that have been proved to statistically fit real spike trains data and age-structured partial differential equations as introduced by Pakdaman, Perthame and Salort. To do so, we focused on a seemingly simple one-neuron model, for which we stated the - nonlinear and strongly coupled - PDE model satisfied in average by its point measure when the process model is a Poisson, a Wold or a Hawkes process [10].

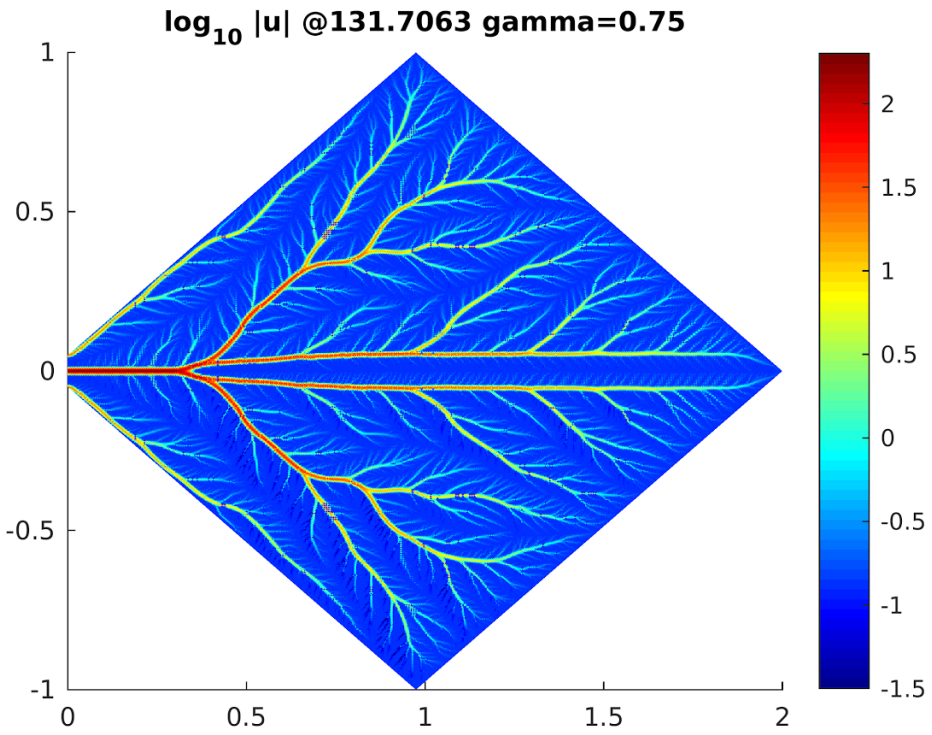

Figure 1. Network formation based on an elliptic-parabolic system of partial differential equations.

7.4.3. Uncertainty propagation

In [42], we study two intrusive methods for uncertainty propagation in scalar conservation laws based on their kinetic formulations. The first one is based on expansions on an orthogonal family of polynomials. The second method uses convolutions based on Jackson kernels. We prove that it satisfies BV bounds and converges to the entropy solution but with a spurious damping phenomenon. Therefore we introduce a second method, which is based on projection on layered Maxellians, and which arises as a minimisation of entropy. This new method satisfies the maximum principle by construction as well as partial entropy inequalities and thus provides an alternative to the standard method of moments which, in general, does not satisfy the maximum principle. Simple numerical simulations for the Burgers equation illustrate these theoretical results.

7.4.4. Simulation of tissue mechanics with agent-based models

In ref. [29] we study and discuss in how far mechanical effects of cells in tissue organisation and growth processes can be captured by agent-based models. We consider a wide range of agent-based models, i.e., lattice base models with one lattice site allowing for many cells or one cell at most, many lattice sites occupied by a single cells (so called Cellular Potts model, Lattice Gas Cellular Automaton approaches, center-based models and vertex models, in which the forces between cells are calculated as forces between the cell centers, as well as deformable cell models in which the cell surface is triangulated. We consider growth of monolayers and multicellular spheroids as reference problems. We also compare in this paper spatial resolution, the capability of the different approaches to represent the physics, cell shape, the computational efficiency and code access. In addition, models evaluating the mechanical effects of growing cell populations in elastic capsules were established and studied.

MATHERIALS Project-Team

6. New Results

6.1. Electronic structure calculations

Participants: Eric Cancès, Virginie Ehrlacher, David Gontier, Claude Le Bris, Gabriel Stoltz.

In electronic structure calculation as in most of our scientific endeavors, we pursue a twofold goal: placing the models on a sound mathematical grounding, and improving the numerical approaches.

E. Cancès and N. Mourad have clarified the mathematical framework underlying the construction of norm-conserving semilocal pseudopotentials for Kohn-Sham models, and have proved the existence of optimal pseudopotentials for a family of optimality criteria [34].

E. Cancès and R. Scott (University of Chicago) have examined a technique of Slater and Kirkwood which provides an exact resolution of the asymptotic behavior of the van der Waals attraction between two hydrogen atoms. They have modified their technique to make the problem more tractable analytically and more easily solvable by numerical methods [35].

In [33], E. Cancès, D. Gontier and G. Stoltz analyze the GW method for finite electronic systems. This method allows to compute excited states. To understand it, a first step is to provide a mathematical framework for the usual one-body operators that appear naturally in many-body perturbation theory. It is then possible to study the GW equations which construct an approximation of the one-body Green's function, and give a rigorous mathematical formulation of these equations. With this framework, results can be established for the well-posedness of the GW_0 equations, a specific instance of the GW model. In particular, the existence of a unique solution to these equations is proved in a perturbative regime.

D. Gontier extended his last-year result on N-representability by including the characterization of representable paramagnetic currents [42]. Together with Salma Lahbabi (former student of E. Cancès, University Hassan II Casablanca, ENSEM), he proved the exponential convergence rates of the uniform sampling of the Brillouin zone for the calculation of crystalline structure properties, in linear and nonlinear settings [43].

A. Bakhta, E. Cancès and V. Ehrlacher have recently been working on the design of an efficient numerical method to solve the inverse band structure problem. The aim of this work is the following: given a set of electronic bands partially characterizing the electronic structure of a crystal, is it possible to recover the structure of a material which could achieve similar electronic properties? The main difficulty in this problem relies in the practical resolution of an associated optimization problem with numerous local optima.

E. Cancès has pursued his long-term collaboration with Y. Maday (Paris 6) on the numerical analysis of electronic structure models. Together with G. Dusson (Paris 6), B. Stamm (Paris 6), and M. Vohralík (Inria), they have designed a new postprocessing method for planewave discretizations of nonlinear Schrödinger equations, and used it to compute sharp *a posteriori* error estimators for both the discretization error and the algorithmic error (convergence threshold in the iterations on the nonlinearity). They have then extended this approach to the Kohn-Sham model. In parallel, they have derived a posteriori error estimates for conforming numerical approximations of the Laplace eigenvalue problem with a homogeneous Dirichlet boundary condition [32]. In particular, upper and lower bounds for the first eigenvalue are given. These bounds are guaranteed, fully computable, and converge with the optimal speed to the exact eigenvalue.

Implicit solvation models aim at computing the properties of a molecule in solution (most chemical reactions take place in the liquid phase) by replacing all the solvent molecules but the few ones strongly interacting with the solute, by an effective continuous medium accounting for long-range electrostatics. E. Cancès, Y. Maday (Paris 6), and B. Stamm (Paris 6) have recently introduced a very efficient domain decomposition method for the simulation of large molecules in the framework of the so-called COSMO implicit solvation models. In collaboration with F. Lipparini (Paris 6), B. Mennucci (Department of Chemistry, University of Pisa) and J.-P.

Picquemal (Paris 6), they have implemented this algorithm in widely used computational software products (Gaussian and Tinker). E. Cancès, Y. Maday, F. Lipparini and B. Stamm have also extended this approach to the more complex polarizable continuum model (PCM).

C. Le Bris, in collaboration with P. Rouchon (École des Mines de Paris) and with J. Roussel, in the context of an internship at École des Ponts, has pursued the study of a new efficient numerical approach, based on a model reduction technique, to simulate high dimensional Lindblad type equations at play in the modelling of open quantum systems. The added value of the most recent contribution with respect to the previous studies lies in two different aspects. First, the rank of the reduced model used as surrogate for the full model can now be dynamically adjusted, in an adaptive strategy. Second, a variance reduction approach based on the technique of control variate has been developed. The noise intrinsically present in the Monte-Carlo simulation of the underlying stochastic dynamics may indeed be reduced by using the deterministic reduced model as control variate. A publication collects these two aspects and reports on the results achieved [19].

6.2. Complex fluids

Participant: Sébastien Boyaval.

The aim of the research performed in the project-team about complex fluids is mainly focused on the mathematical modelling and numerical simulation of i) non-Newtonian rheologies, with application to geophysical fluids such as mudflows, or the solid transport in rivers, and ii) stratified flows, in particular free-surface flows, which naturally occur in the geophysical context under gravity influence.

The need for reduced models is crucial for numerical computations at the large geophysical scale. S. Boyaval has therefore pursued his research about a systematic asymptotic reduction technique for thin-layers of non-Netwonian fluids with a near hydrostatic pressure [11]. On the other hand, accurate numerical simulations (for benchmark purposes at least) require a full 3D model mainly based on Stokes-like equations, and there is a constant need for better computation methods in that field too. With a view to *condensed* high-order approximations of elliptic PDEs like the Stokes equation on generic meshes (obtained by refinement or agglomeration of a simplicial initial mesh), S. Boyaval has participated in a joint work about hybridization of a mixed-dual generic approach [8]. On the hydraulic applications side, the studies initiated at CEMRACS 2013 about a stochastic representation of fluctuations in the transport of river sediments by bed-load have been published [9].

6.3. Homogenization

Participants: Michael Bertin, Ludovic Chamoin, Virginie Ehrlacher, Thomas Hudson, Marc Josien, Claude Le Bris, Frédéric Legoll, Simon Lemaire, François Madiot, William Minvielle.

6.3.1. Deterministic non periodic systems

The homogenization of (deterministic) non periodic systems is a well known topic. Although well explored theoretically by many authors, it has been less investigated from the standpoint of numerical approaches (except in the random setting). In collaboration with X. Blanc (Paris 7) and P.-L. Lions (Collège de France), C. Le Bris has introduced a possible theory, giving rise to a numerical approach, for the simulation of multiscale nonperiodic systems. The theoretical considerations are based on earlier works by the same authors (derivation of an algebra of functions appropriate to formalize a theory of homogenization). The numerical endeavor is completely new. The theoretical results obtained to date are being collected in a series of manuscripts that will be available shortly. The publications [30] and [10] specifically address the issues related to a local perturbation of the periodic problem and the challenging, practically relevant problem of interfaces between periodic structures of different nature (the celebrated "twin boundaries" problem in materials science). Some related problems will now be addressed in the context of the PhD thesis of M. Josien.

6.3.2. Stochastic homogenization

The project-team has pursued its efforts in the field of stochastic homogenization of elliptic equations, aiming at designing numerical approaches that both are practically relevant and keep the computational workload limited.

Using the standard homogenization theory, one knows that the homogenized tensor, which is a deterministic matrix, depends on the solution of a stochastic equation, the so-called corrector problem, which is posed on the *whole* space $\mathbb{R}^d$. This equation is therefore delicate and expensive to solve. In practice, the space $\mathbb{R}^d$ is truncated to some bounded domain, on which the corrector problem is numerically solved. In turn, this yields a converging approximation of the homogenized tensor, which happens to be a *random* matrix.

In [47], C. Le Bris, F. Legoll and W. Minvielle have investigated the possibility to use a variance reduction technique based on computing the corrector equation *only for selected environments*. These environments are chosen based on the fact that their statistics in the finite supercell matches the statistics of the materials in the infinite supercell. This method yields an approximation of the homogenized matrix with an error smaller than standard approximations. The efficiency of the approach has been demonstrated for various types of random materials, including composite materials with randomly located inclusions.

In addition, M. Bertin and F. Legoll, in collaboration with S. Brisard (École des Ponts), have investigated the possibility to use the Hashin-Shtrikman bounds as control variables in a control variate approach. The Hashin-Shtrikman bounds are often used in the computational mechanics community as approximations of the homogenized quantities. Our aim is use them to improve the efficiency of the reference computations, somewhat in the spirit of a preconditionner. Preliminary encouraging numerical results have been obtained.

Over the past years, the project-team has proposed several variance reduction techniques, see e.g. [21] for a method using antithetic variables (in a nonlinear context) and [20] for a control variate approach using a surrogate model based on a defect-type theory. These various approaches have been reviewed and compared to one another in [29].

In collaboration with B. Stamm (Paris 6), E. Cancès, V. Ehrlacher and F. Legoll have proposed in [13] a new approach to approximate the homogenized coefficients of a random stationary material. This method is an alternative to that proposed e.g. by A. Bourgeat and A. Piatniski in [Approximations of effective coefficients in stochastic homogenization, Annales de l'Institut Henri Poincaré 40, 2004] which consists in solving a corrector problem on a bounded domain. The method introduced in [13] is based on a new corrector problem, which is posed on the entire space, but which is simpler than the standard corrector problem in that the coefficients of the equation are uniform outside some ball of finite radius. This implies that, in some cases (including the case of randomly located spherical inclusions), this new corrector problem can be recast as an integral equation posed on the surface of the inclusions. The problem can then be efficiently solved via domain decomposition and using spherical harmonics.

6.3.3. Multiscale Finite Element approaches

From a numerical point of view, the Multiscale Finite Element Method (MsFEM) is a classical strategy to address the situation when the homogenized problem is not known (e.g. in difficult nonlinear cases), or when the scale of the heterogeneities, although small, is not considered to be zero (and hence the homogenized problem cannot be considered as an accurate enough approximation).

The MsFEM has been introduced more than 10 years ago. However, even in simple deterministic cases, there are still some open questions, for instance concerning multiscale advection-diffusion equations. Such problems are possibly advection dominated and a stabilization procedure is therefore required. How stabilization interacts with the multiscale character of the equation is an unsolved mathematical question worth considering for numerical purposes. In that spirit, C. Le Bris, F. Legoll and F. Madiot have studied in [46] several variants of the Multiscale Finite Element Method (MsFEM), specifically designed to address multiscale advection-diffusion problems in the convection-dominated regime. Generally speaking, the idea of the MsFEM is to perform a Galerkin approximation of the problem using specific basis functions, that are precomputed (in an offline stage) and adapted to the problem considered. Several possibilities for the basis functions have been examined (for instance, they may or may not encode the convection field). Depending on how basis functions are defined, stabilization techniques (such as SUPG) may be required. Another option to handle such problems is to use a splitting approach, with two legacy codes, one solving a purely diffusive multiscale equation, the other one solving a single scale, convection-dominated advection-diffusion equation. In [46], these various approaches have been compared in terms of accuracy and computational costs.

In the context of the PhD thesis of F. Madiot, current efforts are focused on the study of an advection-diffusion equation with a dominating convection in a *perforated domain*. The multiscale character of the problem here stems from the geometry of the domain. A paramount difference with the case considered in [46] is that boundary layers may appear throughout the domain (i.e. in the neighborhood of each perforation). The accuracy of the numerical approaches in the boundary layers thus becomes critical.

Most of the numerical analysis studies of the MsFEM are focused on obtaining *a priori* error bounds. In collaboration with L. Chamoin, who is currently in delegation in the project-team for the second year (from ENS Cachan, since September 2014), members of the project-team have been working on *a posteriori* error analysis for MsFEM approaches, with the aim to develop error estimation and adaptation tools. They have extended to the MsFEM case an approach that is classical in the computational mechanics community for single scale problems, and which is based on the so-called Constitutive Relation Error (CRE). Once a numerical solution u_h has been obtained, the approach needs additional computations in order to determine a divergence-free field as close as possible to the exact flux $k\nabla u$. In the context of the MsFEM, it is important to be able to do all the expensive computations in an offline stage, independently of the right-hand side. The standard CRE approach thus needs to be adapted to that context, in order to keep that feature that makes it adapted to a multiscale, multi-query context. The proposed approach yields very interesting results, and provide an accurate and robust estimation of the global error.

Current efforts are targeted towards the design of adaptive algorithms for specific quantities of interest (in the so-called "goal-oriented" setting), and the design of model reduction approaches (such as the Proper Generalized Decomposition, or PGD) in the specific context of multiscale problems.

6.3.4. Coarse approximation of an elliptic problem with oscillatory coefficients

Still another question investigated in the project-team is to find an alternative to standard homogenization techniques when the latter are difficult to use in practice. Consider a linear elliptic equation, say in divergence form, with a highly oscillatory matrix coefficient, and assume that this problem is to be solved for a large number of right-hand sides. If the coefficient oscillations are infinitely rapid, the solution can be accurately approximated by the solution to the homogenized problem, where the homogenized coefficient has been evaluated beforehand by solving the corrector problem. If the oscillations are moderately rapid, one can think instead of MsFEM-type approaches to approximate the solution to the reference problem. However, in both cases, the complete knowledge of the oscillatory matrix coefficient is required, either to build the average model or to compute the multiscale basis. In many practical cases, this coefficient is often only partially known, or merely completely unavailable, and one only has access to the solution of the equation for some loadings. This observation has led to think about alternative methods, in the following spirit. Is it possible to approximate the reference solution by the solution to a problem with a *constant* matrix coefficient? How can this "best" constant matrix approximating the oscillatory problem be constructed in an efficient manner?

A preliminary step, following discussion and interaction with A. Cohen (Paris 6), has been to cast the problem as a convex optimization problem. We have then shown that the "best" constant matrix defined as the solution of that problem converges to the homogenized matrix in the limit of infinitely rapidly oscillatory coefficients. Furthermore, the optimization problem being convex, it can be efficiently solved using standard algorithms. C. Le Bris, F. Legoll and S. Lemaire have comprehensively explored that problem. The algorithm can be made very efficient, and it yields accurate approximation of the homogenized matrix. We have also shown that it is possible to construct, in a second stage, approximations to the correctors, in order to recover an approximation of the *gradient* of the solution.

6.3.5. Optimization of a material microstructure

A project involving V. Ehrlacher and F. Legoll, in collaboration with G. Leugering and M. Stingl (Cluster of Excellence, Erlangen-Nuremberg University), aims at optimizing the shape of some materials (modeled as structurally graded linear elastic materials) in order to achieve the best mechanical response at the minimal cost. As is often the case in shape optimization, the solution tends to be highly oscillatory, hence the need for homogenization techniques. Materials under consideration are being thought of as microstructured materials composed of steel and void and whose microstructure patterns are constructed as the macroscopic deformation

of a reference periodic microstructure. The optimal material (i.e. the best macroscopic deformation) is the deformation achieving the best mechanical response.

For a given deformation, one can first compute the mechanical response using a *homogenized model*. This is the first variant that has been followed. Model reduction techniques are then required, in order to expedite the resolution of the corrector problem needed to identify the homogenized coefficient at each loop of the optimization algorithm. In that context, a PGD-type approach has been proposed.

A second variant is to compute the mechanical response at the *microscale*, using the highly oscillatory model. Preliminary results have been obtained. Current efforts are focused towards choosing an appropriate model reduction strategy.

6.3.6. Discrete systems and their thermodynamic limit

We conclude this section by describing works of the project-team on discrete models with highly oscillatory coefficients.

Dislocations are geometric line defects which interact via long-range stress fields in crystalline solids. In [45], T. Hudson has studied the thermally-driven motion of dislocations in a discrete Monte Carlo model, showing that over long observation times at low temperature in a large body, the most probable trajectory of straight dislocation lines lie close to the solution of an explicit deterministic evolution equation.

Another work is related to the understanding of the origin of hysteresis in rubber-made materials. When submitted to cyclic deformations, the strain-stress curve of these materials indeed shows a hysteresis behavior, which seems to be independent of the speed of loading. Some years ago, members of the project-team have suggested a model, at a mesoscale, to explain this behavior. This model was written in terms of a system made of a finite number of particles. F. Legoll, T. Lelièvre and T. Hudson are currently studying whether a thermodynamic limit of the model previously proposed can be identified. In order to simplify the setting, the reference discrete model has been replaced by a continuum model with highly oscillatory coefficients. This model is nonlinear and time-dependent. The question is now to identify (e.g. using two-scale convergence arguments) its homogenized limit, first in a periodic setting, second in a stochastic setting.

6.4. Computational Statistical Physics

Participants: Giacomo Di Gesù, Thomas Hudson, Dorian Le Peutrec, Frédéric Legoll, Tony Lelièvre, Antoine Levitt, Boris Nectoux, Julien Roussel, Mathias Rousset, Gabriel Stoltz, Pierre Terrier, Pierre-André Zitt.

The work of the project-team in this area is concentrated on two new directions: the sampling of reactive trajectories (where rare events dictate the dynamics of the system), and the computation of average properties of nonequilibrium systems (which complements the more traditional field of techniques to compute free energy differences).

6.4.1. Sampling of reactive trajectories

Finding trajectories for which the system undergoes a significant change is a challenging task since the transition events are typically very rare. Several methods have been proposed in the physics and chemistry literature, and members of the project-team have undertaken their study in the past years.

A first class of techniques are the accelerated dynamics introduced by A. Voter (Los Alamos National Lab) and his collaborators. A short review on the mathematical analysis of these dynamics was written by T. Lelièvre, see [48]. In [23], T. Lelièvre and F. Nier (Paris 13) analyze the low temperature asymptotics for Quasi-Stationary Distributions in a bounded domain. The objective of this analysis is to justify mathematically the validity of hyperdynamics.

Another class of techniques to compute reactive trajectories is based on splitting techniques. After the first result obtained in [12], C.E. Bréhier, T. Lelièvre and M. Rousset pursued their analysis of the Adaptive Multilevel Splitting algorithm, which is a rare event simulation method. In [31], a generalization of the method is proposed, and it is shown how to make the estimator unbiased in a discrete-in-time setting (which is generically the setting encountered in practice). Numerical experiments illustrate the performance of the method.

6.4.2. Nonequilibrium systems and non-reversible dynamics

In [38], T. Lelièvre has studied with A. Duncan and G.A. Pavliotis nonreversible diffusion processes to sample a probability measure. It is shown that nonreversible dynamics are always better in terms of the asymptotic variance (statistical error), but the efficiency of the whole algorithm sensitively depends on the time discretization algorithm, which may induce some bias (deterministic error).

T. Lelièvre together with R. Assaraf, B. Jourdain and R. Roux, have analyzed in [27] the validity of non equilibrium molecular dynamics techniques to compute the derivative of an observable with respect to a parameter-dependent probability measure. The probability measure is defined as the stationary state of a non-reversible stochastic dynamics (in particular no analytical formula for this measure is available). Such computations are at the basis of the numerical approximation of transport coefficients in molecular dynamics.

6.4.3. Numerical analysis of simulation techniques

In [44], G. Stoltz, together with A.-A. Homman (École des Ponts) and J.-B. Maillet (CEA/DAM), present new parallelizable numerical schemes for the integration of Dissipative Particle Dynamics with Energy conservation. So far, no numerical scheme was able to correctly preserve the energy over long times and give rise to small errors on average properties for moderately small timesteps, while being straightforwardly parallelizable. Two new methods are proposed, both of them straightforwardly parallelizable, and allowing to correctly preserve the total energy of the system. The accuracy and performance of these new schemes are illustrated both on equilibrium and nonequilibrium parallel simulations.

The discretization of overdamped Langevin dynamics, through schemes such as the Euler-Maruyama method, may lead to numerical methods which are unstable when the forces are non-globally Lipschitz. One way to stabilize numerical schemes is to superimpose some acceptance/rejection rule, based on a Metropolis-Hastings criterion for instance. However, rejections perturb the dynamical consistency of the resulting numerical method with the reference dynamics. G. Stoltz and M. Fathi (Berkeley) present in [40] some modifications of the standard stabilization of discretizations of overdamped Langevin dynamics by a Metropolis-Hastings procedure, which allow to either improve the strong order of the numerical method, or to decrease the bias in the estimation of transport coefficients characterizing the effective dynamical behavior of the dynamics. The latter approach relies on modified numerical schemes together with a Barker rule for the stabilization.

A. Levitt, in collaboration with C. Ortner (University of Warwick), has worked on the numerical analysis of saddle point search, an important step in the computation of reaction rates. While the convergence theory of minimization algorithms, such as the gradient method, is well-understood and standard, no such theory exists for saddle point algorithms such as the dimer method. Their work reveals a major obstruction to convergence: for some systems, the dimer method can oscillate indefinitely. This shows that there is no Lyapunov function for the associated flow, and highlights the fundamental difference between minimization and saddle search. Further work focuses on improving the reliability and convergence speed of such methods.

6.4.4. Free energy computations

The topic of free energy computations is still a significant research area of the project-team. T. Lelièvre has co-authored a review article [14] on the adaptive biasing force (ABF) method.

In addition, two new results have been obtained on the ABF method by H. Al Rachid (École des Ponts) in collaboration with T. Lelièvre: a numerical result concerning a projected version of the ABF algorithm, which enables to reduce the variance, see [25] ; and a theoretical result on the existence of a solution to the non linear Fokker Planck equation associated to the ABF process, see [49].

T. Lelièvre and G. Stoltz, together with G. Fort (Télécom Paris) and B. Jourdain (École des Ponts), have studied the Self-Healing Umbrella Sampling (SHUS) method in [16]. This method is an adaptive biasing method to compute free energies on the fly by appropriately penalizing already visited regions. The convergence of the method relies on a rewriting as a stochastic approximation method with random steps, and can therefore be seen as a variation of the Wang-Landau method.

6.4.5. Convergence of processes

D. Le Peutrec and G. Di Gesù have studied in [37] the rate of convergence to equilibrium at low temperature of a stochastic interacting large particle system which can be seen as a spatially discrete approximation of the stochastic Allen-Cahn equation on the one-dimensional torus. Upper and lower bounds for the leading term of the associated spectral gap in the small temperature regime are proven, uniformly in the system size. It is also shown that the upper bound is sharp under a suitable control of the growth of the system size by the temperature.

The article [17] by B. Jourdain (École des Ponts), T. Lelièvre and B. Miasojedow (Warsaw) on the mean-field limit for the transient phase of the random walk Metropolis algorithm in the infinite dimension limit has been published in Annals of Applied Probability. In this article, the authors prove that the Metropolis Hastings algorithm converges to a nonlinear stochastic differential equation in the infinite dimensional limit.

6.4.6. Force fields and modeling

In [41], G. Stoltz, together with G. Ferré (École des Ponts) and J.B Maillet (CEA/DAM), has presented a distance between atomic configurations, which is invariant with respect to permutations of the atoms. This distance is defined through a functional representation of atomic positions. It allows to directly compare different atomic environments with an arbitrary number of particles without going through a space of reduced dimensionality (i.e. fingerprints) as an intermediate step. Moreover, this distance is naturally invariant through permutations of atoms and through global rotations. This distance provides an important building block for the construction of accurate force-fields using machine learning techniques.

E. Cancès has contributed to the development of more efficient algorithms for polarizable force field molecular dynamics, which have been implemented and successfully tested on massively parallel computers [18].

During the post-doctoral position of I.G. Tejada, G. Stoltz, F. Legoll and E. Cancès studied in collaboration with L. Brochard (École des Ponts) the derivation of a concurrent coupling technique to model fractures at the atomistic level by combining a reactive potential with a harmonic approximation; see [50].

6.5. Various topics

A. Bakhta (École des Ponts) and V. Ehrlacher [28] have studied a system of PDEs modeling the cross-diffusion of different atomic species in a crystalline solid thin film during a Physical Vapor Deposition process, coupled with the evolution of the domain as external chemical species fluxes are absorbed at the surface of the solid layer. This model leads to a system of degenerate elliptic cross-diffusion equations. They proved the existence of a global weak solution to this system in arbitrary dimension in the case of a constant domain using analysis tools from gradient flow theory. The existence of a global weak solution in a one-dimensional case with external fluxes was also proved. Under the assumption that this solution is unique, the existence of optimal external fluxes in order to achieve desired concentration profiles of the different species in the thickness of the solid layer at the end of the process was also obtained.

Numerical simulations of crystal defects are necessarily restricted to finite computational domains, supplying artificial boundary conditions that emulate the effect of embedding the defect in an effectively infinite crystalline environment. V. Ehrlacher, in a joint work with C. Ortner (U. of Warwick) and A. Shapeev (Skolkovo Institute of Science and Technology) [39] have studied a mathematical framework within which the accuracy of different types of boundary conditions can be precisely assessed.

T. Lelièvre together with F. Casenave (Safran) and A. Ern (École des Ponts) have proposed in the short note [36] an analysis of the Empirical Interpolation Method which highlights the symmetry played by the two variables (parameter and space variable). A variant of the Empirical Interpolation Method is introduced in order to deal with situations where some observations have to be discarded, and the number of observed values is thus different for the two variables.

In collaboration with P.-L. Lions (Collège de France), C. Le Bris has written an extensive set of lecture notes on parabolic equations with irregular data (initial conditions and parameter coefficients). These lecture notes correspond to joint works between the two authors and to an expanded version of the works by P.-L. Lions specifically exposed in his lectures delivered at Collège de France in 2012–2013. The application of the theory to the specific context of stochastic differential equations with irregular coefficients is also examined.

MATHRISK Project-Team

7. New Results

7.1. Liquidity risk

Participants: Aurélien Alfonsi, Pierre Blanc.

A. Alfonsi and P. Blanc are working on the optimal execution problem when many large traders who modify the market prices. In a previous study, they have developed a price impact model that takes into account an exogeneous flow of market orders, in which the optimal execution strategy is known explicitly. This year,they have worked on the practical implementation of this model. Namely, they have proposed an estimation procedure to estimate the model parameters (decay kernel of the price impact and Hawkes kernel for the self excitation of the order flow). They have run this estimation on market data and backtested the optimal execution strategy.

7.2. Backward stochastic (partial) differential equations with jumps, optimal stopping and stochastic control with nonlinear expectation, risk minimization

Participants: Roxana Dumitrescu, Marie-Claire Quenez [(Univ Paris 7)], Arnaud Lionnet, Agnès Sulem.

R. Dumitrescu, M.C. Quenez and A. Sulem have provided a weak dynamic principle for Combined Optimal Stopping/Stochastic Control with $\mathcal{E}^f-$conditional Expectation . They have investigated the links between generalized Dynkin games and double barriers reflected BSDEs with jumps and also studied mixed generalized Dynkin games in a Markovian framework and associated nonlinear HJB equations with barriers .

In the recent paper [43], they study game options in an imperfect market with default. They extend the results obtained by Kifer [68] in a perfect market model to the case of imperfections taken into account via the nonlinearity of the wealth dynamics. In this framework, the pricing system is expressed as a nonlinear g-expectation/evaluation induced by a nonlinear BSDE with jump. They prove that the superhedging price of a game option coincides with the value function of a corresponding *generalized* Dynkin game expressed in terms of the g-evaluation. They also address the case of ambiguity on the model, - for example an ambiguity on the default probability -, and characterize the superhedging price of a game option as the value function of a *mixed generalized* Dynkin game. They prove the existence of a cancellation time and a trading strategy which allows the seller to be super-hedged, whatever the model is. This study is introduced by the analysis of the simpler case of American options.

In collaboration with Jane Bielagt (Humbold Univ.) and Gonąlo Dos Reis (Univ. of Edimburgh), Arnaud Lionnet investigates in the effects of the social interactions of a finite set of agents on an equilibrium pricing mechanism. They consider an incomplete market where agents invest so as to minimize their risk measure. Here, agents assess risk using convex dynamic risk measures expressed by Backward Stochastic Differential Equations (BSDE). Beside the risk associated with their own economic activity, the agents compare their trading gains to that of the others, and factor this relative performance in the evaluation of their risk/satisfaction. When a derivative product is introduced to complete the market and allow agents to trade a non-financial risk factor (such as temperature), the risk of each agent is lowered, as expected. However, agents then find it in their interest to be more concerned with their relative performance. This leads them to behave more like a herd and this destabilizes the previously stable, purely financial market.

7.3. Systemic risk

Participants: Hamed Amini [EPFL], Andreea Minca [Cornell University)], Agnès Sulem, Rui Chen, Romuald Elie.

We study the issue of control of systemic risk in the framework of random graph models. The paper [16] by H. Amini, A. Minca and A. Sulem, provides important insight on the relation between the value of a financial system, connectivity and optimal intervention. More precisely, we consider a core-periphery random financial network in which links lead to the creation of projects in the outside economy but make banks prone to contagion risk. The controller seeks to maximize, under budget constraints, the value of the financial system, defined as the total value of the projects funded. Under partial information on interbank links, revealed in conjunction with the spread of contagion, the optimal control problem is shown to become a Markov decision problem. Our results show that up to a certain connectivity, the value of the financial system increases with connectivity. However, this is no longer the case if connectivity becomes too large. This insight shows that it is far from obvious that connectivity of a core bank should always be brought forward as an argument for priority intervention and it may be sometimes preferable to invest in non-core banks that lend directly to the economy. The natural question remains how to create incentives for the banks to attain an optimal level of connectivity and how to design a guarantee fund that would represent an intervention fund that can be used to maximize the benefits of connectivity. This is under study with the PhD student Rui Chen.

Moreover R. Elie obtained a CVRS PEPS grant on systemic risk modeling with graphs in collaboration with the Inria team COATI and the economic department of Université de Nice.

7.4. Dependence modeling

7.4.1. Estimation of the parameters of a Wishart process

A.Alfonsi with A. Kebaier and C. Rey have studied the Maximum Likelihood Estimator for the Wishart processes and in particular its convergence in the ergodic and in some non ergodic cases. In the non ergodic cases, their analysis rely on refined results on the Laplace transform for Wishart processes. This work also extends a recent paper by Ben Alaya and Kebaier on the maximum likelihood estimation for the CIR process.

7.5. Interest rate modeling

A. Alfonsi, E. Palidda and A. Ahdida extend the Linear Gaussian Model (LGM) by replacing the constant convariation matrix by some Wishart dynamics. This extension allows them to generate smile while keeping the affine structure of the model. They have obtained a price expansion around the LGM for Caplet and Swaption prices. They also present a second order discretization scheme that allow them to compute exotic prices within this model.

7.6. Numerical Probability

A. Alfonsi with A. Kohatsu-Higa and M. Hayashi are investigating how to apply the parametric method recently proposed by V. Bally and A. Kohatsu-Higa for reflected SDEs. This method allows them to obtain an unbiased estimator for expectations of general functions of the process.

7.7. Optimal transport

Participant: Benjamin Jourdain.

With J. Corbetta (postdoc financed by the chair financial risks), A. Alfonsi and B. Jourdain study a general formula for the time-derivative of the wasserstein distance between the time-marginals of two Markov processes. They have checked the validity of this formula for pure-jump Markov processes with a bounded intensity of jumps. They now study the extension to piecewise deterministic Markov processes.

7.8. Multitype sticky particle systems

Participant: Benjamin Jourdain.

B. Jourdain and J. Reygner study multitype sticky particle systems which can be obtained as vanishing noise limits of multitype rank-based diffusions. Rank-based diffusion processes and their multitype generalization permit to reproduce empirical features of stock markets. B. Jourdain and J. Reygner have obtained the optimal rate of convergence as the number of particles grows to infinity of the approximate solution to a diagonal system of hyperbolic conservation laws based on multitype sticky particles.

7.9. Numerical Probability

7.9.1. American option pricing.

Damien Lamberton with M. Pistorius has worked on the approximation of American options by Canadian options, which originated from the work of Peter Carr. This lead them to revise old results on the binomial approximation of the American put. D. Lamberton is also working with M. Zervos on American options involving the maximum of the underlying.

7.9.2. Convergence in total variation of approximation schemes for Markov processes

(V. Bally and PhD student C. Rey [40])

The main issue was to consider very general approximation schemes and to estimate the approximation error for test functions which are just measurable and bounded. It is worth to mention that the input of noise in the approximation schemes is allowed to be quite general, while in the standard approximation schemes for diffusion processes one considers Gaussian input only. In some sense this means that we treat invariance principles as well. We also considered approximation schemes of higher order, as the Victoir Nynomia scheme for example. An important ingredient is an abstract Malliavin calculus for general random variables (which has been settled in previous papers of V. Bally and Lucia Caramellino.

7.9.3. Approximation schemes for Piecewise Deterministic Markov Processes

(V. Bally and PhD student V.Rabiet [39]).

PDMP processes are very popular in many practical fields as biology, chemistry or fiability theory. The main idea is that such a model may present different scales: slow ones and rapid onces. And from a numerical point of view it is extremely difficult to implement algorithms which take care of rapid scales in details. Then the idea is to average the rapid scales (in the spirit of the Central Limit Theorem) and consequently to replace small (and rapid) jumps by a Brownien component. This procedure is already widely used by practitioners. Our work was to derive estimates of the error which is done by this procedure.

7.9.4. Convergence in distribution norms in the Central Limit Theorem

(V. Bally with Lucia Caramellino and Guillaume Poly)

In the classical theory, the convergence which has already been studied is the convergence in total variation (measurable test functions). The main result is the theorem of Prohorov, in the fifties. We have proved that under similar hypothesis (with more finite moments however) one may obtain a much more accurate estimate of the error, in some norms which are close to distribution norms. As a remarkable consequence, we obtained a CLT for the zeros of trigonometric polynomials with random coefficients.

MIMOVE Team

7. New Results

7.1. Introduction

MiMove's research activities in 2015 have focused on a set of areas directly related to the team's research topics. Hence, we have worked on QoS for Emergent Mobile Systems (§ 7.2) in relation to our research topic regarding Emergent Mobile Distributed Systems (§ 3.2). Furthermore, our effort on SoundCity (§ 7.3) is linked to our research on Mobile Social Crowd-sensing (§ 3.4). Still in the context of Mobile Social Crowd-sensing (§ 3.4), we have developed AppCivist-PB (§ 7.4) related to our interest in social applications aiming to actively involve citizens (see § 4.1); this is further linked to our research on composition of Emergent Mobile Distributed Systems (§ 3.2).

7.2. QoS for Emergent Mobile Systems

Participants: Georgios Bouloukakis, Nikolaos Georgantas, Rachit Agarwal, Valérie Issarny, Raphael de Aquino Gomes.

With the emergence of Future Internet applications that connect web services, sensor-actuator networks and service feeds into open, dynamic, mobile choreographies, heterogeneity support of interaction paradigms is of critical importance. Heterogeneous interactions can be abstractly represented by client-server, publish/subscribe, tuple space and data streaming middleware connectors that are interconnected via bridging mechanisms providing interoperability among the choreography peers. We make use of the *eVolution Service Bus (VSB)* (see § 6.2) as the connector enabling interoperability among heterogeneous choreography participants. VSB models interactions among peers through generic *post* and *get* operations that represent peer behavior with varying time/space coupling.

Within this context, we study end-to-end Quality of Service (QoS) properties of choreographies, where in particular we focus on the effect of middleware interactions on QoS. We consider both homogeneous and heterogeneous (via VSB) interactions. We report in the following our results in three complementary directions:

- While VSB ensures functional interoperability of heterogeneous choreography interactions, differences in timing requirements and constraints of such interactions can severely affect their latencies and success rates. To model timeliness, we introduce the *lease* and *timeout* parameters. The former captures data availability and validity in time, while the latter represents intermittent availability of data recipients due to mobility and disconnection. By precisely studying the related timing thresholds using timed automata models, we verify conditions for successful interactions with VSB connectors. Furthermore, we statistically analyze through simulations, the effect of varying lease and timeout periods to ensure higher probabilities of successful interactions. Simulation experiments are compared with experiments run on the VSB implementation testbed to evaluate the accuracy of results. This work can provide application developers with precise design time information when setting these timing thresholds in order to ensure accurate runtime behavior [23].
- Choreography peers deployed in mobile environments are typically characterized by intermittent connectivity and asynchronous reception of data. In such environments, it is essential to guarantee acceptable levels of timeliness between the data sources and mobile users. In order to provide QoS guarantees in different application scenarios and contexts, it is necessary to model the system performance by incorporating the intermittent connectivity. Queueing Network Models (QNMs) offer a simple modeling environment, which can be used to represent various application scenarios, and provide accurate analytical solutions for performance metrics, such as system response time. We provide an analytical solution regarding the end-to-end response time between the users and the

data sources by modeling the intermittent connectivity of mobile users with product-form QNMs. We utilize the publish/subscribe middleware as the underlying communication infrastructure for the mobile users. To represent the subscriber's connections/disconnections, we model and solve analytically an ON/OFF queueing system by applying a mean value approach. Finally, we validate our model using both simulations with real-world workload traces and comparison with an actual implementation of a Java Messaging Service middleware. The deviations between the performance results foreseen by the analytical model and the ones provided by the simulator and the prototype implementation of a real system are shown to be less than 5% for a variety of scenarios.

- Large-scale mobile environments are characterized by, among others, a large number of mobile users, intermittent connectivity and non-homogeneous arrival rate of data to the users, depending on the region's context. Multiple application scenarios in major cities need to address the above situation for the creation of robust mobile systems. Towards this, it is fundamental to enable system designers to tune a communication infrastructure using various parameters depending on the specific context. We take a first step towards enabling an application platform for large-scale information management relying on mobile social crowd-sourcing [26]. To inform the stakeholders of expected loads and costs, we model a large-scale mobile pub/sub system as a queueing network. We introduce additional timing constraints such as (i) mobile user's intermittent connectivity period; and (ii) data validity lifetime period (e.g. that of sensor data). Using our MobileJINQS simulator (http://xsb.inria.fr/d4d#mobilejinqs), we parameterize our model with realistic input loads derived from the D4D CDR (Call Detail Record) dataset (http://www.d4d.orange.com/en/home) and varied lifetime periods in order to analyze the effect on response time. This work provides system designers with coarse grain design time information when setting realistic loads and time constraints [18].

7.3. Urban Civics: An IoT Middleware for Democratizing Crowdsensed Data in Smart Societies

Participants: Valérie Issarny, Fadwa Rebhi, Animesh Pathak, Sara Hachem.

The growth of our cities comes along with the aggravation of urban nuisances (e.g., air pollution), which significantly alters the citizens' quality of life and especially their health. It then becomes essential to ensure the growth of cities is both environmentally and socially sustainable. As computer scientists, it is our vision that ICT shall play a key role in achieving the above sustainability requirements, as already put forward by the smart city/society concept. However, smart cities have mostly emphasized the big data dimension and related knowledge engineering to ease the management of the city's infrastructure and resources. While this is an important part of smart cities, we believe that ICT should be leveraged to promote participatory democracy so that citizens and government can communicate openly about the issues facing their societies as much as about their solutions. Toward that goal, we have introduced the Urban Civics middleware, which addresses three complementary research questions underlying participatory democracy from an ICT perspective [20], [21]:

(RQ1) How to leverage the richness of urban sensors of the new digital era that features the Internet of Things, open data, social networking, and mobile computing to serve both citizens and government with better insights? Our answer lies in connecting those various data sources where probabilistic protocols combined with semantic technology allow for an urban-scale middleware solution.

