

Inria

RESEARCH CENTER
Paris

FIELD

Activity Report 2019

Section New Results

Edition: 2020-03-21

ALGORITHMICS, PROGRAMMING, SOFTWARE AND ARCHITECTURE

1. ANTIQUE Project-Team	4
2. CAMBIUM Project-Team	8
3. CASCADE Project-Team	14
4. KOPERNIC Team	15
5. OURAGAN Project-Team	18
6. PARKAS Project-Team	22
7. PIR2 Project-Team	28
8. POLSYS Project-Team	37
9. PROSECCO Project-Team	41
10. SECRET Project-Team	43

APPLIED MATHEMATICS, COMPUTATION AND SIMULATION

11. CAGE Project-Team	50
12. MATHERIALS Project-Team	59
13. MATHRISK Project-Team	66
14. MOKAPLAN Project-Team	70
15. QUANTIC Project-Team	74
16. SIERRA Project-Team	78

DIGITAL HEALTH, BIOLOGY AND EARTH

17. ANGE Project-Team	83
18. ARAMIS Project-Team	90
19. COMMEDIA Project-Team	96
20. MAMBA Project-Team	99
21. REO Team	105
22. SERENA Project-Team	106

NETWORKS, SYSTEMS AND SERVICES, DISTRIBUTED COMPUTING

23. ALPINES Project-Team	109
24. DELYS Project-Team	112
25. DYOGENE Project-Team	116
26. EVA Project-Team	128
27. GANG Project-Team	137
28. GANG Project-Team	143
29. MIMOVE Project-Team	149
30. WHISPER Project-Team	155

PERCEPTION, COGNITION AND INTERACTION

31. ALMANACH Project-Team	158
32. COML Team	165
33. RITS Project-Team	170
34. VALDA Project-Team	181
35. WILLOW Team	184

ANTIQUÉ Project-Team

7. New Results

7.1. Relational Static Analysis

7.1.1. *Relational abstraction for memory properties*

Participants: Hugo Illous, Matthieu Lemerre, Xavier Rival [correspondant].

Static analyses aim at inferring semantic properties of programs. We can distinguish two important classes of static analyses: state analyses and relational analyses. While state analyses aim at computing an over-approximation of reachable states of programs, relational analyses aim at computing functional properties over the input-output states of programs. Several advantages of relational analyses are their ability to analyze incomplete programs, such as libraries or classes, but also to make the analysis modular, using input-output relations as composable summaries for procedures. In the case of numerical programs, several analyses have been proposed that utilize relational numerical abstract domains to describe relations. On the other hand, designing abstractions for relations over input-output memory states and taking shapes into account is challenging. We have proposed a set of novel logical connectives to describe such relations, which are inspired by separation logic. This logic can express that certain memory areas are unchanged, freshly allocated, or freed, or that only part of the memory was modified. Using these connectives, we have built an abstract domain and design a static analysis that over-approximates relations over memory states containing inductive structures. We implemented this analysis and evaluated it on a basic library of list manipulating functions.

This work was done as part of the Phd of Hugo Illous [10] and a journal paper is currently under submission.

7.2. Static Analysis of Probabilistic Programming Languages

7.2.1. *Towards the verification of semantic assumptions required by probabilistic inference algorithms*

Participants: Wonyeol Lee, Hangeol Wu, Xavier Rival [correspondant], Hongseok Yang.

Probabilistic programming is the idea of writing models from statistics and machine learning using program notations and reasoning about these models using generic inference engines. Recently its combination with deep learning has been explored intensely, which led to the development of so called deep probabilistic programming languages, such as Pyro, Edward and ProbTorch. At the core of this development lie inference engines based on stochastic variational inference algorithms. When asked to find information about the posterior distribution of a model written in such a language, these algorithms convert this posterior-inference query into an optimisation problem and solve it approximately by a form of gradient ascent or descent. We analysed one of the most fundamental and versatile variational inference algorithms, called score estimator or REINFORCE, using tools from denotational semantics and program analysis. We formally expressed what this algorithm does on models denoted by programs, and exposed implicit assumptions made by the algorithm on the models. The violation of these assumptions may lead to an undefined optimisation objective or the loss of convergence guarantee of the optimisation process. We then describe rules for proving these assumptions, which can be automated by static program analyses. Some of our rules use nontrivial facts from continuous mathematics, and let us replace requirements about integrals in the assumptions, such as integrability of functions defined in terms of programs' denotations, by conditions involving differentiation or boundedness, which are much easier to prove automatically (and manually). Following our general methodology, we have developed a static program analysis for the Pyro programming language that aims at discharging the assumption about what we call model-guide support match. Our analysis is applied to the eight representative model-guide pairs from the Pyro webpage, which include sophisticated neural network models such as AIR. It found a bug in one of these cases, and revealed a non-standard use of an inference engine in another, and showed that the assumptions are met in the remaining six cases.

This work has been published in [12].

7.3. Static Analysis of JavaScript Code

7.3.1. *Weakly Sensitive Analysis for Unbounded Iteration over JavaScript Objects*

Participants: Yoonseok Ko, Xavier Rival [correspondant], Sukyoung Ryu.

In [23] and [11], we studied composite object abstraction for the analysis JavaScript.

JavaScript framework libraries like jQuery are widely use, but complicate program analyses. Indeed, they encode clean high-level constructions such as class inheritance via dynamic object copies and transformations that are harder to reason about. One common pattern used in them consists of loops that copy or transform part or all of the fields of an object. Such loops are challenging to analyze precisely, due to weak updates and as unrolling techniques do not always apply. In this work, we observe that precise field correspondence relations are required for client analyses (e.g., for call-graph construction), and propose abstractions of objects and program executions that allow to reason separately about the effect of distinct iterations without resorting to full unrolling. We formalize and implement an analysis based on this technique. We assess the performance and precision on the computation of call-graph information on examples from jQuery tutorials.

7.4. Rule-based Modeling with Arithmetics

7.4.1. *Counters in Kappa: Semantics, Simulation, and Static Analysis.*

Participants: Pierre Boutillier, Ioana Cristescu, Jérôme Feret.

Site-graph rewriting languages, such as Kappa or BNGL, offer parsimonious ways to describe highly combinatorial systems of mechanistic interactions among proteins. These systems may be then simulated efficiently. Yet, the modeling mechanisms that involve counting (a number of phosphorylated sites for instance) require an exponential number of rules in Kappa. In BNGL, updating the set of the potential applications of rules in the current state of the system comes down to the sub-graph isomorphism problem (which is NP-complete).

In [14], we extend Kappa to deal both parsimoniously and efficiently with counters. We propose a single push-out semantics for Kappa with counters. We show how to compile Kappa with counters into Kappa without counters (without requiring an exponential number of rules). We design a static analysis, based on affine relationships, to identify the meaning of counters and bound their ranges accordingly.

7.5. Reduced product

7.5.1. *Sharing Ghost Variables in a Collection of Abstract Domains.*

Participants: Marc Chevalier, Jérôme Feret.

In abstract interpretation, it is often necessary to be able to express complex properties while doing a precise analysis. A way to achieve that is to combine a collection of domains, each handling some kind of properties, using a reduced product. Separating domains allows an easier and more modular implementation, and eases soundness and termination proofs. This way, we can add a domain for any kind of property that is interesting. The reduced product, or an approximation of it, is in charge of refining abstract states, making the analysis precise.

In program verification, ghost variables can be used to ease proofs of properties by storing intermediate values that do not appear directly in the execution.

In [15], we propose a reduced product of abstract domains that allows domains to use ghost variables to ease the representation of their internal state. Domains must be totally agnostic with respect to other existing domains. In particular the handling of ghost variables must be entirely decentralized while still ensuring soundness and termination of the analysis.

7.6. Static Analysis of Neural Networks

7.6.1. *Perfectly Parallel Fairness Certification.*

Participants: Caterina Urban [correspondant], Maria Christakis, Valentin Wüestholz, Fuyuan Zhang.

Recently, there is growing concern that machine-learning models, which currently assist or even automate decision making, reproduce, and in the worst case reinforce, bias of the training data. The development of tools and techniques for certifying fairness of these models or describing their biased behavior is, therefore, critical.

In [19], we propose a perfectly parallel static analysis for certifying causal fairness of feed-forward neural networks used for classification tasks. When certification succeeds, our approach provides definite guarantees, otherwise, it describes and quantifies the biased behavior. We design the analysis to be sound, in practice also exact, and configurable in terms of scalability and precision, thereby enabling pay-as-you-go certification. We implement our approach in an open-source tool and demonstrate its effectiveness on models trained with popular datasets.

7.7. Reductions between synchronous and asynchronous programming abstractions

7.7.1. *Communication closed asynchronous protocols.*

Participants: Andrei Damien, Cezara Drăgoi, Alexandru Militaru, Josef Widder.

Fault-tolerant distributed systems are implemented over asynchronous networks, where performance emerges from the load of the system. Due to asynchronous communication and the occurrence of faults (e.g., process crashes or the network dropping messages) the implementations are hard to understand and analyze. In contrast, synchronous computation models simplify design and reasoning.