(RQ2) How to assimilate urban data so as to generate explanatory city models to inform urban problem solving? Our solution leverages data assimilation (developed by the Inria CLIME team) that has proven successful in geosciences and paves the way to the comprehensive integration of heterogeneous data sources whose accuracy may vary significantly.

(RQ3) How to integrate the solutions to the above into a scalable urban middleware and further ensure citizen participation? Building on our past experience in developing middleware solutions for the mobile environment and especially the – mobile – Internet of Things, we have conceived and introduced the architecture of Urban Civics, a novel IoT middleware solution for democratizing

crowd-sensed data in smart societie. We are in particular confident that, in addition to leveraging existing incentive mechanisms, the citizen participation will also be prompted by the very nature of participatory democracy. However, such an assumption needs to be validated through actual experiments at an urban scale for which we deploy use cases in the Paris and San Francisco Bay areas.

7.4. AppCivist: Engineering Software Assemblies for Participatory Democracy

Participants: Valérie Issarny, Cristhian Parra Trepowski, Animesh Pathak.

Information and communication technologies (ICT) are profoundly changing the nature of human social and environmental interactions. One such change concerns innovations in the way that citizens both interact with government institutions and engage in greater self-government through democratic assembly and collective action. Our research focuses on this transformation of politics, asking how new social media can contribute to new forms of democracy. The pervasive use of ICT suggests that they present an unprecedented opportunity to rethink the constraints of time and space that are generally thought to make the exercise of a more direct and engaging democracy at a large scale practically impossible. In effect, ICT challenge the assumption that citizens of large political units must be content with systems of representative democracy that typically produce a more passive and legalistic citizenship than an active and participatory one.

To consider this challenge, we undertake a pragmatic and modest investigation of how ICT and more precisely software systems can contribute to enabling direct democracy at a large scale. Our research has two immediate objectives. One is to engineer software that leverages the reach of the Internet and the powers of computation to enhance the experience and efficacy of civic participation. The second is to use the ICT software platform to induce the associational forms of a new digitally-inspired citizenship among residents.

Our research is multi-disciplinary in nature, bringing together anthropologists and computer scientists to coinvestigate how to build software systems that promote the development of such digital democratic assemblies and citizens. Our initiative is further rooted in the principles of social activism in that we want to provide citizens with new software systems that help them articulate projects, deliberate directly among themselves, and mobilize activities. A number of digital tools and in particular social networks and web-based content management systems already support aspects of social activism. However, these tools need to be customized as much as composed to become really useful for activists. To that end, we have set the principles of the AppCivist service-oriented software platform in [24]. AppCivist is built around the vision of letting activist users compose their own applications, called Assemblies, using relevant Internet-based components that enable various aspects of democratic assembly and collective action. Starting from a social science perspective, we identified the following high-level categories of functions for AppCivist Assemblies: Mobilizing people, Co-creating proposals, Acting collectively, and Communicating.

Following, we have concentrated on developing the first instance of AppCivist for Participatory Budgeting (PB), as a representative use case of participatory democracy. As a result, we are able to account for various initiatives in citizen participation, including lessons learned from existing PB campaigns worldwide since their emergence in Brazil in the late 1980s. Research contributions more specifically relate to [22]:

- *State of the art survey and analysis of software systems that contribute to enabling participatory democracy*, which lacks an adequate bottom-up approach to digital proposal making. Such an approach would allow groups of citizens to self-assemble on the basis of common interests and enable the resulting citizen assemblies to initiate ideas and elaborate on them using convenient assemblies of software services.
- *State of the art survey and analysis of digital tools oriented towards Participatory Budgeting*, where leveraging ICT to enable truly urban-scale participation in PB campaigns remains unrealized. AppCivist-PB utilizes the concepts of *citizen assembly* and *software assembly* to address this challenge.
- *AppCivist-PB software architecture* enabling citizen and software assemblies, which following the design of AppCivist introduced in [24] strictly adheres to the principles of service orientation. In

that framework, citizen assemblies allow registered users and groups of users to self assemble into higher-level groups to coordinate idea generation and to elaborate proposals through versioning. In a complementary way, software assemblies adhere to the well-known principle of service composition, configuring software services and components oriented towards the implementation of functions supporting participatory democracy.

- *AppCivist-PB prototype* permits an early assessment of the effectiveness of AppCivist-PB in supporting actual urban-scale PB campaigns, such as the one of Paris in 2015. In addition, the prototype provides an opportunity to experiment with developing service wrappers to integrate third-party services (e.g., Etherpad.org) into its software assemblies. In the near future, we intend to automate this integration as much as possible, building on our background in the synthesis of mediators [13], [12].

This research is carried out in collaboration with the Social Apps Lab at CITRIS at UC Berkeley in the context of CityLab@Inria and Inria@SiliconValley.

MOKAPLAN Project-Team

7. New Results

7.1. Numerical methods for JKO Gradient Flows

J-D. Benamou, G. Carlier, M. Laborde, G. Peyré, B. Schmitzer, V. Duval

Taking advantage of the Benamou-Brenier dynamic formulation of optimal transport, we propose in [28], a convex formulation for each step of the JKO scheme for Wasserstein gradient flows which can be attacked by an augmented Lagrangian method which we call the ALG2-JKO scheme. We test the algorithm in particular on the porous medium equation. We also consider a semi implicit variant which enables us to treat nonlocal interactions as well as systems of interacting species. Regarding systems, we can also use the ALG2-JKO scheme for the simulation of crowd motion models with several species.

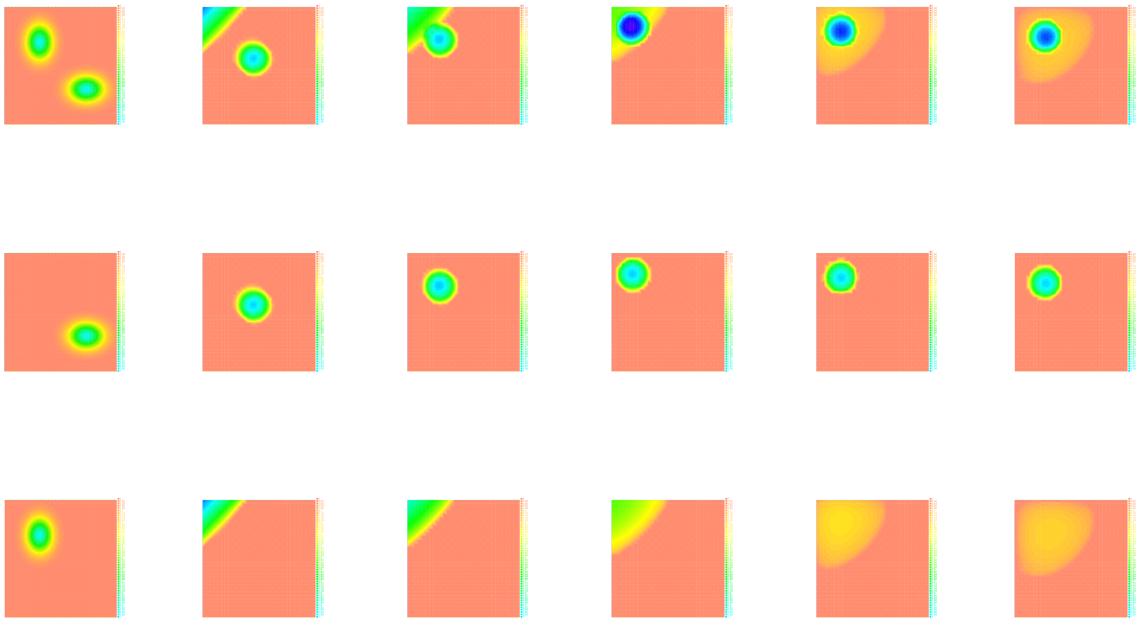

Figure 12. Evolution of two species where the first one is attracted by the other and the second one is repelled by the first one. Top row: display of $\rho_1 + \rho_2$. Middle row: display of ρ_1. Bottom row: display of ρ_2.

We have also investigated the entropy-regularization of the Wasserstein metric to compute gradient flows [19], [34]. This entropic regularization trades the usual Wasserstein fidelity term for a Kullback-Leibler divergence term. Adapting first-order proximal methods to this framework, we have developed numerical schemes which dramatically reduce the computational load needed to simulate the evolution of a mass density through a JKO flow. By construction, the entropy regularization yields an additional diffusion effects to the evolution, but we have proved that a careful choice of the regularization parameter with respect to the timestep yields the convergence of the scheme towards the solutions of the continuous PDE.

A novel Lagrangian method using a discretization of the Monge-Ampère operator for JKO has been developed in [13]. Not only convergence of the scheme has been established but also one advantage of this method is that it makes it possible to use a Newton's method .

7.2. Density Functional Theory

J-D. Benamou Luca Nenna, G. Carlier

In [41] is presented the state of art and recent developments of the optimal transportation theory with many marginals for a class of repulsive cost functions. We introduce some aspects of the Density Functional Theory (DFT) from a mathematical viewpoint, and revisit the theory of optimal transport from its perspective. Moreover, in the last three sections, we describe some recent and new theoretical and numerical results obtained for the Coulomb cost, the repulsive harmonic cost and the determinant.

In [29] we present a numerical method, based on iterative Bregman projections, to solve the optimal transport problem with Coulomb cost. This is related to the strong interaction limit of Density Functional Theory. The first idea is to introduce an entropic regularization of the Kantorovich formulation of the Optimal Transport problem. The regularized problem then corresponds to the projection of a vector on the intersection of the constraints with respect to the Kullback-Leibler distance. Iterative Bregman projections on each marginal constraint are explicit which enables us to approximate the optimal transport plan. We validate the numerical method against analytical test cases.

7.3. Stability for inverse problems with sparsity prior

G. Peyré, V. Duval, Q. Denoyelle,C. Poon

In [42], we have analyzed the recovery performance of two popular finite dimensional approximations of the sparse spikes deconvolution problem over Radon measures, namely the LASSO, and the Continuous Basis-Pursuit. The LASSO is the de-facto standard for the sparse regularization of inverse problems in imaging. It performs a nearest neighbor interpolation of the spikes locations on the sampling grid. The C-BP method, introduced by Ekanadham, Tranchina and Simoncelli, uses a linear interpolation of the locations to perform a better approximation of the infinite-dimensional optimization problem, for positive measures. We have proved that, in the small noise regime, both methods estimate twice the number of original spikes, and we have provided an explicit formula which allows to predict the locations and amplitudes of the spurious spikes. All those properties are in fact connected to an intrinsinc property of the signal: the source condition [16], [24].

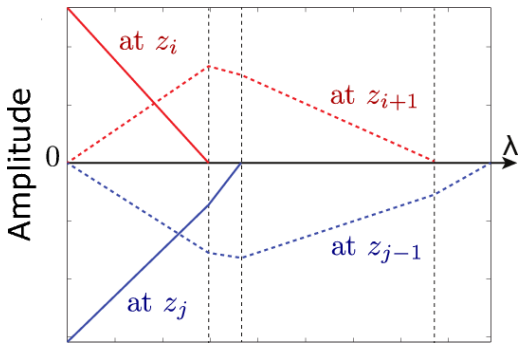

Figure 13. The solution path of the discrete LASSO (as a function of λ) for some discrete measure m_0 (the noise w is set to zero). This shows the amplitudes of the coefficients at $z_i = ih$, resp. $z_j = jh$, (continuous line) and at the next, resp. previous, point of the grid (dashed line) as λ varies.

Those effects are typically due to the use of a discrete grid in the reconstruction process. Several authors have recently proposed algorithms to tackle the problem directly in a continuous setting [75], [92]. As we have shown in [16], the method fails when the distance between spikes with opposite signs are below a certain threshold. However, when all the spikes have the same sign, the LASSO on a continuous domain works for arbitrarily close spikes, being all the more sensitive to noise. In [40], we have given a detailed analysis of the noise sensitivity of the method: if t denotes the minimum separation of the input measure (the minimum distance between two spikes), w refers to the noise and λ is the regularization parameter, when $\|w\|_{L^2}/\lambda$, $\|w\|_{L^2}/t^{2N-1}$ and λ/t^{2N-1} are small enough (where N is the number of spikes), there exists a unique solution to the BLASSO program with exactly the same number of spikes as the original measure. We show that the amplitudes and positions of the spikes of the solution both converge toward those of the input measure when the noise and the regularization parameter drops to zero faster than t^{2N-1}.

7.4. Generalized Solution of Euler

Minimal geodesics along volume preserving maps, through semi-discrete optimal transport

Q. Mérigot and J.-M. Mirebeau introduced a numerical method for extracting minimal geodesics along the group of volume preserving maps, equipped with the L^2 metric, which as observed by Arnold solve Euler's equations of inviscid incompressible fluids. The method relies on the generalized polar decomposition of Brenier, numerically implemented through semi-discrete optimal transport. It is robust enough to extract non-classical, multi-valued solutions of Euler's equations, for which the dimension of the support of the flow is higher than the dimension of the domain, a striking and unavoidable consequence of this model. Our convergence results encompass this generalized model, and our numerical experiments illustrate it for the first time in two space dimensions (see Figure 14).

7.5. Principal Agent Problem

J-D. Benamou, Xavier Dupuis, G. Carlier An alternated projection numerical scheme for the more general c-concavity constraint using Dykstra's algorithm has been recently developed in [33] but being able to handle realistic principal-agent problems remains a challenging issue. Investigating the structure of equilibria in matching problems with non-transferable utilities is also one of our objectives, together with numerical methods in the spirit of the IPFP algorithm.

A semi-discrete approach to the PA problem is investigated. The range of products is discrete and leads to a non convex problem. Non-linear optimization methods are tested. See https://mathmarx.paris.inria.fr:8080.

7.6. Unbalanced Optimal Transport

G. Carlier, F-X. Vialard, B. Schmitzer, L. Chizat Classical optimal transport theory and algorithms assume that the input measures are normalized, i.e. that their total mass is 1. This is an important limitation for many problems in imaging sciences and machine learning, where input data are typically not normalized, and where one should enables local creation or destruction of mass. Handling such "unbalanced" transportation problem is also relevant for applications in biological modeling, for instance to take into account cellular growth through optimal transport gradient flows.

Recently, several researchers of MOKAPLAN made important progress on this problem, by deriving a general framework extending optimal transport to this "unbalanced" setting. In [38] we derived a dynamic optimal transport formulation that enables a source term in the initial formulation of Benamou and Brenier [55]. We proved that it defines a distance on positive measures, enjoy many important properties (dual formulation) and can be computed using fast first order convex optimization methods. We then provided in [39] an even larger class of "unbalanced" optimal transport optimization problems, that are obtained via a static formulation, and show that one can recovers the dynamic formulation in some specific cases. Similar models were derived independently and at the same time by two other international teams [143], [137], which shows the timeliness of our research. We believe these new theoretical and numerical findings will have a strong impact on the developpement of optimal transport methods in imaging sciences and machine learning.

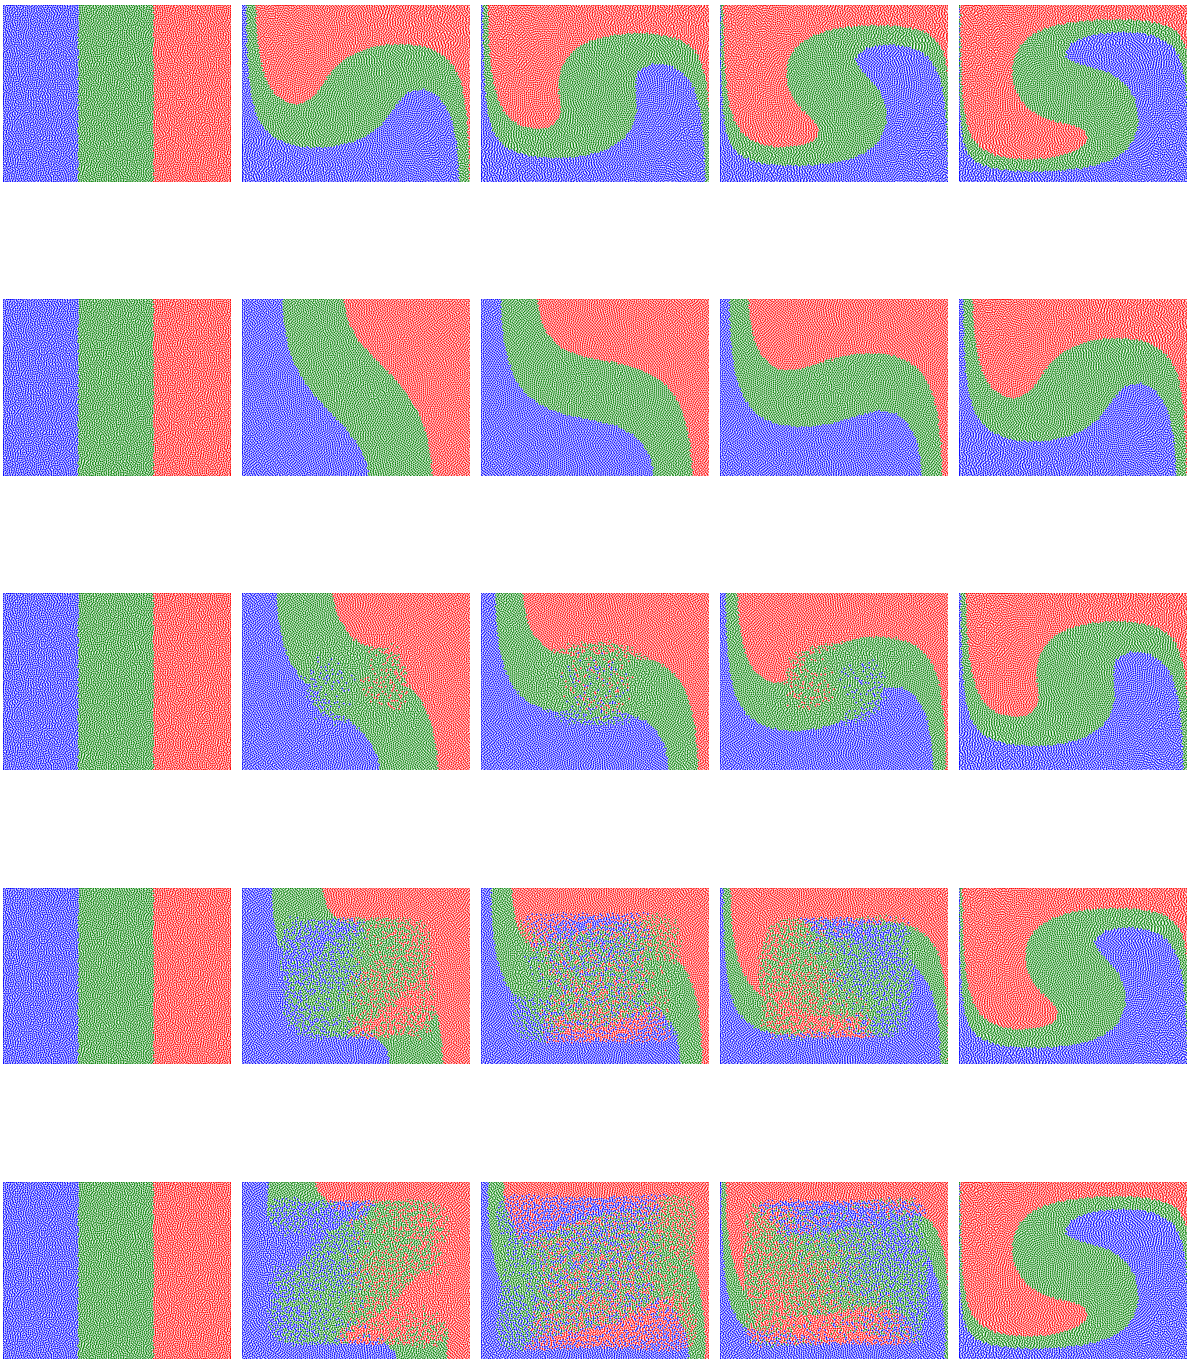

Figure 14. (First row) Beltrami flow in the unit square at various timesteps, a classical solution to Euler's equation. The color of the particles depend on their initial position. (Second to fifth row) Generalized fluid flows that are reconstructed by our algorithm, using boundary conditions displayed in the first and last column. When $t_{\max} < 1$ we recover the classical flow, while for $t_{\max} \geq 1$ the solution is not classical any more and includes some mixing.

MUSE Team

6. New Results

6.1. Home Network or Access Link? Locating Last-mile Downstream Throughput Bottlenecks

Participants: Srikanth Sundaresan (ICSI), Nick Feamster (Princeton), Renata Teixeira

As home networks see increasingly faster downstream throughput speeds, a natural question is whether users are benefiting from these faster speeds or simply facing performance bottlenecks in their own home networks. In our paper recently accepted for publication in PAM'16, we studied whether downstream throughput bottlenecks occur more frequently in their home networks or in their access ISPs. We identified lightweight metrics that can accurately identify whether a throughput bottleneck lies inside or outside a user's home network and developed a detection algorithm that locates these bottlenecks. We validated this algorithm in controlled settings and characterized bottlenecks on two deployments, one of which included 2,652 homes across the United States. We found that wireless bottlenecks are more common than access-link bottlenecks—particularly for home networks with downstream throughput greater than 20 Mbps, where access-link bottlenecks are relatively rare.

6.2. On the Reliability of Profile Matching Across Large Online Social Networks

Participants: Oana Goga and Krishna Gummadi (MPI-SWS), Patrick Loiseau (EURECOM), Robin Sommer (ICSI), Renata Teixeira

Matching the profiles of a user across multiple online social networks brings opportunities for new services and applications as well as new insights on user online behavior, yet it raises serious privacy concerns. Prior literature has showed that it is possible to accurately match profiles, but their evaluation focused only on sampled datasets. In our KDD'15 paper [2], we study the extent to which we can reliably match profiles in practice, across real-world social networks, by exploiting public attributes, i.e., information users publicly provide about themselves. Today's social networks have hundreds of millions of users, which brings completely new challenges as a reliable matching scheme must identify the correct matching profile out of the millions of possible profiles. We first define a set of properties for profile attributes–Availability, Consistency, non-Impersonability, and Discriminability (ACID)–that are both necessary and sufficient to determine the reliability of a matching scheme. Using these properties, we propose a method to evaluate the accuracy of matching schemes in real practical cases. Our results show that the accuracy in practice is significantly lower than the one reported in prior literature. When considering entire social networks, there is a non-negligible number of profiles that belong to different users but have similar attributes, which leads to many false matches. Our paper sheds light on the limits of matching profiles in the real world and illustrates the correct methodology to evaluate matching schemes in realistic scenarios.

6.3. Exploiting crowd sourced reviews to explain movie recommendation

Participants: Sara El Aouad, Christophe Dupuy, Francis Bach, and Renata Teixeira (Inria), Christophe Diot (Technicolor)

Streaming services such as Netflix, M-Go, and Hulu use advanced recommender systems to help their customers identify relevant content quickly and easily. These recommenders display the list of recommended movies organized in sublists labeled with the genre or some more specific labels. Unfortunately, existing methods to extract these labeled sublists require human annotators to manually label movies, which is time-consuming and biased by the views of annotators. In our work [6], we design a method that relies on crowd sourced reviews to automatically identify groups of similar movies and label these groups. Our method takes the content of movie reviews available online as input for an algorithm based on Latent Dirichlet Allocation (LDA) that identifies groups of similar movies. We separate the set of similar movies that share the same combination of genre in sublists and personalize the movies to show in each sublist using matrix factorization. The results of a side-by-side comparison of our method against Technicolor's M-Go VoD service are encouraging.

6.4. Characterizing Home Device Usage From Wireless Traffic Time Series

Participants: Katsiaryna Mirylenka (IBM Research - Zurich), Vassilis Christophides, Themis Palpanas (Paris Descartes University), Ioannis Pefkianakis (Hewlett Packard Labs), Martin May (Technicolor).

The analysis of *temporal behavioral patterns* of home network users can reveal important information to Internet Service Providers (ISPs) and help them to optimize their networks and offer new services (e.g., remote software upgrades, troubleshooting, energy savings). Our study [4] uses time series analysis of continuous traffic data from wireless home networks, to extract traffic patterns recurring within, or across homes, and *assess the impact of different device types (fixed or portable) on home traffic*. Traditional techniques for time series analysis are not suited in this respect, due to the limited stationary and evolving distribution properties of wireless home traffic data. We propose a novel framework that relies on a *correlation-based similarity* measure of time series, as well as a notion of *strong stationarity* to define recurring motifs and dominant devices. Using this framework, we analyze the wireless traffic collected from 196 home gateways over two months. Our framework goes beyond existing application-specific analysis techniques, such as analysis of wireless traffic, which mainly rely on data aggregated across hundreds, or thousands of users. It enables the extraction of recurring patterns from traffic time series of individual homes, leading to a much more fine-grained analysis of the behavior patterns of the users. We also determine the best time aggregation policy w.r.t. to the number and statistical importance of the extracted motifs, as well as the device types dominating these motifs and the overall gateway traffic. Our results show that ISPs can exceed the simple observation of the aggregated gateway traffic and better understand their networks.

6.5. On Continuous Top-k Queries with Real-Time Scoring Functions

Participants: Nelly Vouzoukidou (Google, France), Bernd Amann (LIP6), Vassilis Christophides.

Modern news sharing and social media platforms allow millions of users to *produce and consume information in real-time*. To assess relevance of published information in this new setting, batch scoring based on content similarity, link centrality or page views is no longer sufficient. Instead, streams of events like "replies" (for posting comments), "likes" (for rating content) or "retweets" (for diffusing information) explicitly provided by users represent valuable online feedback on published information that has to be exploited in order to adjust in real-time any available score of information items. Note that in the future Internet of Things (IoT), not only digital, but also physical objects will be expected to be ranked in a fully automated way with respect to real-time human activities (viewing concentration), vital signals (emotional arousal), etc.

Rather than indexing as quickly as possible information items to re-evaluate *snapshot queries*, publish/subscribe systems index *continuous queries* and update on the fly their results each time a new matching item arrives. Existing publish/subscribe systems rely on two alternative continuous filetring semantics, namely *predicate-based* filtering or *similarity-based top-k* filtering. In predicate-based systems, incoming items that match the filtering predicates are simply added to the result list of continuous queries, while in similarity-based top-k publish/subscribe systems, matching items have also to exhibit better relevance w.r.t. the items already appearing as the top-k results of the continuous query. In top-k publish/subscribe systems the relevance of an item remains constant during a pre-specified time window, and once its lifetime exceeds the

item simply expires. Only recently, information recency has become part of the relevance score of continuous queries. Clearly, when information relevance decays as time passes both (a) results lists maintenance and (b) early pruning of the query index traversal are challenged. While these problems have been studied for (textual or spatio-textual) content scoring functions with time decay, non-homogeneous scoring functions accommodating various forms of *query-dependent* and *query-independent* information relevance with time decay is supported only by MeowsReader. In this work we are going beyond this general form of *time-decayed static scores* and consider continuous queries featuring *real-time scoring functions* under the form of *time decaying positive user feedback* for millions of online social media events per minute and millions of user queries.

MUTANT Project-Team

7. New Results

7.1. Weakly-Supervised Discriminative Model for Audio-to-Score Alignment

We consider a new discriminative approach to the problems of segmentation and of audio-to-score alignment. For each musical event, templates have to be built or learnt before performing any alignment. Because annotating a large database music files would be a tedious task, we develop an original approach to learn templates without annotations, but only the knowledge of the music scores associated to music files. We consider the two distinct informations provided by the music scores: (i) an exact ordered list of musical events and (ii) an approximate prior information about relative duration of events. We extend the celebrated Dynamic Time Warping algorithm (DTW) to a convex problem that learns optimal classifiers for all events while jointly aligning files, using this weak supervision only. We show that the relative duration between events can be easily used as a penalization of our cost function and allows us to drastically improve performances of our approach. We describe in details our approach and preliminary results obtained on a large-scale database in [18].

This work was done in collaboration with the SIERRA project-team at Inria Paris.

7.2. Semi-Markov Models for Real-time MIDI-to-Score Alignment

We develop a new stochastic model of symbolic (MIDI) performance of polyphonic scores, based on Semi-Markov models, to align MIDI performances of music scores. In our approach, the evolution of the music performer and the production of performed notes are modeled with a hierarchical extension of hidden semi-Markov models (HSMM). By comparing with a previously studied model based on hidden Markov model (HMM), we give theoretical reasons why the present model is advantageous to deal with complex music event such as trills, tremolos, arpeggios, and other ornaments. This is also confirmed empirically by comparing the accuracy of score following and analysing the errors. We also develop a hybrid of this HSMM-based model and the HMM-based model which is computationally more efficient and retains the advantages of the former model. The present model yields one of the state-of-the-art score following algorithms for symbolic performance and can possibly be applicable for other music recognition problems. Details and results are published in [19].

This work was done in collaboration with Eita Nakamura from the National Institute of Informatics of Tokyo, Japan.

7.3. Real-time Audio-to-Score Alignment of Singing Voice

Singing voice is specific in music: a vocal performance conveys both music (melody/pitch) and lyrics (text/phoneme) content. We develop and original approach that aims at exploiting the advantages of melody and lyric information for real-time audio-to-score alignment of singing voice. First, lyrics are added as a separate observation stream into a template-based hidden semi-Markov model (HSMM), whose observation model is based on the construction of vowel templates. Second, early and late fusion of melody and lyric information are processed during real-time audio-to-score alignment. An experiment conducted with two professional singers (male/female) shows that the performance of a lyrics-based system is comparable to that of melody-based score following systems. Furthermore, late fusion of melody and lyric information substantially improves the alignment performance. Finally, maximum a posteriori adaptation (MAP) of the vowel templates from one singer to the other suggests that lyric information can be efficiently used for any singer. Preliminary results are published in [15].

7.4. Online Methods for Audio Segmentation and Clustering

Audio segmentation is an essential problem in many audio signal processing tasks, which tries to segment an audio signal into homogeneous chunks. Rather than separately finding change points and computing similarities between segments, we focus on joint segmentation and clustering, using the framework of hidden Markov and semi-Markov models. We introduced a new incremental EM algorithm for hidden Markov models (HMMs) and showed that it compares favorably to existing online EM algorithms for HMMs. Early experimental results on musical note segmentation and environmental sound clustering are promising and will be pursued further in 2015.

Theoretical results were published in [11] in collaboration with the SIERRA project-team, and experimental results were further extended in [32]. Early experimental setups show that our algorithms out perform state-of-the-art supervised methods for Percussion Sound classification. In collaboration with IRCyNN (Nantes) we are currently studying algorithmic extensions to complex environmental sounds.

7.5. Adaptive Synchronization Strategies for Automatic Accompaniment

José Echeveste developped several synchonization strategies in the framework of his PhD thesis. Their formalization is based on a dynamic real-time extension of the time map formalism, going beyond state-of-the-art where the largest body of literature on time maps is devoted to static functions, defined and known at all times before any manipulation is done. Only the latest work of Liang and Danneberg (2011) have considered dynamic time map in the synchronization problem. However their approach suffer from a consistency drawback: the convergence of the tempo depends on the events occuring during the catching trajectory. In our approach we have developped a lag-depend formulation of the catching trajectory, which is insensitive to the actual events. This adaptive strategy consider only the deviation in tempo and position and is otherwise context-independant, it ensure convergence both in position and tempo, and it is efficient: there is no need to a fine sampling clock to discretize the time evolution: as long as the prediction time map do not change, delays are computed only once using the accompaniment time map. Our approach is general enough to handle various important issues in automatic accompaniment: latency management, integration of non constant tempo specifications in the score (*accelerando*, *ritardanto*, *rubato*...), handling of missing events, *etc.* Synchronization strategies have been fully formalized in the PhD report of José Echeveste [8] together with a complete Antescofo core including other dynamic constructions.

7.6. Temporal objects for the design of reusable library in Antescofo

Composers develops their own idiosyncratic compositional language through their pieces. In addition, composers and sound engineers have to face drastically different performance set-up for the same piece. This situation advocate for the development of new generic mechanism to simplify the development of generic yet dedicated libraries in Antescofo. In cooperation with various composers (Marco Stroppa, Julia Blondeau, Jason Freeman, Jose Miguel Fernandez, Yann Marez) we have introduced seevral new mechanisms in Antescofo to ease the building of dedicated yet reusable library of compositional pieces: extnesion of the functional language to include new control structure, introduction of *continuation combinators* making possible to start actions at the end of other durative actions, marshalling of Antescofo values, *etc.* The most notable ones are actor-based features to implement *temporal objects*. Object templates are specified and then instantiated at will. A temporal object encapsulate a local state; it can react to logical condition; it offers instantaneous as well as durative methods; reaction to synchronous broadcast can be defined as well as exceptional condition handlers. These new features are currently tested in the development of new pieces and are expected to evolve following the feedbacks from these applications.

7.7. Embedding real-time audio computation in Antescofo

DSP processing in Antescofo is an experimental extension of the language started in 2014 and aimed at driving various DSP processing capabilities directly within Antescofo. DSP processors are defined directly in an Antescofo score, harnessing various signal processing libraries. These DSP processors are then dynamically

connected together using Antescofo audio links. Input and output channels are used to link these processors with the host environment while internal channels connect DSP among themselves. The connections are specified with a new kind of Antescofo actions, the patch. So, the connections can be changed dynamically in response to the events detected by the listening machine and can be synchronized using the expressive repertoire of synchronization strategies available in Antescofo. Ordinary Antescofo variables can be used to control the DSP computations, which add an additional level of dynamicity. Currently, FAUST and and a few specific signal processors (notably FFT) can be defined. Several benefits results of this tight integration. The network of signal processors is heterogeneous, mixing DSP nodes specified with different tools. The network of signal processors can change dynamically in time following the result of a computation. This approach answers the shortcomings of fixed (static) dataflow models of the Max or PureData host environments. Signal processing is controlled at a symbolic level and can be guided, *e.g.* by information available in the augmented score (like position, expected tempo, etc.). The tight integration makes possible to specify, concisely and more effectively, finer and more precise control of the signal processing, at a lower computational cost. One example is the use of symbolic curve specification to specify variations of control parameters at sample rate. It makes it possible to embed sound analysis inside Antescofo as well. At last but not least, signal processing can be done more efficiently. For example, in the *remaking* of Boulez' piece *Antheme 2* there is an improvment of performance in time of 45 % compared to the original version with the audio effects managed in Max.

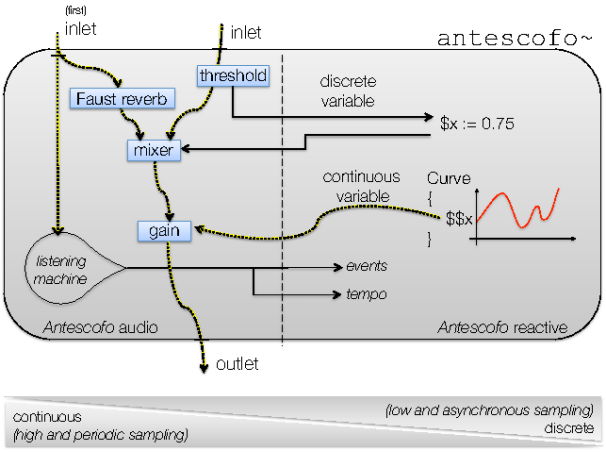

Figure 8. Articulation between DSP and the reactive engine.

The current work focuses on the development of a dedicated type system enabling a finer control of scheduling and audio buffer size, refining results previously developped in the cyclostatic scheduling of synchronous dataflow. Early results are published in [29].

7.8. Visualizing Timed and Hierarchical Code Structures

This work applies an information visualisation perspective to a set of revisions in the timeline-based representation of action items in *AscoGraph*, the dedicated user interface to Antescofo. Our contribution is twofold: (a) a design study of the proposed new model, and (b) a technical, algorithmic component. In the former, we show how our model relates to principles of information coherence and clarity, facility of seeking and navigation, hierarchical distinction and explicit linking. In the latter, we frame the problem of arranging action rectangles in a 2D space as a strip packing problem, with the additional constraint that the (horizontal) time coordinates of each block are fixed. We introduce three algorithms of increasing complexity for automatic arrangement, estimate their packing performance and analyse their strengths and weaknesses. We evaluate the systemic

improvements achieved and their applicability for other time-based datasets. Furthermore, algorithms for efficient automatic stacking of time-overlapping action blocks are developped, as well as mathematical proof for their time-coherency during dynamic visualizations.

Results are implemented in Section 6.2 and reported in [12] and [13].

7.9. Model-based Testing an Interactive Music System

We have been pursuing our studies on the application of model-based timed testing techniques to the interactive music system (IMS) Antescofo, in the context of the Phd of Clément Poncelet and in relation with the developments presented in Section 6.3 .

Several formal methods have been developed for automatic conformance testing of critical embedded software, with the execution of a real implementation under test (IUT, or black-box) in a testing framework, where carefully selected inputs are sent to the IUT and then the outputs are observed and analyzed. In conformance model-based testing (MBT), the input and corresponding expected outputs are generated according to formal models of the IUT and the environment. The case of IMS presents important originalities compared to other applications of MBT to realtime systems. On the one hand, the time model of IMS comprises several time units, including the wall clock time, measured in seconds, and the time of music scores, measured in number of beats relatively to a tempo. This situation raises several new problems for the generation of test suites and their execution. On the other hand, we can reasonably assume that a given mixed score of Antescofo specifies completely the expected timed behavior of the IMS, and compile automatically the given score into a formal model of the IUT's expected behavior, using an intermediate representation. This give a fully automatic test method, which is in contrast with other approaches which generally require experts to write the specification manually.

We have developed online and offline approches to MBT for Antescofo. The offline approach relies on tools of the Uppaal suite [38], [37], using a translation of our models into timed automata. These results have been presented during the 30th ACM/SIGAPP Symposium On Applied Computing, track Software Verification and Testing [21] and an article describing this approach has been accepted for publication in the Journal of New Music Research. The online approach is based on a new virtual machine executing the models of score in intermediate representation (see Section 6.3).

7.10. Representation of Rhythm and Quantization

Rhythmic data are commonly represented by tree structures (rhythms trees) in assisted music composition environments, such as OpenMusic, due to the theoretical proximity of such structures with traditional musical notation. We are studying the application in this context of techniques and tools for processing tree structure, which were originally developed for other areas such as natural language processing, automatic deduction, Web data processing... We are particularly interested in two well established formalisms with solid theoretical foundations: tree automata and term rewriting.

Our first main contribution in that context is the development of a new framework for rhythm transcription, the problem of the generation, from a sequence of timestamped notes, *e.g.* a file in MIDI format, of a score in traditional music notation) – see Section 6.4 . This problem arises immediately as insoluble unequivocally: we shall calibrate the system to fit the musical context, balancing constraints of precision, or of simplicity / readability of the generated scores. We are developing in collaboration with Jean Bresson (Ircam) and Slawek Staworko (LINKS, currently on leave at University of Edinburgh) an approach based on algorithms for the enumeration of large sets of weighted trees (tree series), representing possible solutions to a problem of transcription. The implementation work is performed by Adrien Ycart, under a research engineer contract with Ircam. This work has been presented in [23].

Our second contribution, in collaboration with Prof. Masahiko Sakai (Nagoya University), is a proposal of a structural theory (equational system on rhythm trees) defining equivalence on rhythm notations [14], [16]. This approach can be used for example to generate, by transformation, different notations possible the same rate, with the ability to select in accordance with certain constraints. We have also conducted related work on the theory of term rewriting [17].

MYCENAE Project-Team

6. New Results

6.1. Numerical and theoretical studies of slow-fast systems with complex oscillations

6.1.1. Canard-Mediated (De)Synchronization in Coupled Phantom Bursters

Participants: Elif Köksal Ersöz, Mathieu Desroches, Maciej Krupa, Frédérique Clément.

In [32], we study canard-mediated transitions in mutually coupled phantom bursters. We extend a multiple-timescale model which provides a sequence of dynamic events, i.e. transition from a frequency modulated relaxation cycle to a quasi-steady state and resumption of the relaxation regime through small amplitude oscillations. Folded singularities and associated canard solutions have a particular impact on the dynamics of the original system, which consists of two feedforward coupled FitzHugh-Nagumo oscillators, where the slow subsystem (regulator) controls the periodic behavior of the fast subsystem (secretor). We first investigate the variability in the dynamics depending on the canard mechanism that occurs near the folded singularities of the 4D secretor- regulator configuration. Then, we introduce a second secretor and focus on the slow-fast transitions in the presence of a linear coupling between the secretors. In particular, we explore the impact of the relationship between the canard structures and the coupling on patterns of synchronization and desynchronization of the collective dynamics of the resulting 6D system. We identify two different sources of desynchronization induced by canards, near a folded-saddle singularity and a folded-node singularity, respectively.

Part of these results have also been presented as posters at the *SIAM Conference on Applications of Dynamical Systems* (Snowbird, May 17-21, 2015) and 1st *International Conference on Mathematical Neuroscience* (Antibes Juan les Pins, June 8-10-2015).

6.1.2. Mixed-Mode Oscillations in a piecewise linear system with multiple time scale coupling

Participants: Soledad Fernández García, Maciej Krupa, Frédérique Clément.

We analyze a four dimensional slow-fast piecewise linear system with three time scales presenting Mixed-Mode Oscillations. The system possesses an attractive limit cycle along which oscillations of three different amplitudes and frequencies can appear, namely, small oscillations, pulses (medium amplitude) and one surge (largest amplitude). In addition to proving the existence and attractiveness of the limit cycle, we focus our attention on the canard phenomena underlying the changes in the number of small oscillations and pulses. We analyze locally the existence of secondary canards leading to the addition or subtraction of one small oscillation and describe how this change is globally compensated for or not with the addition or subtraction of one pulse.

6.1.3. Noise-induced canard and mixed-mode oscillations in large stochastic networks with multiple timescales

Participants: Jonathan Touboul, Maciej Krupa, Mathieu Desroches.

We investigate in [28] the dynamics of large stochastic networks with different timescales and nonlinear mean-field interactions. After deriving the limit equations for a general class of network models, we apply our results to the celebrated Wilson-Cowan system with two populations with or without slow adaptation, paradigmatic example of nonlinear mean-field network. This system has the property that the dynamics of the mean of the solution exactly satisfies an ODE. This reduction allows to show that in the mean-field limit and in multiple populations with multiple timescales, noise induces canard explosions and Mixed-Mode Oscillations on the mean of the solution. This sheds new light on the qualitative effects of noise and sensitivity to precise noise values in large stochastic networks. We further investigate finite-sized networks and show that systematic differences with the mean-field limits arise in bistable regimes (where random switches between different attractors occur) or in mixed-mode oscillations, were the finite-size effects induce early jumps due to the sensitivity of the attractor.

6.1.4. Canard explosion in delayed equations with multiple timescales, applications to the delayed Fitzhugh-Nagumo system

Participants: Maciej Krupa, Jonathan Touboul.

In two contributions, we investigated theoretically the presence of canard explosions of delayed differential equations, and have applied these results to the FitzHugh-Nagumo neuronal model.

- In [21] we analyze canard explosions in delayed differential equations with a one-dimensional slow manifold. This study is applied to explore the dynamics of the van der Pol slow-fast system with delayed self-coupling. In the absence of delays, this system provides a canonical example of a canard explosion. We show that as the delay is increased a family of 'classical' canard explosions ends as a Bogdanov-Takens bifurcation occurs at the folds points of the S-shaped critical manifold.
- Motivated by the dynamics of neuronal responses, we analyze in [21] the dynamics of the Fitzhugh-Nagumo slow-fast system with delayed self-coupling. Beyond the regime of small delays, delays significantly enrich the dynamics, leading to mixed-mode oscillations, bursting and chaos. These behaviors emerge from a delay-induced subcritical Bogdanov-Takens instability arising at the fold points of the S-shaped critical manifold. Underlying the transition from canard-induced to delay-induced dynamics is an abrupt switch in the nature of the Hopf bifurcation.

6.1.5. Canard-induced loss of stability across a homoclinic bifurcation

Participants: Mathieu Desroches, Jean-Pierre Françoise, Lucile Megret.

In [16], we investigate the possibility of bifurcations which display a dramatic change in the phase portrait in a very small (on the order of 10^{-7} in the example presented here) change of a parameter. We provide evidence of existence of such a very rapid loss of stability on a specific example of a singular perturbation setting. This example is strongly inspired of the explosion of canard cycles first discovered and studied by E. Benoît, J.-L. Callot, F. Diener and M. Diener. After some presentation of the integrable case to be perturbed, we present the numerical evidences for this rapid loss of stability using numerical continuation. We discuss then the possibility to estimate accurately the value of the parameter for which this bifurcation occurs.

6.1.6. Analysis of Interspike-Intervals for the General Class of Integrate-and-Fire Models with Periodic Drive

Participant: Justyna Signerska-Rynkowska.

In [27], we study one-dimensional integrate-and-fire models of the general type $\dot{x} = F(t, x)$ and analyze properties of the firing map which iterations recover consecutive spike timings. We impose very week constraints for the regularity of the function $F(t, x)$ e.g. often it suffices to assume that F is continuous. If additionally F is periodic in t, using mathematical study of the displacement sequence of an orientation preserving circle homeomorphism, we provide a detailed description of the regularity properties of the sequence of interspike-intervals and behaviour of the interspike-interval distribution.

6.1.7. A geometric mechanism for mixed-mode bursting oscillations in a hybrid neuron model

Participants: Justyna Signerska-Rynkowska, Jonathan Touboul, Alexandre Vidal.

In [35], we exhibit and investigate a new type of mechanism for generating complex oscillations featuring an alternation of small oscillations with spikes (MMOs) or bursts (MMBOs) in a class of hybrid dynamical systems modeling neuronal activity. These dynamical systems, called nonlinear adaptive integrate-and-fire neurons, combine nonlinear dynamics modeling input integration in a nerve cell with discrete resets modeling the emission of an action potential and the subsequent return to reversal potential. We show that presence of complex oscillations in these models relies on a fundamentally hybrid structure of the flow: invariant manifolds of the continuous dynamics govern small oscillations, while discrete resets govern the emission of spikes or bursts. The decomposition into these two mechanisms leads us to propose a purely geometrical interpretation of these complex trajectories, and this relative simplicity allows to finely characterize the MMO patterns through the study of iterates of the adaptation map associated with the hybrid system. This map is however

singular: it is discontinuous and has unbounded left- and right-derivatives. We apply and develop rotation theory of circle maps for this class of adaptation maps to precisely characterize the trajectories with respect to the parameters of the system. In contrast to more classical frameworks in which MM(B)Os were evidenced, the present geometric mechanism neither requires no more than two dimensions, does not necessitate to have separation of timescales nor complex return mechanisms.
Part of these results have also been presented as posters at the *SIAM Conference on Applications of Dynamical Systems* (Snowbird, May 17-21, 2015) and 1st *International Conference on Mathematical Neuroscience* (Antibes Juan les Pins, June 8-10-2015).

6.2. Non conservative transport equations for cell population dynamics

6.2.1. Cell-kinetics based calibration of a multiscale model: application to cell population dynamics in ovarian follicles

Participants: Benjamin Aymard [ICL], Frédérique Clément, Danielle Monniaux [INRA], Marie Postel.

In [30], we present a strategy for tuning the parameters of a multiscale model of structured cell populations in which physiological mechanisms are embedded into the cell scale. This strategy allows one to cope with the technical difficulties raised by such models, that arise from their anchorage in cell biology concepts: localized mitosis, progression within and out of the cell cycle driven by time- and possibly unknown-dependent, and nonsmooth velocity coefficients. We compute different mesoscopic and macroscopic quantities from the microscopic unknowns (cell densities) and relate them to experimental cell kinetic indexes. We study the expression of reaching times corresponding to characteristic cellular transitions in a particle-like reduction of the original model. We make use of this framework to obtain an appropriate initial guess for the parameters and then perform a sequence of optimization steps subject to quantitative specifications. We finally illustrate realistic simulations of the cell populations in cohorts of interacting ovarian follicles.

6.2.2. Dimensional reduction of a multiscale cell population model

Participants: Frédérique Clément, Frédéric Coquel [CMAP], Marie Postel, Kim Long Tran.

We have designed a dimensional reduction of a multiscale structured cell population model, consisting of a system of 2D transport equations, into a system of twice as many 1D transport equations. The reduced model is obtained by computing the moments of the 2D model with respect to one space variable. The 1D solution is defined from the solution of the 2D model starting from an initial condition that is a Dirac mass in the direction removed by reduction. Long time properties of the 1D model solution are obtained in connection with properties of the support of the 2D solution for general case initial conditions. Finite volume numerical approximations of the 1D reduced model can be used to compute the moments of the 2D solution with satisfying accuracy. The numerical robustness is studied in the scalar case and a full scale vector case is presented.

6.3. Macroscopic limits of stochastic neural networks and neural fields

6.3.1. Pinwheel-Dipole configuration in cat visual cortex

Participants: Jérôme Ribot [CIRB], Alberto Romagnoni [CIRB], Chantal Milleret [CIRB], Daniel Bennequin [CIRB], Jonathan Touboul.

One fascinating aspect of the brain is its ability to process information in a fast and reliable manner. The functional architecture is thought to play a central role in this task, by encoding efficiently complex stimuli and facilitating higher level processing. In the early visual cortex of higher mammals, information is processed within functional maps whose layout is thought to underlie visual perception. The possible principles underlying the topology of the different maps, as well as the role of a specific functional architecture on information processing, is however poorly understood.

- In [25], we show that spatial frequency representation in cat areas 17 and 18 exhibits singularities around which the map organizes like an electric dipole potential. These singularities are precisely co-located with singularities of the orientation map: the pinwheel centers. We first show, using high resolution optical imaging, that a large majority (around 80%) of pinwheel centers exhibit in their neighborhood semi-global extrema in the spatial frequency map. These extrema created a sharp gradient that was confirmed with electrophysiological recordings. Based on an analogy with electromagnetism, a mathematical model of a dipolar structure is proposed, that was accurately fitted to optical imaging data for two third of pinwheel centers with semi-global extrema.
- Mathematically, this pinwheel-dipole architecture is fascinating. We demonstrated mathematically in [26] that two natural principles, local exhaustivity of representation and parsimony, would indeed constrain the orientation and spatial frequency maps to display co-located singularities around which the orientation is organized as a pinwheel and spatial frequency as a dipole. Moreover, using a computational model, we showed that this architecture allows a trade-off in the local perception of orientation and spatial frequency, but this would occur for sharper selectivity than the tuning width reported in the literature. We therefore re-examined physiological data and show that indeed the spatial frequency selectivity substantially sharpens near maps singularities, bringing to the prediction that the system tends to optimize balanced detection between different attributes.

These results shed new light on the principles at play in the emergence of functional architecture of cortical maps, as well as their potential role in processing information.

6.3.2. Absorption properties of stochastic equations with Hölder diffusion coefficients

Participants: Jonathan Touboul, Gilles Wainrib [ENS].

In [29], we address the absorption properties of a class of stochastic differential equations around singular points where both the drift and diffusion functions vanish. According to the Hölder coefficient alpha of the diffusion function around the singular point, we identify different regimes. Stability of the absorbing state, large deviations for the absorption time, existence of stationary or quasi-stationary distributions are discussed. In particular, we show that quasi-stationary distributions only exist for alpha < 3/4, and for alpha in the interval (3/4, 1), no quasi-stationary distribution is found and numerical simulations tend to show that the process conditioned on not being absorbed initiates an almost sure exponential convergence towards the absorbing state (as is demonstrated to be true for alpha = 1). Applications of these results to stochastic bifurcations are discussed.

6.3.3. On a kinetic FitzHugh-Nagumo model of neuronal network

Participants: Stéphane Mischler [CEREMADE], Cristóbal Quiñinao [CIRB], Jonathan Touboul.

We investigate in [33] the existence and uniqueness of solutions of a McKean-Vlasov evolution PDE representing the macroscopic behavior of interacting Fitzhugh-Nagumo neurons. This equation is hypoelliptic, nonlocal and has unbounded coefficients. We proved existence of a solution to the evolution equation and non trivial stationary solutions. Moreover, we demonstrated uniqueness of the stationary solution in the weakly nonlinear regime. Eventually, using a semigroup factorisation method, we showed exponential nonlinear stability in the small connectivity regime.

6.4. Modeling of neurogenesis and brain development

6.4.1. Lhx2 regulates the timing of β-catenin-dependent cortical neurogenesis

Participants: Lea-Chia-Ling Hsu [Taipei], Sean Nama [Taipei], Yi Cui, Ching-Pu Chang [Taipei], Chia-Fang Wang [Taipei], Hung-Chih Kuo [Taipei], Jonathan Touboul, Shen-Ju Chou [Taipei].

The timing of cortical neurogenesis has a major effect on the size and organization of the mature cortex. The deletion of the LIM-homeodomain transcription factor Lhx2 in cortical progenitors by Nestin-cre leads to a dramatically smaller cortex. In [19] we report that Lhx2 regulates the cortex size by maintaining the cortical progenitor proliferation and delaying the initiation of neurogenesis. The loss of Lhx2 in cortical progenitors results in precocious radial glia differentiation and a temporal shift of cortical neurogenesis. We further investigated the underlying mechanisms at play and demonstrated that in the absence of Lhx2, the Wnt/β-catenin pathway failed to maintain progenitor proliferation. We developed and applied a mathematical model that reveals how precocious neurogenesis affected cortical surface and thickness. Thus, we concluded that Lhx2 is required for β-catenin function in maintaining cortical progenitor proliferation and controls the timing of cortical neurogenesis.

6.4.2. Competition and boundary formation in heterogeneous media: Application to neuronal differentiation

Participants: Cristóbal Quiñinao [CIRB], Benoît Perthame [LJLL], Jonathan Touboul.

We analyze in [22] an inhomogeneous system of coupled reaction-diffusion equations representing the dynamics of gene expression during differentiation of nerve cells. The outcome of this developmental phase is the formation of distinct functional areas separated by sharp and smooth boundaries. It proceeds through the competition between the expression of two genes whose expression is driven by monotonic gradients of chemicals, and the products of gene expression undergo local diffusion and drive gene expression in neighboring cells. The problem therefore falls in a more general setting of species in competition within a non-homogeneous medium. We show that in the limit of arbitrarily small diffusion, there exists a unique monotonic stationary solution, which splits the neural tissue into two winner-take-all parts at a precise boundary point: on both sides of the boundary, different neuronal types are present. In order to further characterize the location of this boundary, we use a blow-up of the system and define a traveling wave problem parametrized by the position within the monotonic gradient: the precise boundary location is given by the unique point in space at which the speed of the wave vanishes.

6.4.3. Local homeoprotein diffusion can stabilize boundaries generated by graded positional cues

Participants: Cristóbal Quiñinao [CIRB], Alain Prochiantz [CIRB], Jonathan Touboul.

Boundary formation in the developing neuroepithelium decides on the position and size of compartments in the adult nervous system. In [23], we started from the French Flag model proposed by Lewis Wolpert, in which boundaries are formed through the combination of morphogen diffusion and of thresholds in cell responses. In contemporary terms, a response is characterized by the expression of cell-autonomous transcription factors, very often of the homeoprotein family. Theoretical studies suggest that this sole mechanism results in the formation of boundaries of imprecise shapes and positions. Alan Turing, on the other hand, proposed a model whereby two morphogens that exhibit self-activation and reciprocal inhibition, and are uniformly distributed and diffuse at different rates lead to the formation of territories of unpredictable shapes and positions but with sharp boundaries (the 'leopard spots'). Here, we have combined the two models and compared the stability of boundaries when the hypothesis of local homeoprotein intercellular diffusion is, or is not, introduced in the equations. We find that the addition of homeoprotein local diffusion leads to a dramatic stabilization of the positioning of the boundary, even when other parameters are significantly modified. This novel Turing/Wolpert combined model has thus important theoretical consequences for our understanding of the role of the intercellular diffusion of homeoproteins in the developmental robustness of and the changes that take place in the course of evolution.

6.4.4. Designing a mathematical model of the dynamics of progenitor cell populations in the mouse cerebral cortex

Participants: Marie Postel, Alice Karam [UPMC], Mérina Latbi [UPMC], Guillaume Pezeron [UPMC], Kim Long Tran, Frédérique Clément, Sylvie Schneider-Maunoury [UPMC].

The mammalian cortex is a laminar structure in the dorsal telencephalon, composed of distinct cell types with different spatial and temporal origins. Cortical projection neurons display different patterns of layering and connectivity that depend on their birth date. We have designed a multi-scale mathematical model of structured cell populations, taking into account three main cell types: apical progenitors (APs), intermediate progenitors (IPs) and neurons (N). APs self-renew and produce IPs that divide to give Ns. The main originality of this spatio-temporal model is to explicitly represent the different phases of the cell cycle, G1, S, G2 and M. Biological data from the experiments and from the literature provide values for parameters of the model (e.g. duration of each cell cycle phase and division rates for each cell type). The outputs of the model are interpretable in terms of cell kinetics (e.g. mitotic index, labelling index, cell numbers). They are adjusted to experimental observations by numerical simulation.

PARKAS Project-Team

6. New Results

6.1. Reasoning about C11 Program Transformations

Participants: Francesco Zappa Nardelli, Robin Morisset.

We have shown that the weak memory model introduced by the 2011 C and C++ standards does not permit many of common source-to-source program transformations (such as expression linearisation and "roach motel" reordering) that modern compilers perform and that are deemed to be correct. As such it cannot be used to define the semantics of intermediate languages of compilers, as, for instance, LLVM aimed to. We consider a number of possible local fixes, some strengthening and some weakening the model. We have evaluated the proposed fixes by determining which program transformations are valid with respect to each of the patched models. We have provided formal Coq proofs of their correctness or counterexamples as appropriate.

A paper on this work has been accepted in [18]. In collaboration with Viktor Vafeiadis (MPI-SWS, Germany) and Thibaut Balabonski (U. Paris Sud).

6.2. Language design on top of JavaScript

Participant: Francesco Zappa Nardelli.

This research project aims at improving the design of the JavaScript language. We propose a typed extension of JavaScript combining dynamic types, concrete types and like types to let developers pick the level of guarantee that is appropriate for their code. We have implemented our type system in the V8 JavaScirpt engine and we have explored the performance and software engineering benefits.

A paper on this work has been accepted in ECOOP 2015 [21].

With Gregor Richards (Waterloo University) and Jan Vitek (Northeastern University).

6.3. Synchronous Functional Language with Integer Clocks

Participant: Adrien Guatto.

Adrien Guatto defended his PhD thesis on the modular description of space/time tradeoffs at the language level. His thesis work extends the n-synchronous framework proposed by Cohen, Mandel, Plateau, Pouzet and others. Clocks now feature arbitrary positive integers that model bursty communication between subprograms: “integer clocks”. The activation conditions of Lustre are revisited in this new setting to become “local time scales” that allow subprograms to perform several steps atomically relative to their context. The thesis details the integration of these features in a clock type system for a higher-order functional language, giving full formal treatment of its metatheory and compilation to finite-state digital circuits.

6.4. Fidelity in Real-Time Programming

Participants: Guillaume Baudart, Timothy Bourke.

In this work we study embedded systems with a significant mix of discrete reactive behaviours and ‘physical’ timing constraints. The idea is to make the most of the advantages of synchronous languages for precisely specifying discrete behaviours but to adapt or extend them to treat real-time constraints more abstractly, that is, without an *a priori* definition of an eventual sampling interval.

This year we concluded our study of the Loosely Timed-Triggered Architectures (LTTA) by developing simplified models of the underlying implementations and protocols. This enabled us to improve the protocols, simplify the correctness and performance arguments, and compare them to systems built using modern clock synchronization algorithms. We developed our models in the Zélus programming language which enables (instances of) them to be compiled for simulation and contributes to our work on better exploiting synchronous languages for real-time specification and analysis. This work was presented at the EMSOFT conference and a journal article has been submitted.

This year we also concluded our study of the Quasi-synchronous Approach to modelling real-time distributed systems. We formalized the relation between the discrete abstraction proposed by Paul Caspi and the real-time archictectures for which it is intended. This enabled us to precisely state a correctness requirement for the abstraction and to show that it is sound for systems of two nodes (a typical case explored in other publications) but not for general systems of three or more nodes. Our formalization clarifies the relation between the causality of traces of the real-time system and the causality introduced by the synchronous abstraction. This enables us to state and show necessary and sufficient restrictions on the communication topologies and timing characteristics of systems to ensure soundness. A paper explaining this result has been drafted and will be submitted early in 2016.

6.5. Verified compilation of Lustre

Participants: Timothy Bourke, Marc Pouzet.

Synchronous dataflow languages and their compilers are increasingly used to develop safety-critical applications, like fly-by-wire controllers in aircraft and monitoring software for power plants. A striking example is the SCADE Suite tool of ANSYS/Esterel Technologies which is DO-178B/C qualified for the aerospace and defense industries. This tool allows engineers to develop and validate systems at the level of abstract block diagrams that are automatically compiled into executable code.

Formal modelling and verification in an interactive theorem prover can potentially complement the industrial certification of such tools to give very precise definitions of language features and increased confidence in their correct compilation; ideally, right down to the binary code that actually executes.

This year we picked up on previous work in the PARKAS team to develop a verified compiler for a Lustre/SCADE-like synchronous language. We focused on the critical and until now unresolved compiler stage that transforms dataflow equations into imperative code. We developed, in Coq, a prototype compiler for the core language (without modular resets or tuples) and showed its correctness with respect to a dataflow semantics based on functions from natural numbers to present or absent values. This required the development of a novel intermediate model for relating delayed dataflow streams to imperative memories in such a way that a critical induction could be stated and proved. We further showed how to justify a post-transformation optimization that is essential for the efficiency of clock-directed code generation. We are preparing a paper describing these results. Work continues on both semantic questions (existence of a semantics for well-type and well-clocked programs, treatment of resets, etc.) and compilation issues (integration with the verified CompCert compiler).

In collaboration with Pierre-Évariste Dagand (CNRS) and Lionel Reig (Collège de France).

PI.R2 Project-Team

6. New Results

6.1. Effects in proof theory and programming

Participants: Guillaume Claret, Pierre-Louis Curien, Hugo Herbelin, Étienne Miquey, Ludovic Patey, Pierre-Marie Pédrot, Yann Régis-Gianas, Alexis Saurin.

6.1.1. Axiom of dependent choice in classical arithmetic

In 2012, Hugo Herbelin showed that classical arithmetic in finite types extended with strong elimination of existential quantification proves the axiom of dependent choice. To get classical logic and choice together without being inconsistent is made possible first by constraining strong elimination of existential quantification to proofs that are essentially intuitionistic and secondly by turning countable universal quantification into an infinite conjunction of classical proofs evaluated along a call-by-need evaluation strategy so as to extract from them intuitionistic contents that complies to the intuitionistic constraint put on strong elimination of existential quantification. Étienne Miquey has been working on a sequent-calculus version of this system, using Danvy's methodology of semantic artifacts, to progressively reduce the consistency of such a system to the normalisation of Girard-Reynold's system F. To achieve this goal, he incidentally proposed a way to get a dependently-typed sequent calculus, as well as a method to type a state-and-continuation-passing style translation of call-by-need calculus.

6.1.2. The computational contents of completeness proofs

Hugo Herbelin worked on the computational content of Gödel's completeness theorem, developing a proof with side-effects suitable for normalisation-by-evaluation.

6.1.3. Gödel's functional interpretation

Pierre-Marie Pédrot extended the proof-as-program interpretation of Gödel's Dialectica translation to the fully dependent setting, including dependent elimination [17].

6.1.4. Logical foundations of call-by-need evaluation

Alexis Saurin and Pierre-Marie Pédrot extended their reconstruction of call-by-need based on linear head reduction with control. They showed how linear head reduction could be adapted to the $\lambda\mu$-calculus. This classical linear head reduction lifts the usual properties of the intuitionistic one (with respect to σ-equivalence) to the $\lambda\mu$-calculus (and its σ-equivalence already formulated by Olivier Laurent in his PhD thesis). Moreover, they showed that substitution sequences of the $\lambda\mu$-calculus linear head reduction are in correspondence with the classical Krivine abstract machine substitution sequences, validating the known fact that the KAM implements linear head reduction. In a second step, they could lift to the $\lambda\mu$-calculus their three-step transformation from linear head reduction to call-by-need, and study the correspondence with Ariola, Herbelin and Saurin's classical call-by-need. This work appeared as one of the chapters of Pierre-Marie Pédrot's thesis and has been accepted for publication at ESOP'16 [30].

6.1.5. Call-by-name forcing

Pierre-Marie Pédrot studied variants of the forcing construction by decomposing it through call-by-push-value. In particular, the by-name decomposition behaves much more nicely w.r.t. the computational content of proofs and is a candidate for a dependently-typed extension. This work is partially reported on in his PhD [17].

6.1.6. A theory of effects and resources

In joint work with Marcelo Fiore and Guillaume Munch-Maccagnoni, Pierre-Louis Curien considered the Curry-Howard-Lambek correspondence for effectful computation and resource management, specifically proposing polarised calculi together with presheaf-enriched adjunction models as the starting point for a comprehensive semantic theory relating logical systems, typed calculi, and categorical models in this context. Our thesis is that the combination of effects and resources should be considered orthogonally. Model theoretically, this leads to an understanding of our categorical models from two complementary perspectives: (i) as a linearisation of CBPV (Call-by-Push-Value) adjunction models, and (ii) as an extension of linear/non-linear adjunction models with an adjoint resolution of computational effects. When the linear structure is cartesian and the resource structure is trivial, we recover Levy's notion of CBPV adjunction model, while when the effect structure is trivial, we have Benton's linear/non-linear adjunction models. Further instances of our model theory include the dialogue categories with a resource modality of Melliès and Tabareau, and the Enriched Effect Calculus models of Egger, Møgelberg and Simpson. Our development substantiates the approach by providing a lifting theorem of linear models into cartesian ones. To each of our categorical models we systematically associate a typed term calculus, each of which corresponds to a variant of the sequent calculi LJ (Intuitionistic Logic) or ILL (Intuitionistic Linear Logic). The adjoint resolution of effects corresponds to polarisation whereby, syntactically, types locally determine a strict or lazy evaluation order and, semantically, the associativity of cuts is relaxed. In particular, our results show that polarisation provides a computational interpretation of CBPV in direct style. Further, we characterise depolarised models: those where the cut is associative, and where the evaluation order is unimportant. This work will be presented at POPL 2016 [26].

6.1.7. Coq as a programming language with effects

As part of his PhD thesis, Guillaume Claret defined a notion of effectful interactive computation as an embedded DSL in Coq (in the spirit of the works on algebraic effects), and used it to implement a web server. It is equipped with a dual notion of effectful interactive execution context. Using these two notions together, Guillaume Claret is able to specify and reason about interactive programs inside Coq. He submitted several papers about this line of work: one has been published [32], others will be part of his PhD manuscript.

6.2. Reasoning and programming with infinite data

Participants: Amina Doumane, Alexis Saurin, Pierre-Marie Pédrot, Yann Régis-Gianas.

This theme is part of the ANR project Rapido (see the National Initiatives section).

6.2.1. Interactive semantics for logic fixed-points and infinitary logics.

Amina Doumane and Alexis Saurin, in a joint work with David Baelde published at CSL 2015 [24], developed a game-semantics of $\mu MALL$ (Multiplicative Additive Linear Logic with least and greatest fixpoints).

This interactive semantics was worked out in computational ludics, benefitting from both the work by Clairambault on a HO style game semantics for an intuitionistic logic with least and greatest fixpoints and from the flexibility of Terui's computational ludics (in particular its ability to consider designs with cuts).

This framework is built around the notion of design, which can be seen as an analogue of the strategies of game semantics. The infinitary nature of designs makes them particularly well suited for representing computations over infinite data. We provided $\mu MALL$ with a denotational semantics (that is invariant by cut-elimination), interpreting proofs by designs and formulas by particular sets of designs called behaviours. Then a completeness result for a specific class of designs is proved, the class of "essentially finite designs", which are those designs performing a finite computation followed by a copycat. On the way to the previous completeness result, we investigate semantic inclusion, proving its decidability (given two formulas A and B, one can decide whether the semantics of A is included in the semantics of B) and completeness (if semantic inclusion holds, the corresponding implication is provable in $\mu MALL$).

6.2.2. Proof theory of circular proofs

In a collaboration with David Baelde, Amina Doumane and Alexis Saurin developed further the theory of infinite proofs. Studying the proof theory of circular proofs on MALL, they established a result of focalisation for these infinite proofs. The usual result of focalisation for linear logic can actually be extended to circular proofs but, contrarily to $\mu MALL$ where fixed-points operators can be given an arbitrary polarity, the least fixed-point must be set to be a positive construction and the greatest fixed-points to be negative, which is consistent with intuition from programming with inductive and co-inductive datatypes. An interesting phenomenon arising with focalisation is that some infinite but regular proofs may not have any regular focused proofs. This is similar to what happens for cut-elimination of regular proofs.

Works on cut-elimination for circular proofs are still ongoing.

6.2.2.1. Automata theory meets proof theory: proof certificates for Büchi inclusion

In a joint work with David Baelde and Lucca Hirschi, Amina Doumane and Alexis Saurin carried out a proof-theoretical investigation of the linear-time μ-calculus, proposing well-structured proof systems and showing constructively that they are complete for inclusions of Büchi automata suitably encoded as formulas.

They do so in a way that combines the advantages of two lines of previous work: Kaivola gave a proof of completeness for an axiomatisation that amounts to a finitary proof system, but his proof is non-constructive and yields no reasonable procedure. On the other hand, Dax, Hofmann and Lange recently gave a deductive system that is appropriate for algorithmic proof search, but their proofs require a global validity condition and do not have a well understood proof theory.

They work with well-structured proof systems, effectively constructing proofs in a finitary sequent calculus that enjoys local correctness and cut elimination.

This involves an intermediate circular proof system in which one can obtain proofs for all inclusions of parity automata, by adapting Safra's construction. In order to finally obtain finite proofs of Büchi inclusions, a translation result from circular to finite proofs is designed.

6.3. Effective higher dimensional algebra

Participants: Cyrille Chenavier, Pierre-Louis Curien, Yves Guiraud, Maxime Lucas, Philippe Malbos, Jovana Obradović.

6.3.1. Rewriting methods for Artin monoids

With Stéphane Gaussent (ICJ, Univ. Saint-Étienne), Yves Guiraud and Philippe Malbos have used higher-dimensional rewriting methods for the study of Artin monoids, a class of monoids that is fundamental in algebra and geometry. This work formulates in a common language several known results in combinatorial group theory: one by Tits about the fundamental group of a graph associated to an Artin monoid [76], and one by Deligne about the actions of Artin monoids on categories [58], both originally proved by geometrical and non-constructive methods. An improved completion procedure, called the homotopical completion-reduction procedure (see also [8]), is formalised and used to give constructive proofs of (improved versions of) both theorems. This work has been published in Compositio Mathematica [19] and has been implemented in a Python library (http://www.pps.univ-paris-diderot.fr/~guiraud/cox/cox.zip).

6.3.2. Rewriting and Garside theory

Yves Guiraud has collaborated with Patrick Dehornoy (LNO, Univ. Caen) to develop an axiomatic setting for monoids with a special notion of quadratic normalisation map with good computational properties. This theory generalises the normalisation procedure known for monoids that admit a special family of generators called a Garside family [57] to a much wider class that also includes the plactic monoids. It is proved that good quadratic normalisation maps correspond to quadratic convergent presentations, together with a sufficient condition for this to happen, based on the shape of the normalisation paths on length-three words. This work has been submitted for publication to the Journal de l'École Polytechnique — Mathématiques [44].

Building on this last article, Yves Guiraud currently collaborates with Matthieu Picantin (Automates team, LIAFA, Univ. Paris 7) to generalise the main results of [19] to monoids with a Garside family. This will allow an extension of the field of application of the rewriting methods to other geometrically interesting classes of monoids, such as the dual braid monoids.

6.3.3. Higher-dimensional linear rewriting

With Eric Hoffbeck (LAGA, Univ. Paris 13), Yves Guiraud and Philippe Malbos have introduced in [64] the setting of linear polygraphs to formalise a theory of linear rewriting, generalising Gröbner bases. They have adapted the computational method of [7] to compute polygraphic resolutions of associative algebras, with applications to the decision of the Koszul homological property. They are currently engaged into a major overhaul of this work, whose main goal is to ease the adaptation of the results to other algebraic varieties, like commutative algebras or Lie algebras.

6.3.4. Theory of reduction operators

Cyrille Chenavier, supervised by Yves Guiraud and Philippe Malbos, explores the use of Berger's theory of reduction operators [50] to design new rewriting methods in algebra. In [42], he proposed a construction of a contracting homotopy for the Koszul complex of an algebra (a complex characterising the homological property of Koszulness): when an algebra admits a side-confluent presentation (a strong hypothesis of confluence), he gave a candidate for the contracting homotopy, built using specific representations of confluence algebras; when the presentation satisfies an additional condition, called the extra-condition, it turns out that this candidate works.

6.3.5. Rewriting methods for coherence

In [45], Maxime Lucas, supervised by Yves Guiraud and Pierre-Louis Curien, has applied the rewriting techniques of [65] to prove coherence theorems for bicategories and pseudofunctors. He obtained a coherence theorem for pseudonatural transformations thanks to a new theoretical result, improving on the former techniques, that relates the properties of rewriting in 1- and 2-categories.

6.3.6. Wiring structure of operads and operad-like structures

Building on recent ideas of Marcelo Fiore on the one hand, and of François Lamarche on the other hand, Pierre-Louis Curien and Jovana Obradović developed a syntactic approach, using some of the kit of Curien-Herbelin's duality of computation and its polarised versions by Munch and Curien, to the definition of various structures that have appeared in algebra under the names of operads, cyclic operads, dioperads, properads, modular and wheeled operads, permutads, etc. These structures are defined in the literature in different flavours. The goal is to formalise the proofs of equivalence between these different styles of definition. This work is completed for cyclic operads and was presented at the conference Category Theory 2015 in Aveiro [43]. Further work will be to make these proofs modular, so as not to repeat them for each variation of the notion of operad.

6.3.7. A graphical proof of the Bénabou-Roubaud theorem

As a substantial development of reasoning with string diagrams, Jovana Obradović gave a complete proof of the Bénabou-Roubaud monadic descent theorem in [47]. One of the essential points concerning Grothendieck's original approach to descent theory consists of identifying the class of effective descent morphisms for a given fibration. In the special case of a bifibration satisfying Beck-Chevalley condition, Bénabou and Roubaud have given such a characterisation by means of monadicity. Due to the technically complicated calculations involving Grothendieck's cocycle condition, the categorical equivalence which reflects the comparison of the descent in fibered categories with monadic descent is usually not worked out in complete detail in the literature. Jovana Obradović linked the monadic and the original viewpoint via another possible definition of the category of descent data. This intermediate step, due to Janelidze and Tholen, involves constructions in internal categories and it provides an example on how one can stay in the world of string diagrams even when dealing with morphisms which do not have an explicit string diagram definition.

6.4. Incrementality

Participants: Yann Régis-Gianas, Lourdes Del Carmen González Huesca, Thibaut Girka.

An optimisation to perform incremental computations was developed by Lourdes del Carmen González Huesca and Yann Régis-Gianas, providing a mechanism to achieve efficiency. Incrementality as a way to propagate an input change into a corresponding output change is guided by formal change descriptions over terms and dynamic differentiation of functions. The data-changes are represented by displaceable types, a general framework to displace terms directed by types. An extension of the simply-typed lambda-calculus with differentials and partial derivatives offers a language to reason about incrementality. The basic system, λ-diff, was enriched with expressions for fixed-points and data-types together with their corresponding derivatives to analyse incrementality over them. The above results are reported in the second part of Lourdes González Huesca PhD thesis [16].

In collaboration with Paolo Giarrusso and Yufei Cai (Univ Marburg, Allemagne), Yann Régis-Gianas developed a new method to incrementalise higher-order programs using formal derivatives and static caching. A paper is in preparation.

In collaboration with David Mentré (Mitsubishi), Thibaut Girka and Yann Régis-Gianas designed and certified a new algorithm for correlating program generation: such a program is used to characterise the differences between two close programs. (Therefore, a correlating program is a good input for an incremental static analyser.) Before their work, only one algorithm existed in the literature and it was unsound. The new algorithm is sound and certified in Coq. This work has been published in the ATVA conference. Thibaut Girka has presented this work [33] at ATVA 2015.

In collaboration with David Mentré (Mitsubishi), Thibaut Girka and Yann Régis-Gianas are developing a theoretical framework to define a notion of differential operational semantics: a general mathematical object to characterise the difference of behavior of two close programs.

6.5. Metatheory and development of Coq

Participants: Pierre-Louis Curien, Hugo Herbelin, Pierre Letouzey, Yann Régis-Gianas, Matthieu Sozeau.

6.5.1. Models of type theory

Simplicial sets and their extensions as Kan complexes can serve as models of homotopy type theory. Hugo Herbelin extended his concrete type-theoretic formalisation of semi-simplicial sets [20] to simplicial sets.

6.5.2. Unification

Matthieu Sozeau is working in collaboration with Beta Ziliani (PhD at MPI-Saarbrücken, now assistant professor at Cordoba, Argentina) on formalising the unification algorithm used in Coq, which is central for working with advanced type inference features like Canonical Structures. This is the first precise formalisation of all the rules of unification including the ones used for canonical structure resolution and universes. The presentation includes a careful study of the heuristics used in the existing Coq algorithms, which can be added or removed from the new implementation modularly. This work has been presented at the ICFP'15 conference [31].

6.5.3. Nominal techniques

Matthieu Sozeau cosupervised the internship of Gabriel Lewertowski with Nicolas Tabareau (Ascola team, Nantes), on the development of a library for nominal reasoning in Coq/Ssreflect. The goal of this internship was to study the use of nominal sets to ease the formalisation of programming language (meta-)theory. A library based on the Mathematical Components formalisation of finite sets and effective quotients was built, providing generic definitions of substitution and elimination operators for simple descriptions of programming language syntax as a grammar. This work was done in collaboration with Assia Mahboubi (Specfun) and Cyril Cohen (Marelle). It forms the basis for the formalisation of cubical type theory, a new type theory using name abstraction that implements an axiom-free version of Homotopy Type Theory.

POLSYS Project-Team

6. New Results

6.1. Fundamental algorithms and structured polynomial systems

6.1.1. On the complexity of the F5 Gröbner basis algorithm

We study the complexity of Gröbner bases computation, in particular in the generic situation where the variables are in simultaneous Noether position with respect to the system.

We give a bound on the number of polynomials of degree d in a Gröbner basis computed by F_5 algorithm in this generic case for the grevlex ordering (which is also a bound on the number of polynomials for a reduced Gröbner basis, independently of the algorithm used). Next, we analyse more precisely the structure of the polynomials in the Gröbner bases with signatures that F_5 computes and use it to bound the complexity of the algorithm.

Our estimates show that the version of F_5 we analyse, which uses only standard Gaussian elimination techniques, outperforms row reduction of the Macaulay matrix with the best known algorithms for moderate degrees, and even for degrees up to the thousands if Strassen's multiplication is used. The degree being fixed, the factor of improvement grows exponentially with the number of variables.

6.1.2. On the complexity of computing Gröbner bases for weighted homogeneous systems

Solving polynomial systems arising from applications is frequently made easier by the structure of the systems. Weighted homogeneity (or quasi-homogeneity) is one example of such a structure: given a system of weights $W = (w_1, ..., w_n)$, W-homogeneous polynomials are polynomials which are homogeneous w.r.t the weighted degree $\deg(X_1^{\alpha_1} \cdots X_n^{\alpha_n}) = \sum_{i=1}^n w_i \alpha_i$. Gröbner bases for weighted homogeneous systems can be computed by adapting existing algorithms for homogeneous systems to the weighted homogeneous case. In [6], we show that in this case, the complexity estimate for Algorithm F_5 ($\binom{n+d_{\max}-1}{d_{\max}}^{\omega}$ can be divided by a factor $(\prod_{i=1} w_i)^{\omega}$). For zero-dimensional systems, the complexity of Algorithm FGLM nD^{ω} (where D is the number of solutions of the system) can be divided by the same factor $(\prod_{i=1} w_i)^{\omega}$. Under genericity assumptions, for zero-dimensional weighted homogeneous systems of W-degree $(d_1, ..., d_n)$, these complexity estimates are polynomial in the weighted Bézout bound $\prod_{i=1}^n d_i / \prod_{i=1}^n w_i$. Furthermore, the maximum degree reached in a run of Algorithm F_5 is bounded by the weighted Macaulay bound $\sum_{i=1}^n (d_i - w_i) + w_n$, and this bound is sharp if we can order the weights so that $w_n = 1$. For overdetermined semi-regular systems, estimates from the homogeneous case can be adapted to the weighted case. We provide some experimental results based on systems arising from a cryptography problem and from polynomial inversion problems. They show that taking advantage of the weighted homogeneous structure yields substantial speed-ups, and allows us to solve systems which were otherwise out of reach.

6.1.3. Linear Algebra for Computing Gröbner Bases of Linear Recursive Multidimensional Sequences

Sakata generalized the Berlekamp – Massey algorithm to n dimensions in 1988. The Berlekamp – Massey – Sakata (BMS) algorithm can be used for finding a Gröbner basis of a 0-dimensional ideal of relations verified by a table. We investigate this problem using linear algebra techniques, with motivations such as accelerating change of basis algorithms (FGLM) or improving their complexity. In [12], we first define and characterize multidimensional linear recursive sequences for 0-dimensional ideals. Under genericity assumptions, we propose a randomized preprocessing of the table that corresponds to performing a linear change of coordinates on the polynomials associated with the linear recurrences. This technique then essentially reduces our problem to using the efficient 1-dimensional Berlekamp – Massey (BM) algorithm. However, the number of probes to the table in this scheme may be elevated. We thus consider the table in

the *black-box* model: we assume probing the table is expensive and we minimize the number of probes to the table in our complexity model. We produce an FGLM-like algorithm for finding the relations in the table, which lets us use linear algebra techniques. Under some additional assumptions, we make this algorithm adaptive and reduce further the number of table probes. This number can be estimated by counting the number of distinct elements in a multi-Hankel matrix (a multivariate generalization of Hankel matrices); we can relate this quantity with the *geometry* of the final staircase. Hence, in favorable cases such as convex ones, the complexity is essentially linear in the size of the output. Finally, when using the LEX ordering, we can make use of fast structured linear algebra similarly to the Hankel interpretation of Berlekamp – Massey.

6.1.4. Nearly optimal computations with structured matrices

In [9] we estimate the Boolean complexity of multiplication of structured matrices by a vector and the solution of nonsingular linear systems of equations with these matrices. We study four basic and most popular classes, that is, Toeplitz, Hankel, Cauchy and Vandermonde matrices, for which the cited computational problems are equivalent to the task of polynomial multiplication and division and polynomial and rational multipoint evaluation and interpolation. The Boolean cost estimates for the latter problems have been obtained by Kirrinnis in [10], except for rational interpolation. We supply them now as well as the Boolean complexity estimates for the important problems of multiplication of transposed Vandermonde matrix and its inverse by a vector. All known Boolean cost estimates for such problems rely on using Kronecker product. This implies the d-fold precision increase for the d-th degree output, but we avoid such an increase by relying on distinct techniques based on employing FFT. Furthermore we simplify the analysis and make it more transparent by combining the representations of our tasks and algorithms both via structured matrices and via polynomials and rational functions. This also enables further extensions of our estimates to cover Trummer's important problem and computations with the popular classes of structured matrices that generalize the four cited basic matrix classes, as well as the transposed Vandermonde matrices. It is known that the solution of Toeplitz, Hankel, Cauchy, Vandermonde, and transposed Vandermonde linear systems of equations is generally prone to numerical stability problems, and numerical problems arise even for multiplication of Cauchy, Vandermonde, and transposed Vandermonde matrices by a vector. Thus our FFT-based results on the Boolean complexity of these important computations could be quite interesting because our estimates are reasonable even for more general classes of structured matrices, showing rather moderate growth of the complexity as the input size increases.