In [17], we defined the first algorithm that automatically transforms an asynchronous protocol into a synchronous one. The method is sound but not complete. The transformation is based on an axiomatization of the notion of communication closure introduced by Elrad and Frances. If the asynchronous protocol is communication-closed then the translator will successfully compute its synchronous counter-part. Checking communication closure is done locally without considering any interferences between processes. The translator was successfully applied to Multi-Paxos, ViewStamped, and the atomic broadcast of Chandra and Toueg, generating the first synchronous counterparts of these protocols. The transformation from asynchronous to synchronous preserves the local states processes go through and the exchanged messages. The translator has been implemented in a prototype tool called Athos, i.e., Asynchronous To Heard-Of Synchronizer, that is open source. The tool takes as input protocols in an intermediate protocol language that has an asynchronous semantics and it is very close to C. These results have been published in one of the main verification venues Computer Aided Verification, CAV 2019 (acceptance rate <25% out of >250 submissions). The impact of the translator from asynchronous protocols to equivalent synchronous ones is important for the verification community because such a transformation reduces dramatically the state space and the set of traces to explore in order to prove the program correct, independently of the used verification technique.

7.7.2. Executable Rounds: a Programming Abstraction for Fault-Tolerant Protocols.

Participants: Cezara Drăgoi, Josef Widder, Damien Zufferey.

Fault-tolerant distributed systems are notoriously difficult to design and implement. Although programming languages for distributed systems is an active research area, appropriate synchronization primitives for fault-tolerance and group communication remains an important challenge. In [18] we present a new programming abstraction, HSync, for implementing benign and Byzantine distributed protocols. HSync is based on communication-closed rounds. Round models offer a simple abstraction for group communication and communication-closed rounds simplify dealing with faults. Protocols are implemented in a modular way in HSync. The language separates the message reception from the process local computation. It extends classic rounds with language constructs that give to the programmer the possibility to implement network and algorithm-specific policies for message reception. We have implemented an execution platform for HSync that runs on top of commodity hardware. We evaluate experimentally its performance, by comparing consensus implementations in HSync with LibPaxos3 and Bft-SMaRt, two consensus libraries tolerant to benign, resp. Byzantine faults.

7.8. Introduction

Participant: Andreea Beica.

The PhD of Andreea Beica [9] aims at studying two aspects related to the modelling of Biochemical Reaction Networks, in the context of Systems Biology.

In the first part, we analyse how scale-separation in biological systems can be exploited for model reduction. We first argue for the use of rule-based models for prototyping genetic circuits, and then show how the inherent multi-scaleness of such systems can be used to devise a general model approximation method for rule-based models of genetic regulatory networks. The reduction proceeds via static analysis of the rule system. Our method relies on solid physical justifications, however not unlike other scale-separation reduction techniques, it lacks precise methods for quantifying the approximation error, while avoiding to solve the original model. Consequently, we next propose an approximation method for deterministic models of biochemical networks, in which reduction guarantees represent the major requirement. This second method combines abstraction and numerical approximation, and aims at providing a better understanding of model reduction methods that are based on time- and concentration- scale separation.

In the second part of the thesis, we introduce a new re-parametrisation technique for differential equation models of biochemical networks, in order to study the effect of intracellular resource storage strategies on growth, in self-replicating mechanistic models. Finally, we aim towards the characterisation of cellular growth as an emergent property of a novel Petri Net model semantics of Biochemical Reaction Networks.

CAMBIUM Project-Team

6. New Results

6.1. Programming language design and implementation

6.1.1. The OCaml system

Participants: Damien Doligez, Armaël Guéneau, Xavier Leroy, Luc Maranget, David Allsop [Cambridge University], Florian Angeletti, Frédéric Bour [Facebook, until Sep 2019], Stephen Dolan [Cambridge University], Alain Frisch [Lexifi], Jacques Garrigue [Nagoya University], Sébastien Hinderer [SED], Nicolás Ojeda Bär [Lexifi], Gabriel Radanne, Thomas Refis [Jane Street], Gabriel Scherer [Inria team Parsifal], Mark Shinwell [Jane Street], Leo White [Jane Street], Jeremy Yallop [Cambridge University].

This year, we released four versions of the OCaml system: versions 4.08.0, 4.08.1, 4.09.0, and 4.09.1. Versions 4.08.1 and 4.09.1 are minor releases that respectively fix 6 and 5 issues. Versions 4.08.0 and 4.09.0 are major releases that introduce new language features, improve performance and usability, and fix about 50 issues. The main novelties are:

- User-defined binding operators are now supported, with syntax similar to `let*`, `let+`, and `and*`. These operators make it much easier to write OCaml code in monadic style or using applicative structures.
- The `open` construct now applies to arbitrary module expressions in structures and to applicative paths in signatures.
- A new notion of user-defined “alerts” generalizes the “deprecated” warning.
- New modules were added to the standard library: `Fun`, `Bool`, `Int`, `Option`, `Result`.
- Many floating-point functions were added, including fused multiply-add, as well as a new `Float.Array` submodule.
- Many error messages were improved, as well as error and warning reporting mechanisms.
- Pattern-matching constructs that correspond to affine functions are now optimized into arithmetic computations.

6.1.2. Evolution of the OCaml type system

Participants: Florian Angeletti, Jacques Garrigue, Thomas Refis [Jane Street], Didier Rémy, Gabriel Radanne, Gabriel Scherer [Inria team Parsifal], Leo White [Jane Street].

In addition to the work done on the above releases, efforts have been done to improve the type system and its implementation. Those include:

- Formalizing the typing of the pattern-matching of generalized algebraic data types (GADTs).
- Fixing some issues related to the incompleteness of the treatment of GADTs.
- Proposing extensions of the type system to reduce this incompleteness in concrete cases, by refining the information on abstract types.
- Exploring practical ways to obtain more polymorphism for functions whose soundness does not rely on the value restriction.
- Improving the readability of the type-checker code.
- Making the module layer of the type-checker more incremental, in order to improve efficiency and to facilitate integration with documentation tools.

6.1.3. Refactoring with ornaments in ML

Participants: Didier Rémy, Thomas Williams [Google Paris].

Thomas Williams, Lucas Baudin, and Didier Rémy have been working on refactoring and other transformations of ML programs based on mixed ornamentation and disornamentation. Ornaments have been introduced as a way of describing changes in data type definitions that can reorganize or add pieces of data. After a new data structure has been described as an ornament of an older one, the functions that operate on the bare structure can be partially or sometimes totally lifted into functions that operate on the ornamented structure.

This year, Williams and Rémy improved the formalization of the lifting framework. In particular, they introduced an intermediate language, in which nonexpansive expressions can be marked on source terms and traced during reduction. This allows to treat the nonexpansive part of expansive expressions as nonexpansive and use equational reasoning on nonexpansive parts of terms that appear in types. This approach significantly simplifies the metatheory of ornaments. This calculus could also have some interest in itself, beyond ornaments, to study languages with side effects.

6.1.4. A better treatment of type abbreviations during type inference

Participants: Didier Rémy, Carine Morel.

During her M2 internship under the supervision of Didier Rémy, Carine Morel revisited the treatment of type abbreviations in type inference for ML-like type systems, using a modern approach based on typing constraints [24]. Instead of expanding type abbreviations prior to unification, both the original abbreviated view and all expanded views are kept during unification, so as to avoid unnecessary expansions and use the least-expanded view whenever possible in the result of unification.

6.2. Software specification and verification

6.2.1. The CompCert formally-verified compiler

Participants: Xavier Leroy, Jacques-Henri Jourdan [CNRS], Michael Schmidt [AbsInt GmbH], Bernhard Schommer [AbsInt GmbH].

In the context of our work on compiler verification, 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 ARM, PowerPC, RISC-V and x86 architectures [8]. This compiler comprises a back-end part, which translates the Cminor intermediate language to PowerPC assembly and which is reusable for source languages other than C [7], 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 100000-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 added a new optimization to CompCert: “if-conversion”, that is, the replacement of conditional statements and expressions by conditional move operations and similar branchless instruction sequences. As a consequence, fewer conditional branch instructions are generated. This replacement usually improves worst-case execution time (WCET), because mispredicted conditional branches tremendously increase execution time. This replacement is also interesting for cryptographic code and other programs that manipulate secret data: conditional branches over secret data take time that depends on the data, leaking some information, while conditional move instructions are constant-time and do not leak. The new if-conversion optimization plays a role in the ongoing work of Inria team Celtique on compilation that preserves constant-time properties. Its proof of semantic preservation is nontrivial and prompted the development of a new kind of simulation diagram.

Other recent improvements to the CompCert C compiler include:

- a new code generator targeting the AArch64 instruction set, that is, the 64-bit mode of the ARMv8 architecture;
- the ability to specify the semantics of certain built-in functions, making them amenable to optimizations such as constant propagation and common subexpression elimination;
- improvements to the verified C parser generated by Menhir, including fewer run-time checks, faster validation, and the removal of all axioms from the proof.

We released two versions of CompCert incorporating these improvements: version 3.5 in February 2019 and version 3.6 in September 2019.

6.2.2. *Time credits and time receipts in Iris*

Participants: Glen Mével, François Pottier, Jacques-Henri Jourdan [CNRS].

From March to August 2018, Glen Mével did an M2 internship at Gallium, where he was co-advised by Jacques-Henri Jourdan (CNRS) and François Pottier. Glen extended the program logic Iris with time credits and time receipts.