6.2. Solving Polynomial Systems over the Reals and Applications

6.2.1. Probabilistic Algorithm for Computing the Dimension of Real Algebraic Sets

Let $f \in \mathbb{Q}[X_1, ..., X_n]$ be a polynomial of degree D. We consider the problem of computing the real dimension of the real algebraic set defined by $f = 0$. Such a problem can be reduced to quantifier elimination. Hence it can be tackled with Cylindrical Algebraic Decomposition within a complexity that is doubly exponential in the number of variables. More recently, denoting by d the dimension of the real algebraic set under study, deterministic algorithms running in time $D^{O(d(n-d))}$ have been proposed. However, no implementation reflecting this complexity gain has been obtained and the constant in the exponent remains unspecified. In [11], we design a probabilistic algorithm which runs in time which is essentially cubic in $D^{d(n-d)}$. Our algorithm takes advantage of genericity properties of polar varieties to avoid computationally difficult steps of quantifier elimination. We also report on a first implementation. It tackles examples that are out of reach of the state-of-the-art and its practical behavior reflects the complexity gain.

6.2.2. Real root finding for determinants of linear matrices

Let $A_0, A_1, ..., A_n$ be given square matrices of size m with rational coefficients. The paper [7] focuses on the exact computation of one point in each connected component of the real determinantal variety $\{x \in \mathbb{R}^n : \det(A_0 + x_1A_1 + \cdots + x_nA_n) = 0\}$. Such a problem finds applications in many areas such as control theory, computational geometry, optimization, etc. Using standard complexity results this problem can be solved using $m^{O(n)}$ arithmetic operations. Under some genericity assumptions on the coefficients of the

matrices, we provide in an algorithm solving this problem whose runtime is essentially quadratic in $\binom{n+m}{n}^3$. We also report on experiments with a computer implementation of this algorithm. Its practical performance illustrates the complexity estimates. In particular, we emphasize that for subfamilies of this problem where m is fixed, the complexity is polynomial in n.

6.2.3. Real root finding for rank defects in linear Hankel matrices

Let $H_0, ..., H_n$ be $m \times m$ matrices with entries in $\mathbb{Q}$ and Hankel structure, i.e. constant skew diagonals. We consider the linear Hankel matrix $H(X) = H_0 + X_1H_1 + \cdots + X_nH_n$ and the problem of computing sample points in each connected component of the real algebraic set defined by the rank constraint $\mathsf{rank}(H(X)) \leq r$, for a given integer $r \leq m-1$. Computing sample points in real algebraic sets defined by rank defects in linear matrices is a general problem that finds applications in many areas such as control theory, computational geometry, optimization, etc. Moreover, Hankel matrices appear in many areas of engineering sciences. Also, since Hankel matrices are symmetric, any algorithmic development for this problem can be seen as a first step towards a dedicated exact algorithm for solving semi-definite programming problems, i.e. linear matrix inequalities. Under some genericity assumptions on the input (such as smoothness of an incidence variety), we design in [18] a probabilistic algorithm for tackling this problem. It is an adaptation of the so-called critical point method that takes advantage of the special structure of the problem. Its complexity reflects this: it is essentially quadratic in specific degree bounds on an incidence variety. We report on practical experiments and analyze how the algorithm takes advantage of this special structure. A first implementation outperforms existing implementations for computing sample points in general real algebraic sets: it tackles examples that are out of reach of the state-of-the-art.

6.2.4. Optimizing a Parametric Linear Function over a Non-compact Real Algebraic Variety

In [17], we consider the problem of optimizing a parametric linear function over a non-compact real trace of an algebraic set. Our goal is to compute a representing polynomial which defines a hypersurface containing the graph of the optimal value function. Rostalski and Sturmfels showed that when the algebraic set is irreducible and smooth with a compact real trace, then the least degree representing polynomial is given by the defining polynomial of the irreducible hypersurface dual to the projective closure of the algebraic set. First, we generalize this approach to non-compact situations. We prove that the graph of the opposite of the optimal value function is still contained in the affine cone over a dual variety similar to the one considered in compact case. In consequence, we present an algorithm for solving the considered parametric optimization problem for generic parameters' values. For some special parameters' values, the representing polynomials of the dual variety can be identically zero, which give no information on the optimal value. We design a dedicated algorithm that identifies those regions of the parameters' space and computes for each of these regions a new polynomial defining the optimal value over the considered region.

6.2.5. Bounds for the Condition Number of Polynomials Systems with Integer Coefficients

Polynomial systems of equations are a central object of study in computer algebra. Among the many existing algorithms for solving polynomial systems, perhaps the most successful numerical ones are the homotopy algorithms. The number of operations that these algorithms perform depends on the condition number of the roots of the polynomial system. Roughly speaking the condition number expresses the sensitivity of the roots with respect to small perturbation of the input coefficients. A natural question to ask is how can we bound, in the worst case, the condition number when the input polynomials have integer coefficients? In [19] we address this problem and we provide effective bounds that depend on the number of variables, the degree and the maximum coefficient bitsize of the input polynomials. Such bounds allows to estimate the bit complexity of the algorithms that depend on the separation bound, like the homotopy algorithms, for solving polynomial systems.

6.2.6. Nearly Optimal Refinement of Real Roots of a Univariate Polynomial

In [10] we assume that a real square-free polynomial A has a degree d, a maximum coefficient bitsize τ and a real root lying in an isolating interval and having no nonreal roots nearby (we quantify this assumption). Then, we combine the *Double Exponential Sieve* algorithm (also called the *Bisection of the Exponents*), the

bisection, and Newton iteration to decrease the width of this inclusion interval by a factor of $t = 2^{-L}$. The algorithm has Boolean complexity $\widetilde{O}_B(d^2\tau + dL)$. Our algorithms support the same complexity bound for the refinement of r roots, for any $r \leq d$.

6.2.7. Accelerated Approximation of the Complex Roots and Factors of a Univariate Polynomial

The known algorithms approximate the roots of a complex univariate polynomial in nearly optimal arithmetic and Boolean time. They are, however, quite involved and require a high precision of computing when the degree of the input polynomial is large, which causes numerical stability problems. We observe that these difficulties do not appear at the initial stages of the algorithms, and in [8] we extend one of these stages, analyze it, and avoid the cited problems, still achieving the solution within a nearly optimal complexity estimates, provided that some mild initial isolation of the roots of the input polynomial has been ensured. The resulting algorithms promise to be of some practical value for root-finding and can be extended to the problem of polynomial factorization, which is of interest on its own right. We conclude with outlining such an extension, which enables us to cover the cases of isolated multiple roots and root clusters.

6.2.8. Polynomial Interrupt Timed Automata

Interrupt Timed Automata (ITA) form a subclass of stopwatch automata where reachability and some variants of timed model checking are decidable even in presence of parameters. They are well suited to model and analyze real-time operating systems. Here we extend ITA with polynomial guards and updates, leading to the class of polynomial ITA (polITA). In [13], we prove that reachability is decidable in 2EXPTIME on polITA, using an adaptation of the cylindrical algebraic decomposition algorithm for the first-order theory of reals using symbolic computation. Compared to previous approaches, our procedure handles parameters and clocks in a unified way. We also obtain decidability for the model checking of a timed version of CTL and for reachability in several extensions of polITA.

6.3. Solving Systems in Finite Fields, Applications in Cryptology and Algebraic Number Theory

6.3.1. Polynomial-Time Algorithms for Quadratic Isomorphism of Polynomials: The Regular Case

Let $\mathbf{f} = (f_1, ..., f_m)$ and $\mathbf{g} = (g_1, ..., g_m)$ be two sets of $m \geq 1$ nonlinear polynomials in $\mathbb{K}[x_1, ..., x_n]$ ($\mathbb{K}$ being a field). In [3], we consider the computational problem of finding – if any – an invertible transformation on the variables mapping $\mathbf{f}$ to $\mathbf{g}$. The corresponding equivalence problem is known as *Isomorphism of Polynomials with one Secret* (`IP1S`) and is a fundamental problem in multivariate cryptography. Amongst its applications, we can cite Graph Isomorphism (`GI`) which reduces to equivalence of cubic polynomials with respect to an invertible linear change of variables, according to Agrawal and Saxena. The main result is a randomized polynomial-time algorithm for solving `IP1S` for quadratic instances, a particular case of importance in cryptography. To this end, we show that `IP1S` for quadratic polynomials can be reduced to a variant of the classical module isomorphism problem in representation theory. We show that we can essentially *linearize* the problem by reducing quadratic-`IP1S` to test the orthogonal simultaneous similarity of symmetric matrices; this latter problem was shown by Chistov, Ivanyos and Karpinski (ISSAC 1997) to be equivalent to finding an invertible matrix in the linear space $\mathbb{K}^{n\times n}$ of $n \times n$ matrices over $\mathbb{K}$ and to compute the square root in a certain representation in a matrix algebra. While computing square roots of matrices can be done efficiently using numerical methods, it seems difficult to control the bit complexity of such methods. However, we present exact and polynomial-time algorithms for computing a representation of the square root of a matrix in $\mathbb{K}^{n\times n}$, for various fields (including finite fields), as a product of two matrices. Each coefficient of these matrices lie in an extension field of $\mathbb{K}$ of polynomial degree. We then consider `#IP1S`, the counting version of `IP1S` for quadratic instances. In particular, we provide a (complete) characterization of the automorphism group of homogeneous quadratic polynomials. Finally, we also consider the more general *Isomorphism of Polynomials* (`IP`) problem where we allow an invertible linear transformation on the variables *and* on the set

of polynomials. A randomized polynomial-time algorithm for solving IP when $\mathbf{f} = (x_1^d, ..., x_n^d)$ is presented. From an algorithmic point of view, the problem boils down to factoring the determinant of a linear matrix (*i.e.* a matrix whose components are linear polynomials). This extends to IP a result of Kayal obtained for PolyProj.

6.3.2. Factoring $N = p^r q^s$ for Large r and s

Boneh *et al.* showed at Crypto 99 that moduli of the form $N = p^r q$ can be factored in polynomial time when $r \simeq \log p$. Their algorithm is based on Coppersmith's technique for finding small roots of polynomial equations. In [15] we show that $N = p^r q^s$ can also be factored in polynomial time when r or s is at least $(\log p)^3$; therefore we identify a new class of integers that can be efficiently factored. We also generalize our algorithm to moduli with k prime factors $N = \prod_{i=1}^{k} p_i^{r_i}$; we show that a non-trivial factor of N can be extracted in polynomial-time if one of the exponents r_i is large enough.

6.3.3. On the Complexity of the BKW Algorithm on LWE

This work [1] presents a study of the complexity of the Blum–Kalai–Wasserman (BKW) algorithm when applied to the Learning with Errors (LWE) problem, by providing refined estimates for the data and computational effort requirements for solving concrete instances of the LWE problem. We apply this refined analysis to suggested parameters for various LWE-based cryptographic schemes from the literature and compare with alternative approaches based on lattice reduction. As a result, we provide new upper bounds for the concrete hardness of these LWE-based schemes. Rather surprisingly, it appears that BKW algorithm outperforms known estimates for lattice reduction algorithms starting in dimension $n \approx 250$ when LWE is reduced to SIS. However, this assumes access to an unbounded number of LWE samples.

6.3.4. Structural Cryptanalysis of McEliece Schemes with Compact Keys

A very popular trend in code-based cryptography is to decrease the public-key size by focusing on subclasses of alternant/Goppa codes which admit a very compact public matrix, typically quasi-cyclic (QC), quasi-dyadic (QD), or quasi-monoidic (QM) matrices. In [5], we show that the very same reason which allows to construct a compact public-key makes the key-recovery problem intrinsically much easier. The gain on the public-key size induces an important security drop, which is as large as the compression factor p on the public-key. The fundamental remark is that from the $k \times n$ public generator matrix of a compact McEliece, one can construct a $k/p \times n/p$ generator matrix which is - from an attacker point of view - as good as the initial public-key. We call this new smaller code the folded code. Any key-recovery attack can be deployed equivalently on this smaller generator matrix. To mount the key-recovery in practice, we also improve the algebraic technique of Faugère, Otmani, Perret and Tillich (FOPT). In particular, we introduce new algebraic equations allowing to include codes defined over any prime field in the scope of our attack. We describe a so-called "structural elimination" which is a new algebraic manipulation which simplifies the key-recovery system. As a proof of concept, we report successful attacks on many cryptographic parameters available in the literature. All the parameters of CFS-signatures based on QD/QM codes that have been proposed can be broken by this approach. In most cases, our attack takes few seconds (the harder case requires less than 2 hours). In the encryption case, the algebraic systems are harder to solve in practice. Still, our attack succeeds against several cryptographic challenges proposed for QD and QM encryption schemes, but there are still some parameters that have been proposed which are out of reach for the methods given here. However, regardless of the key-recovery attack used against the folded code, there is an inherent weakness arising from Goppa codes with QM or QD symmetries. It is possible to derive from the public key a much smaller public key corresponding to the folding of the original QM or QD code, where the reduction factor of the code length is precisely the order of the QM or QD group used for reducing the key size. To summarize, the security of such schemes are not relying on the bigger compact public matrix but on the small folded code which can be efficiently broken in practice with an algebraic attack for a large set of parameters.

6.3.5. A Polynomial-Time Key-Recovery Attack on MQQ Cryptosystems

In [16], we investigate the security of the family of MQQ public key cryptosystems using multivariate quadratic quasigroups (MQQ). These cryptosystems show especially good performance properties. In particular, the

MQQ-SIG signature scheme is the fastest scheme in the ECRYPT benchmarking of cryptographic systems (eBACS). We show that both the signature scheme MQQ-SIG and the encryption scheme MQQ-ENC, although using different types of MQQs, share a common algebraic structure that introduces a weakness in both schemes. We use this weakness to mount a successful polynomial time key-recovery attack that finds an equivalent key using the idea of so-called good keys. In the process we need to solve a MinRank problem that, because of the structure, can be solved in polynomial-time assuming some mild algebraic assumptions. We highlight that our theoretical results work in characteristic 2 which is known to be the most difficult case to address in theory for MinRank attacks and also without any restriction on the number of polynomials removed from the public-key. This was not the case for previous MinRank like-attacks against $\mathcal{MQ}$ schemes. From a practical point of view, we are able to break an MQQ-SIG instance of 80 bits security in less than 2 days, and one of the more conservative MQQ-ENC instances of 128 bits security in little bit over 9 days. Altogether, our attack shows that it is very hard to design a secure public key scheme based on an easily invertible MQQ structure.

6.3.6. Algebraic Cryptanalysis of a Quantum Money Scheme The Noise-Free Case

In [14], we investigate the Hidden Subspace Problem (HSP_q) over $\mathbb{F}_q$ which is as follows:

Input : $p_1, ..., p_m, q_1, ..., q_m \in \mathbb{F}_q[x_1, ..., x_n]$ of degree $d \geq 3$ (and $n \leq m \leq 2n$).
Find : a subspace $A \subset {\mathbb{F}_q}^n$ of dimension $n/2$ (n is even) such that

$$p_i(A) = 0 \ \forall i \in \{1, ..., m\} \text{ and } q_j(A^{\perp}) = 0 \ \forall j \in \{1, ..., m\},$$

where $A^{\perp}$ denotes the orthogonal complement of A with respect to the usual scalar product in $\mathbb{F}_q$.

This problem underlies the security of the first public-key quantum money scheme that is proved to be cryptographically secure under a non quantum but classic hardness assumption. This scheme was proposed by S. Aaronson and P. Christiano at STOC'12. In particular, it depends upon the hardness of HSP_2. More generally, Aaronson and Christiano left as an open problem to study the security of the scheme for a general field $\mathbb{F}_q$. We present a randomized polynomial-time algorithm that solves the HSP_q for $q > d$ with success probability $\approx 1 - 1/q$. So, the quantum money scheme extended to $\mathbb{F}_q$ is not secure for big q. Finally, based on experimental results and a structural property of the polynomials that we prove, we conjecture that there is also a randomized polynomial-time algorithm solving the HSP_2 with high probability. To support our theoretical results we also present several experimental results confirming that our algorithms are very efficient in practice. We emphasize that S. Aaronson and P. Christiano proposes a non-noisy and a noisy version of the public-key quantum money scheme. The noisy version of the quantum money scheme remains secure.

6.3.7. Folding Alternant and Goppa Codes with Non-Trivial Automorphism Groups

The main practical limitation of the McEliece public-key encryption scheme is probably the size of its key. A famous trend to overcome this issue is to focus on subclasses of alternant/Goppa codes with a non trivial automorphism group. Such codes display then symmetries allowing compact parity-check or generator matrices. For instance, a key-reduction is obtained by taking quasi-cyclic (QC) or quasi-dyadic (QD) alternant/Goppa codes. We show that the use of such symmetric alternant/Goppa codes in cryptography introduces a fundamental weakness. It is indeed possible to reduce the key-recovery on the original symmetric public-code to the key-recovery on a (much) smaller code that has not anymore symmetries. This result [4] is obtained thanks to a new operation on codes called folding that exploits the knowledge of the automorphism group. This operation consists in adding the coordinates of codewords which belong to the same orbit under the action of the automorphism group. The advantage is twofold: the reduction factor can be as large as the size of the orbits, and it preserves a fundamental property: folding the dual of an alternant (resp. Goppa) code provides the dual of an alternant (resp. Goppa) code. A key point is to show that all the existing constructions of alternant/Goppa codes with symmetries follow a common principal of taking codes whose support is globally invariant under the action of affine transformations (by building upon prior works of T. Berger and A. Dür). This enables not only to present a unified view but also to generalize the construction of QC, QD and even

quasi-monoidic (QM) Goppa codes. All in all, our results can be harnessed to boost up any key-recovery attack on McEliece systems based on symmetric alternant or Goppa codes, and in particular algebraic attacks.

6.3.8. Improved Sieving on Algebraic Curves

The best algorithms for discrete logarithms in Jacobians of algebraic curves of small genus are based on index calculus methods coupled with large prime variations. For hyperelliptic curves, relations are obtained by looking for reduced divisors with smooth Mumford representation (Gaudry); for non-hyperelliptic curves it is faster to obtain relations using special linear systems of divisors (Diem, Diem and Kochinke). Recently, Sarkar and Singh have proposed a sieving technique, inspired by an earlier work of Joux and Vitse, to speed up the relation search in the hyperelliptic case. In [20], we give a new description of this technique, and show that this new formulation applies naturally to the non-hyperelliptic case with or without large prime variations. In particular, we obtain a speed-up by a factor approximately 3 for the relation search in Diem and Kochinke's methods.

PROSECCO Project-Team

7. New Results

7.1. Verification of Security Protocols in the Symbolic Model

Participants: Bruno Blanchet, Miriam Paiola.

The applied pi calculus is a widely used language for modeling security protocols, including as a theoretical basis of ProVerif. However, the seminal paper that describes this language [24] does not come with proofs, and detailed proofs for the results in this paper were never published. This year, Martín Abadi, Bruno Blanchet, and Cédric Fournet finished the detailed proofs of all results of this paper, started last year, and added a new example on a symbolic analog of indifferentiability of hash functions. This work is submitted to a journal.

Previously [37], Bruno Blanchet and Miriam Paiola presented an automatic technique for proving secrecy and authentication properties for security protocols that manipulate lists of unbounded length, for an unbounded number of sessions. That work relies on an extension of Horn clauses, generalized Horn clauses, designed to support unbounded lists, and on a resolution algorithm on these clauses. However, in that previous work, they had to model protocols manually with generalized Horn clauses, which is unpractical. They recently extended the input language of ProVerif to model protocols with lists of unbounded length. They give the formal meaning of this extension, translate it automatically to generalized Horn clauses, and prove that this translation is sound. This work appears as a research report [21].

We implemented several extensions of ProVerif: Bruno Blanchet and Vincent Cheval improved the algorithm for proving observational equivalence between two processes, by merging them into a single biprocess that encodes the two processes. Bruno Blanchet also introduced a new construct **new** $a[x_1, ..., x_n]$ in ProVerif which allows to specify the arguments $x_1, \cdots, x_n$ used in the internal representation of the fresh name a. This extension allows one to tune the precision and speed of the analysis performed by ProVerif. The extended tool is available at http://proverif.inria.fr, and deposited to the APP (*Agence pour la Protection des Programmes*).

Stéphanie Delaune, Mark Ryan, and Ben Smyth [42] introduced the idea of swapping data in order to prove observational equivalence. For instance, ballot secrecy in electronic voting is formalized by saying that A voting a and B voting b is observationally equivalent to (indistinguishable from) A voting b and B voting a. Proving such an equivalence typically requires swapping the votes. However, Delaune et al's approach was never proved correct. Bruno Blanchet and Ben Smyth filled this gap by formalizing the approach and providing a detailed soundness proof. They plan to submit this work to a conference.

7.2. Verification of Security Protocols in the Computational model

Participant: Bruno Blanchet.

Bruno Blanchet implemented several extensions of his computational protocol verifier CryptoVerif. In particular, he improved the global dependency analysis, used in order to show that the result of all tests is independent from some random values. He improved the proof of secrecy properties, in particular to prove forward secrecy properties. He also improved the merging of branches of tests, in particular to be able to merge the two branches of **if** b **then** P_1 **else** P_2 even when variables are renamed between P_1 and P_2. Finally, he added the display of an explanation of why a cryptographic transformation fails, to make the tool easier to use. The extended tool is available at http://cryptoverif.inria.fr.

Within the ANR project AnaStaSec, Bruno Blanchet verified an air-ground avionic security protocol (International Civil Aviation Organization (ICAO) Document 9880: Manual on Detailed Technical Specifications for the Aeronautical Telecommunication Network (ATN) using ISO/OSI standards and protocols, Part IV) using CryptoVerif. He proved entity authentication and message authenticity for the main protocol, in the computational model of cryptography, and made comments on some points that should be clarified in the protocol specification. He presented this work at a meeting of the secure dialog service working group of ICAO, in Toulouse, September 2015. The working group was strongly interested by the presentation and welcomed the proposal to apply these modelling and formal verification techniques as part of its validation activities.

7.3. The F* programming language

Participants: Nikhil Swamy [Microsoft Research], Catalin Hritcu, Chantal Keller [LRI], Aseem Rastogi [Univ of Maryland], Antoine Delignat-Lavaud, Simon Forest, Karthikeyan Bhargavan, Cedric Fournet [Microsoft Research], Pierre-Yves Strub [IMDEA], Markulf Kohlweiss [Microsoft Research], Jean Karim Zinzindohoue, Santiago Zanella Beguelin [Microsoft Research, MSR-Inria].

F* is a new higher order, effectful programming language (like ML) designed with program verification in mind. Its type system is based on a core that resembles System Fω (hence the name), but is extended with dependent types, refined monadic effects, refinement types, and higher kinds. Together, these features allow expressing precise and compact specifications for programs, including functional correctness properties. The F* type-checker aims to prove that programs meet their specifications using an automated theorem prover (usually Z3) behind the scenes to discharge proof obligations. Programs written in F* can be translated to OCaml, F#, or JavaScript for execution. We published a paper on the design, implementation, and formal core of F* at POPL 2016. F* is being developed as an open-source project at GitHub: https://github.com/FStarLang and the official webpage is at http://fstar-lang.org. We released several beta versions of the software this year.

7.4. Micro-Policies and Secure Compilation

Participants: Catalin Hritcu, Arthur Azevedo de Amorim, Zoi Paraskevopoulou, Nikolaos Giannarakis.

Following on from previous work on the *micro-policy* framework, Catalin Hritcu and his collaborators published new work on applications and efficient implementations of micro-policies. They published work on low-level implementations of micro-policies at ASPLOS 2015 [18]. At IEEE S&P, they published a paper how to write formally verified reference monitors using micro-policies [26].

Other than these published works, Hritcu and his colleagues also worked on using micro-policies to enforce secure information flow at the hardware level [25], and a secure compiler for a high-level language that relies on micro-policies to enforce programming language abstractions [45].

7.5. Dependable Property-Based Testing

Participants: Catalin Hritcu, Zoi Paraskevopoulou.

Catalin Hritcu and his student, Zoi Paraskevopoulou, worked on a methodology for formally verified property-based testing and implemented it as a foundational verification framework for QuickChick, a port of QuickCheck to Coq. This work was published at ITP 2015 [19]. Catalin Hritcu also worked with a number of co-authors on a new technique for creating random generators for property-based testing. This work is currently under submission [46].

7.6. Attacks and Proofs for Transport Layer Security

Participants: Benjamin Beurdouche, Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Cedric Fournet [Microsoft Research], Markulf Kohlweiss [Microsoft Research], Alfredo Pironti, Pierre-Yves Strub [IMDEA], Jean Karim Zinzindohoue.

As a countermeasure to our earlier work on the triple handshake attack, we proposed a TLS extension called *session hash* which has now been published as an Internet standard (IETF RFC 7627). We also formally analyzed various protocols such as TLS, IKE, and SSH for key synchronization and triple handshake attacks, and proved that our session hash countermeasure prevents such attacks on TLS. This work appeared at NDSS 2015 [15].

We discovered and reported an important class of *state machine attacks* on implementations of the Transport Layer Security (TLS) protocol. These attacks appear when TLS implementations incorrectly accept messages which are forbidden by the TLS state machine. We built a test framework for such attacks and analyzed a number of open source implementations. Our analysis uncovered critical vulnerabilities such as the SKIP attack on Java and the FREAK attack on almost all mainstream web browsers. The research results were published at IEEE S&P where our paper won a distinguished paper award [14]. Our work also led to security updates and CVEs for many web browsers, TLS libraries, and web servers.

Along with colleagues at several other institutions, we discovered the Logjam vulnerability on protocols that still support weak Diffie-Hellman groups in their key exchange. We showed that the attack could be used for online and offline attacks on real-world TLS clients and servers. We also showed how the vulnerability could weaken the security of IPsec and SSH connections. Our research led to widespread changes to the configurations of web servers, mail servers, web browsers, and TLS libraries. The research was published at ACM CCS 2015 [12] where it won a Best Paper award.

Antoine Delignat-Lavaud showed how the unsafe sharing of certificates across multiple HTTPS websites could be exploited to fully compromise the same origin policy for websites, using a vulnerability called *virtual host confusion*. A research paper on these attacks appeared at WWW 2015 [17].

7.7. Privacy, Electronic Voting, and Auctions

Participants: Benjamin Smyth [correspondant], Elizabeth Quaglia.

Benjamin Smyth worked on a formal analysis of privacy in Direct Anonymous Attestation schemes [50]. He also showed how to verify commitment protocols in ProVerif without False attacks [39].

Apart from these published works, Benjamin Smyth and Elizabeth Quaglia worked on formal security analyses of electronic auction schemes based on existing models for electronic voting [48]. Benjamin Smyth worked on developing new formal definitions for secrecy and independence in election schemes [51], and on applying such definitions to the security analysis of real-world voting protocols such as Helios and JCJ [49].

7.8. Computationally Complete Symbolic Attacker Models

Participants: Gergei Bana, Hubert Comon-Lundh [ENS Cachan], Rohit Chadha [University of Missouri].

In previous work, Bana and Comon-Lunch proposed a new approach to computational verification of cryptographic protocols, by defining a *computationally complete* symbolic attacker, so that a symbolic proof against this attacker can be shown to imply a computational proof of security [27], [28].

Following on from this work, Bana and Chadha fully developed the core parts of the computationally complete symbolic attacker based on indistinguishability. This covers both trace properties and equivalence properties and can be proved partially complete. They evaluated their method by applying it to several classic protocols. This work is currently under submission.

Bana, Comon-Lundh, and Koutsos also worked on a decision procedure for the computationally complete symbolic attacker based on indistinguishability.

QUANTIC Project-Team

6. New Results

6.1. Entanglement between stationary and propagating modes

Participants: B. Huard and F. Mallet.

The results of this section were published in [14].

Entanglement being instrumental in quantum machines, we have shown how a Josephson mixer can generate and distribute entangled microwave radiations on separated transmission lines and different frequencies by spontaneous parametric down-conversion in 2012. Using two Josephson mixers, we have provided the first demonstration of entanglement between spatially separated propagating fields in the microwave domain. Therefore, a new variety of entangled states, the so-called EPR states (after Einstein, Podolsky and Rosen), which are encoded on continuous variables, is now available in this frequency range.

In 2015, we have shown that it could constitute the central component of a potential quantum network based on continuous-variable entanglement. The device essentially acts as a regular mixer performing frequency conversion but without adding extra noise. Used as a switch, it is able to open and close the coupling to a high-quality factor cavity in a time-controlled way. We have demonstrated how this feature leads to a new kind of quantum memory. Coupled to its ability to generate entanglement, we have demonstrated the time-controlled generation, storage and on-demand release of an entangled state, which is the prerequisite for the node of a quantum network.

Several implementations of quantum memories for microwave radiation have been realized in the past few years. In order to store the state of microwave signals, some use spin ensembles [81], [130], [71], or mechanical oscillators [98], while others use superconducting cavities with tunable input coupling [124], [102]. Our own implementation is sketched in Fig. 3 b, where the Josephson Mixer allows an on-demand access to the long lived 3D cavity based on noiseless frequency conversion. Its main advantage consists in the ability to generate entanglement between the memory and the output port.

Noiseless frequency conversion is another regime of the Josephson mixer. The frequency of the pump tone is now chosen to be at the difference between the frequencies of the modes $\widehat{a}$ and $\widehat{b}$, $\Omega = |\omega_a - \omega_b|$. In the rotating frame, the effective Hamiltonian reduces to a beam-splitter Hamiltonian with an implicit frequency conversion:

$$H = \hbar\chi(\widehat{a}^\dagger\widehat{b}\widehat{c} + \widehat{a}\widehat{b}^\dagger\widehat{c}^\dagger).$$

The elementary process corresponds to the conversion of photons between the mode a and b mediated by the pump at a rate $\chi|\langle\widehat{c}\rangle|$ as sketched in Fig. 3 c. Therefore, the noiseless frequency conversion generates a coupling between the long lived cavity mode $\widehat{b}$ and the propagating modes at the input of mode $\widehat{a}$. This pump field can then be varied in time to switch on and off the coupling.

A first measurement consists in the capture, storage and retrieval of a microwave pulse. The protocol is quite simple, we turn the pump tone on when the incoming pulse reaches the memory input. The signal pulse has been designed such that it is optimally absorbed by the memory. The pump tone is turned off after the absorption and turned back on at a later time τ to retrieve the pulse in the transmission line. The measured output amplitude in time shown in Fig. 3 d demonstrate that this protocol can be performed with a great efficiency for a few microseconds.

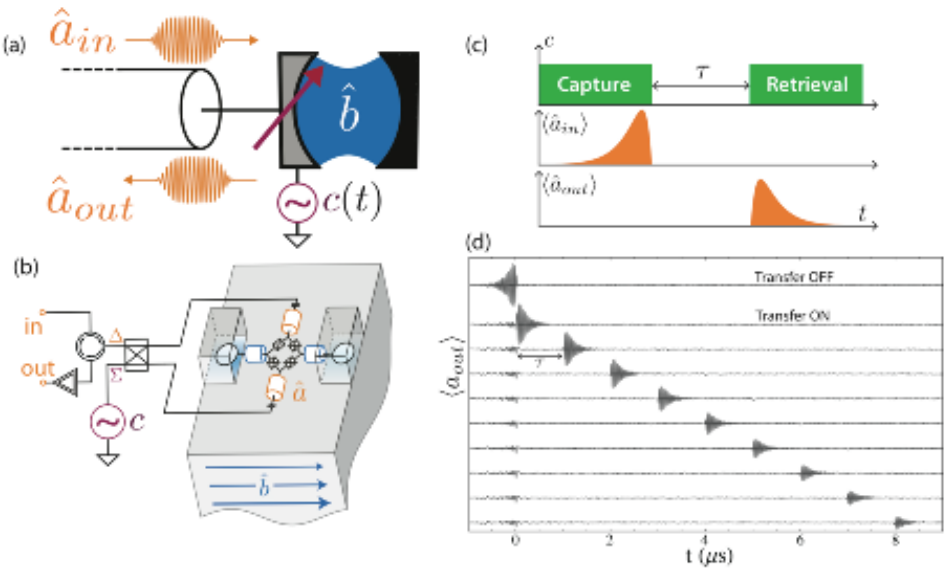

Figure 3. (a) Simplified schematics of the quantum memory. When the pump is driven at $\Omega = |\omega_a - \omega_b|$, the JRM behaves as a beam splitter with an implicit frequency conversion whose transparency depends on the pump amplitude. (b) Schematics of the device. The core of the device is similar to the usual design [107] excepted that one of the two transmission lines is replaced by a superconducting 3D cavity that defines the memory mode. (c) Protocol of the capture, storage and release of a incoming microwave pulse. (d) Measured output amplitude as a function of time. In the first trace, the pump is always turned off and the measured amplitude corresponds to the reflected incoming pulse. In the following traces, the pump is turned on and varied in time as indicated in (c). The storage time is varied from 0 μs to 8 μs.

However, the unique ability of this device lies in the possibility to combine this storage operation with the entanglement generation demonstrated previously. A second measurement consists in the generation, storage and characterization of an EPR state distributed between the memory and the transmission line. The protocol is sketched in Fig. 4 b. The pump is first applied at $\Omega = \omega_a + \omega_b$ to generate an EPR state shared between the memory and the propagating mode. The propagating mode complex amplitude is measured and at a later time, the pump is turned on again at $\Omega = |\omega_a - \omega_b|$ to activate the noiseless conversion. The memory mode is then retrieved in the transmission line and its complex amplitude is measured. By analyzing the cross-correlations between these two measurements, we have been able to show that the memory preserves the entanglement of the EPR state. Furthermore, the contours of the EPR state Wigner function have been inferred from this correlation measurement (Fig. 4 c) and the entanglement quantified.

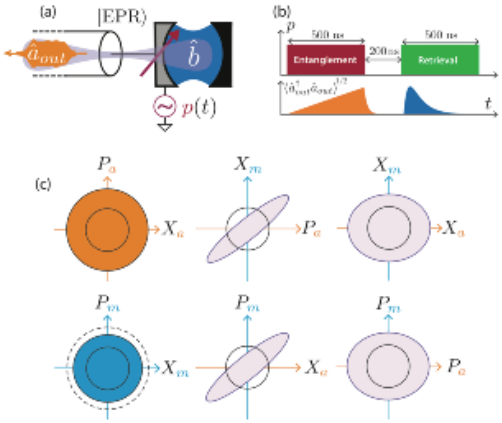

Figure 4. (a) When the pump is shined at $\Omega = \omega_a + \omega_b$, an EPR state is distributed between the transmission line and the memory. (b) Protocol for the entanglement distribution, storage and retrieval. (c) Contour of the marginal Wigner distributions reconstructed from the correlation measurements corresponding to the protocol (b).

6.2. Wideband Josephson mixer

Participants: B. Huard and F. Mallet.

The results of this section were published in [22].

For nearly a decade, the superconducting circuits community develops microwave amplifiers in the quantum regime, i.e. adding only a noise comparable to the vacuum fluctuations of the signal. We participated in this effort in 2012 [107] by adding frequency tunability to the only non-degenerate existing amplifier: the Josephson Parametric Converter (JPC) invented by the group of Michel Devoret at Yale.

However, this amplifier showed the defect of being limited to a few MHz bandwidth for a gain of 20 dB and a dynamic range (maximum input power before changing the gain) capable of amplifying signals typical of circuit-QED. We conducted a theoretical study to understand the various constraints involved in the manufacture of such an amplifier. This study has allowed us to make the first lumped element version of the JPC with bandwidth only limited by the mismatch between the characteristic impedance of the resonators and that of the transmission line.

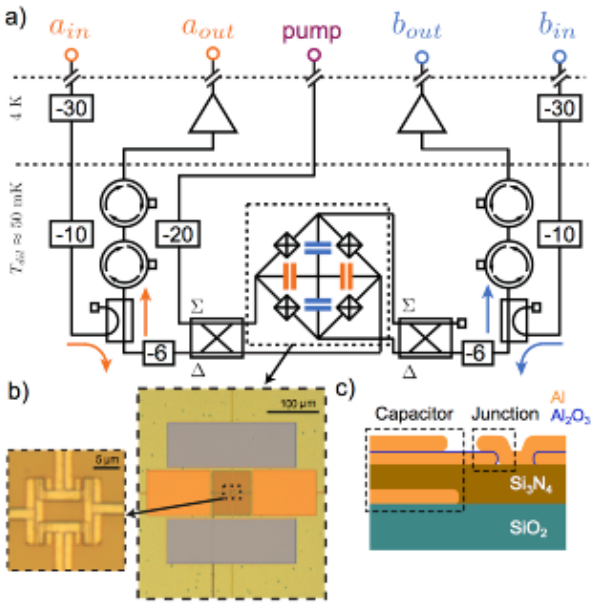

Figure 5. (a) Simplified schematic of the experimental setup. Differential a and b modes of the Josephson mixer are addressed in reflection through two 180 degree hybrid couplers. All input lines are filtered and attenuated (partially shown). Output signals are separated from input signals by a directional coupler and amplified by a low noise HEMT amplifier at 4K. (b) Optical microscope picture of the device showing the planar capacitors (right) and the Josephson junction ring (left). (c) Side view of the device. The thickness of the bottom plate of the capacitors is 35 nm and buried below 200 nm of silicon nitride, the top plate of the capacitors and the Josephson junctions are obtained by double angle deposition of 100 nm and 120 nm of aluminium with an intermediate oxidation.

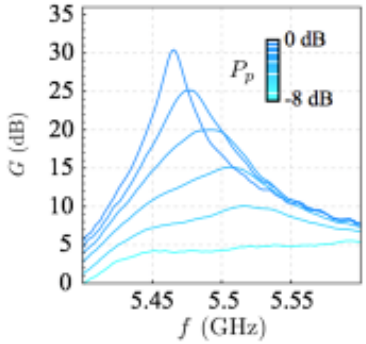

Figure 6. Gain in reflection as a function of frequency for various pump powers. The color bar encodes the pump power referred to the parametric oscillation threshold.

Finally we have measured the quantum efficiency of this amplifier and obtained almost 70%, which means that only 30% of the noise power observed at the end of line comes from technical noise while 70% is the signal, including quantum noise.

6.3. Quantum Zeno dynamics

Participants: B. Huard, L. Bretheau, P. Campagne-Ibarcq, F. Mallet.

The results of this section were published in [13].

Electromagnetic modes are instrumental for realizing quantum physics experiments and building quantum machines. Their manipulation usually involves the tailoring of their Hamiltonian in time. An alternative control scheme, called Quantum Zeno Dynamics (QZD), consists in restricting the evolution of a mode to a subset of possible states. This promising control scheme had been implemented in 2014 on atomic levels of Rb and of a Rydberg atom.

We have made the first observation of QZD of light, using superconducting circuits. By preventing the access to a single energy level, the dynamics of the field is dramatically changed. In this experiment, it was indeed possible to avoid a number of photons N, which was arbitrarily chosen between 2 and 5. Under this constraint, and starting in its ground state, a resonantly driven mode is confined to levels 0 to $N-1$. The level occupation is then found to oscillate in time, similarly to an N-level system. Performing a direct Wigner tomography of the field reveals its non-classical features. In particular, at half period in the evolution, it resembles a "Schrödinger cat state".

In its original definition, the quantum Zeno effect corresponds to the inhibition of coherent transitions from, or to, the pointer states of a strong measurement or dissipative process. Instead of freezing the dynamics, one can restrict it to a given subspace by choosing a measurement with degenerate eigenvalues.

Similar behavior can also be induced by rapid unitary "kicks", leaving the subspace to protect unaffected. It can be understood considering a model for the original Zeno measurement as a series of coherent interactions with ancillary systems. When the interactions are strong enough, departure from the subspace is perfectly suppressed, so that the outcome of the detector is always the same. Therefore, the ancillas are all left in the same state after the interaction and they do not need to be reset. One can then enforce Zeno dynamics by performing repeatedly unitary operations controlling the state of an auxiliary degree of freedom. This amounts

to re-using the same ancilla, at the condition that the unitary evolutions are fast enough to effectively randomize the phase of coherences created with the system. In that sense, QZD is a coherent feedback, which engineers the energy level landscape of a system or its environment by coherent coupling with an ancillary degree of freedom.

In the experiment, a qubit in the resolved photon number regime plays the role of the ancillary system. A strong Rabi drive is applied on its transition conditioned on the cavity mode hosting N photons (N = 3 on Fig. 7). The drive hybridizes the levels $|N, g\rangle$ and $|N, e\rangle$ that repel each other. The level $|N\rangle$ is then moved out from the harmonic ladder of the cavity mode. When starting in the vacuum and applying a coherent drive at ω_r, the generated state cannot contain N photons so that it is restricted to N levels.

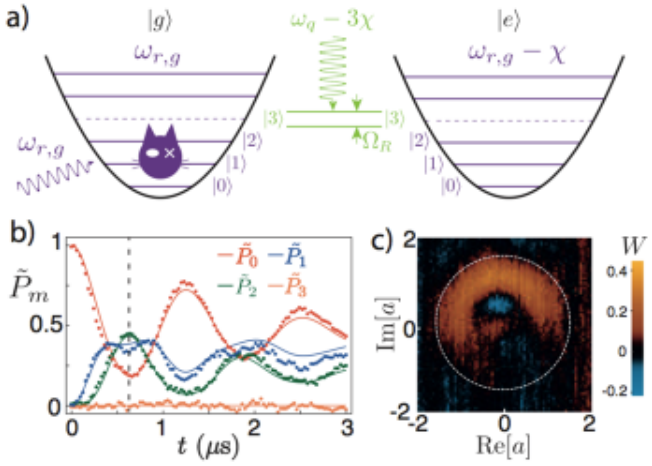

Figure 7. a) Combined energy level diagram for the qubit and cavity. By applying a strong Rabi drive on the $|3, g\rangle \longleftrightarrow |3, e\rangle$ transition, the $|2\rangle \longleftrightarrow |3\rangle$ transition of the cavity becomes off resonant at $\omega_{r,g}$. b) Oscillations of the Fock state occupation when driving the cavity mode from the vacuum and blocking $|3\rangle$. c) Wigner tomography of the field at half period of oscillation (dashed line in b). The quasi-probability density is confined within a circular barrier of radius $\sqrt{3}$ (white circle). Negativities (in blue) reveal a non classical state.

When measuring the Fock state occupation probabilities as a function of time for this effective driven N-level system, characteristic oscillations appear (see Fig. 7 b). Quantum coherence of the field is revealed by direct Wigner tomography (see Fig. 7 c). At half-period of the oscillations, fringes with negativities can be observed. This non classical state is similar to a "Schrödinger cat state", confined in phase space within a circular barrier of radius $\sqrt{N}$.

All these observations are well captured by a model based on N levels only. Our results demonstrate that QZD allows the direct control of the field state in its phase space. This experiment paves the way to the realization of various protocols, such as phase space tweezers, generation and protection of entanglement, and quantum logic operations.

6.4. Efficient quantum filtering for quantum feedback control

Participants: Pierre Rouchon

The results of this section were published in [23].

We discuss an efficient numerical scheme for the recursive filtering of diffusive quantum stochastic master equations. We show that the resulting quantum trajectory is robust and may be used for feedback based on inefficient measurements. The proposed numerical scheme is amenable to approximation, which can be used to further reduce the computational burden associated with calculating quantum trajectories and may allow real-time quantum filtering. We provide a two-qubit example where feedback control of entanglement may be within the scope of current experimental systems.

6.5. Adaptive low-rank approximation and denoised Monte-Carlo approach for high-dimensional Lindblad equations

Participants: Pierre Rouchon

The results of this section were published in [17].

We present a twofold contribution to the numerical simulation of Lindblad equations. First, an adaptive numerical approach to approximate Lindblad equations using low-rank dynamics is described: a deterministic low-rank approximation of the density operator is computed, and its rank is adjusted dynamically, using an on-the-fly estimator of the error committed when reducing the dimension. On the other hand, when the intrinsic dimension of the Lindblad equation is too high to allow for such a deterministic approximation, we combine classical ensemble averages of quantum Monte Carlo trajectories and a denoising technique. Specifically, a variance reduction method based upon the consideration of a low-rank dynamics as a control variable is developed. Numerical tests for quantum collapse and revivals show the efficiency of each approach, along with the complementarity of the two approaches.

This work results from a collaboration with Claude Le Bris of the Matherials project-team and in the framework of the ANR-project EMAQS entitled "Evaluation and Manipulation At Quantum Scale" coordinated by Karine Beauchard from ENS-Rennes.

6.6. Stabilization of photon-number states via single-photon corrections: a first convergence analysis under an ideal set-up

Participants: Pierre Rouchon

The results of this section were published in [33].

This work presents a first mathematical convergence analysis of a Fock states feedback stabilization scheme via single-photon corrections. This measurement-based feedback has been developed and experimentally tested in 2012 by the cavity quantum electrodynamics group of Serge Haroche and Jean-Michel Raimond. Here, we consider the infinite-dimensional Markov model corresponding to the ideal set-up where detection errors and feedback delays have been disregarded. In this ideal context, we show that any goal Fock state can be stabilized by a Lyapunov-based feedback for any initial quantum state belonging to the dense subset of finite rank density operators with support in a finite photon-number sub-space. Closed-loop simulations illustrate the performance of the feedback law.

Paulo Sergio Pereira da Silva and Pierre Rouchon are participants to the Inria associate Team CDSS with principal Inria investigator, François Dufour of the Inria Team Project CQFD on the topic "Control of dynamic systems subject to stochastic jumps".

6.7. Convergence and adiabatic elimination for a driven dissipative quantum harmonic oscillator

Participants: Rémi Azouit, Alain Sarlette, Pierre Rouchon

The results of this section were published in [30].

We prove that a harmonic oscillator driven by Lindblad dynamics where the typical drive and loss channels are two-photon processes instead of single-photon ones, converges to a protected subspace spanned by two coherent states of opposite amplitude. We then characterize the slow dynamics induced by a perturbative single-photon loss on this protected subspace, by performing adiabatic elimination in the Lindbladian dynamics.

6.8. Parameter estimation from measurements along quantum trajectories

Participants: Pierre Six, Ph. Campagne-Ibarcq, Benjamin Huard, Pierre Rouchon

The results of this section were published in [34].

The dynamics of many open quantum systems are described by stochastic master equations. In the discrete-time case, we recall the structure of the derived quantum filter governing the evolution of the density operator conditioned to the measurement outcomes. We then describe the structure of the corresponding particle quantum filters for estimating constant parameter and we prove their stability. In the continuous-time (diffusive) case, we propose a new formulation of these particle quantum filters. The interest of this new formulation is first to prove stability, and also to provide an efficient algorithm preserving, for any discretization step-size, positivity of the quantum states and parameter classical probabilities. This algorithm is tested on experimental data to estimate the detection efficiency for a superconducting qubit whose fluorescence field is measured using a heterodyne detector.

6.9. Adding a single state memory optimally accelerates symmetric linear maps

Participants: Alain Sarlette

The results of this section are to be published in IEEE Trans. Automatic Control [24].

This work is exploring the context and benefits of so-called “non-Markovian” dynamics, where the dynamics implied by hidden variables modifies the behavior of an iterative procedure. Such mechanisms appear in both classical and quantum systems, and one of our future goals is to better characterize the benefits of engineered non-Markovianity in terms of stabilizing power in very constrained systems. The precise setting here is a discrete-time linear map, which is unknown except for a lower and upper bound on its eigenvalues. By adding one memory slot to each coordinate, this map can be accelerated quadratically. We prove that by adding more memory slots, this cannot be further improved. This is reminiscent of the acceleration of random walks by lifting them or by quantizing them, which we are currently exploring.

6.10. A common symmetrization framework for iterative (linear) maps

Participants: Alain Sarlette

The results of this section were presented at [29].

We review a “symmetrization” abstraction of iterative consensus algorithms, which allows to generalize them to general discrete group operations including those acting on quantum systems and on sequences of control actions. We highlight a few new applications of the framework including: consensus networks with antagonistic interactions; sub-stochastic matrix iterations; and coordinate descent on (locally) quadratic functions. The purpose is to show which types of iterative dynamics can be covered by this group-theoretic framework, and potentially operationally generalized to non-classical systems.

6.11. Deterministic hidden coordinate for a qubit under fluorescence measurement

Participants: Alain Sarlette, Pierre Rouchon

The experimentalists in the group have set up an experiment with continuous heterodyne measurement of an energy loss operator on a superconducting qubit. We have observed that in the associated mathematical model, due to the degeneracy of the diffusion operator, the resulting quantum trajectories are supported not in the entire Bloch sphere, but instead they belong to the surface of a *deterministically* evolving ellipsoid. We have entirely characterized this fact and highlighted that such behavior is not generic. A paper comparing this to the experimental data and a more general theory about deterministic evolutions in quantum stochastic differential equations are being finalized. This work has been presented at [28].

6.12. Relations between quantum walks, open quantum walks, and lifted walks: the cycle graph

Participants: Alain Sarlette

The convergence time of a random walk on a graph towards its stationary distribution is an important indication of the efficiency of random algorithms based on it. Quantum random walks have been shown to allow quadratically accelerated convergence for large graphs, at least in some cases. The famous Grover search algorithm has been shown to actually fit this framework in an abstracted setting (it is doing the opposite of a random walk: converging from the uniform distribution towards a particular identified element). Yet also with classical dynamics, simple mechanisms have been proposed which allow to quadratically accelerate the convergence with respect to a standard random walk. Some basic principles have been conjectured to cause this acceleration, basically transforming a diffusion-like behavior into a more transport-like behavior, but with remaining trail. We are working towards formally characterizing the effect of these principles, and extracting similar principles in the quantum walks. This should help identify key effects to be protected in the associated quantum algorithms. We currently have worked out the equivalence of all these accelerating settings for the simplest example of the cycle graph. Quantum coherences turn out to play no major role and a classical feedback structure can be identified. We are now working towards other graphs, where the convergence effect of quantum coherences might be hidden in propagating classical information. This work has been presented at [35].

6.13. Confining the state of light to a quantum manifold by engineered two-photon loss

Participants: Zaki Leghtas and Mazyar Mirrahimi

Physical systems usually exhibit quantum behavior, such as superpositions and entanglement, only when they are sufficiently decoupled from a lossy environment. Paradoxically, a specially engineered interaction with the environment can become a resource for the generation and protection of quantum states. This notion can be generalized to the confinement of a system into a manifold of quantum states, consisting of all coherent superpositions of multiple stable steady states. In a collaboration with the team of Michel H. Devoret at Yale university, we have confined the state of a superconducting resonator to the quantum manifold spanned by two coherent states of opposite phases and have observed a Schrödinger cat state spontaneously squeeze out of vacuum before decaying into a classical mixture. As suggested by our earlier work [93], this experiment points toward robustly encoding quantum information in multidimensional steady-state manifolds and should lead to significant hardware shortcuts for quantum error correction and fault-tolerant quantum computation.

This experimental work was published in Science [18].

6.14. Single-Photon-Resolved Cross-Kerr Interaction for Autonomous Stabilization of Photon-Number States

Participants: Zaki Leghtas and Mazyar Mirrahimi

Quantum states can be stabilized in the presence of intrinsic and environmental losses by either applying active feedback conditioned on an ancillary system or through reservoir engineering. Reservoir engineering maintains a desired quantum state through a combination of drives and designed entropy evacuation. In a collaboration with the team of Robert J. Schoelkopf at Yale university, we propose and implement a quantum reservoir engineering protocol that stabilizes Fock states in a microwave cavity. This protocol is realized with a circuit quantum electrodynamics platform where a Josephson junction provides direct, nonlinear coupling between two superconducting waveguide cavities. The nonlinear coupling results in a single photon resolved cross-Kerr effect between the two cavities enabling a photon number dependent coupling to a lossy environment. The quantum state of the microwave cavity is discussed in terms of a net polarization and is analyzed by a measurement of its steady state Wigner function.

This work was published in Physical Review Letters [15].

6.15. Characterizing entanglement of an artificial atom and a cavity cat state with Bell's inequality

Participants: Zaki Leghtas and Mazyar Mirrahimi

The Schrödinger's cat thought experiment highlights the counterintuitive concept of entanglement in macroscopically distinguishable systems. The hallmark of entanglement is the detection of strong correlations between systems, most starkly demonstrated by the violation of a Bell inequality. No violation of a Bell inequality has been observed for a system entangled with a superposition of coherent states, known as a cat state. In a collaboration with the team of Robert J. Schoelkopf at Yale university, we use the Clauser-Horne-Shimony-Holt formulation of a Bell test to characterize entanglement between an artificial atom and a cat state, or a Bell-cat. Using superconducting circuits with high-fidelity measurements and real-time feedback, we detect correlations that surpass the classical maximum of the Bell inequality. We investigate the influence of decoherence with states up to 16 photons in size and characterize the system by introducing joint Wigner tomography. Such techniques demonstrate that information stored in superpositions of coherent states can be extracted efficiently, a crucial requirement for quantum computing with resonators.

This work was published in Nature Communications [25].

RAP Project-Team

4. New Results

4.1. Random Graphs

Participant: Nicolas Broutin.

And/Or trees for random Boolean functions

For some time, a number of teams have tried to devise natural probability distributions on Boolean functions. Indeed, the most natural one, the uniform one, is not quite satisfactory: almost all Boolean functions have maximal complexity, while it is extremely difficult to construct some with high complexity. One approach consists in generating functions by seeing them as "expressions" encoded as a tree of computation. We generalize and unify the previous approaches that are restricted to very specific cases by looking at the distributions induced on the Boolean function by large computation trees that are arbitrary, except for the fact they the neighborhoods of the root (where the computation concentrates) stabilizes in distribution as the sizes of the tree increases [12].

4.2. Resource Allocation in Large Data Centres

Participants: Christine Fricker, Philippe Robert, Guilherme Thompson.

With the exponential increase in internet data transmission volume over the past years, efficient bandwidth allocation in large data centres has become crucial. Illustrating examples are the rapid spread of cloud computing technology, as well as the growth of the demand for video streaming, both of which were quasi non-existent 10 years ago.

Currently, most systems operate under decentralised policies due to the complexity of managing data exchange on large scales. In such systems, customer demands are served respecting their initial service requirements (a certain video quality, amount of memory or processing power etc.) until the system reaches saturation, which then leads to the blockage of subsequent customer demands. Strategies that rely on the scheduling of tasks are often not suitable to address this load balancing problem as the users expect instantaneous service usage in real time applications, such as video transmission and elastic computation. Our research goal is to understand and redesign its algorithms in order to develop decentralised policies that can improve global performance using local instantaneous information. This research is made in collaboration with Fabrice Guillemin, from Orange Labs.

In a first approach to this problem, we examined offloading schemes in fog computing context, where one data centres are installed at the edge of the network. We analyse the case with one data centre close to user which is backed up by a central (bigger) data centre. When a request arrives at an overloaded data centre, it is forwarded to the other data centre with a given probability, in order to help coping with saturation and reducing the rejection of requests. In [16], we have been able to show that the performance of such a system can be expressed in terms of the invariant distribution of a random walk in the quarter plane. As a consequence we have been able to assess the behaviour and performance of these systems, proving the effectiveness of such an offloading arrangement.

In a second step, we investigated allocation schemes which consist in reducing the bandwidth of arriving requests to a minimal value when the system is close to saturation. We analysed the effectiveness of such a downgrading policy, which, if the system is correctly designed, will reduce the fraction of rejected transmissions. We developed a mathematical model which allows us to predict system behaviour under such a policy and calculate the ideal threshold (in the same scale as the resource) after which downgrading should be initiated, given system parameters. We proved the existence of a unique equilibrium point, around which we have been able to determine the probability of the system being above or under the threshold. We found that system blockage can be almost surely eliminated. This policy finds a natural application in the context of video streaming services and other real time applications, such as MPEG-DASH. A document is being written to further publication.

Finally, with those results, we now try to extend our research towards more complex systems, investigating the behaviour of multiple resource systems (such as a Cloud environment, where computational power is provided using unities of CPU and GB of RAM) and other offloading schemes, such as the compulsory forwarding of a request when it's blocked at the edge server, but keeping a trunk reservation to protect the service originally assigned to the big data centre.

4.3. Resource allocation in vehicle sharing systems

Participants: Christine Fricker, Plinio Santini Dester, Hanene Mohamed, Yousra Chabchoub.

This is a collaboration with Danielle Tibi, Université Denis Diderot.

Vehicle sharing systems are becoming an urban mode of transportation, and launched in many cities, as Velib' and Autolib' in Paris. One of the major issues is the avail ability of the resources: vehicles or free slots to return them. These systems became a hot topic in Operation Research and now the importance of stochasticity on the system behavior is commonly admitted. The problem is to understand the system behavior and how to manage these systems in order to provide both resources to users. Our stochastic model is the first one taking into account the finite number of spots at the stations.

With Danielle Tibi, we use limit local theorems to obtain the asymptotic stationary joint distributions of several station states when the system is large (both numbers of stations and bikes), in the case of finite capacities of the stations. This gives an asymptotic independence property for node states. This widely extends the existing results on heterogeneous bike-sharing systems.

Recently we investigate some network load balancing algorithms to improve the bike sharing system behavior. We focus on the choice of the least loaded station among two t o return the bike. A problem is the influence of the delay between the choice time (the beginning of the trip) and the time the station is joined (the end of the trip). However the main challenge is to deal with the choice between two neighboring stations. For that, a system of infinite queues is studied in light traffic. For a bike-shar ing homogeneous model, we restrict our study to a deterministic cooperation of two by two stations. It relies on new results for the classical system of two queues under the join-the-shortest-queue policy.

JC Decaux provides us data describing Velib' user trips. These data are useful to measure the system parameters, validate our models and test our algorithms. Indeed, we use these data to investigate load balancing algorithms such as two-choice policies.

4.4. Scaling Methods

Participants: Philippe Robert, Wen Sun.

4.4.1. Fluid Limits in Wireless Networks

This is a collaboration with Amandine Veber (CMAP, École Polytechnique). The goal is to investigate the stability properties of wireless networks when the bandwidth allocated to a node is proportional to a function of its backlog: if a node of this network has x requests to transmit, then it receives a fraction of the capacity proportional to $\log(1+x)$, the logarithm of its current load. This year we completed the analysis of a star network topology with multiple nodes. Several scalings were used to describe the fluid limit behaviour.

4.4.2. The Time Scales of a Transient Network

A large distributed system where users' files are duplicated on unreliable data servers is investigated. Due to a server breakdown, a copy of a file can be lost, it can be retrieved if another copy of the same file is stored on other servers. In the case where no other copy of a given file is present in the network, it is definitely lost. In order to have multiple copies of a given file, it is assumed that each server can devote a fraction of its processing capacity to duplicate files on other servers to enhance the durability of the system.

A trade-off is necessary between the bandwidth and the memory used for this back-up mechanism and the data loss rate. Back-up mechanisms already exist and have been studied thanks to simulation. To our knowledge, no theoretical study exists on this topic. With a very simple centralized model, we have been able to emphasise a trade-off between capacity and life-time with respect to the duplication rate. From a mathematical point of view, we are currently studying different time scales of the system with an averaging phenomenon.

We have used scaling methods with different time scales to derive some asymptotic results on the decay of a simplified network: it is assumed that any copy of a given file is lost at some fixed rate and the total processing capacity of the system is devoted to duplicate the file with least number of copies. We start from the optimal initial state: each file has the maximum number of copies. Due to random losses, the state of the network is transient and all files will be eventually lost. There is a stability assumption for the system having a critical time scale of decay. When the stability condition is not satisfied, i.e. when it is initially overloaded, we have shown that the state of the network converges to an interesting local equilibrium. We are currently studying a more general case which the duplication depends on the structure of the system. See [7].

4.5. Stochastic Models of Biological Networks

Participants: Renaud Dessalles, Sarah Eugene, Philippe Robert.

4.5.1. Stochastic Modelling of self-regulation in the protein production system of bacteria

This is a collaboration with Vincent Fromion from INRA Jouy-en-Josas, which started on December 2014.

In prokaryotic cells (e.g. E. Coli. or B. Subtilis) the protein production system has to produce in a cell cycle (i.e. less than one hour) more than 10^6 molecules of more than 2500 kinds, each having different level of expression. The bacteria uses more than 85% of its resources to the protein production. Gene expression is a highly stochastic process: bacteria sharing the same genome, in a same environment will not produce exactly the same amount of a given protein. Some of this stochasticity can be due to the system of production itself: molecules, that take part in the production process, move freely into the cytoplasm and therefore reach any target in the cell after some random time; some of them are present in so much limited amount that none of them can be available for a certain time; the gene can be deactivated by repressors for a certain time, etc. We study the integration of several mechanisms of regulation and their performances in terms of variance and distribution. As all molecules tends to move freely into the cytoplasm, it is assumed that the encounter time between a given entity and its target is exponentially distributed.

4.5.1.1. Feedback model

We have also investigated the production of a single protein, with the transcription and the translation steps, but we also introduced a direct feedback on it: the protein tends to bind on the promoter of its own gene, blocking therefore the transcription. The protein remains on it during an exponential time until its detachment caused by thermal agitation.

The mathematical analysis aims at understanding the nature of the internal noise of the system and to quantify it. We tend to test the hypothesis usually made that such feedback permits a noise reduction of protein distribution compared to the “open loop” model. We have made the mathematical analysis of the model (using a scaling to be able to have explicit results), it appeared that reduction of variance compared to an “open loop” model is limited: the variance cannot be reduced for more than 50%.

We proposed another possible effect of the feedback loop: the return to equilibrium is faster in the case of a feedback model compared to the open loop model. Such behaviour can be beneficial for the bacteria to change of command for a new level of production of a particular protein (due, for example, to a radical change in the environment) by reducing the respond time to reach this new average. This study has been mainly performed by simulation and it has been shown that the feedback model can go 50% faster than the open loop results. See [13]..

4.5.1.2. Transcription-translation model for all proteins

The other model that has been studied integrates the production of all the proteins. Each gene has to be transcribed in mRNA (using RNA-Polymerase molecules) and each mRNA has to be translated in protein (using ribosome molecules). Experiments (as the one from Taniguchi et al. (2010)) have shown that protein production is subject to high variability especially for highly expressed proteins. Our goal is to determine what in the protein production mechanism is responsible for the noise.

We already made simulations that takes into amount of RNA-Polymerases and Ribosomes and that genes and mRNAs sequester these molecules during the whole the time of elongation. This global sharing of Ribosomes/RNA-Polymerases reproduce only a part of the unknown noise experimentally seen. We are developing Python simulations that extends this model and take into account other feature that might be responsible for the noise in protein production. This new simulation will include new features such as:

- The volume of the cell. We consider it as proportional to the total number of proteins, and will increase as the cell grows. Transcription and translation initiation are then depending on the concentration of respectively free RNA-polymerase and free ribosomes.
- The division of the cell. At division, all components have an equal chance to go in either one of the two daughter cell.
- DNA replication. At some point in the cell cycle, the genome duplicates, doubling therefore the copy number of each gene

The simulation parameters will be fit with the data of Taniguchi et al. (2010) and the goal is to compare our result to see if which aspects of the protein production are responsible for the noise of the proteins.

4.5.2. Stochastic Modelling of Protein Polymerization

This is a collaboration with Marie Doumic, Inria MAMBA team.

The first part of our work focuses on the study of the polymerization of protein. This phenomenon is involved in many neurodegenerative diseases such as Alzheimer's and Prion diseases, e.g. mad cow. In this context, it consists in the abnormal aggregation of proteins. Curves obtained by measuring the quantity of polymers formed in in vitro experiments are sigmoids: a long lag phase with almost no polymers followed by a fast consumption of all monomers. Furthermore, repeating the experiment under the same initial conditions leads to somewhat identical curves up to translation. After having proposed a simple model to explain this fluctuations, we studied a more sophisticated model, closer to the reality. We added a conformation step: before being able to polymere, proteins have to misfold. This step is very quick and remains at equilibrium during the whole process. Nevertheless, this equilibrium depends on the polymerization which is happening on a slower time scale. The analysis of these models involves stochastic averaging principles.

The second part concerns the study of telomeres. This work is made in collaboration with Zhou Xu, Teresa Teixeira, from IBCP in Paris.
In eukaryotic cells, at each mitosis, chromosomes are shortened, because the DNA polymerase is not able to duplicate one ending of the chromosome. To prevent loss of genetic information- which could be catastrophic for the cell-chromosomes are equipped with telomeres at their endings. These telomeres do not contain any genetic information; they are a repetition of the sequence T-T-A-G-G-G thousands times. At each mitosis, there is therefore a loss of telomere. As it has a finite length, when the telomeres are too short, the cell cannot divide anymore: they enter in replicative senescence. Our model tries to captures the two phases of the shortening of telomeres: first, the initial state of the cells, when the telomerase is still active to repair the telomeres. Second, when the telomerase is inhibited, we try to estimate the senescence threshold, when the replication of the cells stops.

REGAL Project-Team

6. New Results

6.1. Distributed algorithms for dynamic networks

Participants: Luciana Bezerra Arantes [correspondent], Marjorie Bournat, Swan Dubois, Denis Jeanneau, Mohamed Hamza Kaaouachi, Sébastien Monnet, Franck Petit [correspondent], Pierre Sens, Julien Sopena.

Nowadays, distributed systems are more and more heterogeneous and versatile. Computing units can join, leave or move inside a global infrastructure. These features require the implementation of dynamic systems, that is to say they can cope autonomously with changes in their structure in terms of physical facilities and software. It therefore becomes necessary to define, develop, and validate distributed algorithms able to managed such dynamic and large scale systems, for instance mobile *ad hoc* networks, (mobile) sensor networks, P2P systems, Cloud environments, robot networks, to quote only a few.

We have obtained results both on fundamental aspects of distributed algorithms and on specific emerging large-scale applications.

We study various key topics of distributed algorithms: agreement, failure detection, data dissemination and data finding in large scale systems, self-stabilization and self-* services.

6.1.1. Agreement and failure detection in dynamic Distributed Systems

Distributed systems should provide reliable and continuous services despite the failures of some of their components. A classical way for a distributed system to tolerate failures is to detect them and then to recover. It is now well recognized that the dominant factor in system unavailability lies in the failure detection phase. In 2015, we obtain the following results on failure detection:

Assuming a message-passing environment with a majority of correct processes, the necessary and sufficient information about failures for implementing a general state machine replication scheme ensuring consistency is captured by the Ω failure detector. We show in [46] that in such a message-passing environment, Ω is also the weakest failure detector to implement an eventually consistent replicated service, where replicas are expected to agree on the evolution of the service state only after some (a priori unknown) time.

We also study the k-set agreement problem is a generalization of the consensus problem where processes can decide up to k different values. Very few papers have tackled this problem in dynamic networks. Exploiting the formalism of the Time Varying Graph model, we propose in [70]a new quorum-based failure detector for solving k-set agreement in dynamic networks with asynchronous communications. We present two algorithms that implement this new failure detector using graph connectivity and message pattern assumptions. We also provide an algorithm for solving k-set agreement using our new failure detector.

We propose several algorithms to implement efficient failure detection services. We introduce in [60] the Two Windows Failure Detector (2WFD), an algorithm that provides QoS and is able to react to sudden changes in network conditions, a property that currently existing algorithms do not satisfy. We ran tests on real traces and compared the 2W-FD to state-of-the-art algorithms. Our results show that our algorithm presents the best performance in terms of speed and accuracy in unstable scenarios. In [62], we propose a new approach towards the implementation of failure detectors for large and dynamic networks: we study reputation systems as a means to detect failures. The reputation mechanism allows efficient node cooperation via the sharing of views about other nodes. Our experimental results show that a simple prototype of a reputation-based detection service performs better than other known adaptive failure detectors, with improved flexibility. It can thus be used in a dynamic environment with a large and variable number of nodes.

6.1.2. Probabilistic Byzantine Tolerance allocation strategies in Hybrid Cloud Environments

We explore the node allocation challenges in providing probabilistic Byzantine fault tolerance in a hybrid cloud environment, consisting of nodes with varying reliability levels, compute power, and monetary cost. We consider hybrid computing architectures that combine edge nodes with cloud hosted computing. In such a system, a large fraction of the computation is performed by donated machines at the edge of the network, which significantly reduces the cost to the owner of the computation.

Considering "bag of tasks" (BoT) applications where a large computational problem is broken into a large number of independent tasks, the probabilistic Byzantine fault tolerance guarantee refers to the confidence level that the result of a given computation is correct despite potential Byzantine failures. In [36] we explore probabilistic Byzantine tolerance, in which computation tasks are replicated on dynamic replication sets whose size is determined based on ensuring probabilistic thresholds of correctness.

6.1.3. Covering problems in dynamic systems

We study covering problems (such as minimal dominating set or maximal matching) in the context of highly dynamic distributed systems. We first obtain some general results. In [48], we first propose a new definition of this family of problems since classical ones are meaningless in such systems. We generalize the classical definition of time complexity (for static systems) to our setting. We also provided in [40] a generic tool to help the writing of impossibility proofs in dynamic distributed systems. Then, we focus on the particular case of the minimal dominating set problem. We characterize the necessary and sufficient condition to construct deterministically a minimal dominating set in a dynamic system according to our definition.

6.1.4. Self-Stabilization

Self-stabilization is a generic paradigm to tolerate transient faults (*i.e.*, faults of finite duration) in distributed systems. Results obtained in this area by Regal members in 2015 follow.

Spanning tree construction is a well-studied problem in distributed computing for its numerous applications like routing, broadcast...Properties of the obtained trees, efficiency of the construction, and fault-tolerance guarantees are naturally at the heart of many researches. In this context, we propose in [39] a new self-stabilizing algorithm for the minimum diameter spanning tree that achieves better time and space complexity than existing solutions. Moreover, our solution tolerates a fully asynchronous adversary.

A classical way to endowed self-stabilization with (permanent) fault tolerance is *confinement*. That is, we ensure that the self-stabilizing system moreover ensures that the effect of permanent faults is limited to some topological areas of the system. In [27], we propose a characterization of optimal confinement areas for a large set of spanning tree metrics in presence of Byzantine faults. In [24], we propose a stabilizing implementation of an atomic register in presence of crash faults. By avoiding the propagation of fault effects further than a given radius, confinement is clearly a *spatial* approach. Another approach, called *temporal*, consists in recovering as quick as possible to a configuration from which some forms of safety are satisfied.

In [68], we introduce the notion of *gradual stabilization* and provide a gradually self-stabilizing algorithm that solves the *unison* problem, *i.e.*, the problem that consists in synchronizing logical clocks locally maintained by the processes.

6.1.5. Team of Mobile Robots

Swarm of autonomous mobile sensor devices (or, robots) recently emerged as an attractive issue in the study of dynamic distributed systems permits to assess the intrinsic difficulties of many fundamentals tasks, such as exploring or gathering in a discrete space. We consider autonomous robots that are endowed with visibility sensors (but that are otherwise unable to communicate) and motion actuators. The robots we consider are weak, *i.e.*, they are anonymous, uniform, unable to explicitly communicate, and oblivious (they do not remember any of their past actions). Despite their weakness, those robots must collaborate to solve a collective tasks such as exploration, gathering, flocking, to quote only a few.

In [45], we first show that it is impossible to explore any simple torus of arbitrary size with (strictly) less than four robots, even if the algorithm is probabilistic. Next, we propose an optimal (*w.r.t.* the number of robots) solution for the terminating exploration of torus-shaped networks by a team of k such robots in the SSYNC model. The proposed algorithm is probabilistic and works for any simple torus of size $\ell \times L$, where $7 \leq \ell \leq L$. Since the optimal number of robots is also four in rings, our result shows that increasing the number of possible symmetries in the network (due to increasing dimensions) does not necessarily come at an extra cost *w.r.t.* the number of robots that are necessary to solve the problem.

6.2. Management of distributed data

Participants: Rudyar Cortes, Mesaac Makpangou, Olivier Marin, Sébastien Monnet [correspondent], Pierre Sens.

6.2.1. Long term durability and storage load distribution

In 2014, we had proposed SPLAD (for Scattering and PLAcing Data replicas to enhance long-term durability), a model that allows us to vary the data scattering degree by tuning a selection range width. We have enhanced our model [57] and we have focused on the study of the policy used while choosing a storing node within the selection range. Some policies may lead to heavily unbalanced storage load distribution which can be harmful for the system. Simple policies to balance the load (e.g. storing new blocks on least loaded nodes) may induce network congestion and thus data losses. We have shown that the "power of two choices" policy (choosing the least loaded node among two random ones) brings good results both in terms of storage load distribution and fault tolerance.

6.2.2. Management of dynamic big data

Managing and processing Dynamic Big Data, where multiple sources produce new data continuously, is very complex. Static cluster- or grid-based solutions are prone to induce bottleneck problems, and are therefore ill-suited in this context. Our objective in this domain is to design and implement a Reliable Large Scale Distributed Framework for the Management and Processing of Dynamic Big Data. In 2015, we focused on Spatio-temporal range queries over Big Location Data aim to extract and analyze relevant data items generated around a given location and time. They require concurrent processing of massive and dynamic data flows. We proposed a scalable architecture for continuous spatio-temporal range queries built by coalescing multiple computing nodes on top of a Distributed Hash Table. The key component of our architecture is a distributed spatio-temporal indexing structure which exhibits low insertion and low index maintenance costs. We assessed our solution with a public data set released by Yahoo! which comprises millions of geotagged multimedia files [43].

6.3. CISE Logic and tool for proving invariants in distributed databases

Participants: Marc Shapiro [correspondent], Mahsa Najafzadeh, Alexey Gotsman, Carla Ferreira.

We have developed a new sound logic for proving the correctness of a distributed database under concurrent updates, showing whether the application maintains the database's *integrity invariants*. An operation of the application is specified as a *preparator*, which checks the operation's precondition at an origin replica and generates an *effector*. The effector abstracts the update to be applied to every replica. The application also specifies which operations are allowed to take place concurrently. In summary, the logic shows that the application maintains the invariant if the three following rules are satisfied:

- Each operation individually maintains the invariant. It follows that operations' preconditions are sufficiently strong to ensure correctness in a sequential execution.
- The effectors of any two operations that can execute concurrently commute. This implies that the database replicas all converge to the same state.
- For any pair of operations u and v that can execute concurrently, the precondition of u is stable under the effector of v, and vice-versa.

This result is published at POPL 2016 [50].

We have implemented a tool (based on the Z3 SMT solver) that implements these rules. A demo of the tool is available online [78]. If the application passes the tool, it is correct. If not, the tool returns a counter-example, which the application developer can inspect to find the source of the error. Generally speaking, the developer can either weaken the invariants or the effects of operations, or strengthen consistency by disallowing concurrency. By choosing one or the other, the developer performs a co-design of the application with its consistency protocol, in order to have the highest possible concurrency that still ensures correctness.

For instance, consider a database of bank accounts, with the invariant that an account's balance must be positive. The banking application has operations $credit(acct, amt)$, $debit(acct, amt)$, and $accrue - interest(acct)$. The first rule dictates that $debit$ has the precondition $amt = balance$. The second rule dictates that $accrue - interest$ computes the amount of interest according to the state at the origin, not at every replica. The third rule is violated if concurrent $debit$s are allowed; if the bank wishes to uphold the invariant, the only correct solution is to disallow concurrent $debit$s.

6.4. Memory management for big data

Participants: Antoine Blin, Damien Carver, Maxime Lorrillere, Sébastien Monnet, Julien Sopena [correspondent].

6.4.1. Automated file cache pooling

Some applications, like online sales servers, intensively use disk I/Os. Their performance is tightly coupled with I/Os efficiency. To speed up I/Os, operating systems use free memory to offer caching mechanisms. Several I/O intensive applications may require a large cache to perform well. However, nowadays resources are virtualized. In clouds, for instance, virtual machines (VMs) offer both isolation and flexibility. This is the foundation of cloud elasticity, but it induces fragmentation of the physical resources, including memory. This fragmentation reduces the amount of available memory a VM can use for caching I/Os. Previously, we proposed Puma (for Pooling Unused Memory in Virtual Machines) which allows I/O intensive applications running on top of VMs to benefit of large caches. This was realized by providing a remote caching mechanism that provides the ability for any VM to extend its cache using the memory of other VMs located either in the same or in a different host.

We have performed an extensive evaluation of Puma [53] and we have enhanced our solution: Puma adapts automatically the amount a memory that a VM offers to another VM. Furthermore, if the network becomes overloaded, Puma detects a performance degradation and stops using a remote cache.

REO Project-Team

7. New Results

7.1. Mathematical and numerical analysis of fluid-structure interaction problems

Participants: Matteo Aletti, Faisal Amlani, Benoit Fabrèges, Miguel Ángel Fernández Varela, Jean-Frédéric Gerbeau, Mikel Landajuela Larma, Damiano Lombardi, Marina Vidrascu.

In [55] we present a numerical study in which several partitioned solution procedures for incompressible fluid-structure interaction are compared and validated against the results of an experimental FSI benchmark. The numerical methods discussed cover the three main families of coupling schemes: strongly coupled, semi-implicit and loosely coupled. Very good agreement is observed between the numerical and experimental results. The comparisons confirm that strong coupling can be efficiently avoided, via semi-implicit and loosely coupled schemes, without compromising stability and accuracy.

In [14] we introduce a Nitsche-XFEM method for fluid-structure interaction problems involving a thin-walled elastic structure (Lagrangian formalism) immersed in an incompressible viscous fluid (Eulerian formalism). The fluid domain is discretized with an unstructured mesh not fitted to the solid mid-surface mesh. Weak and strong discontinuities across the interface are allowed for the velocity and pressure, respectively. The fluid-solid coupling is enforced consistently using a variant of Nitsche's method with cut-elements. Robustness with respect to arbitrary interface intersections is guaranteed through suitable stabilization. Several coupling schemes with different degrees of fluid-solid time splitting (implicit, semi-implicit and explicit) are investigated. A series of numerical tests in 2D, involving static and moving interfaces, illustrates the performance of the different methods proposed.

In [15] we investigated the autoregulation in the retinal haemodynamics by means of three-dimensional simulations. The autoregulation is a key phenomenon from a physiological standpoint, consisting in the ability of the vasculature to control the flow in different pressure conditions. A simplified fluid-structure interaction method was devised in order to render the vessels wall contraction in a large network, with an affordable computational cost. Several test cases were performed on a patient-specific arteriolar network, whose geometry was reconstructed by using fundus camera images. The tests were in agreement with experimental trends and confirm the ability of the approach to reproduce the phenomena involved.

In [33] we study an unsteady nonlinear fluid-structure interaction problem which is a simplified model to describe blood flow through viscoelastic arteries. We consider a Newtonian incompressible two-dimensional flow described by the Navier-Stokes equations set in an unknown domain depending on the displacement of a structure, which itself satisfies a linear viscoelastic beam equation. The fluid and the structure are fully coupled via interface conditions prescribing the continuity of the velocities at the fluid-structure interface and the action-reaction principle. We prove that strong solutions to this problem are global-in-time. We obtain in particular that contact between the viscoleastic wall and the bottom of the fluid cavity does not occur in finite time. To our knowledge, this is the first occurrence of a no-contact result, but also of existence of strong solutions globally in time, in the frame of interactions between a viscous fluid and a deformable structure.

In [27] and [45] we study the effect of wall bending resistance on the motion of an initially spherical capsule freely suspended in shear flow or in a a planar hyperbolic flow. We consider a capsule with a given thickness made of a three–dimensional homogeneous elastic material. A numerical method is used to model the coupling of a boundary integral method for the fluids with a shell finite element method for the capsule envelope. For a given wall material, the capsule deformability strongly decreases when the wall bending resistance increases. In addition, if one expresses the same results as a function of the two–dimensional mechanical properties of the mid–surface, which is how the capsule wall is modeled in the thin–shell model, the capsule deformed shape is identical to the one predicted for a capsule devoid of bending resistance. The bending rigidity is found to have

a negligible influence on the overall deformation of an initially spherical capsule, which therefore depends only on the elastic stretching of the mid–surface. Still, the bending resistance of the wall must be accounted for to model the buckling phenomenon, which is observed locally at low flow strength and persist at steady state. We show that the wrinkle wavelength only depends on the bending number, which compares the relative importance of bending and shearing phenomena, and provide the correlation law. Such results can then be used to infer values of the bending modulus and wall thickness from experiments on spherical capsules in simple shear flow.

In [57] we consider the motion of an elastic structure represented by the nonlinear Saint-Venant Kirchhoff model immersed in a compressible fluid modeled by the compressible Navier-Stokes equations. Existence and uniqueness of a regular solution defined locally in time is proved.

7.2. Numerical methods for biological flows

Participants: Chloé Audebert, Benoit Fabrèges, Miguel Ángel Fernández Varela, Jean-Frédéric Gerbeau, Céline Grandmont, Sanjay Pant, Marc Thiriet, Irène Vignon-Clementel.

In [37] we present a closed-loop global lumped parameter model for pre stage-II single-ventricle physiology. This model, which is built on a fibre mechanics based description of the heart chambers, benefits from a novel method to describe regurgitant valves. As many as 33 model parameters are estimated from uncertain clinical measurements in two patients—with and without atrioventricular valve regurgitation—through the method of data assimilation. Results are validated qualitatively through measurements and clinical estimates that were not included in the parameter estimation procedure. The methods are shown to successfully capture patient-specific clinical observations such as double peaked nature of valvular flows and abnormalities in electrocardiogram readings.

In [39] we propose a methodology for full propagation of uncertainty from clinical data to model results that enables estimation of the confidence associated with model predictions. We illustrate this problem in a pre stage-II single-ventricle physiology, for which coherence of simulations and clinical data indicated that the flow split to the right lung was highly uncertain. We want to assess here how such uncertainty translates into surgical planning of removing the stenosis or not. Taking into account the effect of the rest of the circulation is also studied in the uncertainty propagation.

In [21] 3D blood flow simulations are carried out for the design of a stented valve reducer in enlarged ventricular outflow tracts. Different device designs are built and compared with the initial device-free state, or with the reducer alone. Results suggest that pressure loss is higher for the reducer alone than for the full device, and that the latter successfully restores hemodynamics to a healthy state. Pressure forces on the reducer and on the valve have the same magnitudes. Migration would occur towards the right ventricle rather than the pulmonary arteries.

In [44] we aim at developing a mathematical model in order to reproduce hemodynamics changes due to liver ablation surgeries. First, a 0D closed-loop model is developed, to simulate hepatectomy and compute post-operative average values. Due to the closed loop, the surgery impact both on and from the whole circulation can be captured, including bleeding and infusion. Then, a one-dimensional artery model is implemented to improve the closed-loop model and simulate better the changes in arterial waveforms due to surgery.

In [54] we investigate the spatial and time discretization of the transient Oseen equations. Finite elements with symmetric stabilization in space are combined with several time-stepping schemes (monolithic and fractional-step). Quasi-optimal (in space) and optimal (in time) error estimates are established for smooth solutions in all flow regimes. We first analyze monolithic time discretizations using the Backward Differentation Formulas of order 1 and 2 (BDF1 and BDF2). We derive a new estimate on the time-average of the pressure error featuring the same robustness with respect to the Reynolds number as the velocity estimate. Then, we analyze fractional-step pressure-projection methods using BDF1. The stabilization of velocities and pressures can be treated either implicitly or explicitly. Numerical results illustrate the main theoretical findings.

In [26] we study the effects of inserted needle on the subcutaneous interstitial flow. A goal is to describe the physical stress affecting cells during acupuncture treatment. The model consists of the convective Brinkman equations to describe the flow through a fibrous medium. Numerical studies in FreeFem++ are performed to illustrate the acute physical stress developed by the implantation of a needle that triggers the physiological reactions of acupuncture. We emphasize the importance of numerical experiments for advancing in modeling in acupuncture. In [40] we show that the acupoint must contain a highly concentrated population of mastocytes (e.g., very-high–amplitude, small-width Gaussian distribution) to get an initial proper response. Permanent signaling is provided by chemotaxis and continuous recruitment of mastocytes. Therefore, the density and distribution of mastocytes are crucial factors for efficient acupuncture as well as availability of circulating and neighboring pools of mastocytes.

In [61] we carry out a three-dimensional blood flow simulation through a complete macrovascular circuit, the cerebral venous network, rather than using reduced order simulation and partial vascular network. The bio-mechanical modeling step is carefully performed and leads to the description of the flow governed by the Navier-Stokes equations for an incompressible viscous fluid. We then numerically solve the equations with a free finite element software in five meshes of a realistic geometry obtained from medical images to prove the feasibility of the pipeline. Some particularities of the venous network, as asymmetry for example, are discussed.

7.3. Numerical methods for cardiac electrophysiology

Participants: Muriel Boulakia, Jean-Frédéric Gerbeau, Damiano Lombardi.

In [58] we investigate the monodomain equation which describes the evolution of the cardiac electrical potential and which corresponds to a coupled system involving a reaction-diffusion equation and an ordinary differential equation. Lipschitz stability inequalities are shown for the identification of some parameters of the model from measurements on the cardiac potential and the ionic variable.