Time credits are a well-understood concept, and have been used in several papers already by Armaël Guéneau, Arthur Charguéraud, and François Pottier. However, because Iris is implemented and proved sound inside Coq, extending Iris with time credits requires a nontrivial proof, which Glen carried out, based on a program transformation which inserts “tick” instructions into the code. As an application of time credits, Glen verified inside Iris the correctness of Okasaki’s notion of “debits”, which allows reasoning about the time complexity of programs that use thunks.

Time receipts are a new concept, which allows proving that certain undesirable events, such as integer overflows, cannot occur until a very long time has elapsed. Glen extended Iris with time receipts and proved the soundness of this extension. As an application of time credits and receipts together, Jacques-Henri Jourdan updated Charguéraud and Pottier’s earlier verification of the Union-Find data structure [12] and proved that integer ranks cannot realistically overflow, even if they are stored using only $\log W$ bits, where W is the number of bits in a machine word.

This work carried out in 2018 has been published at ESOP 2019 [16].

6.2.3. *A program logic for Multicore OCaml*

Participants: Glen Mével, François Pottier, Jacques-Henri Jourdan [CNRS].

Glen Mével, who is co-advised by Jacques-Henri Jourdan and François Pottier, has been working on designing a mechanized program logic for Multicore OCaml.

One of the key challenges is to enable deductive reasoning under a weak memory model. In such a model, the behaviors of a program are no longer described by a naive interleaving semantics. Thus, the operational semantics that describes a weak memory model often feels unnatural to the programmer, and is difficult to reason about.

This year, Glen designed and implemented a proof system on top of Iris, a modular separation logic framework whose implementation and soundness proof are both expressed in Coq. This system allows mechanized program verification for a fragment of the Multicore OCaml language. It provides a certain degree of abstraction over the low-level operational semantics, in the hope of simplifying reasoning. This abstraction includes an abstract concept of “local view” of the shared memory; views are exchanged between threads via atomic locations.

A few simple concurrent data structures have been proven correct using the system. They include several variants of locks and mutual exclusion algorithms.

Glen presented preliminary results at the Iris Workshop in October 2019.

6.2.4. *Verifying a generic local solver in Iris*

Participants: Paulo Emílio de Vilhena, Jacques-Henri Jourdan [CNRS], François Pottier.

From March to August 2019, Paulo Emílio de Vilhena did an M2 internship in our team, where he was advised by François Pottier, with precious help from Jacques-Henri Jourdan (CNRS).

Paulo verified a short but particularly subtle piece of code, namely a “local generic solver”, that is, an on-demand, incremental, memoizing least fixed point computation algorithm. This algorithm is a slightly simplified version of `Fix0`, an OCaml library published by François Pottier in 2009.

⁰<https://gitlab.inria.fr/fpottier/fix>

The specification of this algorithm is simple: the solver computes the optimal least fixed point of a system of monotone equations. Although the solver relies on mutable internal state for memoization and for “spying”, a form of dynamic dependency discovery, no side effects are mentioned in the specification. The challenge is precisely to formally justify why it is permitted to hide these side effects from the user.

The verification is carried out in Iris, a modern breed of concurrent separation logic. Iris is embedded in Coq, so the proof is machine-checked. The proof makes crucial use of prophecy variables, a novel feature of Iris. Auxiliary contributions include a restricted infinitary conjunction rule for Iris and a specification and proof of Longley’s “modulus” function, an archetypical example of spying.

This paper [13] has been accepted for presentation at the conference POPL 2020, which will take place in New Orleans in January 2020.

6.2.5. Formal reasoning about asymptotic complexity

Participants: Armaël Guéneau, Arthur Charguéraud [Inria team Camus], François Pottier, Jacques-Henri Jourdan [CNRS].

For several years, Armaël Guéneau, Arthur Charguéraud, François Pottier have been investigating the use of Separation Logic, extended with Time Credits, as an approach to the formal verification of the time complexity of OCaml programs. In 2018 and 2019, in collaboration with Jacques-Henri Jourdan, Armaël has worked on a more ambitious case study, namely a state-of-the-art incremental cycle detection algorithm, whose amortized complexity analysis is nontrivial. Armaël has proposed an improved and simplified algorithm and has carried out a machine-checked proof of its complexity. Furthermore, the verified algorithm has been released and is now used in production inside the Dune build system for OCaml. A paper has been published and presented at the International Conference on Interactive Theorem Proving (ITP 2019) [15]. A more detailed version of these results appears in Armaël Guéneau’s dissertation [11], which was defended on December 16, 2019.

6.2.6. TLA+

Participants: Damien Doligez, Leslie Lamport [Microsoft Research], Ioannis Filippidis, Stephan Merz [Inria team VeriDis].

Damien Doligez is the 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 [25], and to build tools for writing TLA+ specifications and mechanically checking the proofs.

We have made a bug-fix release of TLAPS (version 1.4.4). In parallel, we are working on adding features for dealing with temporal properties, that is, fairness and liveness. We have implemented support for the `ENABLED` operator and the action composition operator in TLA+ proofs. This support is still experimental, but we hope to release a new version of TLAPS next year with these features.

6.3. Shared-memory concurrency

6.3.1. Instruction fetch in the ARMv8 architecture

Participants: Luc Maranget, Peter Sewell [University of Cambridge], Ben Simmer [University of Cambridge].

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 is taking 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 initial focus on devices.

Luc Maranget participates in project REMS, for *Rigorous Engineering for Mainstream Systems*, an EPSRC project led by Peter Sewell. This year Luc Maranget took part in a research effort that resulted in a paper entitled *ARMv8-A system semantics: instruction fetch in relaxed architectures*. This paper has been accepted for presentation at ESOP 2020. This paper introduces a robust model of instruction fetch and cache maintenance, a central aspect of a processor system’s semantics, for ARMv8-A. Luc Maranget specifically extended the **litmus** and **diy** test generators so as to account for self-modifying code. He also performed part of the experiments that support the instruction fetch model.

6.3.2. An ARMv8 mixed-size memory model

Participants: Luc Maranget, Jade Alglave [ARM Ltd & University College London].

Jade Alglave and Luc Maranget have completed their work on a mixed-size version of the ARMv8 memory model. This model builds on the `aarch64.cat` model authored by Will Deacon (ARM Ltd). The model is now ready, and a paper has been written. They hope to work around certain intellectual property restrictions and to submit this paper for publication next year.

6.3.3. Work on diy

Participants: Luc Maranget, Jade Alglave [ARM Ltd & University College London], Antoine Hacquard.

The **diy** suite (for “Do It Yourself”) provides a set of tools for testing shared memory models: the **litmus** tool for running tests on hardware, various generators for producing tests from concise specifications, and **herd**, a memory model simulator. Tests are small programs written in x86, Power, ARM, generic (LISA) assembler, or a subset of the C language that can thus be generated from concise specifications, run on hardware, or simulated on top of memory models. Test results can be handled and compared using additional tools. On distinctive feature of our system is Cat, a domain-specific language for memory models.

This year, new synchronisation primitives and instructions were added to various models. Some sizable developments occurred that facilitate the integration of mixed-size models into **herd**: a default definition of the same-instruction relation, which allows using mixed-size models on all tests; an automatic adjustment of the machine’s elementary granularity, which facilitates massive testing; and the addition of equivalence classes and relations on them as basic values, which extends the expressiveness of Cat to some abstract mixed-size models.

During a 3-month internship, Antoine Hacquard (an EPITA second-year student) extended the complete tool suite to handle a new target, namely X86_64. The addition of this new target significantly enhances the **diy** tool suite, as X86_64 is a very popular architecture. Moreover, Antoine Hacquard implemented all memory access instructions for all sizes (from byte to quadword), which enabled us to design a mixed-size TSO model for this very popular architecture.

6.3.4. Unifying axiomatic and operational weak memory models

Participants: Quentin Ladeveze, Jean-Marie Madiot, Jade Alglave [ARM Ltd & University College London], Simon Castellan [Imperial College London].

Modern multi-processors optimize the running speed of programs using a variety of techniques, including caching, instruction reordering, and branch speculation. While those techniques are perfectly invisible to sequential programs, such is not the case for concurrent programs that execute several threads and share memory: threads do not share at every point in time a single consistent view of memory. A *weak memory model* offers only weak consistency guarantees when reasoning about the permitted behaviors of a program. Until now, there have been two kinds of such models, based on different mathematical foundations: axiomatic models and operational models.

Axiomatic models explicitly represent the dependencies between the program and memory actions. These models are convenient for causal reasoning about programs. They are also well-suited to the simulation and testing of *hardware* microprocessors.

Operational models represent program states directly, thus can be used to reason on programs: program logics become applicable, and the reasoning behind nondeterministic behavior is much clearer. This makes them preferable for reasoning about *software*.

Jean-Marie Madiot has been collaborating with weak memory model expert Jade Alglave and concurrent game semantics researcher Simon Castellan in order to unify these styles, in a way that attempts to combine the best of both approaches. The first results are a formalisation of TSO-style architectures using partial-order techniques similar to the ones used in game semantics, and a proof of a stronger-than-state-of-art “data-race freedom” theorem: well-synchronised programs can assume a strong memory model.

Since October 2019, Luc Maranget and Jean-Marie Madiot are advising a PhD candidate, Quentin Ladeveze. His goal is to further generalize and formalize weak memory models. This involves reasoning about linearizations of interdependent acyclic relations.

This is a first step towards tractable verification of concurrent programs, combining software verification using concurrent program logics, in the top layer, and hardware testing using weak memory models, in the bottom layer. Our hope is to leave no unverified gap between software and hardware, even (and especially) in the presence of concurrency.

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 to the research program (see before) and the research projects (see after):

- Advanced primitives for privacy in the cloud
- Efficient functional encryption
- Attribute and predicate encryption schemes
- New primitives for efficient anonymous authentication
- Applications to machine learning
- Blockchain protocols
- Searchable Encryption

KOPERNIC Team

7. New Results

7.1. Uniprocessor Mixed-Criticality Real-Time Scheduling

In the context of the FUI CEOS project 8.1.1.1, last two years we transformed the free software program PX4, which performs the autopilot of the CEOS drone, in a graph of hard real-time tasks. This transformation was intended to achieve a schedulability analysis guaranteeing the autopilot is able to perform safety critical missions since its behaviour is deemed to be hard real-time, i.e., all deadlines of all tasks are satisfied. It is worth noting that the autopilot is one of the most important programs of the drone since it maintains its stability not only during hover phases but also during automatic flight missions from one GPS point to another. This transformation resulted in a "real-time autopilot" that we called PX4-RT.

For the first version of PX4-RT we chose, as periods, the periods used in the original version of PX4 which was not hard real-time as we shown last year. Then, since these periods was inherited from an automatic control analysis achieved by initial designers of PX4 in a non hard real-time context, we had to determine the right combination of periods of tasks, allowing on the one hand to correctly control the drone, and on the other hand, using a schedulability analysis, to satisfy all the deadlines. In order to achieve this goal, we used a hardware in the loop simulation (HitL) which simulates only the sensors and the actuators, whereas the PX4-RT program runs on the Pixhawk board based on an ARM Cortex-M4 uniprocessor. Eventually, we determined some period combinations that fit our needs, other combinations did not allow the drone to follow correctly the given mission or resulted in a crash. Moreover, we verified that all the right combinations led to a schedulable set of tasks, meaning that corresponding versions of PX4-RT were hard real-time. Finally, we used with success the best combination of periods to run PX4-RT on the real drone of CEOS during a simple flight. Of course, we plan to achieve numerous realistic flights planned in the three industrial use cases of the CEOS project.

In addition to this study intended to determine the right combination of periods, we addressed two other issues. In the first one we tried to decrease the worst case execution times (WCET) of tasks in order to increase the schedulability ratio. Such decrease allows to add on the same processor new tasks presently executed on other processors, e.g., mission planning, fault tolerance, etc. Since we found out that the Kalman filter had the largest measured execution time of all the tasks, we studied the Kalman filter algorithm implemented in PX4 to decrease its WCET. We suppressed the two states of the Kalman Filter corresponding to the wind speed estimation since our drone do not have a sensor measuring this speed. Then, we suppressed the three states of the Kalman filter corresponding to the accelerometer bias whose standard deviation was close to zero. Each of these modifications brought an improvement of 15 percent in term of largest measured execution time without decreasing the performances of the drone. In the second issue we started a theoretical study about relations between the stability of a set of automatic control laws and the schedulability of the corresponding set of real-time tasks. In the literature some results exist about one control law corresponding to one real-time task. To the best of our knowledge there is no result for a set of control laws that exchange data.

Finally, we deeply studied NuttX the real-time operating system used presently to support PX4 and PX4-RT autopilot programs. Indeed, we plan to modify the scheduler of this operating system in order to manage real-time tasks more safely. In order to do that we will draw inspiration from the technique proposed in our time triggered offline scheduler that accounts for the preemption and scheduler cost [14].

7.2. Multicore processor graph tasks scheduling

Due to widespread of multicore processors on embedded and real-time systems, we concentrate our work on the study of the schedulability of real-time tasks with precedence constraints on such processors. We consider preemptive fixed-priority scheduling policies. First, we have proposed a response time analysis

for directed acyclic graphs task model with non-probabilistic execution time and preemptive fixed-priority scheduling policy [10]. Our response time analysis improves importantly the state of the art analyses, while allowing scalable extensions for response time analysis of tasks with worst case execution times described by probability distributions. We extend this response time analysis to similar task model with probabilistic worst case execution time with the advantage of providing efficient results also for task model with non-probabilistic worst case execution times. Our response time analysis is based on iterative equations which offer run-time enhancement compared to existing work [21] requesting the resolution of complex MILP optimization problem. In addition, we have defined priority on sub-task level enhancing the schedulability and reducing the worst-case response time. The proposed priority assignment algorithm is adapted for the studied task model and it outperforms several state-of-the art methods. We have also proposed a partitioning heuristic that assigns each sub-task to a given core. This heuristic takes into consideration communication delays between sub-tasks inside the same graph in order to minimize the communication while balancing different cores load and maximizing possible parallelism. The proposed heuristics and response time analysis (RTA) are validated on randomly generated task sets and on the PX4-RT drone autopilot programs developed by Kopernic team in FR FUI21 CEOS project.

7.3. Power consumption of probabilistic real-time systems

Energy consumption on real-time systems is a crucial problem nowadays as these systems are becoming complex and are expected to deliver more and more functionalities. At the same time, while the processing demand increases, the vast majority of these systems are powered by batteries and are deployed in hazardous environments making their maintenance difficult and impractical. Existing works on energy consumption and real-time systems are often based on a technique called Dynamic Voltage and Frequency Scaling (DVFS). The principle of this technique is to reduce the frequency of the processor in order to lower its input voltage, consequently reducing the energy required to power the processor. Nevertheless, by reducing the frequency of the processor, programs tend to take more time to complete their execution. In the context of real-time systems, programs need to finish their execution before a given deadline. Therefore, the goal of DVFS techniques is to derive proper frequencies that minimize energy consumption and still ensure that all deadlines of all the programs will be respected. Works carried during this postdoc are twofold. The first contribution consisted in observing how the Worst-Case Execution Time (WCET) of programs varies with regards to the frequency of the processor. Many existing works have considered that the WCET is completely scalable, i.e., a simple factor can be applied to derive a new WCET under a different frequency setup. Nevertheless, researchers have recognize that this hypothesis may be too optimistic since other components, that do not run at the same speed as the processor, e.g., the memory, are used by programs. We derived an experimental setup to observe how the execution of programs varied by setting different frequencies on the processor and the memory. We measured CPU cycles and execution times and it was clear from our experiments that the theoretical speedup bound that should be achieved when the processor is running at its maximum speed is never achieved. We also observed, that DVFS techniques could also be applied to the memory of the system, since some programs do not perform many memory request. Our experiments led to a short paper accepted for the Work-in-Progress session of the 40th Real-Time System Symposium. The paper also introduced the task model that will be used as a basis of the next contribution of the postdoc. This next contribution consists in developing RTA techniques for probabilistic real-time systems in order to derive hardware frequency setups. The inclusion of probabilistic real-time system is motivated by the ever-increasing demand of functionalities for this type of systems. To the best of our knowledge, DVFS techniques in conjunction with probabilistic real-time systems have never been studied. The solution to this optimization problem is ongoing work while preparing the submission of first results beginning of February 2020.

7.4. Data-oriented scheduling approaches

We consider the scheduling problem of tasks using an inter-task communication model based on a circular buffer, which eases the data consistency between tasks [13], [12]. The tasks are scheduled on one processor by a fixed priority preemptive scheduling algorithm and they have implicit deadlines. We provide a formal

method calculating the optimal size for each of the buffers while ensuring data consistency, i.e., it is required that a buffer slot is accessed for reading the input data. This later slot will never be used by the producer task to write new data before the execution completion of the instances of all consumers that are currently reading from this slot. As a second contribution, we provide an analytical characterization of the temporal validity and reachability properties of the data flowing in between communicating tasks. These two properties are characterized by considering both tasks execution and data propagation orders. Moreover, we assume that a task instance reads all its inputs data at its activation time and writes back the output data at the completion time where this data becomes immediately available for consumption. Given that, they may be several data samples available in the buffer, we say that a data sample is fresh or temporal valid if, since the time instant it is produced, its producer has not completed another execution. Given that, we use buffers whose size may be larger than one, it is obvious that the consumer task will not implicitly know which data is temporally valid. In order to use the data that reflects the current status of the system environment (valid data), we introduce a novel parameter; the sub-sampling rate used within two scheduling algorithms. These scheduling algorithms ensure the data consistency and temporal validity, while deadlines are met.

OURAGAN Project-Team

7. New Results

7.1. Certified non-conservative tests for the structural stability of discrete multidimensional systems

In [18], we present new computer algebra based methods for testing the structural stability of n -D discrete linear systems (with n at least 2). More precisely, we show that the standard characterization of the structural stability of a multivariate rational transfer function (namely, the denominator of the transfer function does not have solutions in the unit polydisc of \mathbb{C}^n) is equivalent to the fact that a certain system of polynomials does not have real solutions. We then use state-of-the-art computer algebra algorithms to check this last condition, and thus the structural stability of multidimensional systems.

7.2. Computing period matrices and the Abel-Jacobi map of superelliptic curves

In [24], we present an algorithm for the computation of period matrices and the Abel-Jacobi map of complex superelliptic curves given by an equation $y^m = f(x)$. It relies on rigorous numerical integration of differentials between Weierstrass points, which is done using Gauss method if the curve is hyperelliptic ($m = 2$) or the Double-Exponential method. The algorithm is implemented and makes it possible to reach thousands of digits accuracy even on large genus curves.