In [32] we studied the application of a Reduced-Order Modeling method (Approximated Lax Pairs) to the solution of the partial differential equations describing the polarisation of tissues. Due to the complexity of the scenarios involved and the presence of propagating waves, the performances of the standard methods proposed in the literature to provide a low computational cost solution are not always satisfactory. The ALP method consists of the construction of an adaptive time dependent basis that diagonalises, at each time, a Schrödinger-type operator. Its application to several 2D and 3D test-cases on the equations arising in electrophysiology was investigated, showing that the performances of the method in terms of speed-up and accuracy are promising.

In [62] we considered the simulation of full cycles of the electrical activity of the heart and the corresponding body surface potential. The model is based on a realistic torso and heart anatomy, including ventricles and atria. One of the specificities of our approach is to model the atria as a surface, which is the kind of data typically provided by medical imaging for thin volumes. The bidomain equations are considered in their usual formulation in the ventricles, and in a surface formulation on the atria. Two ionic models are used: the Courtemanche-Ramirez-Nattel model on the atria, and the " Minimal model for human Ventricular action potentials " (MV) by Bueno-Orovio, Cherry and Fenton in the ventricles. The heart is weakly coupled to the torso by a Robin boundary condition based on a resistor-capacitor transmission condition. Various ECGs are simulated in healthy and pathological conditions (left and right bundle branch blocks, Bachmann's bundle block, Wolff-Parkinson-White syndrome). To assess the numerical ECGs, we use several qualitative and quantitative criteria found in the medical literature. Our simulator can also be used to generate the signals measured by a vest of electrodes. This capability is illustrated at the end of the article.

In [24] we address the inverse problem of electrocardiography from a new perspective, by combining electrical and mechanical measurements. Our strategy relies on the definition of a model of the electromechanical contraction which is registered on ECG data but also on measured mechanical displacements of the heart tissue typically extracted from medical images. In this respect, we establish in this work the convergence of a sequential estimator which combines for such coupled problems various state of the art sequential data

assimilation methods in a unified consistent and efficient framework. Indeed we aggregate a Luenberger observer for the mechanical state and a Reduced Order Unscented Kalman Filter applied on the parameters to be identified and a POD projection of the electrical state. Then using synthetic data we show the benefits of our approach for the estimation of the electrical state of the ventricles along the heart beat compared with more classical strategies which only consider an electrophysiological model with ECG measurements. Our numerical results actually show that the mechanical measurements improve the identifiability of the electrical problem allowing to reconstruct the electrical state of the coupled system more precisely. Therefore, this work is intended to be a first proof of concept, with theoretical justifications and numerical investigations, of the advantage of using available multi-modal observations for the estimation and identification of an electromechanical model of the heart.

7.4. Lung and respiration modeling

Participants: Laurent Boudin, Muriel Boulakia, Céline Grandmont, Jessica Oakes, Nicolas Pozin, Irène Vignon-Clementel.

In silico models of flow and transport in the lung are increasingly being used to predict regional deposition in healthy and diseased lungs. However, very few models have been validated with in vivo human or animal experimental data. In [36], we create a physiologically-based simulation of airflow and particle transport in healthy and emphysematous rat lungs. Excellent agreement between the numerical predictions and experimental data is found for the healthy lungs. However, the numerical predictions are unable to predict the experimental findings of enhanced deposition in the normal regions of the emphysematous lungs and thus more sophisticated models of transport in the deep regions of the lung are needed. This is what is being explored in [42], where interactions of flow and transport between 3D upper-parts and 1D downstream respiratory trees are captured for inspiration and expiration for the first time.

While several groups have investigated detailed flow and particle transport in the acinar regions of the healthy lung, little is currently known about diseased acini. In [35] we perform numerical simulations of flow and transport in healthy and emphysematous acini. As the alveolar septa is deteriorated in emphysema there is less surface area available for particles to deposit on. Therefore, fewer particles deposit in the diseased models. In addition, we find that particle deposition is more heterogeneously distributed in emphysema, a phenomenon that was also found in the in vivo animal experiments.

7.5. Methods for the interaction data - simulation

Participants: Jean-Frédéric Gerbeau, Damiano Lombardi, Sanjay Pant, Irène Vignon-Clementel.

In [38] we proposed an information theoretical framework to study the practical identifiability of dynamical systems. The fundamental question arising in parameter estimation problems is whether, given a set of observations of the system, it is possible to retrieve the parameters values. The method proposed exploits a database of direct numerical simulations and study the parameters-to-observables map by means of differential entropies. Contrary to other approaches proposed in the literature it is not restricted to ordinary differential equations and take the experimental noise into account. Several test cases were performed on a large spectrum of bio-physical systems, providing promising results.

In [60] we studied a differential entropy estimator based on $kp-$neighbours, aiming at applying a Bayesian framework and some information-theoretic ideas to inverse problems. The goal of this work is to estimate the Shannon differential entropy in high dimensional settings, in possible presence of functional or nearly functional dependences. A modification of the Kozachenko-Leonenko estimator is proposed, consisting of introducing a local gaussian approximation to the probability measure. Test-cases were performed to assess the properties of the method and to compare its performances with other methods proposed in the literature.

The articles [37], presented in the section about biological flows, and [24], presented in the section about electrophysiology, also present methods concerning the interaction data - simulation.

7.6. Miscellaneous

Participants: Laurent Boudin, Irène Vignon-Clementel.

In [34] we develop a quantitative single cell-based model for multi-cellular tumor spheroids of a specific lung cancer cell line, growing under various nutrient conditions: we confront the simulations performed with this model with data on the growth kinetics and spatial labeling patterns for cell proliferation, extracellular matrix, cell distribution and cell death. We stepwise arrive at a model that mimics the spheroid growth under two conditions, and can predict two other ones. The number of mechanisms the model contains is necessary and sufficient to explain the data.

In [19] we consider a kinetic model describing some mechanisms of opinion formation in the framework of referendums, where the individuals, who can interact between themselves and modify their opinion by means of spontaneous self-thinking, are moreover under the influence of mass media. We study, at the numerical level, both the transient and the asymptotic regimes. In particular, we point out that a plurality of media, with different orientations, is a key ingredient to allow pluralism and prevent consensus. The forecasts of the model are compared to some surveys related to the Scottish independence referendum of 2014.

In [56] we review various results on the compactness of the linearized Boltzmann collision operator and of its generalization to mixtures of non-reactive monatomic gases.

RITS Project-Team

7. New Results

7.1. 2D Laser Based Road Obstacle Classification for Road Safety Improvement

Participants: Pierre Merdrignac, Evangeline Pollard, Fawzi Nashashibi.

Vehicle and pedestrian collisions often result in fatality to the vulnerable road users (VRU), indicating a strong need to protect such persons. Laser sensors have been extensively used for moving obstacles detection and tracking. Laser impacts are produced by reflection on these obstacles which suggests an information is available to recognize multiple road obstacles classes (pedestrian, cyclists, vehicles,...). We introduce a new system to address this problem that is divided in three parts: definition of geometric features describing road obstacles, multi-class object classification from an adaboost trained classifier and Bayesian estimation of the obstacle class. This approach benefits from consecutive observations of a single obstacle to estimate its class more precisely. We tested our system on some laser sequences and showed that it can estimate the class of some road obstacles around the vehicle with an accuracy of 87.4%. The vehicle class is determined with more than 97% of success. However, the main source of confusion is for static obstacles (posts and trees) for which 15% are classified as pedestrians. More detail can be fund in [36], [16].

7.2. On line Mapping and Global Positioning technique based on evidential SLAM

Participants: Guillaume Trehard, Evangeline Pollard, Fawzi Nashashibi.

Locate a vehicle in an urban environment remains a challenge for the autonomous driving community. By fusing information from a LIDAR, a Global Navigation by Satellite System (GNSS) and the vehicle odometry, we introduced and developed an original solution based on evidential grids and a particle filter to map the static environment and simultaneously estimate the position in a global reference at a high rate and without any prior knowledge (see [39]).

7.3. PML-SLAM

Participants: Zayed Alsayed, Fawzi Nashashibi, Anne Verroust-Blondet.

Our goal is to improve localization systems performances in order to be able to navigate in large-scale urban environments. In this context, we first optimized CPU and memory consumption of a SLAM laser-based technique [52] by introducing a map manager system. This strategy allows a smooth navigation while saving and loading probabilities-grid submaps into/from a hard-disc in a transparent way (cf. [27]). This work was validated and extended in the context of ITS Bordeaux demonstrations (VEDECOM demonstrator), where GPS information was integrated into SLAM environment Maps.

7.4. Motion planning techniques

Participants: David Gonzalez Bautista, Fernando Garrido Carpio, Joshué Pérez Rastelli, Vicente Milanés Montero, Fawzi Nashashibi.

Intelligent vehicles have increased their capabilities for highly, and even fully, automated driving under controlled environments. Scene information is received using on-board sensors and communication network systems–i.e. infrastructure and other vehicles. Considering the available information, different motion planning techniques have been implemented to autonomously driving on complex environments. The main goal is focused on executing strategies to improve safety, comfort and energy optimization. However, research challenges such as navigation in urban dynamic environments with obstacle avoidance capabilities–i.e. Vulnerable Road Users (VRU) and vehicles–and cooperative maneuvers among automated and semi-automated vehicles still need further efforts for a real environment implementation. We have recently carried out a deep state-of-the-art review to find the gaps in this hot topic into the autonomous vehicle field, paying special attention to overtaking and obstacle avoidance maneuvers.

Based on this review, we have mainly identified two main gaps: trajectory and speed planning with dynamics obstacle avoidance capabilities and real-time performance of the algorithms in the sense of significantly reducing the computational time, moving the system closer to what a vehicle should be able to provide in the real world.

According to this review, a speed planner has been designed with specific considerations on computing time efficiency, with an optimal comfort and avoiding to exceed speed and acceleration limits [31]. The comfort is evaluated as the minimization and smoothness of acceleration and jerk profiles, while maintaining a coherent speed profile with respect to traffic rules, the geometry of the path and the lateral accelerations associated to it. Specifically, this speed planner uses fifth order polynomial curves. These curves are C2 continuous and smooth, meaning that the jerk profile is also continuous and smooth. The method proposed computes the velocity in terms of the length of the path, instead of time, greatly reducing the errors. Specific targets for the speed planner are:

- Compute a smooth and continuous speed profile accounting for acceleration limits (longitudinal and lateral) according to ISO 2631-1 standard.
- Minimize distance error problems by associating the speed profile in the path speed planner instead of the time.

This speed planner was tested against other techniques providing better results in terms of computational time and smoothness (cf. [32]).

Additionally, a novel trajectory planning with a significant reduction on the computational time with respect to prior implementations from the team has been implemented. Our approach is mainly affected by vehicle's kinematics and physical road constraints. Based on these assumptions, computational time for path planning can be significantly reduced by creating a database containing already optimized versions of all the potential trajectories in each curve the vehicle can carry out. Therefore, this algorithm generates a database of smooth and continuous curves considering a big set of different intersection scenarios, taking into account the constraints of the infrastructure and the physical limitations of the vehicle. According to the real scenario, the local planner selects from the database the appropriate curves, searching for the ones that fit with the intersections defined on it. The path planning algorithm has been tested in simulation against the previous control architecture. The results obtained show path generation improvements in terms of smoothness and to continuity. Next steps on this algorithm is to test its performance in real platform and add the dynamics obstacle avoidance capabilities, establishing the link with the perception algorithms research line currently open in the team.

7.5. Control techniques

Participants: Francisco Navas Matos, Carlos Eduardo Flores Pino, David Gonzalez Bautista, Joshué Pérez Rastelli, Vicente Milanés Montero.

The final stage for automating a vehicle relies on the control algorithms. They are in charge of providing the proper behavior and performance to the vehicle, leading to provide the fully automated capabilities. Having this in mind, there are two research lines currently open in the time: the first one is mainly related to what we call "naturalistic driving" in the sense of adding the human reasoning to the vehicle. We are mainly focusing

our effort on artificial intelligent algorithms as neuro-fuzzy techniques. The main reason is the growing interest of the car makers in adding sharing control capabilities (between the vehicle and the driver) to the automated car. Our initial results show a big potential of using this approach and we already achieved some simulations results that were well-accepted by the scientific community and will be shown in mid-December at the final event of the EU project DESERVE.

On the other hand, we are also further investigating robust control algorithms for providing stability not only to an automated vehicle but also to a chain of automated vehicles that should be able to cooperate intelligently. This work is mainly divided in two main research lines:
1) Controllability and stability of dynamic complex systems are the key aspects when it comes to design intelligent control algorithms for vehicles. Current advances in the field are mainly oriented to advanced multi-sensor fusion toward multi-target decision-making systems. These artificial intelligence-based algorithms are able to provide reasonable responses under controlled environments (i.e. highly-detailed maps). However, new trends are proposing intelligent algorithms able to handle any unexpected circumstances as unpredicted uncertainties or even fully outages from sensors. The goal of this new research line at RITS is to further investigate control algorithms able to provide stability responses for autonomous vehicles under uncontrolled circumstances, including modifications on the input/output sensors. Dynamic plant models where different inputs/outputs can be added or subtracted in real-time during its operation is one of the hot topics in the control research arena. This system has to provide stable enough response when these operations occur. This is especially true on high-risk environments as autonomous driving; and
2) Data-driven control techniques based on model-free algorithms. Vehicles exhibit a highly non-linear behavior, especially at low speeds (as occur in urban environments). The research on novel data-driven techniques that are independent of the plant model provides huge benefits when applying them to automated vehicles. This novel research line in the team tries to further investigate on stable algorithm that doesn't need an accurate model of the vehicle dynamic, leading to compensate the effects of nonlinear dynamics, disturbances, or uncertainties in the parameters. [35]

7.6. Study on Perception and Communication Sytems for Safety

Participants: Pierre Merdrignac, Oyunchimeg Shagdar, Ines Ben Jemaa, Fawzi Nashashibi.

The existing R&D efforts for protecting vulnerable road users (VRU) are mainly based on perception techniques, which aim to detect VRUs utilizing vehicle embedded sensors. The efficiency of such a technique is largely affected by the sensor's visibility condition. Vehicle-to-Pedestrian (V2P) communication can also contribute to the VRU safety by allowing vehicles and pedestrians to exchange information. This solution is, however, largely affected by the reliability of the exchanged information, which most generally is the GPS data. Since perception and communication have complementary features, we can expect that a combination of such approaches can be a solution to the VRU safety. This is the motivation of this work. We develop theoretical models to present the characteristics of perception and communications systems. Experimental studies are conducted to compare the performances of these techniques in real-world environments. Our results show that the perception system reliably detects pedestrians and other objects within 50 m of range in the line-of-sight (LOS) condition. In contrast, the V2P communication coverage is approximately 340 and 200 meters in LOS and non-LOS (NLOS) conditions, respectively. However, the communication-based system fails to correctly position the VRU w.r.t the vehicle, preventing the system from meeting the safety requirement. Finally, we propose a cooperative system that combines the outputs of the communication and perception systems. More detail can be fund in [37], [16].

7.7. Asynchronous Reactive Distributed Congestion Control Algorithms for the ITS G5 Vehicular Communications

Participant: Oyunchimeg Shagdar.

The IEEE 802.11p is the technology dedicated to vehicular communications to support road safety, efficiency, and comfort applications. A large number of research activities have been carried out to study the characteristics of the IEEE 802.11p. The key weakness of the IEEE 802.11p is the channel congestion issue, where the wireless channel gets saturated when the road density increases. The European Telecommunications Standardization Institute (ETSI) is in the progress of studying the channel congestion problem and proposed so-called Reactive Distributed Congestion Control (DCC) algorithm as a solution to the congestion issue. In this work we investigate the impacts of the Reactive DCC mechanism in comparison to the conventional IEEE 802.11p with no congestion control. Our study shows that the Reactive DCC scheme creates oscillation on channel load that consequently degrades communication performance. The results reveal that the channel load oscillation is due to the fact that in the Reactive DCC, the individual CAM (Cooperative Awareness Message) controllers react to the channel congestion in a synchronized manner. To reduce the oscillation, we propose a simple extension to Reactive DCC, Asynchronous Reactive DCC, in which the individual CAM controllers adopt randomized rate setting, which can significantly reduce the oscillation and improve the network performance. See [45] for more detail.

7.8. Vehicle to vehicle visible light communication

Participants: Mohammad Abu Alhoul, Oyunchimeg Shagdar, Fawzi Nashashibi.

Visible Light Communication (VLC) technology utilizes the light spectral range between 380 nm and 750 nm, which enables the dual functionality of lightning and information delivery. A use of VLC for the ITS domain has many benefits including that it can be a complementary technology to the IEEE 802.11p, which is the radio communications technology dedicated to the V2X communication but suffers from its channel congestion problem.

This year, we conducted theoretical and experimental studies on the optical channel characteristics. Based on our studies and the previous contributions, we developed a transmitter and receiver VLC prototype to be integrated to the vehicle lightning systems dedicated to platooning applications. Using the low-cost Arduino micro-controller, a transmitter broadcasts the vehicle status information including the vehicle identity, velocity, orientation, acceleration through the vehicle rear Light Emitting Diodes (LED). The receiver is based on a simple Photo Diode (PD) with an accurate 635 nm optical filtering stage to overcome the saturation and the unwanted ambient noise issues. Experimental studies show that the system can provide 8.5 Kbps of information delivery between vehicles with up to 30 meters of bumper to bumper distance.

7.9. Analysis of broadcast strategies in IEEE 802.11p VANETs

Participants: Younes Bouchaala, Oyunchimeg Shagdar, Paul Muhlethaler.

We analyze different broadcast strategies in IEEE 802.11p Vehicular Ad-hoc NETworks (VANETs). The first strategy is the default IEEE 802.11p strategy. Using a model derived from the Bianchi model, we provide the network performance in terms of throughput and success rate. The second strategy consists in using an acknowledgment technique similar to the acknowledgment with point-to-point traffic. A node will send its broadcast packet as in the default case, but it requires an acknowledgment from a neighbor node. This node may be a random neighbor or may be selected according to precise rules. We analyze this second strategy in terms of throughput and success rate. Somewhat surprisingly, we show that this second strategy improves the delivery ratio of the transmitted packets but reduces the overall throughput. This means that if the CAM messages (Cooperative Awareness Messages) are broadcasted, the total number of packets actually delivered will be greater with the default strategy than with the improved strategy. We propose a third strategy which consists in using the default strategy for normal packets, but we add random redundant transmissions to ensure greater reliability for very important packets. We show that with this simple technique, not only do we obtain suitable reliability, but we also achieve larger global throughput than with the acknowledgment-oriented technique. This is described in [26]. Another contribution of this paper is to compute network performance in terms of throughput and success rate with respect to the network parameters and to analyze their impact on performances.

7.10. Multicast communications for cooperative vehicular systems

Participants: Ines Ben Jemaa, Oyunchimeg Shagdar, Paul Muhlethaler, Arnaud de La Fortelle.

With the advancement of wireless communications technologies, users can now have multicast services while they are driving. Majority of the multicast services require Internet-to-vehicle multicast message dissemination. Conventional group management approaches in Internet is relatively simple because it is performed on the local networks of the multicast members which are usually a priori configured to receive the service. In addition to this, multicast packets flows follow a fixed routing structure that is built between the source and the destinations. These approaches could not be applied to vehicular networks (VANET) due to their dynamic and distributed nature. In order to enable such multicasting, our work deals with two aspects. First, reachability of the moving vehicles to the multicast service and second, multicast message dissemination in the VANET. Regarding the first issue, we find that neither current multicast addressing nor existing mobility management mechanisms are suitable for VANET. We introduce first a self-configuring multicast addressing scheme that allows the vehicles to auto-configure a dynamic multicast address without a need to exchange signalling messages with the Internet. Second, we propose a simplified approach that extends Mobile IP and Proxy Mobile IP. About message dissemination, we first propose to revisit traditional multicast routing techniques that rely on a tree structure. In particular, as vehicular networks are known to have changing topology, we present a theoretical study of the link lifetime between vehicles in urban environments. Then, we propose then Motion-MAODV, an improved version of a tree-based routing mechanism (MAODV) that aims at guaranteeing longer route lifetime. Finally, we also propose a geographic routing protocol Melody that provides a geocast dissemination in urban environments. Through simulations, we show that Melody ensures more reliable and efficient packet delivery to a given geographic area compared to traditional geo-brodcasting schemes in highly dense scenarios. More detail can be found in [28], [41], [47].

7.11. Context Awareness and Priority Control for ITS based on Automatic Speech Recognition

Participants: Oyunchimeg Shagdar, Sakriani Watiasri Sakti.

Bringing rapid assistance to motorists involved in a traffic accident is an important service to be provided by Intelligent Transportation System (ITS). Existing proposals to automatic accident detection are based on the vehicle's perception point of view. In [38] we introduce situational awareness based on the "understanding" of conversational speech of drivers/passengers using an automatic speech recognition (ASR) system. Context-aware priority control and congestion control schemes are presented to ensure coexistence of ASR-triggered applications and cooperative awareness messages (CAM) in the IEEE 802.11p system. Finally, application risk analysis and performance evaluations of ASR and V2X communications are carried out.

7.12. Emergent Behaviors and Traffic Density among Heuristically-Driven Intelligent Vehicles using V2V Communication

Participants: Oyunchimeg Shagdar, Fawzi Nashashibi.

We study the global traffic density and emergent traffic behavior of several hundreds of intelligent vehicles, as a function of V2V communication (for the ego vehicle to perceive traffic) and path-finding heuristics (for the ego vehicle to reach its destination), in urban environments. Ideal/realistic/no V2V communication modes are crossed with straight-line/towards-most-crowded/towards-least-crowded pathfinding heuristics to measure the average trip speed of each vehicle. The behaviors of intelligent vehicles are modeled by a finite state automaton. The V2V communication model is also built based on signal propagation models in an intersection scenario and a Markov-chain based MAC model. Our experiments in simulation over up to 400 vehicles exhibit attractive insights: 1) communication's impact is positive for the performance of the emergent vehicles' behavior, however, 2) the path-finding heuristics may not obtain their expected collective behavior due to the communications errors in realistic road environment (cf. [43]).

7.13. Time-bounded message dissemination in strings

Participant: Gérard Le Lann.

In 2015, besides reviewing prominent open issues regarding safety in IVNs (see [42]), we have investigated coordination problems that arise in string formations. Since the inception of the platoon concept (1977), a number of solutions have been proposed for achieving string control (platoons are a particular case of ad hoc/open string). String control must be exercised in order to avoid rear-end collisions, string instability, and for coping with emergency situations. The cyber components essential for string control have not been fully identified yet. For example, considering the cooperative adaptive cruise control paradigm, data collected in recent platooning experiments show that it is inappropriate to rely on V2V broadcast from a lead vehicle, thus the quest for other approaches. In strings, one can take advantage of short-range directional antennas which enable fast messaging among consecutive string neighbors, leading to the concept of neighbor-to-neighbor (N2N) communications and the cohort construct (a cohort is a string with a specification). String control problems translate into communication protocol issues and distributed algorithmic problems, notably:

- Time-bounded string-wide acknowledged message delivery and dissemination (TBMD),
- Bounded channel access delay (BCAD), a MAC-level problem,
- Time-bounded message acknowledgment (TBMA).

Acceptable solutions shall achieve small non-stochastic worst-case channel access time bounds (BCAD) and bounded delays for successful message delivery (TBMA and TBMD), under worst-case conditions regarding channel contention and message/acknowledgment losses. Non-stochastic worst-case bounds can only be established analytically (obviously, simulations cannot be considered). The importance of the TBMD problem can be exposed simply as follows: would TBMD be solved, then the string instability problem vanishes. Rather than resting solely on stepwise detection-and-reaction strategies based on radars/lasers, every string member adjusts its acceleration/deceleration rate according to observed motions of its predecessor, TBMD delivers a N2N message carrying the newly string-wide targeted velocity, in less than 100 milliseconds in strings comprising in the order of 20 members, in the presence of message/acknowledgment losses. The TBMD problem has been solved (see [34]). The solution rests on assuming that TBMA and TBMD have solutions. Both problems have been solved (solutions are under review). Contrary to strings, groups are ad hoc/open multilane formations. It turns out that solutions aimed at the 3 problems referenced above are instrumental in solving problems arising with multilane SC scenarios. For example, the 3-way handshakes at the core of safe lane changes published previously now achieve significantly better performance figures. Work in progress also includes:

- conflicting concurrent lane changes at high velocities,

- fully automated zipper merging at high velocities, in non-line-of-sight conditions (radio communications), in line-of-sight conditions (optical communications).

7.14. Broadcast Transmission Networks with Buffering

Participants: Guy Fayolle, Paul Muhlethaler.

We analyzed the so-called back-off technique of the IEEE 802.11 protocol in broadcast mode with waiting queues. In contrast to existing models, packets arriving when a station (or node) is in back-off state are not discarded, but are stored in a buffer of infinite capacity. As in previous studies, the key point of our analysis hinges on the assumption that the time on the channel is viewed as a random succession of transmission slots (whose duration corresponds to the length of a packet) and mini-slots during which the back-off of the station is decremented. These events occur independently, with given probabilities. The state of a node is represented by a two-dimensional Markov chain in discrete-time, formed by the back-off counter and the number of packets at the station. Two models are proposed both of which are shown to cope reasonably well with the physical principles of the protocol. Stability (ergodicity) conditions are obtained and interpreted in terms of maximum throughput. Several approximations related to these models are also discussed in [44].

7.15. Belief propagation inference for traffic prediction

Participants: Cyril Furtlehner, Jean-Marc Lasgouttes.

This work [51] deals with real-time prediction of traffic conditions in a setting where the only available information is floating car data (FCD) sent by probe vehicles. The main focus is on finding a good way to encode some coarse information (typically whether traffic on a segment is fluid or congested), and to decode it in the form of real-time traffic reconstruction and prediction. Our approach relies in particular on the belief propagation algorithm.

These studies have been done in particular in the framework of the projects Travesti and Pumas.

This year, the work about the theoretical aspects of encoding real valued variables into a binary Ising model has been accepted for publication in Annals of Mathematics and Artificial Intelligence [23]. Moreover, an informal collaboration has been started with the company SISTeMA ITS, in order to assess the performance of our techniques in real-world city networks.

7.16. Random Walks in Orthants

Participant: Guy Fayolle.

7.16.1. Explicit criterion for the finiteness of the group in the quarter plane

In the book [3], original methods were proposed to determine the invariant measure of random walks in the quarter plane with small jumps, the general solution being obtained via reduction to boundary value problems. Among other things, an important quantity, the so-called *group of the walk*, allows to deduce theoretical features about the nature of the solutions. In particular, when the *order* of the group is finite, necessary and sufficient conditions have been given in [3] for the solution to be rational or algebraic. When the underlying algebraic curve is of genus 1, we propose, in collaboration with R. Iasnogorodski (St-Petersburg, Russia), a concrete criterion ensuring the finiteness of the group. It turns out that this criterion is always tantamount to the cancellation of a single constant, which can be expressed as the determinant of a matrix of order 3 or 4, and depends in a polynomial way on the coefficients of the walk [20].

7.16.2. Second Edition of the Book Random walks in the Quarter Plane

In collaboration with R. Iasnogorodski (St-Petersburg, Russia) and V. Malyshev, we prepared the second edition of the book [3], which will be published by Springer, in the collection *Probability Theory and Stochastic Processes*. Part II of this second edition borrows specific case-studies from queueing theory, and enumerative combinatorics. Five chapters will be added, including examples and applications of the general theory to enumerative combinatorics. Among them:

- Explicit criterions for the finiteness of the group, both in the genus 0 and genus 1 cases.
- Chapter *Coupled-Queues* shows the first example of a queueing system analyzed by reduction to a BVP in the complex plane.
- Chapter *Joining the shorter-queue* analyzes a famous model, where maximal homogeneity conditions do not hold, hence leading to a system of functional equations.
- Chapter *Counting Lattice Walks* concerns the so-called *enumerative combinatorics*. When counting random walks with small steps, the nature (rational, algebraic or holonomic) of the generating functions can be found and a precise classification is given for the basic (up to symmetries) 79 possible walks.

7.17. Global optimization for online resource allocation

Participant: Jean-Marc Lasgouttes.

As part of the Mobility 2.0 FP7 project, we have considered the possibility to allocate charging stations to Full Electric Vehicle (FEV) users in a way that, instead of merely minimizing their travel time, tries to improve the travel time for the whole community.

Our setting can be seen as a resource allocation problem, known as the Transportation Problem in Operations Research literature. It is solvable using several algorithms, among which the simplex algorithm or the Hungarian algorithm. Unfortunately, these algorithms are not well-adapted here for two reasons:

- The allocation of slots to users is done on-line, when the user does a request. It is not possible to wait until all the users are known before doing the allocation;
- The complexity of these algorithms is very high, especially since, due to the effect of range limitations, each request has different characteristics, which is equivalent to increasing the types of customers.

We therefore present a simple heuristic approach, which is fast enough for systems with thousands of stations. Its principle is to penalize the cost for the user with an approximation of the extra cost incurred to future users who compete for the same resource (a charging or parking slot).

This work has been presented at the ITSC'2015 conference [33].

SECRET Project-Team

6. New Results

6.1. Symmetric cryptology

Participants: Anne Canteaut, Pascale Charpin, Sébastien Duval, Virginie Lallemand, Gaëtan Leurent, Nicky Mouha, María Naya Plasencia, Joëlle Roué, Yann Rotella.

6.1.1. Block ciphers

Most of our work on block ciphers is related to an ANR Project named BLOC. Our recent results mainly concern either the analysis and design of lightweight block ciphers.

Recent results:

- Design and study of a new construction for low-latency block ciphers, named *reflection ciphers*, which generalizes the so-called α-reflection property exploited in PRINCE. This construction aims at reducing the implementation overhead of decryption on top of encryption [15], [60].
- Formalization and generic improvements of impossible differential cryptanalysis: our work provides a general framework for impossible differential cryptanalysis including a generic complexity analysis of the optimal attack [36].
- Cryptanalysis of several recently proposed block ciphers which offer an optimal resistance against side-channel attacks in the sense that the cost of Boolean masking is minimized. This includes an attack against Zorro and its variants [39], and an attack against Picaro in the related-key model [44].
- Cryptanalysis of Feistel constructions with secret Sboxes [42].
- Study of the security of the Even-Mansour construction in the multi-key setting [56].

6.1.2. Authenticated encryption

A limitation of all classical block ciphers is that they aim at protecting confidentiality only, while most applications need both encryption and authentication. These two functionalities are provided by using a block cipher like the AES together with an appropriate mode of operation. However, it appears that the most widely-used mode of operation for authenticated encryption, AES-GCM, is not very efficient for high-speed networks. Also, the security of the GCM mode completely collapses when an IV is reused. These severe drawbacks have then motivated an international competition named CAESAR, partly supported by the NIST, which has been recently launched in order to define some new authenticated encryption schemes [0]. Our work related to this competition is then two-fold: G. Leurent and N. Mouha have participated to the design of some CAESAR candidates; Also, the project-team is involved in a national cryptanalytic effort led by the BRUTUS project funded by the ANR.

Recent results:

- Design of new authenticated encryption schemes submitted to the CAESAR competition: SCREAM v3.0 [72] and PRIMATES 2[58]
- Cryptanalysis of the CAESAR candidates: collision attacks [49] against several candidates including AEZ and Marble, attack against LAC [53].

6.1.3. Stream ciphers

Stream ciphers provide an alternative to block-cipher-based encryption schemes. They are especially well-suited in applications which require either extremely fast encryption or a very low-cost hardware implementation.

[0] http://competitions.cr.yp.to/caesar.html

Recent results:

- Cryptanalysis of the recently proposed lightweight stream cipher Sprout [52], [71].
- New types of correlation attacks against filter generators exploiting the approximation of the filtering function composed with non-bijective monomial mappings [63], [87].
- Design of encryption schemes for efficient homomorphic-ciphertext compression: in order to avoid the (extremely) high expansion rate of homomorphic encryption, a solution consists in transmitting to the server the ciphertext c obtained by encrypting m with a symmetric scheme (the corresponding secret key encrypted by the homomorphic cipher is also transmitted). The server then needs to compute m encrypted with the homomorphic scheme from c, i.e. the server needs to homomorphically evaluate the decryption circuit of the symmetric cipher. A. Canteaut, M. Naya-Plasencia together with their coauthors have investigated the constraints on the symmetric cipher imposed by this application and they have proposed some solutions based on additive IV-based stream ciphers [78].

6.1.4. Hash functions and MACS

The international research effort related to the selection of the new hash function standard SHA-3 has led to many important results and to a better understanding of the security offered by hash functions. However, hash functions are used in a huge number of applications with different security requirements, and also form the building-blocks of some other primitives, like MACs. In this context, we have investigated the security of some of these constructions, in order to determine whether some particular constructions for hash functions may affect the security of the associated MACs.

Recent results:

- Improved generic attacks against hash-based MAC [30], [31]
- Cryptanalysis of 7 (out of 8) rounds of the Chaskey MAC [32]. This work has led the designers of Chaskey to increase the number of rounds [80].
- Attack against the XOR of two hash functions, using complex structures build from collisions [54]. This work by G. Leurent and L. Wang shows that, surprisingly, the construction $H_1(M) \oplus H_2(M)$ with common hash functions H_1 and H_2 (e.g. SHA-256 and BLAKE-256) is actually be less secure than each function on their own.

6.1.5. Security of Internet protocols

Hash functions are used to in key-exchange protocols such as TLS, IKE and SSH, to verify the integrity of the exchange. Most practitioners believe that the hash function only need to resist preimage attacks for this use. However, K. Bhargavan and G. Leurent have shown that collisions in the hash function are sufficient to break the integrity of these protocols, and to impersonate some of the parties [41]. Since many protocols still allow the use of MD5 or SHA-1 (for which collision attacks are known), this result in some practical attacks, and extends the real-world impact of the collision attacks against MD5 and SHA-1. This work has already influenced the latest TLS 1.3 draft, and the main TLS libraries are removing support of MD5 signatures

6.1.6. Cryptographic properties and construction of appropriate building blocks

The construction of building blocks which guarantee a high resistance against the known attacks is a major topic within our project-team, for stream ciphers, block ciphers and hash functions. The use of such optimal objects actually leads to some mathematical structures which may be at the origin of new attacks. This work involves fundamental aspects related to discrete mathematics, cryptanalysis and implementation aspects. Actually, characterizing the structures of the building blocks which are optimal regarding to some attacks is very important for finding appropriate constructions and also for determining whether the underlying structure induces some weaknesses or not. For these reasons, we have investigated several families of filtering functions and of S-boxes which are well-suited for their cryptographic properties or for their implementation characteristics.

Recent results:

- Definition of an extended criterion for estimating the resistance of a block cipher to differential attacks. This work emphasizes the role played by the affine permutation of the set of 8-bit words which follows the inverse function in the AES [45], [25], [26], [64], [24] (see Section 5.1.1).
- Construction of new Sboxes for lightweight ciphers: A. Canteaut, S. Duval and G. Leurent have investigated several constructions for obtaining good cryptographic Sboxes (especially 8-bit Sboxes) with a low implementation cost [43], [62], [84].
- P. Charpin, together with S. Mesnager and S. Sarkar, has provided a rigorous study of involutions over the finite field of order 2^n which are relevant primitives for cryptographic designs [47]. Most notably, they have focused on the class of involutions defined by Dickson polynomials [70], [79].

6.2. Code-based cryptography

Participants: Rodolfo Canto Torres, Julia Chaulet, Adrien Hauteville, Irene Márquez Corbella, Aurélie Phesso, Nicolas Sendrier, Jean-Pierre Tillich.

The first cryptosystem based on error-correcting codes was a public-key encryption scheme proposed by McEliece in 1978; a dual variant was proposed in 1986 by Niederreiter. We proposed the first (and only) digital signature scheme in 2001. Those systems enjoy very interesting features (fast encryption/decryption, short signature, good security reduction) but also have their drawbacks (large public key, encryption overhead, expensive signature generation). Some of the main issues in this field are

- security analysis, implementation and practicality of existing solutions,
- reducing the key size, *e.g.*, by using rank metric instead of Hamming metric, or by using particular families of codes,
- addressing new functionalities, like hashing or symmetric encryption.

Recent results:

- Structural attacks against some variants of the McEliece cryptosystem based on subclasses of alternant/Goppa codes which admit a very compact public matrix, typically quasi-cyclic, quasi-dyadic, or quasi-monoidic matrices [20]. This result is obtained thanks to a new operation on codes called folding that exploits the knowledge of the automorphism group of the code [19].
- Cryptanalysis of a variant of McEliece cryptosystem based on polar codes [40], [59].
- Cryptanalysis of a code-based encryption scheme proposed by Baldi *et al.* in the *Journal of Cryptology* [48].
- Cryptanalysis of a code-based signature scheme proposed at PQCrypto 2013 by Baldi at al. [57].
- Improved algorithm for decoding in the rank metric when some additional information about the targeted codeword is provided [51]; this algorithm used together with a folding technique leads to a feasible attack on the LRPC cryptosystem.
- Design on a new code-based stream cipher, named RankSynd, variant of Synd for the rank metric [50].
- In-depth analysis of the complexity of generic decoding algorithms for linear codes [37]. Most notably, R. Canto Torres and N. Sendrier have investigated the information-set decoding algorithms applied to the case where the number of errors is sub-linear in the code length [46]. This situation appears in the analysis of the McEliece based in quasi-cyclic Moderate Density Parity Check (MDPC) codes.

6.3. Quantum Information

Participants: Kaushik Chakraborty, André Chailloux, Anthony Leverrier, Jean-Pierre Tillich.

6.3.1. Quantum codes

Protecting quantum information from external noise is an issue of paramount importance for building a quantum computer. It also worthwhile to notice that all quantum error-correcting code schemes proposed up to now suffer from the very same problem that the first (classical) error-correcting codes had: there are constructions of good quantum codes, but for the best of them it is not known how to decode them in polynomial time.

Recent results:

- A. Leverrier and JP. Tillich, together with G. Zémor, proposed a new class of quantum LDPC codes, "Quantum expander codes", which feature a simple and very efficient decoding algorithm which can correct arbitrary patterns of errors of size scaling as the square-root of the length of the code. These are the first codes with constant rate for which such an efficient decoding algorithm is known (see Section 5.1.3) [55], [35], [73].
- Error analysis for Boson Sampling, a simplified model for quantum computation [21]

6.3.2. Quantum cryptography

A recent approach to cryptography takes into account that all interactions occur in a physical world described by the laws of quantum physics. These laws put severe constraints on what an adversary can achieve, and allow for instance to design provably secure key distribution protocols. We study such protocols as well as more general cryptographic primitives such as coin flipping with security properties based on quantum theory.

Recent results:

- A. Leverrier gave the first composable security proof for a continuous-variable quantum key distribution protocol with coherent states [22]. This essentially completes the security analysis of continuous-variable protocols with coherent states, which are by far the most practical protocols relying on continuous variables.
- A. Leverrier and E. Diamanti reviewed the state-of-the-art concerning quantum key distribution with continuous variables [18].
- A. Leverrier and M. Tomamichel gave the most complete security proof of the BB84 protocol to date, including all finite-size effects and a full description of the protocol [89].
- K. Chakraborty and A. Leverrier studied a general family of quantum protocols for position verification and present a new class of attacks based on the Clifford hierarchy that outperform previously known attacks [17].

6.3.3. Quantum correlations and nonlocality

Since the seminal work from Bell in the 60's, it has been known that classical correlations obtained via shared randomness cannot reproduce all the correlations obtained by measuring entangled quantum systems. This impossibility is for instance witnessed by the violation of a Bell inequality and is known under the name of "Quantum Nonlocality". In addition to its numerous applications for quantum cryptography, the study of quantum nonlocality and quantum games has become a central topic in quantum information theory, with the hope of bringing new insights to our understanding of quantum theory.

Recent results:

- Development of a general framework for the study of quantum correlations with combinatorial tools [14]

6.3.4. Relativistic cryptography

(see Section 5.1.2).

6.3.5. Quantum cryptanalysis of symmetric primitives

Symmetric cryptography seems at first sight much less affected in the post-quantum world than asymmetric cryptography: its main known threat is Grover's algorithm, which allows for an exhaustive key search in the square root of the normal complexity. For this reason, it is usually believed that doubling key lengths suffices to maintain an equivalent security in the post-quantum world. However, a lot of work is certainly required in the field of symmetric cryptography in order to "quantize" the classical families of attacks in an optimized way. G. Leurent, A. Leverrier and M. Naya Plasencia have recently started working in this area in collaboration with M. Kaplan, especially on differential cryptanalysis. Some preliminary results show that counter-intuitive and surprising cases appear: in general, it is not sufficient to consider the best classical attacks and try to "quantize" them if one wants to find the best post-quantum attack [34], [85].

6.4. Reverse-engineering of communication systems

Participants: Nicolas Sendrier, Jean-Pierre Tillich, Audrey Tixier.

Our activity within this domain, whose first aim is to establish the scientific and technical foundations of a discipline which does not exist yet at an academic level, has been supported by some industrial contracts driven by the Ministry of Defense.

Recent results:

- Efficient algorithm for recovering the block interleaver and the convolutional code when several noisy interleaver codewords are given [76], [13].

SERENA Team

7. New Results

7.1. Guaranteed bounds for Laplace eigenvalues and eigenvectors

In [22], we have derived a posteriori error estimates for the Laplace eigenvalue problem. Guaranteed, fully computable, and optimally convergent upper and lower bounds for the first eigenvalue are given. They are valid under explicit, a posteriori conditions on the computational mesh and on the approximate solution. Guaranteed, fully computable, and polynomial-degree robust bounds for the energy error in the approximation of the first eigenvector are derived as well, under the same conditions. Remarkably, all the constants in our theory can be fully estimated, and no convexity/regularity assumption on the computational domain/exact eigenvector(s) is needed. This general result can still be improved when an elliptic regularity assumption is satisfied (with known constants), typically for convex two-dimensional domains. The application of our framework to conforming finite element approximations of arbitrary polynomial degree is provided, along with a numerical illustration on a set of test problems.

SIERRA Project-Team

7. New Results

7.1. On the Global Linear Convergence of Frank-Wolfe Optimization Variants

Participant: Simon Lacoste-Julien [correspondent].

Collaboration with Martin Jaggi (ETH Zurich).

The Frank-Wolfe (FW) optimization algorithm has lately re-gained popularity thanks in particular to its ability to nicely handle the structured constraints appearing in machine learning applications. However, its convergence rate is known to be slow (sublinear) when the solution lies at the boundary. A simple less-known fix is to add the possibility to take 'away steps' during optimization, an operation that importantly does not require a feasibility oracle. In this paper [17], we highlight and clarify several variants of the Frank-Wolfe optimization algorithm that have been successfully applied in practice: away-steps FW, pairwise FW, fully-corrective FW and Wolfe's minimum norm point algorithm, and prove for the first time that they all enjoy global linear convergence, under a weaker condition than strong convexity of the objective. The constant in the convergence rate has an elegant interpretation as the product of the (classical) condition number of the function with a novel geometric quantity that plays the role of a 'condition number' of the constraint set. We provide pointers to where these algorithms have made a difference in practice, in particular with the flow polytope, the marginal polytope and the base polytope for submodular optimization.

7.2. Barrier Frank-Wolfe for Marginal Inference

Participant: Simon Lacoste-Julien [correspondent].

Collaboration with Rahul G. Krishnan [correspondent] and David Sontag (NYU).

In [16], we introduce a globally-convergent algorithm for optimizing the tree-reweighted (TRW) variational objective over the marginal polytope. The algorithm is based on the conditional gradient method (Frank-Wolfe) and moves pseudomarginals within the marginal polytope through repeated maximum a posteriori (MAP) calls. This modular structure enables us to leverage black-box MAP solvers (both exact and approximate) for variational inference, and obtains more accurate results than tree-reweighted algorithms that optimize over the local consistency relaxation. Theoretically, we bound the sub-optimality for the proposed algorithm despite the TRW objective having unbounded gradients at the boundary of the marginal polytope. Empirically, we demonstrate the increased quality of results found by tightening the relaxation over the marginal polytope as well as the spanning tree polytope on synthetic and real-world instances.

7.3. Sequential Kernel Herding: Frank-Wolfe Optimization for Particle Filtering

Participants: Simon Lacoste-Julien [correspondent], Francis Bach.

Collaboration with Fredrik Lindsten (University of Cambridge).

Recently, the Frank-Wolfe optimization algorithm was suggested as a procedure to obtain adaptive quadrature rules for integrals of functions in a reproducing kernel Hilbert space (RKHS) with a potentially faster rate of convergence than Monte Carlo integration (and "kernel herding" was shown to be a special case of this procedure). In this paper [18], we propose to replace the random sampling step in a particle filter by Frank-Wolfe optimization. By optimizing the position of the particles, we can obtain better accuracy than random or quasi-Monte Carlo sampling. In applications where the evaluation of the emission probabilities is expensive (such as in robot localization), the additional computational cost to generate the particles through optimization can be justified. Experiments on standard synthetic examples as well as on a robot localization task indicate indeed an improvement of accuracy over random and quasi-Monte Carlo sampling.

7.4. Variance Reduced Stochastic Gradient Descent with Neighbors

Participant: Simon Lacoste-Julien [correspondent].

Collaboration with Thomas Hofmann [correspondent], Aurelien Lucchi and Brian McWilliams (ETH Zurich).

Stochastic Gradient Descent (SGD) is a workhorse in machine learning, yet its slow convergence can be a computational bottleneck. Variance reduction techniques such as SAG, SVRG and SAGA have been proposed to overcome this weakness, achieving linear convergence. However, these methods are either based on computations of full gradients at pivot points, or on keeping per data point corrections in memory. Therefore speed-ups relative to SGD may need a minimal number of epochs in order to materialize. This paper [15] investigates algorithms that can exploit neighborhood structure in the training data to share and re-use information about past stochastic gradients across data points, which offers advantages in the transient optimization phase. As a side-product we provide a unified convergence analysis for a family of variance reduction algorithms, which we call memorization algorithms. We provide experimental results supporting our theory.

7.5. Rethinking LDA: Moment Matching for Discrete ICA

Participants: Anastasia Podosinnikova [correspondent], Francis Bach, Simon Lacoste-Julien.

In [21], we consider moment matching techniques for estimation in latent Dirichlet allocation (LDA). By drawing explicit links between LDA and discrete versions of independent component analysis (ICA), we first derive a new set of cumulant-based tensors, with an improved sample complexity. Moreover, we reuse standard ICA techniques such as joint diagonalization of tensors to improve over existing methods based on the tensor power method. In an extensive set of experiments on both synthetic and real datasets, we show that our new combination of tensors and orthogonal joint diagonalization techniques outperforms existing moment matching methods.

7.6. Tensorizing Neural Networks

Participant: Anton Osokin [correspondent].

Collaboration with Alexander Novikov, Dmitry Podoprikhin and Dmitry Vetrov.

Deep neural networks currently demonstrate state-of-the-art performance in several domains. At the same time, models of this class are very demanding in terms of computational resources. In particular, a large amount of memory is required by commonly used fully-connected layers, making it hard to use the models on low-end devices and stopping the further increase of the model size. In this paper [20], we convert the dense weight matrices of the fully-connected layers to the Tensor Train format such that the number of parameters is reduced by a huge factor and at the same time the expressive power of the layer is preserved. In particular, for the Very Deep VGG networks we report the compression factor of the dense weight matrix of a fully-connected layer up to 200000 times leading to the compression factor of the whole network up to 7 times.

7.7. Context-Aware CNNs for Person Head Detection

Participant: Anton Osokin [correspondent].

Collaboration with Tuan-Hung Vu [correspondent] and Ivan Laptev from the Willow project-team.

Person detection is a key problem for many computer vision tasks. While face detection has reached maturity, detecting people under a full variation of camera view-points, human poses, lighting conditions and occlusions is still a difficult challenge. In this work [23], we focus on detecting human heads in natural scenes. Starting from the recent local R-CNN object detector, we extend it with two types of contextual cues. First, we leverage person-scene relations and propose a Global CNN model trained to predict positions and scales of heads directly from the full image. Second, we explicitly model pairwise relations among objects and train a Pairwise CNN model using a structured-output surrogate loss. The Local, Global and Pairwise models are combined into a joint CNN framework. To train and test our full model, we introduce a large dataset composed of 369,846 human heads annotated in 224,740 movie frames. We evaluate our method and demonstrate improvements of person head detection against several recent baselines in three datasets. We also show improvements of the detection speed provided by our model.

7.8. Unsupervised Learning from Narrated Instruction Videos

Participants: Jean-Baptiste Alayrac [correspondent], Simon Lacoste-Julien.

Collaboration with Piotr Bojanowski, Josef Sivic and Ivan Laptev from the Willow project-team, and Nishant Agrawal.

In [29], we address the problem of automatically learning the main steps to complete a certain task, such as changing a car tire, from a set of narrated instruction videos. The contributions of this paper are three-fold. First, we develop a new unsupervised learning approach that takes advantage of the complementary nature of the input video and the associated narration. The method solves two clustering problems, one in text and one in video, applied one after each other and linked by joint constraints to obtain a single coherent sequence of steps in both modalities. Second, we collect and annotate a new challenging dataset of real-world instruction videos from the Internet. The dataset contains about 800,000 frames for five different tasks that include complex interactions between people and objects, and are captured in a variety of indoor and outdoor settings. Third, we experimentally demonstrate that the proposed method can automatically discover, in an unsupervised manner, the main steps to achieve the task and locate the steps in the input videos.

7.9. On Pairwise Cost for Multi-Object Network Flow Tracking

Participant: Simon Lacoste-Julien.

Collaboration with Visesh Chari, Ivan Laptev [correspondent] and Josef Sivic from the Willow project-team.

Multi-object tracking has been recently approached with the min-cost network flow optimization techniques. Such methods simultaneously resolve multiple object tracks in a video and enable modeling of dependencies among tracks. Min-cost network flow methods also fit well within the "tracking-by-detection" paradigm where object trajectories are obtained by connecting per-frame outputs of an object detector. Object detectors, however, often fail due to occlusions and clutter in the video. To cope with such situations, we propose in [13] an approach that regularizes the tracker by adding second order costs to the min-cost network flow framework. While solving such a problem with integer variables is NP-hard, we present a convex relaxation with an efficient rounding heuristic which empirically gives certificates of small suboptimality. Results are shown on real-world video sequences and demonstrate that the new constraints help selecting longer and more accurate tracks improving over the baseline tracking-by-detection method.

7.10. Multi-utility Learning: Structured-Output Learning with Multiple Annotation-Specific Loss Functions

Participant: Anton Osokin [correspondent].

Collaboration with Roman Shapovalov, Dmitry Vetrov and Pushmeet Kohli.

Structured-output learning is a challenging problem; particularly so because of the difficulty in obtaining large datasets of fully labelled instances for training. In this paper [22], we try to overcome this difficulty by presenting a multi-utility learning framework for structured prediction that can learn from training instances with different forms of supervision. We propose a unified technique for inferring the loss functions most suitable for quantifying the consistency of solutions with the given weak annotation. We demonstrate the effectiveness of our framework on the challenging semantic image segmentation problem for which a wide variety of annotations can be used. For instance, the popular training datasets for semantic segmentation are composed of images with hard-to-generate full pixel labellings, as well as images with easy-to-obtain weak annotations, such as bounding boxes around objects, or image-level labels that specify which object categories are present in an image. Experimental evaluation shows that the use of annotation-specific loss functions dramatically improves segmentation accuracy compared to the baseline system where only one type of weak annotation is used.

7.11. Convex Optimization for Parallel Energy Minimization

Participants: K. S. Sesh Kumar [correspondent], Francis Bach.

Collaboration with Alvaro Barbero, Stefanie Jegelka and Suvrit Sra.

Energy minimization has been an intensely studied core problem in computer vision. With growing image sizes (2D and 3D), it is now highly desirable to run energy minimization algorithms in parallel. But many existing algorithms, in particular, some efficient combinatorial algorithms, are difficult to parallelize. By exploiting results from convex and submodular theory, we reformulate in [47] the quadratic energy minimization problem as a total variation denoising problem, which, when viewed geometrically, enables the use of projection and reflection based convex methods. The resulting min-cut algorithm (and code) is conceptually very simple, and solves a sequence of TV denoising problems. We perform an extensive empirical evaluation comparing state-of-the-art combinatorial algorithms and convex optimization techniques. On small problems the iterative convex methods match the combinatorial max-flow algorithms, while on larger problems they offer other flexibility and important gains: (a) their memory footprint is small; (b) their straightforward parallelizability fits multi-core platforms; (c) they can easily be warm-started; and (d) they quickly reach approximately good solutions, thereby enabling faster "inexact" solutions. A key consequence of our approach based on submodularity and convexity is that it is allows to combine *any arbitrary combinatorial or convex methods as subroutines*, which allows one to obtain hybrid combinatorial and convex optimization algorithms that benefit from the strengths of both.

7.12. Active-set Methods for Submodular Optimization

Participants: K. S. Sesh Kumar [correspondent], Francis Bach.

In [46], we consider submodular optimization problems such as submodular function minimization (SFM) and quadratic problems regularized by the Lovász extension; for cut functions, this corresponds respectively to graph cuts and total variation (TV) denoising. Given a submodular function with an SFM oracle, we propose a new active-set algorithm for total variation denoising, which is more flexible than existing ones; the algorithm may be seen as a local descent algorithm over ordered partitions with explicit convergence guarantees. For functions that decompose into the sum of two functions F_1 and F_2 with efficient SFM oracles, we propose a new active-set algorithm for total variation denoising (and hence for SFM by thresholding the solution at zero). This algorithm also optimizes over ordered partitions and improves over existing ones based on TV or SFM oracles for F_1 and F_2.

7.13. Spectral Norm Regularization of Orthonormal Representations for Graph Transduction

Participant: Francis Bach [correspondent].

Collaboration with the Indian Institute of Science, Bangalore, India.

Recent literature suggests that embedding a graph on an unit sphere leads to better generalization for graph transduction. However, the choice of optimal embedding and an efficient algorithm to compute the same remains open. In this paper [25], we show that orthonormal representations, a class of unit-sphere graph embeddings are PAC learnable. Existing PAC-based analysis do not apply as the VC dimension of the function class is infinite. We propose an alternative PAC-based bound, which do not depend on the VC dimension of the underlying function class, but is related to the famous Lovasz function. The main contribution of the paper is SPORE, a SPectral regularized ORthonormal Embedding for graph transduction, derived from the PAC bound. SPORE is posed as a non-smooth convex function over an elliptope. These problems are usually solved as semi-definite programs (SDPs) with time complexity $O(n^6)$. We present, Infeasible Inexact proximal (IIP): an Inexact proximal method which performs subgradient procedure on an approximate projection, not necessarily feasible. IIP is more scalable than SDP, has an $O(1/\sqrt{T})$ convergence, and is generally applicable whenever a suitable approximate projection is available. We use IIP to compute SPORE where the approximate projection step is computed by FISTA, an accelerated gradient descent procedure. We show that the method has a convergence rate of $O(1/\sqrt{T})$. The proposed algorithm easily scales to 1000's of vertices, while the standard SDP computation does not scale beyond few hundred vertices. Furthermore, the analysis presented here easily extends to the multiple graph setting.

7.14. On the Equivalence between Quadrature Rules and Random Features

Participant: Francis Bach [correspondent].

In [31], we show that kernel-based quadrature rules for computing integrals can be seen as a special case of random feature expansions for positive definite kernels, for a particular decomposition that always exists for such kernels. We provide a theoretical analysis of the number of required samples for a given approximation error, leading to both upper and lower bounds that are based solely on the eigenvalues of the associated integral operator and match up to logarithmic terms. In particular, we show that the upper bound may be obtained from independent and identically distributed samples from a specific non-uniform distribution, while the lower bound if valid for any set of points. Applying our results to kernel-based quadrature, while our results are fairly general, we recover known upper and lower bounds for the special cases of Sobolev spaces. Moreover, our results extend to the more general problem of full function approximations (beyond simply computing an integral), with results in L_2- and L_infty-norm that match known results for special cases. Applying our results to random features, we show an improvement of the number of random features needed to preserve the generalization guarantees for learning with Lipschitz-continuous losses.

7.15. Preconditioning of a Generalized Forward-Backward Splitting and Application to Optimization on Graphs

Participant: Loïc Landrieu [correspondent].

Collaboration with Hugo Raguet.

In [41], we present a preconditioning of a generalized forward-backward splitting algorithm for finding a zero of a sum of maximally monotone operators $\sum_{i=1}^{n} A_i + B$ with B cocoercive, involving only the computation of B and of the resolvent of each A_i separately. This allows in particular to minimize functionals of the form $\sum_{i=1}^{n} g_i + f$ with f smooth, using only the gradient of f and the proximity operator of each g_i separately. By adapting the underlying metric, such preconditioning can serve two practical purposes: first, it might accelerate the convergence, or second, it might simplify the computation of the resolvent of A_i for some i. In addition, in many cases of interest, our preconditioning strategy allows the economy of storage and computation concerning some auxiliary variables. In particular, we show how this approach can handle large-scale, non-smooth, convex optimization problems structured on graphs, which arises in many image processing or learning applications, and that it compares favourably to alternatives in the literature.

7.16. A Riemannian Low-Rank Method for Optimization over Semidefinite Matrices with Block-Diagonal Constraints

Participant: Nicolas Boumal [correspondent].

In [34], we propose a new algorithm to solve optimization problems of the form $\min f(X)$ for a smooth function f under the constraints that X is positive semidefinite and the diagonal blocks of X are small identity matrices. Such problems often arise as the result of relaxing a rank constraint (lifting). In particular, many estimation tasks involving phases, rotations, orthonormal bases or permutations fit in this framework, and so do certain relaxations of combinatorial problems such as Max-Cut. The proposed algorithm exploits the facts that (1) such formulations admit low-rank solutions, and (2) their rank-restricted versions are smooth optimization problems on a Riemannian manifold. Combining insights from both the Riemannian and the convex geometries of the problem, we characterize when second-order critical points of the smooth problem reveal KKT points of the semidefinite problem. We compare against state of the art, mature software and find that, on certain interesting problem instances, what we call the staircase method is orders of magnitude faster, is more accurate and scales better. Code is available.

7.17. Tightness of the Maximum Likelihood Semidefinite Relaxation for Angular Synchronization

Participant: Nicolas Boumal [correspondent].

Collaboration with Afonso S. Bandeira and Amit Singer.

Many maximum likelihood estimation problems are, in general, intractable optimization problems. As a result, it is common to approximate the maximum likelihood estimator (MLE) using convex relaxations. Semidefinite relaxations are among the most popular. Sometimes, the relaxations turn out to be tight. In this paper [33], we study such a phenomenon. The angular synchronization problem consists in estimating a collection of n phases, given noisy measurements of some of the pairwise relative phases. The MLE for the angular synchronization problem is the solution of a (hard) non-bipartite Grothendieck problem over the complex numbers. It is known that its semidefinite relaxation enjoys worst-case approximation guarantees. In this paper, we consider a stochastic model on the input of that semidefinite relaxation. We assume there is a planted signal (corresponding to a ground truth set of phases) and the measurements are corrupted with random noise. Even though the MLE does not coincide with the planted signal, we show that the relaxation is, with high probability, tight. This holds even for high levels of noise. This analysis explains, for the interesting case of angular synchronization, a phenomenon which has been observed without explanation in many other settings. Namely, the fact that even when exact recovery of the ground truth is impossible, semidefinite relaxations for the MLE tend to be tight (in favorable noise regimes).

7.18. Coherent Diffractive Imaging Using Randomly Coded Masks

Participant: Alexandre d'Aspremont [correspondent].

Collaboration with Matthew H. Seaberg and Joshua J. Turner.

Coherent diffractive imaging (CDI) provides new opportunities for high resolution X-ray imaging with simultaneous amplitude and phase contrast. Extensions to CDI broaden the scope of the technique for use in a wide variety of experimental geometries and physical systems. Here [44], we experimentally demonstrate a new extension to CDI that encodes additional information through the use of a series of randomly coded masks. The information gained from the few additional diffraction measurements removes the need for typical object-domain constraints; the algorithm uses prior information about the masks instead. The experiment is performed using a laser diode at 532.2 nm, enabling rapid prototyping for future X-ray synchrotron and even free electron laser experiments. Diffraction patterns are collected with up to 15 different masks placed between a CCD detector and a single sample. Phase retrieval is performed using a convex relaxation routine known as "PhaseCut" followed by a variation on Fienup's input-output algorithm. The reconstruction quality is judged via calculation of phase retrieval transfer functions as well as by an object-space comparison between reconstructions and a lens-based image of the sample. The results of this analysis indicate that with enough masks (in this case 3 or 4) the diffraction phases converge reliably, implying stability and uniqueness of the retrieved solution.

7.19. Renegar's Condition Number and Compressed Sensing Performance

Participants: Vincent Roulet, Nicolas Boumal, Alexandre d'Aspremont [correspondent].

Renegar's condition number is a data-driven computational complexity measure for convex programs, generalizing classical condition numbers in linear systems. In [42], we provide evidence that for a broad class of compressed sensing problems, the worst case value of this algorithmic complexity measure taken over all signals matches the restricted eigenvalue of the observation matrix, which controls compressed sensing performance. This means that, in these problems, a single parameter directly controls computational complexity and recovery performance.

7.20. Supervised Clustering in the Data Cube

Participants: Vincent Roulet [correspondent], Fajwel Fogel, Alexandre d'Aspremont, Francis Bach.

In [43], we study a supervised clustering problem seeking to cluster either features, tasks or sample points using losses extracted from supervised learning problems. We formulate a unified optimization problem handling these three settings and derive algorithms whose core iteration complexity is concentrated in a k-means clustering step, which can be approximated efficiently. We test our methods on both artificial and realistic data sets extracted from movie reviews and 20NewsGroup.

7.21. Convex Relaxations for Permutation Problems

Participants: Fajwel Fogel [correspondent], Francis Bach, Alexandre d'Aspremont.

Collaboration with Rodolphe Jenatton.

Seriation seeks to reconstruct a linear order between variables using unsorted similarity information. It has direct applications in archeology and shotgun gene sequencing for example. In [4], we prove the equivalence between the seriation and the combinatorial 2-sum problem (a quadratic minimization problem over permutations) over a class of similarity matrices. The seriation problem can be solved exactly by a spectral algorithm in the noiseless case and we produce a convex relaxation for the 2-sum problem to improve the robustness of solutions in a noisy setting. This relaxation also allows us to impose additional structural constraints on the solution, to solve semi-supervised seriation problems. We present numerical experiments on archeological data, Markov chains and gene sequences.

7.22. Phase Recovery, MaxCut and Complex Semidefinite Programming

Participant: Alexandre d'Aspremont [correspondent].

Collaboration with Irène Waldspurger and Stéphane Mallat.

Phase retrieval seeks to recover a signal x from the amplitude $|Ax|$ of linear measurements. We cast the phase retrieval problem as a non-convex quadratic program over a complex phase vector and formulate a tractable relaxation (called PhaseCut) similar to the classical MaxCut semidefinite program. In [10], we solve this problem using a provably convergent block coordinate descent algorithm whose structure is similar to that of the original greedy algorithm in Gerchberg-Saxton, where each iteration is a matrix vector product. Numerical results show the performance of this approach over three different phase retrieval problems, in comparison with greedy phase retrieval algorithms and matrix completion formulations.

7.23. Choice of V for V-Fold Cross-Validation in Least-Squares

Participant: Sylvain Arlot [correspondent].

Collaboration with Matthieu Lerasle.

The paper [30] studies V-fold cross-validation for model selection in least-squares density estimation. The goal is to provide theoretical grounds for choosing V in order to minimize the least-squares loss of the selected estimator. We first prove a non-asymptotic oracle inequality for V-fold cross-validation and its bias-corrected version (V-fold penalization). In particular, this result implies that V-fold penalization is asymptotically optimal in the nonparametric case. Then, we compute the variance of V-fold cross-validation and related criteria, as well as the variance of key quantities for model selection performance. We show that these variances depend on V like $1 + 4/(V - 1)$, at least in some particular cases, suggesting that the performance increases much from $V = 2$ to $V = 5$ or 10, and then is almost constant. Overall, this can explain the common advice to take $V = 5$—at least in our setting and when the computational power is limited—, as supported by some simulation experiments. An oracle inequality and exact formulas for the variance are also proved for Monte-Carlo cross-validation, also known as repeated cross-validation, where the parameter V is replaced by the number B of random splits of the data.

7.24. Gains and Losses are Fundamentally Different in Regret Minimization: The Sparse Case

Participant: Vianney Perchet [correspondent].

Collaboration with Joon Kwon.

In [38], we demonstrate that, in the classical non-stochastic regret minimization problem with d decisions, gains and losses to be respectively maximized or minimized are fundamentally different. Indeed, by considering the additional sparsity assumption (at each stage, at most s decisions incur a nonzero outcome), we derive optimal regret bounds of different orders. Specifically, with gains, we obtain an optimal regret guarantee after T stages of order $\sqrt{T \log s}$, so the classical dependency in the dimension is replaced by the sparsity size. With losses, we provide matching upper and lower bounds of order $\sqrt{Ts \log(d)/d}$, which is decreasing in d. Eventually, we also study the bandit setting, and obtain an upper bound of order $\sqrt{Ts \log(d/s)}$ when outcomes are losses. This bound is proven to be optimal up to the logarithmic factor $\sqrt{\log(d/s)}$.

7.25. Batched Bandit Problems

Participant: Vianney Perchet [correspondent].

Collaboration with Philippe Rigollet, Sylvain Chassang and Erik Snowberg.

Motivated by practical applications, chiefly clinical trials, we study in [39] the regret achievable for stochastic bandits under the constraint that the employed policy must split trials into a small number of batches. Our results show that a very small number of batches gives close to minimax optimal regret bounds. As a byproduct, we derive optimal policies with low switching cost for stochastic bandits.

7.26. Online Learning in Repeated Auctions

Participant: Vianney Perchet [correspondent].

Collaboration with Jonathan Weed and Philippe Rigollet.

Motivated by online advertising auctions, in [40] we consider repeated Vickrey auctions where goods of unknown value are sold sequentially and bidders only learn (potentially noisy) information about a good's value once it is purchased. We adopt an online learning approach with bandit feedback to model this problem and derive bidding strategies for two models: stochastic and adversarial. In the stochastic model, the observed values of the goods are random variables centered around the true value of the good. In this case, logarithmic regret is achievable when competing against well behaved adversaries. In the adversarial model, the goods need not be identical and we simply compare our performance against that of the best fixed bid in hindsight. We show that sublinear regret is also achievable in this case and prove matching minimax lower bounds. To our knowledge, this is the first complete set of strategies for bidders participating in auctions of this type.

SMIS Project-Team

6. New Results

6.1. Embedded Data Management

Participants: Nicolas Anciaux, Saliha Lallali, Philippe Pucheral, Iulian Sandu Popa [correspondent].

Embedded keyword indexing: In this work, we revisit the traditional problem of information retrieval queries over large collections of files in an embedded context. A file can be any form of document, picture or data stream, associated with a set of terms. A query can be any form of keyword search using a ranking function (e.g., TF-IDF) identifying the top-k most relevant files. The proposed search engine can be used in sensors to search for relevant objects in their surroundings, in cameras to search pictures by using tags, in personal smart dongles to secure the querying of documents and files hosted in an untrusted Cloud, or in a personal cloud securely managed using a tamper resistant smart object. A search engine is usually based on a (large) inverted index and queries are traditionally evaluated by allocating one container in RAM per document to aggregate its score, making the RAM consumption linear with the size of the document corpus. To tackle this issue, we designed a new form of inverted index which can be accessed in a pure pipeline manner to evaluate search queries without materializing any intermediate result. Successive index partitions are written once in Flash and maintained in the background by timely triggering merge operations while files are inserted or deleted from the index. By combining this new index and the corresponding evaluation techniques, our embedded search engine is capable of reconciling high insert/delete/update rate and query scalability. We have demonstrated the search engine on a secure USB token in the context of a personal cloud, and have conducted in depth performance evaluations on a development board representative for different smart objects characteristics. The experimental results demonstrate the scalability of the approach and its superiority compared to state of the art methods. This work was published at VLDB'15 [21] and demonstrated at SIGMOD'15 [24]. It constitutes the main contribution of the PhD thesis of Saliha Lallali

Spatio-temporal indexing in Flash storage: The convergence of mobile computing, wireless communications and sensors has raised the development of many applications exploiting massive flows of spatio-temporal data such as in location-based services, participatory sensing, or traffic management [15]. Spatio-temporal data indexing is among the most active research topics in this area. Nevertheless, since a few years a new fundamental parameter has made its entry on the database scene: the NAND flash storage. The peculiar characteristics of flash memory require redesigning the existing data storage and indexing techniques that were devised for magnetic hard-disks. TRIFL, proposed in [16] is an efficient and generic TRajectory Index for FLash, designed around the key requirements of both trajectory indexing and flash storage. TRIFL is generic in the sense that it is efficient for both simple flash storage devices such as the SD cards and more powerful devices such as the solid state drives. In addition, TRIFL includes an online self tuning algorithm that allows adapting the index structure to the workload and the technical specifications of the flash storage device to maximize the index performance. Moreover, TRIFL achieves good performance with relatively low memory requirements, making it appropriate for many application scenarios. The experimental evaluation shows that TRIFL outperforms the representative indexing methods on flash disks but also on magnetic disks. This work [15] [16] is part of Dai Hai Ton That's Ph.D. thesis, co-supervised by Iulian Sandu Popa.

6.2. Secure Global Computing on Asymmetric Architecture

Participants: Benjamin Nguyen [correspondent], Philippe Pucheral, Quoc Cuong To.

Asymmetric Architecture Computing: This research direction studies the secure execution of various algorithms on data stored in an unstructured network of Trusted Cells (i.e., personal trusted device) so that each user can keep control over her data. The data could be stored locally in a trusted cell or encrypted on some external cloud. Execution takes place on a specific infrastructure called the Asymmetric Architecture: the network of trusted cells, supported by an untrusted cloud supporting IaaS or PaaS. Our objective is to show that many different algorithms and computing paradigms can be executed on the Asymmetric Architecture, thus achieving secure and private computation. Our first contribution in this area was to study the execution of Privacy Preserving Data Publishing (PPDP) algorithms on such an architecture, and provided generic protocols to deal with all kinds of PPDP algorithms, which are robust against honest-but-curious and malicious adversaries [2][3]. Our second contribution was to study general SQL queries in this same execution context. We concentrated on the subset of SQL queries without joins, but including Group By and aggregates, and show how to secure their execution in the presence of honest-but-curious attackers [9]. This work was part of Quoc-Cuong To's Ph.D defended in 2015 [13]. We are extending this general framework through a collaboration with INSA Centre Val de Loire, LIFO Lab and University of Paris Nord, LIPN lab, to study the secure execution of Map/Reduce on the Asymmetric Architecture. Computing MapReduce processes on the Asymmetric Architecture means maintaining the flexibility and efficiency of MapReduce, while adding security into the mix. We have shown in [25] that it is possible to achieve seamless integration of distributed MapReduce processing using trusted cells, while maintaining reasonable performance.

Secure spatio-temporal distributed processing: Mobile participatory sensing could be used in many applications such as vehicular traffic monitoring, pollution tracking, or even health surveying (e.g., to allow measuring in real-time the individual exposure to environmental risk factors or the propagation of an epidemic). However, its success depends on finding a solution for querying a large number of users which protects user location privacy and works in real-time. We addressed these issues and proposed PAMPAS, a privacy-aware mobile distributed system for efficient data aggregation in mobile participatory sensing. In PAMPAS, mobile devices enhanced with secure hardware, called secure probes, perform distributed query processing, while preventing users from accessing other users' data. Secure probes exchange data in encrypted form with help from an untrusted supporting server infrastructure. PAMPAS uses two efficient, parallel, and privacy-aware protocols for location-based aggregation and adaptive spatial partitioning of secure probes. Our experimental results and security analysis demonstrate that these protocols are able to collect, aggregate and share statistics or derived data in real-time, without any privacy leakage. This work is part of Dai Hai Ton That's Ph.D. thesis, co-supervised by Iulian Sandu Popa. The system implementation was demonstrated in [26], and a paper describes the technical details of the system [31].

6.3. Personal Cloud

Participants: Nicolas Anciaux [correspondent], Luc Bouganim, Athanasia Katsouraki, Benjamin Nguyen, Philippe Pucheral, Iulian Sandu Popa, Paul Tran Van.

We are witnessing an exponential increase in the acquisition of personal data about the individuals or produced by them. Today, this information is managed using Web applications, centralizing this data in cloud data servers, under the control of few Web majors [4]. However, it has now become clear that (1) centralizing millions of personal records exposes the data to very sophisticated attacks, linked to a very high potential benefit in case of success (millions of records being revealed), and (2) delegating the management of personal records without any tangible guarantee for the individuals leads to privacy violations, the data being potentially made accessible to other organizations (e.g., governments, commercial partners) and being subject to lucrative secondary usages (not advertised to the individuals). To face this situation, many recent initiatives push towards the emergence of the Personal Cloud paradigm. A personal cloud can be viewed as a personal server, owned by a given individual, which gives to its owner the ability to store her complete digital environment, synchronize it among various devices and share it with other individuals and applications under control. In the SMIS team, we claim the need of a Secure Personal Cloud, and promote the introduction of a secure (tamper resistant) data engine in the architecture [1]. On this basis, we investigate new data sharing and dissemination models, where usage and access control rules endorsed by the individuals could be enforced and have presented this

vision at EDBT'14 and at ADBIS'15 [18]. We have started a cooperation with the startup CozyCloud at the end of 2014. A contract was signed at the end of 2014 to integrate PlugDB in a CozyCloud instance and the PhD of Paul Tran Van (CIFRE SMIS-CozyCloud) has started to explore new data sharing techniques which could be enforced in the secure personal cloud model. A second PhD CIFRE SMIS-CozyCloud is being submitted to explore privacy-preserving distributed computations over personal clouds. Athanasia Katsouraki is working on privacy issues and on adoption of the secure data engine [29] in cooperation with the economists (CERDI) in the context of the Digital Society Institute (DSI). A paper written by jurists, economists and computer scientists from DSI has been invited for publication in Legicom'2016 to present our common vision of Privacy-by-Design principles in the context of Open Data and Internet of Things.

6.4. Applications

Participants: Nicolas Anciaux [correspondent], Luc Bouganim, Philippe Pucheral.

In 2014, we proposed a new paradigm, that we call Folk-enabled Information System (Folk-IS), based on a fully decentralized and participatory approach, where each individual implements a small subset of a complete information system without the need for a shared networked infrastructure [5]. Folk-IS builds upon the emergence of highly secure, portable and low-cost storage and computing devices, called hereafter Smart Tokens. Here however, the focus is on low-cost of ownership, deployment and maintenance, and on the absence of a networked infrastructure. With Folk-IS and thanks to their smart tokens, people will transparently and opportunistically perform data management and networking tasks as they physically move, so that IT services are truly delivered by the crowd. Following this work, we collaborate with researchers and doctors from Cameroon to study the specific case of diabete follow-up. Indeed, there are currently more than half a million diabetes cases in Cameroon and the deaths caused by diabetes complications will double before 2030. Diabetes complications mostly occur due to a bad follow-up of patients. Based on an analysis of the current situation, we proposed a new IT architecture for diabetes follow-up and introduce the bases of a new distributed computation protocol for this architecture. Our approach does not require any preexisting support communication infrastructure, can be deployed at low cost, and provides strong privacy and security guarantees. This work, published in AFRICOM [20] envisions an experiment in the field we plan to conduct under the authority of the Cameroonian National Center for Diabetes and Hypertension, with a potential for generalization to other diseases.

WHISPER Project-Team

7. New Results

7.1. Software engineering for infrastructure software

Tracking code fragments of interest is important in monitoring a software project over multiple versions. Various approaches, including our previous work on Herodotos, exploit the notion of Longest Common Subsequence, as computed by readily available tools such as GNU Diff, to map corresponding code fragments. Nevertheless, the efficient code differencing algorithms are typically line-based or word-based, and thus do not report changes at the level of language constructs. Furthermore, they identify only additions and removals, but not the moving of a block of code from one part of a file to another. Code fragments of interest that fall within the added and removed regions of code have to be manually correlated across versions, which is tedious and error-prone. When studying a very large code base over a long time, the number of manual correlations can become an obstacle to the success of a study. In a paper published at the IEEE International Conference on Software Analysis, Evolution, and Reengineering (SANER) [14], we investigate the effect of replacing the current line-based algorithm used by Herodotos by tree-matching, as provided by the algorithm of the differencing tool GumTree. In contrast to the line-based approach, the tree-based approach does not generate any manual correlations, but it incurs a high execution time. To address the problem, we propose a hybrid strategy that gives the best of both approaches.

Understanding the severity of reported bugs is important in both research and practice. In particular, a number of recently proposed mining-based software engineering techniques predict bug severity, bug report quality, and bug-fix time, according to this information. Many bug tracking systems provide a field "severity" offering options such as "severe", "normal", and "minor", with "normal" as the default. However, there is a widespread perception that for many bug reports the label "normal" may not reflect the actual severity, because reporters may overlook setting the severity or may not feel confident enough to do so. In many cases, researchers ignore "normal" bug reports, and thus overlook a large percentage of the reports provided. On the other hand, treating them all together risks mixing reports that have very diverse properties. In a study published at the Working Conference on Mining Software Repositories (MSR) 2015 [16], we investigate the extent to which "normal" bug reports actually have the "normal" severity. We find that many "normal" bug reports in practice are not normal. Furthermore, this misclassification can have a significant impact on the accuracy of mining-based tools and studies that rely on bug report severity information.

Software is continually evolving, to fix bugs and add new features. Industry users, however, often value stability, and thus may not be able to update their code base to the latest versions. This raises the need to selectively backport new features to older software versions. Traditionally, backporting has been done by cluttering the backported code with preprocessor directives, to replace behaviors that are unsupported in an earlier version by appropriate workarounds. This approach however involves writing a lot of error-prone backporting code, and results in implementations that are hard to read and maintain. In a paper published at the 2015 European Dependable Computing Conference (EDCC) [15], we consider this issue in the context of the Linux kernel, for which older versions are in wide use. We present a new backporting strategy that relies on the use of a backporting compatability library and on code that is automatically generated using the program transformation tool Coccinelle. This approach reduces the amount of code that must be manually written, and thus can help the Linux kernel backporting effort scale while maintaining the dependability of the backporting process.

Logging is a common and important programming practice, but choosing how to log is challenging, especially in a large, evolving software code base that provides many logging alternatives. Insufficient logging may complicate debugging, while logging incorrectly may result in excessive performance overhead and an overload of trivial logs. The Linux kernel has over 13 million lines of code, over 1100 different logging functions, and the strategies for when and how to log have evolved over time. To help developers log correctly

we propose, in a paper published at BEneVol 2015 [18], a framework that will learn existing logging practices from the software development history, and that will be capable of identifying new logging strategies, even when the new strategies just start to be adopted.

7.2. Java runtime support

Java class loaders are commonly used in application servers to load, unload and update a set of classes as a unit. However, unloading or updating a class loader can introduce stale references to the objects of the outdated class loader. A stale reference leads to a memory leak and, for an update, to an inconsistency between the outdated classes and their replacements. To detect and eliminate stale references, in a paper published at DSN 2015 [12], we propose Incinerator, a Java virtual machine extension that introduces the notion of an outdated class loader. Incinerator detects stale references and sets them to null during a garbage collection cycle. We evaluate Incinerator in the context of the OSGi framework and show that Incinerator correctly detects and eliminates stale references, including a bug in Knopflerfish. We also evaluate the performance of Incinerator with the DaCapo benchmark on VMKit and show that Incinerator has an overhead of at most 3.3%.

7.3. Parallel and Distributed Computing

The scalability of multithreaded applications on current multicore systems is hampered by the performance of lock algorithms, due to the costs of access contention and cache misses. In an article published in ACM Transactions on Computer Systems (TOCS), we present a new locking technique, Remote Core Locking (RCL) [10], that aims to accelerate the execution of critical sections in legacy applications on multicore architectures. The idea of RCL is to replace lock acquisitions by optimized remote procedure calls to a dedicated server hardware thread. RCL limits the performance collapse observed with other lock algorithms when many threads try to acquire a lock concurrently and removes the need to transfer lock-protected shared data to the hardware thread acquiring the lock because such data can typically remain in the server's cache. Other contributions presented in this article include a profiler that identifies the locks that are the bottlenecks in multithreaded applications and that can thus benefit from RCL, and a reengineering tool that transforms POSIX lock acquisitions into RCL locks. Eighteen applications were used to evaluate RCL: the nine applications of the SPLASH-2 benchmark suite, the seven applications of the Phoenix 2 benchmark suite, Memcached, and Berkeley DB with a TPC-C client. Eight of these applications are unable to scale because of locks and benefit from RCL on an x86 machine with four AMD Opteron processors and 48 hardware threads. By using RCL instead of Linux POSIX locks, performance is improved by up to 2.5 times on Memcached, and up to 11.6 times on Berkeley DB with the TPC-C client. On a SPARC machine with two Sun Ultrasparc T2+ processors and 128 hardware threads, three applications benefit from RCL. In particular, performance is improved by up to 1.3 times with respect to Solaris POSIX locks on Memcached, and up to 7.9 times on Berkeley DB with the TPC-C client.

Software Transactional Memory (STM) is an optimistic concurrency control mechanism that simplifies parallel programming. Still, there has been little interest in its applicability for reactive applications in which there is a required response time for certain operations. In an article published in ACM Transactions on Parallel Computing (TOPC) [11], we propose supporting such applications by allowing programmers to associate time with atomic blocks in the forms of deadlines and QoS requirements. Based on statistics of past executions, we adjust the execution mode of transactions by decreasing the level of optimism as the deadline approaches. In the presence of concurrent deadlines, we propose different conflict resolution policies. Execution mode switching mechanisms allow meeting multiple deadlines in a consistent manner, with potential QoS degradations being split fairly among several threads as contention increases, and avoiding starvation. Our implementation consists of extensions to a STM runtime that allow gathering statistics and switching execution modes. We also propose novel contention managers adapted to transactional workloads subject to deadlines. The experimental evaluation shows that our approaches significantly improve the likelihood of a transaction meeting its deadline and QoS requirement, even in cases where progress is hampered by conflicts and other concurrent transactions with deadlines.

A challenge in designing a peer-to-peer (P2P) system is to ensure that the system is able to tolerate selfish nodes that strategically deviate from their specification whenever doing so is convenient. In a paper published at SRDS 2015 [13], we propose *RACOON*, a framework for the design of P2P systems that are resilient to selfish behaviours. While most existing solutions target specific systems or types of selfishness, *RACOON* proposes a generic and semi-automatic approach that achieves robust and reusable results. Also, *RACOON* supports the system designer in the performance-oriented tuning of the system, by proposing a novel approach that combines Game Theory and simulations. We illustrate the benefits of using *RACOON* by designing two P2P systems: a live streaming and an anonymous communication system. In simulations and a real deployment of the two applications on a testbed comprising 100 nodes, the systems designed using *RACOON* achieve both resilience to selfish nodes and high performance.

7.4. From Sets to Bits in Coq

Sets form the building block of mathematics, while finite sets are a fundamental data structure of computer science. In the world of mathematics, finite sets enjoy appealing mathematical properties, such as a proof-irrelevant equality and extensionality of functions. Computer scientists, on the other hand, have devised efficient algorithms for set operations based on the representation of finite sets as bit vectors and on bit twiddling, exploiting the hardware's ability to efficiently process machine words.

With interactive theorem provers, sets are reinstituted as mathematical objects. While there are several finite set libraries in Coq, these implementations are far removed from those used in efficient code. Recent work on modeling low-level architectures, such as x86 [41] processors, however, have brought the world of bit twiddling within reach of our proof assistants. We are now able to specify and reason about low-level programs.

In this work, we have implemented bitsets and their associated operations in the Coq proof assistant, thus allowing us to transparently navigate between the concrete world of bit vectors and the abstract world of finite sets. This work grew from a puzzled look at the first page of Warren's *Hacker's Delight* [77], where lies the cryptic formula $x\&(x-1)$ to turn off the rightmost bit in a word. How do we translate the English specification given in the book into a formal definition? How do we prove that this formula meets its specification? Could Coq generate efficient and trustworthy code from it? And how efficiently could we simulate it within Coq itself?

In our work, we have established a bijection between bitsets and sets over finite types. Following a refinement approach, we have shown that a significant part of SSReflect*finset* library can be refined to operations manipulating bitsets. We have also developed a trustworthy extraction of bitsets down to OCaml's machine integers. While we were bound to axiomatize machine integers, we adopted a methodology based on exhaustive testing to gain greater confidence in our model. Finally, we have demonstrateed the usefulness of our library through two applications, a certified implementation of Bloom filters and a verified implementation of the n-queens algorithm.

WILLOW Project-Team

7. New Results

7.1. 3D object and scene modeling, analysis, and retrieval

7.1.1. The joint image handbook

Participants: Matthew Trager, Martial Hebert, Jean Ponce.