7.3. Voronoi diagram of orthogonal polyhedra in two and three dimensions

Voronoi diagrams are a fundamental geometric data structure for obtaining proximity relations. In [28], we consider collections of axis-aligned orthogonal polyhedra in two and three-dimensional space under the max-norm, which is a particularly useful scenario in certain application domains. We construct the exact Voronoi diagram inside an orthogonal polyhedron with holes defined by such polyhedra. Our approach avoids creating full-dimensional elements on the Voronoi diagram and yields a skeletal representation of the input object. We introduce a complete algorithm in 2D and 3D that follows the subdivision paradigm relying on a bounding-volume hierarchy; this is an original approach to the problem. The complexity is adaptive and comparable to that of previous methods. Under a mild assumption it is $O(n/\Delta)$ in 2D or $O(n\alpha^2/\Delta^2)$ in 3D, where n is the number of sites, namely edges or facets resp., Δ is the maximum cell size for the subdivision to stop, and α bounds vertex cardinality per facet. We also provide a numerically stable, open-source implementation in Julia, illustrating the practical nature of our algorithm.

7.4. A symbolic computation approach towards the asymptotic stability analysis of differential systems with commensurate delays

In [30], the work aims at studying the asymptotic stability of retarded type linear differential systems with commensurate delays. Within the frequency-domain approach, it is well-known that the asymptotic stability of such a system is ensured by the condition that all the roots of the corresponding quasipolynomial have negative real parts. A classical approach for checking this condition consists in computing the set of critical zeros of the quasipolynomial, i.e., the roots (and the corresponding delays) of the quasipolynomial that lie on the imaginary axis, and then analyzing the variation of these roots with respect to the variation of the delay. Following this approach, based on solving algebraic systems techniques, we propose a certified and efficient symbolic-numeric algorithm for computing the set of critical roots of a quasipolynomial. Moreover, using recent algorithmic results developed by the computer algebra community, we present an efficient algorithm for the computation of Puiseux series at a critical zero which allows us to finely analyze the stability of the system with respect to the variation of the delay. Explicit examples are given to illustrate our algorithms.

7.5. On the computation of stabilizing controllers of multidimensional systems

In [25], we consider the open problem consisting in the computation of stabilizing controllers of an internally stabilizable MIMO multidimensional system. Based on homological algebra and the so-called *Polydisk Nullstellensatz*, we propose a general method towards the explicit computation of stabilizing controllers. We show how the homological algebra methods over the ring of structurally stable SISO multidimensional transfer functions can be made algorithmic based on standard Gröbner basis techniques over polynomial rings. The problem of computing stabilizing controllers is then reduced to the problem of obtaining an effective version of the Polydisk Nullstellensatz which, apart from a few cases, stays open and will be studied in forthcoming publications.

7.6. Algebraic aspects of the exact signal demodulation problem

In [29], we introduce a general class of problems originating from gearbox vibration analysis. Based on a previous work where demodulation was formulated as a matrix approximation problem, we study the specific case applicable to amplitude and phase demodulation. This problem can be rewritten as a polynomial system. Based on algebraic methods such as linear algebra and homological algebra, we focus on the characterization of the problem and solve it in the noise-free case.

7.7. General closed-form solutions of the position self-calibration problem

The work in [36] investigates the anchors and sources position self-calibration problem in the 3D space based on range measurements and without any prior restriction on the network configuration. Using a well known low-rank property of Euclidean distance matrices, we first reduce the problem to finding 12 unknowns ascribed in a 3×3 transformation matrix and a 3×1 translation vector. In order to estimate them, we then introduce a polynomial parametrization with 9 unknowns that are estimated by solving a linear system. Afterwards, we identify an intrinsic matrix polynomial system that encodes the solution set of the problem and provide a direct method for solving it. The resulting procedure is simple and straightforward to implement using standard numerical tools. We also show that closed-form solutions can always be obtained when the reference frame is fixed. This is illustrated by adopting reference frames from the literature and by introducing a triangular reference frame whose constraints are imposed only on one position set (anchor or source). Experimental results on synthetic and real sound data show that the proposed closed-form solutions efficiently solve the position self-calibration problem.

7.8. Certified lattice reduction

Quadratic form reduction and lattice reduction are fundamental tools in computational number theory and in computer science, especially in cryptography. The celebrated Lenstra–Lenstra–Lovász reduction algorithm (so-called LLL) has been improved in many ways through the past decades and remains one of the central methods used for reducing integral lattice basis. In particular, its floating-point variants—where the rational arithmetic required by Gram–Schmidt orthogonalization is replaced by floating-point arithmetic—are now the fastest known. However, the systematic study of the reduction theory of real quadratic forms or, more generally, of real lattices is not widely represented in the literature. When the problem arises, the lattice is usually replaced by an integral approximation of (a multiple of) the original lattice, which is then reduced. While practically useful and proven in some special cases, this method doesn't offer any guarantee of success in general. In [22], we present an adaptive-precision version of a generalized LLL algorithm that covers this case in all generality. In particular, we replace floating-point arithmetic by Interval Arithmetic to certify the behavior of the algorithm. We conclude by giving a typical application of the result in algebraic number theory for the reduction of ideal lattices in number fields.

7.9. Using Maple to analyse parallel robots

In [27], we present the SIROPA Maple Library which has been designed to study serial and parallel manipulators at the conception level. We show how modern algorithms in Computer Algebra can be used to study the workspace, the joint space but also the existence of some physical capabilities w.r.t. to some design parameters left as degree of freedom for the designer of the robot.

7.10. On the effective computation of stabilizing controllers of 2D systems

In [26], we show how stabilizing controllers for 2D systems can effectively be computed based on computer algebra methods dedicated to polynomial systems, module theory and homological algebra. The complete chain of algorithms for the computation of stabilizing controllers, implemented in Maple, is illustrated with an explicit example.

7.11. Updating key size estimations for pairings

Recent progress on NFS imposed a new estimation of the security of pairings. In [15], we study the best attacks against some of the most popular pairings. It allows us to propose new pairing-friendly curves of 128 bits and 192 bits of security.

7.12. Matrix formulae for Resultants and Discriminants of Bivariate Tensor-product Polynomials

The construction of optimal resultant formulae for polynomial systems is one of the main areas of research in computational algebraic geometry. However, most of the constructions are restricted to formulae for unmixed polynomial systems, that is, systems of polynomials which all have the same support. Such a condition is restrictive, since mixed systems of equations arise frequently in many problems. Nevertheless, resultant formulae for mixed polynomial systems is a very challenging problem. In [19], we present a square, Koszul-type, matrix, the determinant of which is the resultant of an arbitrary (mixed) bivariate tensor-product polynomial system. The formula generalizes the classical Sylvester matrix of two univariate polynomials, since it expresses a map of degree one, that is, the elements of the corresponding matrix are up to sign the coefficients of the input polynomials. Interestingly, the matrix expresses a primal-dual multiplication map, that is, the tensor product of a univariate multiplication map with a map expressing derivation in a dual space. In addition we prove an impossibility result which states that for tensor-product systems with more than two (affine) variables there are no universal degree-one formulae, unless the system is unmixed. Last but not least, we present applications of the new construction in the efficient computation of discriminants and mixed discriminants.

7.13. Separation bounds for polynomial systems

In [21], we rely on aggregate separation bounds for univariate polynomials to introduce novel worst-case separation bounds for the isolated roots of zero-dimensional, positive-dimensional, and over-terminated polynomial systems. We exploit the structure of the given system, as well as bounds on the height of the sparse (or toric) resultant, by means of mixed volume, thus establishing adaptive bounds. Our bounds improve upon Canny's Gap theorem [9]. Moreover, they exploit sparseness and they apply without any assumptions on the input polynomial system. To evaluate the quality of the bounds, we present polynomial systems whose root separation is asymptotically not far from our bounds. We apply our bounds to three problems. First, we use them to estimate the bitsize of the eigenvalues and eigenvectors of an integer matrix; thus we provide a new proof that the problem has polynomial bit complexity. Second, we bound the value of a positive polynomial over the simplex: we improve by at least one order of magnitude upon all existing bounds. Finally, we asymptotically bound the number of steps of any purely subdivision-based algorithm that isolates all real roots of a polynomial system.

7.14. On the maximal number of real embeddings of minimally rigid graphs in \mathbb{R}^2 , \mathbb{R}^3 and S^2

Rigidity theory studies the properties of graphs that can have rigid embeddings in a euclidean space \mathbb{R}^d or on a sphere and other manifolds which in addition satisfy certain edge length constraints. One of the major open problems in this field is to determine lower and upper bounds on the number of realizations with respect to a given number of vertices. This problem is closely related to the classification of rigid graphs according to their

maximal number of real embeddings. In [17], we are interested in finding edge lengths that can maximize the number of real embeddings of minimally rigid graphs in the plane, space, and on the sphere. We use algebraic formulations to provide upper bounds. To find values of the parameters that lead to graphs with a large number of real realizations, possibly attaining the (algebraic) upper bounds, we use some standard heuristics and we also develop a new method inspired by coupler curves. We apply this new method to obtain embeddings in \mathbb{R}^3 . One of its main novelties is that it allows us to sample efficiently from a larger number of parameters by selecting only a subset of them at each iteration. Our results include a full classification of the 7-vertex graphs according to their maximal numbers of real embeddings in the cases of the embeddings in \mathbb{R}^2 and \mathbb{R}^3 , while in the case of S^2 we achieve this classification for all 6-vertex graphs. Additionally, by increasing the number of embeddings of selected graphs, we improve the previously known asymptotic lower bound on the maximum number of realizations.

7.15. Multilinear Polynomial Systems: Root Isolation and Bit Complexity

In [20], we exploit structure in polynomial system solving by considering polynomials that are linear in subsets of the variables. We focus on algorithms and their Boolean complexity for computing isolating hyperboxes for all the isolated complex roots of well-constrained, unmixed systems of multilinear polynomials based on resultant methods. We enumerate all expressions of the multihomogeneous (or multigraded) resultant of such systems as a determinant of Sylvester-like matrices, aka generalized Sylvester matrices. We construct these matrices by means of Weyman homological complexes, which generalize the Cayley-Koszul complex. The computation of the determinant of the resultant matrix is the bottleneck for the overall complexity. We exploit the quasi-Toeplitz structure to reduce the problem to efficient matrix-vector multiplication, which corresponds to multivariate polynomial multiplication, by extending the seminal work on Macaulay matrices of Canny, Kaltofen, and Yagati [9] to the multi-homogeneous case. We compute a rational univariate representation of the roots, based on the primitive element method. In the case of 0-dimensional systems we present a Monte Carlo algorithm with probability of success $1 - 1/2^n$, for a given $\eta \geq 1$, and bit complexity $O_B(n^2 D^{4+\epsilon}(n^{N+1} + \tau) + n D^{2+\epsilon} \eta (D + \eta))$ for any $\epsilon > 0$, where n is the number of variables, D equals the multilinear Bézout bound, N is the number of variable subsets, and τ is the maximum coefficient bitsize. We present an algorithmic variant to compute the isolated roots of overdetermined and positive-dimensional systems. Thus our algorithms and complexity analysis apply in general with no assumptions on the input.

PARKAS Project-Team

6. New Results

6.1. Efficiently Subtyping Union Types

Participant: Francesco Zappa Nardelli.

Julia is a programming language recently designed at MIT to support the needs of the scientific community. Julia occupies a unique position in the design landscape, it is a dynamic language with no type system, yet it has a surprisingly rich set of types and type annotations used to specify multimethod dispatch. The types that can be expressed in function signatures include parametric union types, covariant tuple types, parametric user-defined types with single inheritance, invariant type application, and finally types and values can be reified to appear in signatures. In 2017 with Vitek we started a research project to study the design and the pragmatic use of the Julia language, and formalised the Julia subtyping algorithm. In 2018 we have pursued this study, and we have proved correct the clever and space efficient algorithm relied upon by the Julia runtime. This has been published in [17].

6.2. Fast and reliable unwinding via DWARF tables

Participants: Theophile Bastian, Rémy Oudin, Francesco Zappa Nardelli.

DWARF is a widely-used debugging data format. DWARF is obviously relied upon by debuggers, but it plays an unexpected role in the runtime of high-level programming languages and in the implementation of program analysis tools. The debug information itself can be pervaded by subtle bugs, making the whole infrastructure unreliable. In this project we are investigating techniques and tools to perform validation and synthesis of the DWARF stack unwinding tables, to speedup DWARF-based unwinding, as well as exploring adventurous projects that can be built on top of reliable DWARF information.

We have built a tool that can validate DWARF unwind tables generated by mainstream compilers; the approach is effective, we found a problem in Clang table generation and several in GLIBC inline-assembly snippets. We also designed and implemented a tool that can synthesise DWARF unwind tables from binary that lacks them (e.g. because the compiler did not generate them - immediate applications: JITs assembly, inline assembly, ...). Additionally we have designed and implemented an ahead-of-time compiler of DWARF unwind tables to assembly, and an ad-hoc unwinder integrated with the defacto standard unwinder libuwind. It can speed up unwinding by a factor between 12x and 25x (depending on application), with a 2.5x size overhead for unwind information.

This work has been published in [13].

6.3. Verified compilation of Lustre

Participants: Timothy Bourke, Lélío Brun, Paul Jeanmaire, Marc Pouzet.

Vélus is a compiler for a subset of LUSTRE and SCADE that is specified in the Coq [28] Interactive Theorem Prover (ITP). It integrates the CompCert C compiler [34], [29] to define the semantics of machine operations (integer addition, floating-point multiplication, etcetera) and to generate assembly code for different architectures. The research challenges are to

- to mechanize, i.e., put into Coq, the semantics of the programming constructs used in modern languages for MBD;
- to implement compilation passes and prove them correct;
- to interactively verify source programs and guarantee that the obtained invariants also hold of the generated code.

This year we created a website for the project (<https://velus.inria.fr>) and made an initial release under an Inria non-commercial license (<https://github.com/Inria/velus>). T. Bourke’s JCJC (“Jeune Chercheuse Jeune Chercheur”) project *FidelR* was accepted for funding by the ANR: it aims to develop ITP-based techniques for treating state machines and interactive program verification. We also made progress on the compilation of the modular reset construct, the treatment of (non-normalized) Lustre, and our longer term goal of strengthening the main correctness theorem. These results are detailed below.

6.3.1. *Compiling the modular reset construct.*

In the original LUSTRE language, the only way to reset the internal state of an instantiated function is to propagate and test explicit reset signals. Later languages, like LUCID SYNCHRONE and SCADE, provide a construct for resetting an instance modularly (it works for any function) and efficiently (testing occurs only at the point of instantiation). Last year we showed how to encode the semantics of this construct in Coq. This year we focused on its compilation and the associated proof of correctness. We designed and implemented a new intermediate language that exposes different *step* and *reset* actions on node instances. This language facilitates the optimization of conditional statements in the generated code and permits the transformation to imperative code and its proof of correctness to be treated in two steps: one to introduce named memories and another to fix the sequential order of execution. This work forms the core of L. Brun’s thesis, to be defended early next year, and an article accepted at the ACM SIGBED international conference on Principles of Programming Languages (POPL 2020).

6.3.2. *Non-normalized Lustre.*

Our previous work has focused on a subset of “normalized” programs where the form of expressions and equations is constrained to facilitate the compilation. We have generalized the definitions of syntax and semantics in our prototype compiler to accept non-normalized programs. This included simplifying and generalizing the formalization of clocks presented in [20]. With P. Jeanmaire (M2 internship), we have implemented a compilation pass to translate normalized programs from one syntactical form to another. The main challenge was to formally prove an alignment property (signals are present iff their clocks are true) that had been assumed until now. The proof is finished except for the inductive case for the reset construct which we hope to complete soon.

6.3.3. *Strengthening the correctness theorem.*

The current correctness theorem assumes that an accepted program can be given a semantics in terms of the mechanized model. It should be possible to prove this fact for programs that pass the initial type-checking and clock-checking algorithms, that can be scheduled, and which never invoke an undefined operation (such as a division by zero). We made good progress on this problem by defining an interpreter for normalized Lustre programs and showing that the results it calculates satisfy the semantic predicates. This initial work gives some useful insights into how to proceed. We presented it at the Synchron 2019 workshop.

6.4. Specifying multi-clock Lustre programs

Participants: Timothy Bourke, Guillaume Iooss, Baptiste Pauget, Marc Pouzet.

It is sometimes desirable to compile a single synchronous language program into multiple tasks for execution by a real-time operating system. We have been investigating this question from three different perspectives.

6.4.1. *Harmonic clocks*

We studied the extension of a synchronous language with periodic harmonic clocks based on the work of Mandel et al. [31], [37], [32], [35], [36] on n-synchrony and the extension proposed by Forget et al. [33]

Mandel et al. considered a language with periodic clocks expressed as ultimately periodic binary sequences. The decision procedures (equality, inclusion, precedence) for such an expressive language can be very costly. It is thus sometimes useful to apply an envelope-based abstraction, that is, one where sets of clocks are represented by a rational slope and an interval. Forget considered simpler “harmonic” clocks. His decision procedures coincide with those for the envelope-based abstraction but without any loss of information. During his M2 internship, B. Pauget continued this line of work by extending the input language of the Vélus Lustre compiler with harmonic clocks. This work was the starting point for the proposal of a new intermediate language for a synchronous compiler that is capable of exploiting clock information to apply aggressive optimizations and generate parallel code.

6.4.2. *New Intermediate Language MObc (Multi Object Code)*

This intermediate language is reminiscent of the intermediate Obc language used in the Vélus and Heptagon compiler, but with some important differences and new features. MObc permits a synchronous function to be represented as a set of named state variables and possibly nested blocks with a partial ordering which express the way blocks can and must be called. In comparison, Obc represents a synchronous function as a set of state variables and a transition function that is itself written in a sequential language. Each block comprises a set of equations in Single Static Assignment (SSA) form, that is, exactly one equation per variable, so as to simplify the implementation of a number of classic optimizations (for example, constant propagation, inlining, common sub-expression elimination, code specialisation). Then, every block is translated into a step function (e.g., a C function). This intermediate language has been designed to facilitate the generation of code for a real-time OS and a multi-core target. This work exploits two older results: the article of Caspi et al. [30] that introduces an object representation for synchronous nodes and a “scheduling policy” that specifies how their methods may be called, and; the work of Pouzet et al. [38] on the calculation of input/output relations to merge calculations. We are preparing an article on this subject.