Given multiple perspective photographs, point correspondences form the "joint image", effectively a replica of three-dimensional space distributed across its two-dimensional projections. This set can be characterized by multilinear equations over image coordinates, such as epipolar and trifocal constraints. In this work, we revisit the geometric and algebraic properties of the joint image, and address fundamental questions such as how many and which multilinearities are necessary and/or sufficient to determine camera geometry and/or image correspondences. Our new theoretical results answer these questions in a very general setting, and our work, published ICCV 2015 [17], is intended to serve as a "handbook" reference about multilinearities for practitioners.

7.1.2. Trinocular Geometry Revisited

Participants: Jean Ponce, Martial Hebert, Matthew Trager.

When do the visual rays associated with triplets of point correspondences converge, that is, intersect in a common point? Classical models of trinocular geometry based on the fundamental matrices and trifocal tensor associated with the corresponding cameras only provide partial answers to this fundamental question, in large part because of underlying, but seldom explicit, general configuration assumptions. In this project, we use elementary tools from projective line geometry to provide necessary and sufficient geometric and analytical conditions for convergence in terms of transversals to triplets of visual rays, without any such assumptions. In turn, this yields a novel and simple minimal parameterization of trinocular geometry for cameras with non-collinear or collinear pinholes, which can be used to construct a practical and efficient method for trinocular geometry parameter estimation. This work has been published at CVPR 2014, and a revised version that includes numerical experiments using synthetic and real data has been submitted to IJCV [25].

7.1.3. 24/7 place recognition by view synthesis

Participants: Akihiko Torii, Relja Arandjelović, Josef Sivic, Masatoshi Okutomi, Tomas Pajdla.

We address the problem of large-scale visual place recognition for situations where the scene undergoes a major change in appearance, for example, due to illumination (day/night), change of seasons, aging, or structural modifications over time such as buildings built or destroyed. Such situations represent a major challenge for current large-scale place recognition methods. This work has the following three principal contributions. First, we demonstrate that matching across large changes in the scene appearance becomes much easier when both the query image and the database image depict the scene from approximately the same viewpoint. Second, based on this observation, we develop a new place recognition approach that combines (i) an efficient synthesis of novel views with (ii) a compact indexable image representation. Third, we introduce a new challenging dataset of 1,125 camera-phone query images of Tokyo that contain major changes in illumination (day, sunset, night) as well as structural changes in the scene. We demonstrate that the proposed approach significantly outperforms other large-scale place recognition techniques on this challenging data. This work has been published at CVPR 2015 [16]. Figure 1 shows examples of the newly collected Tokyo 24/7 dataset.

7.1.4. NetVLAD: CNN architecture for weakly supervised place recognition

Participants: Relja Arandjelović, Petr Gronat, Akihiko Torii, Tomas Pajdla, Josef Sivic.

Figure 1. Example query images from the newly collected 24/7 Tokyo dataset. Each place in the query set is captured at different times of day: (a) daytime, (b) sunset, and (c) night. For comparison, the database street-view image at a close-by position is shown in (d). Note the major changes in appearance (illumination changes in the scene) between the database image (d) and the query images (a,b,c)

In [21], we tackle the problem of large scale visual place recognition, where the task is to quickly and accurately recognize the location of a given query photograph. We present the following three principal contributions. First, we develop a convolutional neural network (CNN) architecture that is trainable in an end-to-end manner directly for the place recognition task. The main component of this architecture, NetVLAD, is a new generalized VLAD layer, inspired by the "Vector of Locally Aggregated Descriptors" image representation commonly used in image retrieval. The layer is readily pluggable into any CNN architecture and amenable to training via backpropagation. Second, we develop a training procedure, based on a new weakly supervised ranking loss, to learn parameters of the architecture in an end-to-end manner from images depicting the same places over time downloaded from Google Street View Time Machine. Finally, we show that the proposed architecture obtains a large improvement in performance over non-learnt image representations as well as significantly outperforms off-the-shelf CNN descriptors on two challenging place recognition benchmarks. This work is under review. Figure 2 shows some qualitative results.

Figure 2. Our trained NetVLAD descriptor correctly recognizes the location (b) of the query photograph (a) despite the large amount of clutter (people, cars), changes in viewpoint and completely different illumination (night vs daytime).

7.2. Category-level object and scene recognition

7.2.1. Is object localization for free? – Weakly-supervised learning with convolutional neural networks

Participants: Maxime Oquab, Leon Bottou [MSR New York], Ivan Laptev, Josef Sivic.

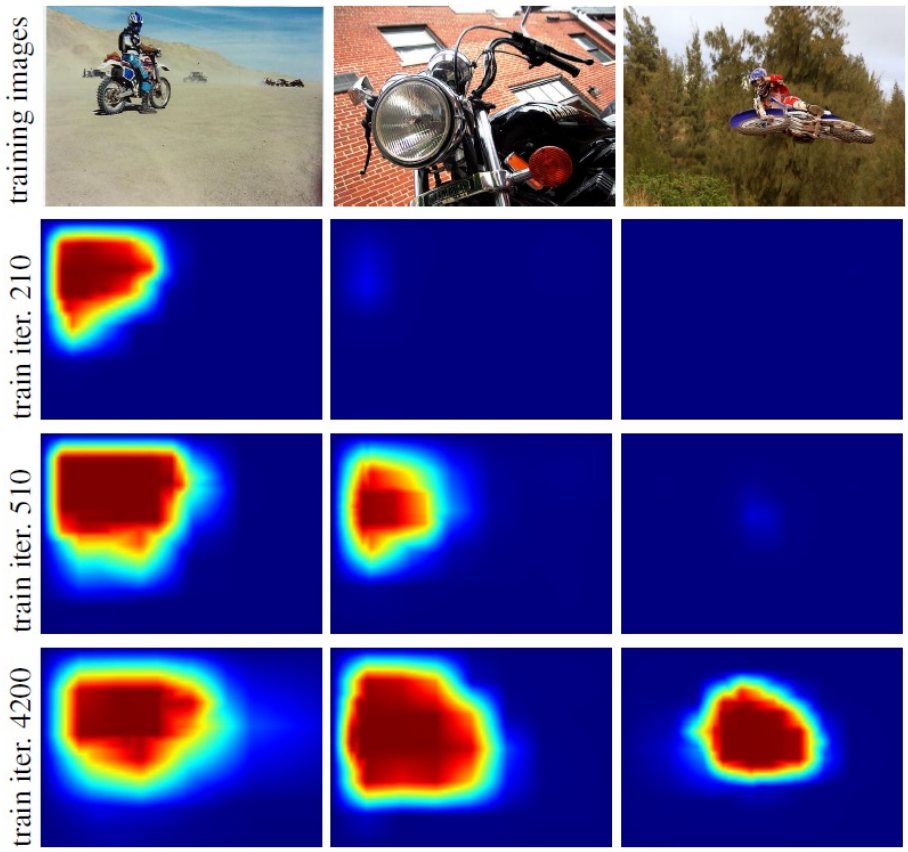

Figure 3. Evolution of localization score maps for the motorbike class over iterations of our weakly-supervised CNN training. Note that locations of objects with more usual appearance are discovered earlier during training.

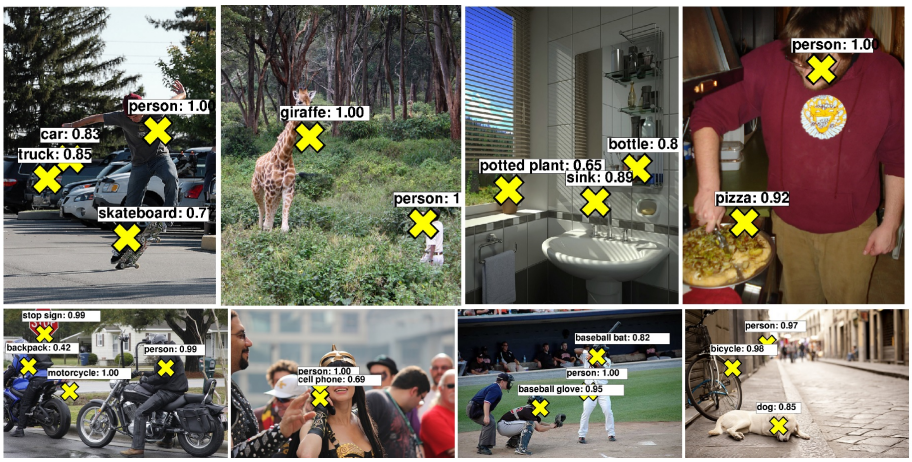

Figure 4. Example location predictions for images from the Microsoft COCO validation set obtained by our weakly-supervised method. Note that our method does not use object locations at training time, yet can predict locations of objects in test images (yellow crosses). The method outputs the most confident location for most confident object classes.

Successful methods for visual object recognition typically rely on training datasets containing lots of richly annotated images. Detailed image annotation, e.g. by object bounding boxes, however, is both expensive and often subjective. We describe a weakly supervised convolutional neural network (CNN) for object classification that relies only on image-level labels, yet can learn from cluttered scenes containing multiple objects (see Figure 3). We quantify its object classification and object location prediction performance on the Pascal VOC 2012 (20 object classes) and the much larger Microsoft COCO (80 object classes) datasets. We find that the network (i) outputs accurate image-level labels, (ii) predicts approximate locations (but not extents) of objects, and (iii) performs comparably to its fully-supervised counterparts using object bounding box annotation for training. This work has been published at CVPR 2015 [14] . Illustration of localization results by our method in Microsoft COCO dataset is shown in Figure 4 .

7.2.2. *Unsupervised Object Discovery and Localization in the Wild: Part-based Matching with Bottom-up Region Proposals*

Participants: Minsu Cho, Suha Kwak, Cordelia Schmid, Jean Ponce.

In [8], we address *unsupervised* discovery and localization of dominant objects from a noisy image collection of multiple object classes. The setting of this problem is fully unsupervised (Fig. 5), without even image-level annotations or any assumption of a single dominant class. This is significantly more general than typical colocalization, cosegmentation, or weakly-supervised localization tasks. We tackle the unsupervised discovery and localization problem using a part-based region matching approach: We use off-the-shelf region proposals to form a set of candidate bounding boxes for *objects* and *object parts*. These regions are efficiently matched across images using a probabilistic Hough transform that evaluates the confidence for each candidate correspondence considering both appearance similarity and spatial consistency. Dominant objects are discovered and localized by comparing the scores of candidate regions and selecting those that stand out over other regions containing them. Extensive evaluations on standard benchmarks (e.g., Object Discovery and PASCAL VOC 2007 datasets) demonstrate that the proposed approach significantly outperforms the current state of the art in colocalization, and achieves robust object discovery even in a fully unsupervised setting. This work has been published in CVPR 2015 [8] as oral presentation.

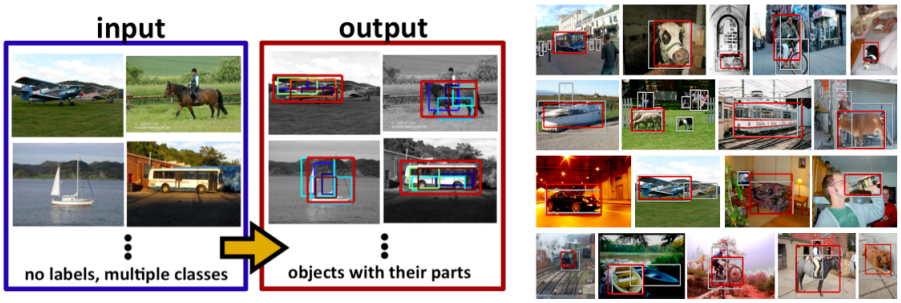

Figure 5. Unsupervised object discovery in the wild. We tackle object localization in an unsupervised scenario without any type of annotations, where a given image collection may contain multiple dominant object classes and even outlier images. The proposed method discovers object instances (red bounding boxes) with their distinctive parts (smaller boxes).

7.2.3. Unsupervised Object Discovery and Tracking in Video Collections

Participants: Suha Kwak, Minsu Cho, Ivan Laptev, Jean Ponce, Cordelia Schmid.

In [11], we address the problem of automatically localizing dominant objects as spatio-temporal tubes in a noisy collection of videos with minimal or even no supervision. We formulate the problem as a combination of two complementary processes: discovery and tracking (Figure 6). The first one establishes correspondences bet ween prominent regions across videos, and the second one associates similar object regions within the same video. It is empirically demonstrated that our method can handle video collections featuring multiple object classes, and substantially outperforms the state of the art in colocalization, even though it tackles a broader problem with much less supervision. This work has been published in ICCV 2015.

7.2.4. Linking Past to Present: Discovering Style in Two Centuries of Architecture

Participants: Stefan Lee, Nicolas Maisonneuve, David Crandall, Alexei A. Efros, Josef Sivic.

With vast quantities of imagery now available online, researchers have begun to explore whether visual patterns can be discovered automatically. Here we consider the particular domain of architecture, using huge collections of street-level imagery to find visual patterns that correspond to semantic-level architectural elements distinctive to particular time periods. We use this analysis both to date buildings, as well as to discover how functionally similar architectural elements (e.g. windows, doors, balconies, etc.) have changed over time due to evolving styles. We validate the methods by combining a large dataset of nearly 150,000 Google Street View images from Paris with a cadastre map to infer approximate construction date for each facade. Not only could our analysis be used for dating or geo-localizing buildings based on architectural features, but it also could give architects and historians new tools for confirming known theories or even discovering new ones. The work was published in [13] and the results are illustrated in figure 7 .

7.2.5. Proposal Flow

Participants: Bumsub Ham, Minsu Cho, Cordelia Schmid, Jean Ponce.

Finding image correspondences remains a challenging problem in the presence of intra-class variations and large changes in scene layout, typical in scene flow computation. In [22], we introduce a novel approach to this problem, dubbed proposal flow, that establishes reliable correspondences using object proposals. Unlike prevailing scene flow approaches that operate on pixels or regularly sampled local regions, proposal flow benefits from the characteristics of modern object proposals, that exhibit high repeatability at multiple scales,

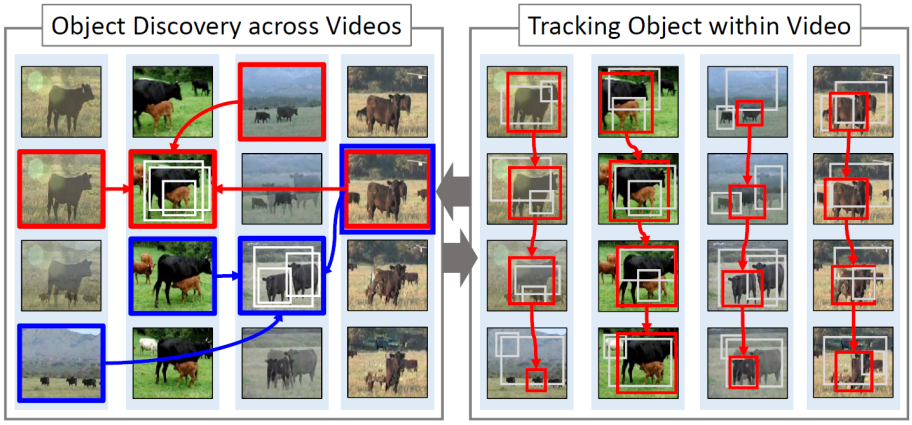

Figure 6. Dominant objects in a video collection are discovered by analyzing correspondences between prominent regions across videos (left). Within each video, object candidates, discovered by the former process, are temporally associated and a smooth spatio-temporal localization is estimated (right). These processes are alternated until convergence or up to a fixed number of iterations.

and can take advantage of both local and geometric consistency constraints among proposals. We also show that proposal flow can effectively be transformed into a conventional dense flow field. We introduce a new dataset that can be used to evaluate both general scene flow techniques and region-based approaches such as proposal flow. We use this benchmark to compare different matching algorithms, object proposals, and region features within proposal flow with the state of the art in scene flow. This comparison, along with experiments on standard datasets, demonstrates that proposal flow significantly outperforms existing scene flow methods in various settings. This work is under review. The proposed method and its qualitative result are illustrated in Figure 8 .

7.3. Image restoration, manipulation and enhancement

7.3.1. Learning a Convolutional Neural Network for Non-uniform Motion Blur Removal

Participants: Jian Sun, Wenfei Cao, Zongben Xu, Jean Ponce.

In this work, we address the problem of estimating and removing non-uniform motion blur from a single blurry image. We propose a deep learning approach to predicting the probabilistic distribution of motion blur at the patch level using a convolutional neural network (CNN). We further extend the candidate set of motion kernels predicted by the CNN using carefully designed image rotations. A Markov random field model is then used to infer a dense non-uniform motion blur field enforcing the motion smoothness. Finally the motion blur is removed by a non-uniform deblurring model using patch-level image prior. Experimental evaluations show that our approach can effectively estimate and remove complex non-uniform motion blur that cannot be well achieved by the previous approaches. This work has been published at CVPR 2015[15].

7.3.2. Robust Image Filtering Using Joint Static and Dynamic Guidance

Participants: Bumsub Ham, Minsu Cho, Jean Ponce.

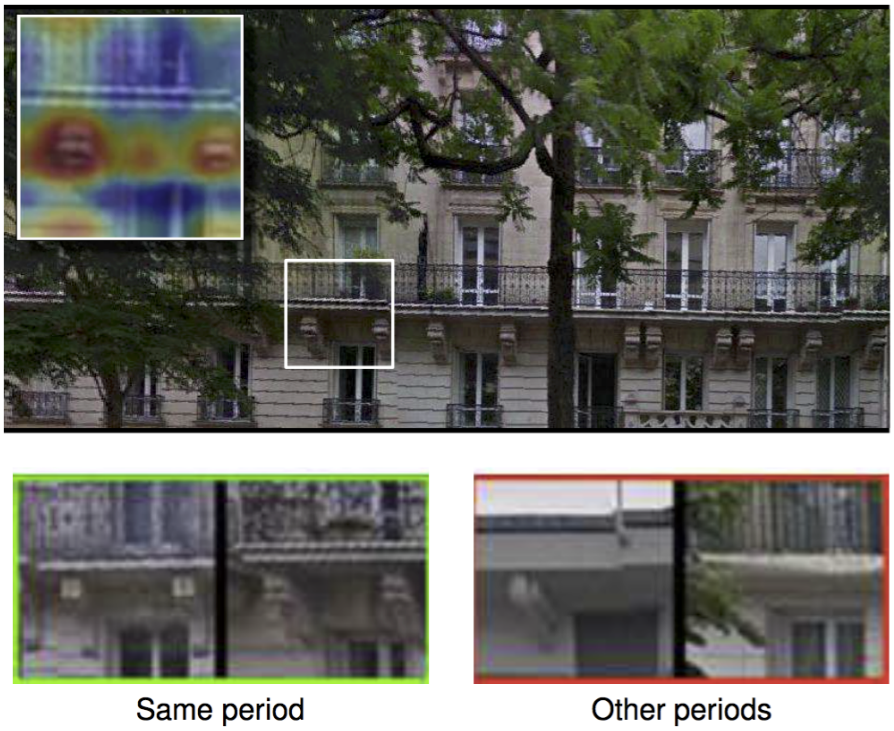

Figure 7. Using thousands of Street View images aligned to a cadastral map, we automatically find visual elements distinctive to particular architectural periods. For example, the patch in white above was found to be distinctive to the Haussmann period (late 1800's) in Paris, while the heat map (inset) reveals that the ornate balcony supports are the most distinctive features. We can also find functionally-similar elements fromthe same and different time periods (bottom).

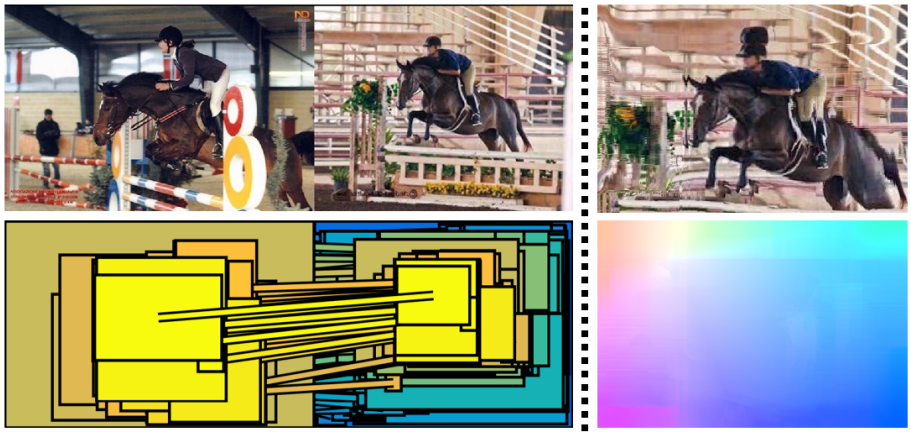

Figure 8. Proposal flow generates a reliable scene flow between similar images by establishing geometrically consistent correspondences between object proposals. (Left) Region-based scene flow by matching object proposals. (Right) Color-coded dense flow field generated from the region matches, and image warping using the flow.

Filtering images using a guidance signal, a process called joint or guided image filtering, has been used in various tasks in computer vision and computational photography, particularly for noise reduction and joint upsampling. The aim is to transfer the structure of the guidance signal to an input image, restoring noisy or altered image structure. The main drawbacks of such a data-dependent framework are that it does not consider differences in structure between guidance and input images, and it is not robust to outliers. We propose a novel SD (for static/dynamic) filter to address these problems in a unified framework by jointly leveraging structural information of guidance and input images. Joint image filtering is formulated as a nonconvex optimization problem, which is solved by the majorization-minimization algorithm. The proposed algorithm converges quickly while guaranteeing a local minimum. The SD filter effectively controls the underlying image structure at different scales and can handle a variety of types of data from different sensors. It is robust to outliers and other artifacts such as gradient reversal and global intensity shifting, and has good edge-preserving smoothing properties. We demonstrate the flexibility and effectiveness of the SD filter in a great variety of applications including depth upsampling, scale-space filtering, texture removal, flash/non-flash denoising, and RGB/NIR denoising. This has been published at CVPR 2015 [10]. The SD filter is illustrated in Figure 9 .

7.3.3. PCS-Net: A Deep learning approach to image restoration

Participants: Jian Sun, Jean Ponce.

This work introduces a novel framework for image restoration casting this problem as a joint classification and regression task. This is a learning-based approach, which first classifies degraded image patches into different categories, then restores these patches using category-specific models. We implement this idea by designing a novel convolutional neural network (dubbed PCS-Net), combining a CNN-based patch classification subnet with a novel patch category switched CNN architecture for category-specific restoration. The proposed PCS-Net learns different weights for different patch categories in a common network structure. Experiments on standard benchmarks show that our approach matches or improves upon the state of the art in image super-resolution and denoising. This work is under review.

7.4. Human activity capture and classification

7.4.1. P-CNN: Pose-based CNN Features for Action Recognition

Participants: Guilhem Chéron, Ivan Laptev, Cordelia Schmid.

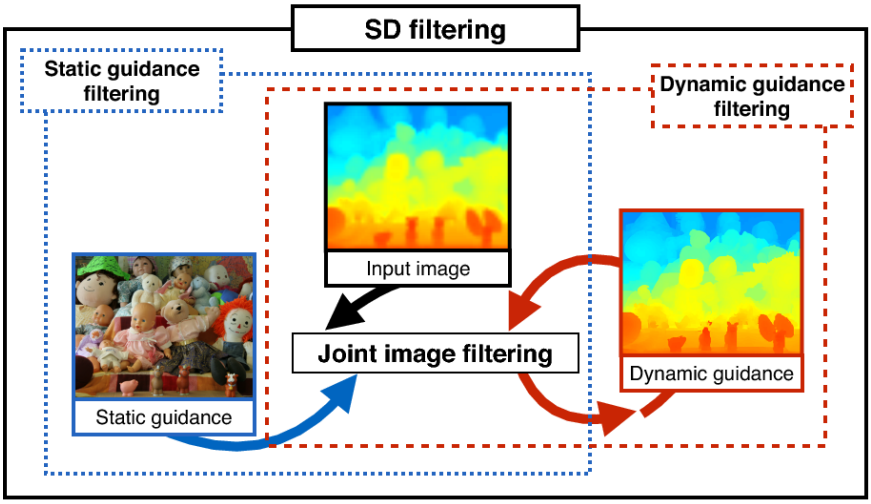

Figure 9. Sketch of joint image filtering and SD filtering: Static guidance filtering convolves an input image with a weight function computed from static guidance, as in the dotted blue box. Dynamic guidance filtering uses weight functions that are repeatedly obtained from regularized input images, as in the dotted red box. We have observed that static and dynamic guidance complement each other, and exploiting only one of them is problematic, especially in the case of data from different sensors (e.g., depth and color images). The SD filter takes advantage of both, and addresses the problems of current joint image filtering.

This work [9] targets human action recognition in video. We argue for the importance of a representation derived from human pose. To this end we propose a new Pose-based Convolutional Neural Network descriptor (P-CNN) for action recognition. The descriptor aggregates motion and appearance information along tracks of human body parts as shown in Figure 10 . We experiment with P-CNN features obtained both for automatically estimated and manually annotated human poses. We evaluate our method on JHMDB and MPII Cooking datasets. For both datasets our method shows consistent improvement over the state of the art. This work has been published at ICCV 2015 [9], and P-CNN code (Matlab) is available online at http://www.di.ens.fr/willow/research/p-cnn/.

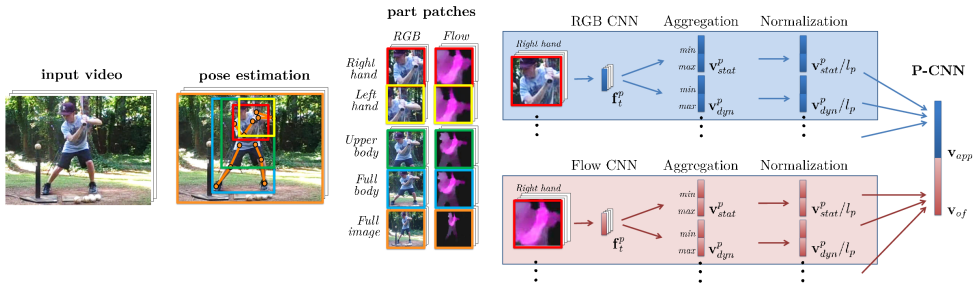

Figure 10. P-CNN features. From left to right: Input video. Human pose. Patches of appearance and optical flow for human body parts. One RGB and one flow CNN descriptor is extracted per frame and per part. Frame descriptors are aggregated over time to obtain the video descriptor. Video descriptors are normalized and concatenated into appearance features and flow features. The final P-CNN feature is the concatenation of appearance and flow.

7.4.2. Context-aware CNNs for person head detection

Participants: Tuan-Hung Vu, Anton Osokin, Ivan Laptev.

Person detection is a key problem for many computer vision tasks. While face detection has reached maturity, detecting people under a full variation of camera view-points, human poses, lighting conditions and occlusions is still a difficult challenge. In this work we focus on detecting human heads in natural scenes. Starting from the recent local R-CNN object detector, we extend it with two types of contextual cues. First, we leverage person-scene relations and propose a Global CNN model trained to predict positions and scales of heads directly from the full image. Second, we explicitly model pairwise relations among objects and train a Pairwise CNN model using a structured-output surrogate loss. The Local, Global and Pairwise models are combined into a joint CNN framework. To train and test our full model, we introduce a large dataset composed of 369,846 human heads annotated in 224,740 movie frames. We evaluate our method and demonstrate improvements of person head detection against several recent baselines in three datasets. We also show improvements of the detection speed provided by our model. This work has been published at ICCV 2015 [18]. The code and the new dataset developed in this work are available online at http://www.di.ens.fr/willow/research/headdetection/.

7.4.3. On Pairwise Costs for Network Flow Multi-Object Tracking

Participants: Visesh Chari, Simon Lacoste-Julien, Ivan Laptev, Josef Sivic.

Multi-object tracking has been recently approached with the min-cost network flow optimization techniques. Such methods simultaneously resolve multiple object tracks in a video and enable modeling of dependencies among tracks. Min-cost network flow methods also fit well within the "tracking-by-detection" paradigm where object trajectories are obtained by connecting per-frame outputs of an object detector. Object detectors, however, often fail due to occlusions and clutter in the video. To cope with such situations, we propose an

approach that regularizes the tracker by adding second order costs to the min-cost network flow framework. While solving such a problem with integer variables is NP-hard, we present a convex relaxation with an efficient rounding heuristic which empirically gives certificates of small suboptimality. Results are shown on real world video sequences and demonstrate that the new constraints help selecting longer and more accurate tracks improving over the baseline tracking-by-detection method. This work has been published at CVPR 2015 [7].

7.4.4. Pose Estimation and Segmentation of Multiple People in Stereoscopic Movies

Participants: Guillaume Seguin, Karteek Alahari, Josef Sivic, Ivan Laptev.

We describe a method to obtain a pixel-wise segmentation and pose estimation of multiple people in stereoscopic videos, illustrated in Figure 11 . This task involves challenges such as dealing with unconstrained stereoscopic video, non-stationary cameras, and complex indoor and outdoor dynamic scenes with multiple people. We cast the problem as a discrete labelling task involving multiple person labels, devise a suitable cost function, and optimize it efficiently. The contributions of our work are two-fold: First, we develop a segmentation model incorporating person detections and learnt articulated pose segmentation masks, as well as colour, motion, and stereo disparity cues. The model also explicitly represents depth ordering and occlusion. Second, we introduce a stereoscopic dataset with frames extracted from feature-length movies “StreetDance 3D" and “Pina". The dataset contains 587 annotated human poses, 1158 bounding box annotations and 686 pixel-wise segmentations of people. The dataset is composed of indoor and outdoor scenes depicting multiple people with frequent occlusions. We demonstrate results on our new challenging dataset, as well as on the H2view dataset from (Sheasby et al.’s ACCV 2012). This work has been published at PAMI [4].

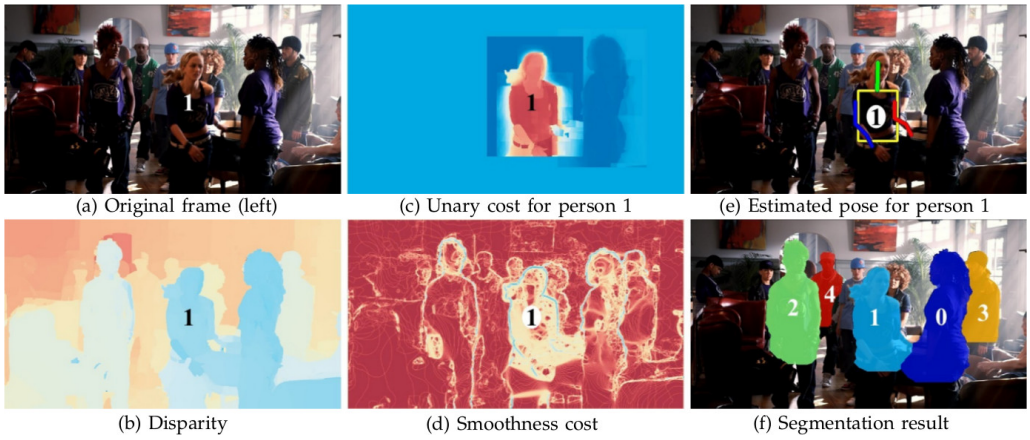

Figure 11. Starting from a stereo pair (a), we estimate disparity maps (b). Using both appearance and disparity cues, we detect persons and estimate their poses (e). We combine pose information with disparity information and occlusion reasonning to compute the unary potentials of a CRF (c) and use standard color and motion cues to compute the binary terms (d). We optimize the CRF problem to produce the final, layered segmentation (f).

7.4.5. Weakly-Supervised Alignment of Video with Text

Participants: Piotr Bojanowski, Rémi Lajugie, Edouard Grave, Francis Bach, Ivan Laptev, Jean Ponce, Cordelia Schmid.

In this work [6], we design a method for aligning natural language sentences with a video stream. Suppose that we are given a set of videos, along with natural language descriptions in the form of multiple sentences (e.g., manual annotations, movie scripts, sport summaries etc.), and that these sentences appear in the same temporal order as their visual counterparts. We propose here a method for aligning the two modalities, i.e., automatically providing a time stamp for every sentence (see Fig. 12). Given vectorial features for both video and text, we propose to cast this task as a temporal assignment problem, with an implicit linear mapping between the two feature modalities. We formulate this problem as an integer quadratic program, and solve its continuous convex relaxation using an efficient conditional gradient algorithm. Several rounding procedures are proposed to construct the final integer solution. After demonstrating significant improvements over the state of the art on the related task of aligning video with symbolic labels, we evaluate our method on a challenging dataset of videos with associated textual descriptions, using both bag-of-words and continuous representations for text. This work has been published at CVPR 2015 [6].

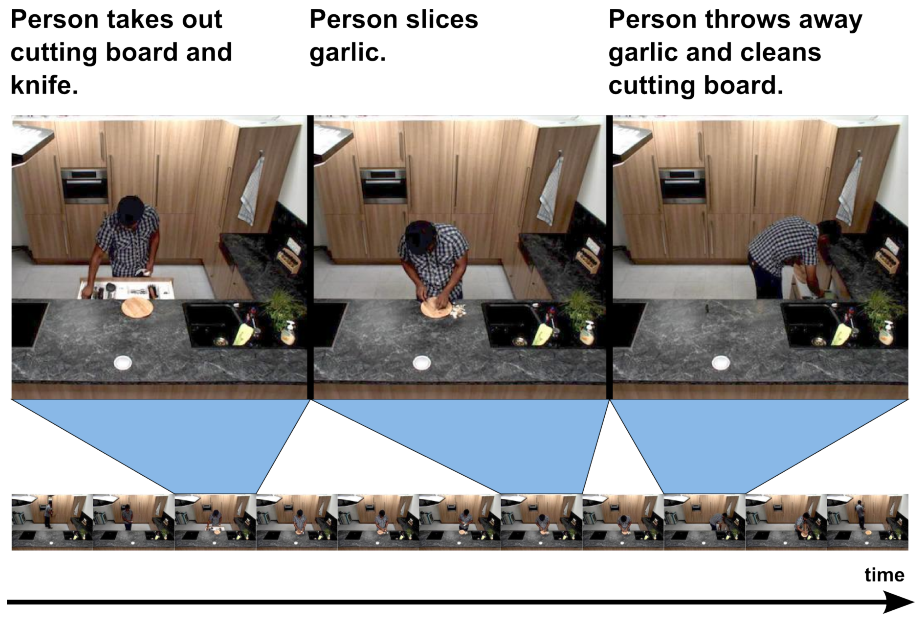

Figure 12. Illustration of the text to video alignment problem. As an output, our model provides a temporal location for every sentence.

7.4.6. Unsupervised learning from narrated instruction videos

Participants: Jean-Baptiste Alayrac, Piotr Bojanowski, Nishant Agrawal, Josef Sivic, Ivan Laptev, Simon Lacoste-Julien.

In [20], we address the problem of automatically learning the main steps to complete a certain task, such as changing a car tire, from a set of narrated instruction videos. The contributions of this paper are three-fold. First, we develop a new unsupervised learning approach that takes advantage of the complementary nature of the input video and the associated narration. The method solves two clustering problems, one in text and one in video, applied one after each other and linked by joint constraints to obtain a single coherent sequence of steps in both modalities. Second, we collect and annotate a new challenging dataset of real-world instruction videos from the Internet. The dataset contains about 800,000 frames for five different tasks that include complex interactions between people and objects, and are captured in a variety of indoor and outdoor settings. Third, we

experimentally demonstrate that the proposed method can automatically discover, in an unsupervised manner, the main steps to achieve the task and locate the steps in the input videos. This work is under review.

7.4.7. Long-term Temporal Convolutions for Action Recognition

Participants: Gül Varol, Ivan Laptev, Cordelia Schmid.

Typical human actions such as hand-shaking and drinking last several seconds and exhibit characteristic spatio-temporal structure. Recent methods attempt to capture this structure and learn action representations with convolutional neural networks. Such representations, however, are typically learned at the level of single frames or short video clips and fail to model actions at their full temporal scale. In [27], we learn video representations using neural networks with long-term temporal convolutions. We demonstrate that CNN models with increased temporal extents improve the accuracy of action recognition despite reduced spatial resolution. We also study the impact of different low-level representations, such as raw values of video pixels and optical flow vector fields and demonstrate the importance of high-quality optical flow estimation for learning accurate action models. We report state-of-the-art results on two challenging benchmarks for human action recognition UCF101 and HMDB51. This work is under review. The results for the proposed method are illustrated in Figure 13 .

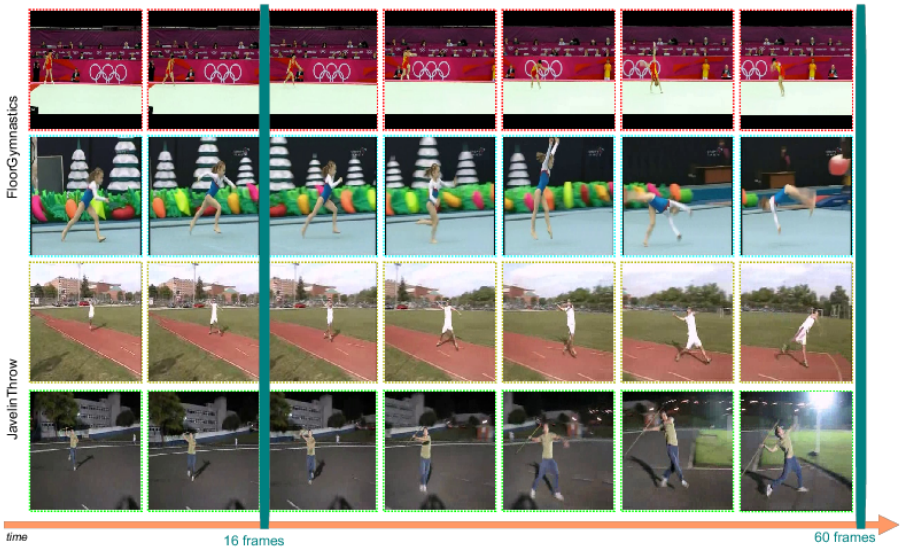

Figure 13. The highest improvement of long-term temporal convolutions in terms of class accuracy is for "JavelinThrow". For 16-frame network, it is mostly confused with "FloorGymnastics" class. We visualize sample videos with 7 frames extracted at every 8 frames. The intuitive explanation is that both classes start by running for a few seconds and then the actual action takes place. Long-term temporal convolutions with 60 frames can capture this interval, whereas 16-frame networks fail to recognize such long-term activities.

7.4.8. Thin-Slicing forPose: Learning to Understand Pose without Explicit Pose Estimation

Participants: Suha Kwak, Minsu Cho, Ivan Laptev.

In [23], we address the problem of learning a pose-aware, compact embedding that projects images with similar human poses to be placed close-by in the embedding space (Figure 14). The embedding function is built on a deep convolutional network, and trained with a triplet-based rank constraint on real image data. This

architecture allows us to learn a robust representation that captures differences in human poses by effectively factoring out variations in clothing, background, and imaging conditions in the wild. For a variety of pose-related tasks, the proposed pose embedding provides a cost-efficient and natural alternative to explicit pose estimation, circumventing challenges of localizing body joints. We demonstrate the efficacy of the embedding on pose-based image retrieval and action recognition problems. This work is under review.

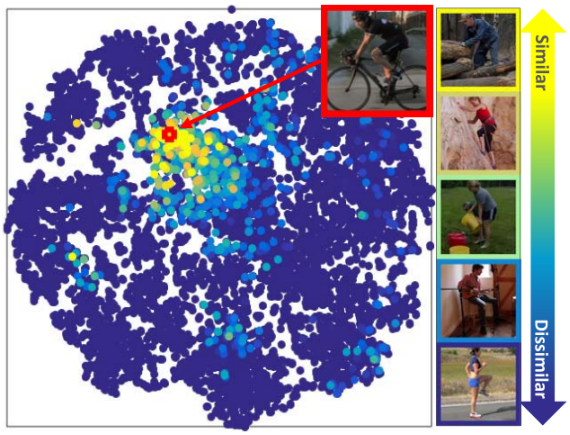

Figure 14. The manifold of our pose embedding visualized using t-SNE. Each point represents a human pose image. To better show correlation between the pose embedding and annotated pose, we color-code pose similarities in annotation between an arbitrary target image (red box) and all the other images. Selected examples of color-coded images are illustrated in the right-hand side. Images similar with the target in annotated pose are colored in yellow, otherwise in blue. As can be seen, yellow images lie closer by the target in general, which indicates that a position on the embedding space implicitly represents a human pose.

7.4.9. Instance-level video segmentation from object tracks

Participants: Guillaume Seguin, Piotr Bojanowski, Rémi Lajugie, Ivan Laptev.

In [26], we address the problem of segmenting multiple object instances in complex videos. Our method does not require manual pixel-level annotation for training, and relies instead on readily-available object detectors or visual object tracking only. Given object bounding boxes at input as shown in Figure 15 , we cast video segmentation as a weakly-supervised learning problem. Our proposed objective combines (a) a discriminative clustering term for background segmentation, (b) a spectral clustering one for grouping pixels of same object instances, and (c) linear constraints enabling instance-level segmentation. We propose a convex relaxation of this problem and solve it efficiently using the Frank-Wolfe algorithm. We report results and compare our method to several baselines on a new video dataset for multi-instance person segmentation. This work is under review.

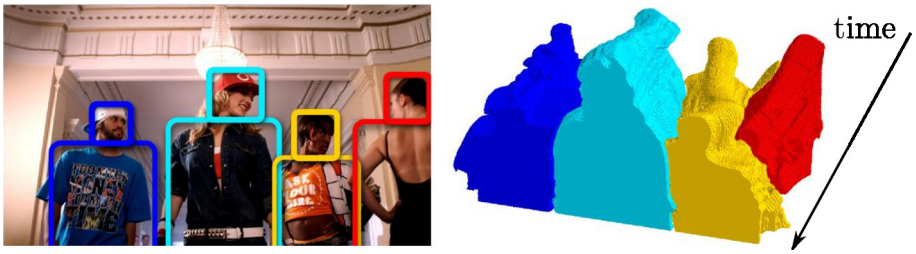

Figure 15. Results of our method applied to multi-person segmentation in a sample video from our database. Given an input video together with the tracks of object bounding boxes (left), our method finds pixel-wise segmentation for each object instance across video frames (right).