6.4.3. *Clocking constraints, communication latencies, and constraint solving*

In this approach, the top-level node of a Lustre program is distinguished from inner nodes. It may contain special annotations to specify the triggering and other details of node instances from which separate “tasks” are to be generated. Special operators are introduced to describe the buffering between top-level instances. Notably, different forms of the `fby` and `current` operators are provided. Some of the operators are under-specified and a constraint solver is used to determine their exact meaning, that is, whether the signal is delayed by zero, one, or more cycles of the receiving clock, which depends on the scheduling of the source and destination nodes. Scheduling is formalized as a constraint solving problem based on latency constraints between some pairs of input/outputs that are specified by the designer. G. Iooss has been prototyping these ideas in the academic Heptagon compiler.

This work is funded by a direct industrial contract with Airbus. In collaboration (this year) with Michel Angot Vincent Bregeon Jean Souyris (Airbus, R&D) and Matthieu Boitrel (Airbus BE).

6.5. The Zelus Language

Participants: Timothy Bourke, Ismail Lakhim-Bennani, Marc Pouzet.

Zelus is our laboratory to experiment our research on programming languages for hybrid systems. It is devoted to the design and implementation of systems that may mix discrete-time/continuous-time signals and systems between those signals. It is essentially a synchronous language reminiscent of Lustre and Lucid Synchronic but with the ability to define functions that manipulate continuous-time signals defined by Ordinary Differential Equations (ODEs). The language is functional in the sense that a system is a function from signals to signals (not a relation). It provides some features from ML languages like higher-order and parametric polymorphism as well as dedicated static analyses.

This year, we have pursued our work on the design, semantics and implementation of hybrid modeling language, in particular the treatment of Differential Algebraic Equations (DAEs) [23].

6.5.1. Compiler Internals: Static Typing and Compiler Organisation

The distribution with manual and examples is distributed at <http://zelus.di.ens.fr> (only Version 1, in binary form). Version 2 (the current active branch) is available in source form on Inria GitLab <https://gitlab.inria.fr/parkas/zelus>, on simple demand.

Several new experimentations have been done this year, in particular on the type system and an extensive rewriting of some compilation internals to simplify the code and make the generated code more shorter (in size) and more efficient.

6.5.2. Co-simulation as Function Lifting

Hybrid models in Zelus (that is, programs that mix discrete and continuous-time signals) are simulated using a single ODE and zero-crossing solvers only. All hybrid modeling languages (e.g., Simulink, Modelica, Ptolemy) act the same way, at least, single solver simulation is the default mechanism.

Its weaknesses are well known: any change of the dynamic, even local, calls for a global reset of the solver, making it slower for that later steps; the mix of a slow and fast signals slows down the whole simulation. Co-simulation is about running several solvers (or instances of the same) at the same time.

We proposed a limited (but useful) manner, by proving a way to internalize the solver to obtain, from a continuous-time function, a synchronous stream function. A preliminary experiment done this year was surprisingly and pleasingly simple to implement in Zelus. It consisted in defining a (higher-order) function *solve* that, given a continuous-time function f returns a stream function $solve\ f$. Given an input stream x and an increasing stream of time horizons h , $solve\ f(x, t)$ returns the stream of approximated values. This function internalizes the ODE solver and the zero-crossing detection mechanism. The overall model is then a purely discrete-time, synchronous model. In particular, classical synchronization protocols between solvers can be programmed in the language itself, hence benefiting from the static checks that track typing, causality and initialization errors, properties that would be more difficult to ensure if programmed directly in C, for example. We think that it is even possible to write a formal synchronous specification of the simulation engine itself, that is, to program the function *solve* directly in Zelus. This experiment on co-simulation gives new insight on the semantics based on non standard analysis that we proposed and, more interestingly, to relate it to the proven and more classical semantics based on super-dense time studied and exploited by Edward Lee.

6.5.3. QSS-based Simulation

Quantized State Systems simulation (QSS) was introduced in the early 2000's by F. Cellier and E. Kofman as an alternative to time-based simulation, which is the dominant approach to ODE/DAE systems simulation.

Rather than linking QSS to Discrete Event Simulation, we have made a preliminary experiment to relate it to Synchronous Programming and its continuous time extension Zelus. Zelus is used to give a formal description of the QSS method that can be executed. We have described the very basic scheme called QSS (or QSS1) for which we can give a Zelus (hence executable) specification. Higher order schemes QSS2, 3, etc. can also be given an Zelus specification. Implicit schemes were also proposed by Kofman for a better handling of stiff systems. Higher order versions of BQSS are nontrivial; they are called LIQSS1, 2, 3, etc. and they can also be specified in Zelus. This preliminary work is done in collaboration with Albert Benveniste (Inria Hycomes, Rennes) and funded by the Modeliscale FUI project.

6.5.4. Property Based Testing of Hybrid Programs

Property-based program testing involves checking an executable specification by running many tests. We build on the work of Georgios Fainekos and Alexandre Donzé, and take inspiration from earlier work by Nicolas Halbwachs, to write a Zelus library of synchronous observers with a quantitative semantics that can be used to specify properties of a system under test. We implemented several optimization algorithms for producing test cases, some of which are gradient-based. To compute the gradients, we use Automatic Differentiation (AD) of the system under test and its specification. Together with François Bidet, we ported

the well-known FADBAD++ library for AD written by Ole Stauning in 1997 to OCaml—the target language of Zélus. Our port is called FADBADml and is now released under an Inria license⁰ and is available on opam.

6.6. Reactive Probabilistic Programming

Participant: Marc Pouzet.

Synchronous languages were introduced to design and implement real-time embedded systems with a (justified) emphasis on determinacy. Yet, they interact with a physical environment that is partially known and are implemented on architectures subject to failures and noise (e.g., channels, variable communication delays or computation time). Dealing with uncertainties is useful online monitoring, learning, statistical testing or to build simplified models for faster simulation. Actual synchronous and languages provide limited support for modeling the non-deterministic behaviors that are omnipresent in embedded systems.

In 2019, we started a new topic on *reactive probabilistic programming* under the initiative of Guillaume Baudart and Louis Mandel (IBM Research, Watson); in collaboration with Erik Atkinson, Michael Carbin and Benjamin Sherman (MIT). We have designed ProbZelus, an extension of Zelus with probabilistic programming constructs. The language makes it possible to describe probabilistic models in interaction with an observable environment. At runtime, a set of inference techniques can be used to learn the distributions of model parameters from observed data. The main results are (1) the design of the ProbZelus compiler, the formalization of the static and dynamic semantics of the language that mixes deterministic and probabilistic components, (2) the design and implementation of inference methods which can be executed with bounded resources, (3) the evaluation of ProbZelus on a set of examples and case studies.

For the moment, ProbZelus is mainly a library of Zelus, with minor changes of the language itself (essentially the type system)⁰ It exploits heavily the higher-order nature of Zelus. E.g., if `f` is a stream function (with type `f: 'a -D-> 'b`) and `x` is a stream (with type `'a`), `Particule.infer f x: 'a -D-> 'b Distribution.t` implements an inference algorithm which computes a distribution for the result of type `'b` with a particule filter algorithm.

A preliminary report describes a part of this work [24] and a presentation at JFLA will be given in January 2020. ProbZelus is available in open source at <https://github.com/IBM/probzelus> since december 2019. Our purpose is to go beyond the library approach with a closer integration of probabilistic constructs and reactive constructs, with dedicated static analyses and compilation techniques to give static guaranties on the result of inference, efficient inference techniques tuned for reactive applications and that ensure execution in bounded time and space; efficient dedicated compilation techniques for probabilistic programs. Finally, the treatment of both discrete-time and continuous-time signals and systems must be investigated (only discrete-time is considered at the moment).

6.7. Identification of matrix operations for Compute-In-Memory architectures from a high-level Machine Learning framework

Participant: Andi Drebes.

Compute-In-Memory (CIM) architectures are capable of performing certain performance-critical operations directly in memory (e.g., matrix multiplications) and represent a promising approach to partially eliminate the bottleneck of traditional von Neumann-based architectures resulting from long-distance communication between main memory and processing units.

In order for applications to benefit from such architectures, their operations must be divided into highly parallel, uniform operations eligible for in-memory computation and control logic that cannot benefit from CIM and that must be carried out by conventional computing devices. It is crucial for this process that as many eligible operations as possible are identified and effectively processed in memory, resulting only in as few computations as possible carried out on the conventional cores.

⁰<https://fadbadml-dev.github.io/FADBADml/>

⁰Yet, higher-order was only used occasionally since then; hence an important implementation effort has been spent this year to make it work well.

The programmability of CIM architectures is a key factor for its overall success. Manual identification of eligible operations and mapping to hardware resources is tedious, error-prone and requires detailed knowledge of the target architecture and therefore does not represent a viable approach to program CIM architectures.

With our partners from the MNEMOSENE project, we have developed a compilation toolchain that unburdens programmers from technical details of CIM architectures by allowing them to express algorithms at a high level of abstraction and that automates parallelization, orchestration and the mapping of operations to the CIM architecture. The solution integrates the Loop Tactics [40] declarative polyhedral pattern recognition and transformation framework into Tensor Comprehensions [39], a framework generating highly optimized kernels for accelerators from an abstract, mathematical notation for tensor operations. The compilation flow performs a set of dedicated optimizations aiming at enabling the reliable detection of computational patterns and their efficient mapping to CIM accelerators.

The results of this work have been submitted to the 10th International Workshop on Polyhedral Compilation Techniques (IMPACT).

6.8. Applying reinforcement learning to improve a branch-and-bound optimizing compiler

Participant: Basile Clement.

Frameworks for image processing, deep learning, etc., work with Directed Acyclic Graphs (DAGs) of computational operators that wrap high-performance libraries. The production of highly optimized, target-specific implementations of the library functions come at a high engineering cost: languages and compilers have failed to deliver performances competitive with expert written code, notably on GPUs and other hardware accelerators. Moreover, library implementations may not offer optimal performance for a specific use case. They may lack inter-operator optimizations and specializations to specific data sizes and shapes.

In his thesis, Ulysse Beaugnon, a former PhD student in the team, proposed to formulate this compilation problem as an optimization research problem using a combination of analytical modeling, experimental search, constraint programming and branch-and-bound optimization techniques. Basile Clement started a PhD to extend this idea, exploring the improvements required to make it fully competitive with handwritten code. In 2019, he evaluated reinforcement learning techniques such as multi-armed bandit schemes to improve the performance and efficiency of the search procedure; extended the analytical model with generic sizes, making it more precise before selecting tiling parameters; and made various improvements to the code generation procedure.

PI.R2 Project-Team

6. New Results

6.1. Effects in proof theory and programming

Participants: Kostia Chardonnet, Emilio Jesús Gallego Arias, Hugo Herbelin, Yann Régis-Gianas, Alexis Saurin, Exequiel Rivas Gadda.

6.1.1. *A theory of effects and resources*

In collaboration with Thomas Letan (ANSSI), Yann Régis-Gianas developed and proved several properties of a simple web server implemented in Coq using FreeSpec. This work will be presented at CPP 2020.

6.1.2. *Call-by-need with probabilistic effects*

As a follow up of Chardonnet's Master 1 internship, Kostia Chardonnet and Alexis Saurin continued investigating call-by-need calculi extended with probabilistic choice and started preliminary discussions with Claudia Faggian.

6.1.3. *Proof-search, algebraically and graphically*

Alexis Saurin worked on proof search in a proof-net scenario, that is proof-net search. A key aspect of proof construction is a management of non-determinism in bottom-up sequent-proof construction, be it when the search succeeds or when facing a failure and the need for backtracking. This is partially dealt with by focussing proof-construction, which reduces drastically the search space while retaining completeness of the resulting proof space (both at the provability level and at the denotational level).

His approach consists in viewing proof-search and sequentialisation as dual aspects of partial proof structures (that is proof nets with open premisses). In particular, he builds on Lafont's parsing criterion to obtain a proof-construction algorithm in which the proof space is not a search tree, as in sequent-calculus, but a dag allowing to share proof-construction paths.

Emilio Jesús Gallego Arias collaborates with Jim Lipton from Wesleyan University on the development of algebraic models for proof search.

6.2. Reasoning and programming with infinite data

Participants: Kostia Chardonnet, Lucien David, Abhishek De, Farzad Jafar-Rahmani, Luc Pellissier, Yann Régis-Gianas, Alexis Saurin.

This theme is part of the ANR project Rapido (see the National Initiatives section) which ended octobre 1st 2019.

6.2.1. *Proof theory of non-wellfounded and circular proofs*

6.2.1.1. *Validity conditions of infinitary and circular proofs*

In collaboration with David Baelde, Amina Doumane and Denis Kuperberg, Alexis Saurin extended the proof theory of infinite and circular proofs for fixed-point logics in various directions by relaxing the validity condition necessary to distinguish sound proofs from invalid ones. The original validity condition considered by Baelde, Doumane and Saurin in CSL 2016 rules out lots of proofs which are computationally and semantically sound and does not account for the cut-axiom interaction in sequent proofs. In the setting of sequent calculus, Alexis Saurin studied together with David Baelde, Amina Doumane and Denis Kuperberg a relaxed validity condition to allow infinite branches to be supported by threads which may leave the infinite branch, visiting other parts of the proofs and bouncing on axioms and cuts. This allows for a much more flexible criterion, inspired from Girard's geometry of interaction. The most general form of this criterion does

not ensure productivity in the sequent calculus due to a discrepancy between the sequential nature of proofs in sequent calculus and the parallel nature of threads. David Baelde, Amina Doumane, Denis Kuperberg and Alexis Saurin provided a slight restriction of the full bouncing validity which grants productivity and validity of the cut-elimination process. This restriction still strictly extends previous notions of validity and is actually expressive enough to be undecidable.

Several directions of research have therefore been investigated from that point:

- Decidability can be recovered by constraining the shapes of bounces (bounding the depth of bounces). They actually exhibited a hierarchy of criteria, all decidable and satisfying the fact that their union corresponds to bouncing validity (which is therefore semi-decidable)
- While the result originally held only for the multiplicative fragment of linear logic, the result was extended to multiplicative and additive linear logic.

Those results are currently submitted.

6.2.1.2. *On the complexity of the validity condition of circular proofs*

Alexis Saurin, together with Rémi Nollet and Christine Tasson, characterised the complexity of deciding the validity of circular proofs. While deciding validity was known to be in PSPACE, they proved that, for $\mu MALL$ proof, it is in fact a PSPACE-complete problem.

The proof is based on a deeper exploration of the connection between thread-validity and the size-change termination principle, a standard tool to prove program termination.

This result has been presented and published at TABLEAUX 2019 [41].

6.2.1.3. *Proof nets for non-wellfounded proofs*

Abhishek De and Alexis Saurin set the basis of the theory of non-wellfounded and circular proofs nets (in the multiplicative setting). Non-wellfounded proof nets, aka infinets, were defined extending Curien's presentation of proof nets allowing for a smooth extension to fixed point logics. The aim of this work is to provide a notion of canonical proof objects for circular proofs free from the irrelevant details of the syntax of the sequent calculus. The first results were published in TABLEAUX 2019 [38] and provide a correctness condition for an infinnet to be sequentialisable in a sequent proof.

The results of the TABLEAUX paper are limited in that they only address the case of proofs with finitely many cuts inferences. Abhishek De and Alexis Saurin are currently investigating, with Luc Pellissier, the general case of infinitely many cut in order to then lift the results from straight thread validity to bouncing thread validity.

6.2.2. *On the denotational semantics of non-wellfounded proofs*

Farzad Jafar-Rahmani started his PhD under the supervision of Thomas Ehrhard and Alexis Saurin in October 2019. His PhD work will focus on the denotational semantics of circular proofs of linear logic with fixed points. After working on the denotational semantics of finitary proofs for linear logic with fixed points (with Kozen rules) during his master, he is currently working at understanding the denotational counterpart of the validity condition of circular proofs.

6.2.3. *Towards inductive and coinductive types in quantum programming languages*

Kostia Chardonnet started his PhD under the supervision of Alexis Saurin and Benoît Valiron in November 2019. Previously, he did his MPRI Master internship under their joint supervision on designing a calculus of reversible programs with inductive and coinductive types. His research focused on extending a languages of type isomorphisms with inductive and coinductive types and understanding the connections of those reversible programs with $\mu MALL$ type isomorphisms and more specifically with $\mu MALL$ focused circular proof isomorphisms. In his PhD, he shall extend this to the case of a quantum programming language with inductive and coinductive data types.

6.2.4. Theory of fixed points in the lambda-calculus

The results of Alexis Saurin in collaboration with Giulio Manzonetto, Andrew Polonsky and Jacob Grue Simonsen, on two long-standing conjectures on fixed points in the λ -calculus – the “fixpoint property” and the “double-fixpoint conjecture” – have now appeared in the Journal of Logic and Computation [34]. The former asserts that every λ -term admits either a unique or an infinite number of β -distinct fixpoints while the second, formulated by Statman, says that there is no fixpoint satisfying $Y\delta = Y$ for $\delta = \lambda y, x.x(yx)$. They proved the first conjecture in the case of open terms and refute it in the case of sensible theories (instead of β). Moreover, they provide sufficient conditions for both conjectures in the general case. Concerning the double-fixpoint conjecture, they propose a proof technique identifying two key properties from which the results would follow, while they leave as conjecture to prove that those actually hold.

6.3. Effective higher-dimensional algebra

Participants: Antoine Allieux, Pierre-Louis Curien, Alen Durić, Eric Finster, Yves Guiraud, Amar Hadzihanović, Cédric Ho Thanh, Matthieu Sozeau.

6.3.1. Rewriting methods in higher algebra

Yves Guiraud has completed a four-year collaboration with Eric Hoffbeck (Univ. Paris 13) and Philippe Malbos (Univ. Lyon 1), whose aim was to develop a theory of rewriting in associative algebras, with a view towards applications in homological algebra. They adapted the known notion of polygraph [64] to higher-dimensional associative algebras, and used these objects to develop a rewriting theory on associative algebras that generalises the two major tools for computations in algebras: Gröbner bases [63] and Poincaré-Birkhoff-Witt bases [100]. Then, they transposed the construction of [14], based on an extension of Squier’s theorem [104] in higher dimensions, to compute small polygraphic resolutions of associative algebras from convergent presentations. Finally, this construction has been related to the Koszul homological property, yielding necessary or sufficient conditions for an algebra to be Koszul. The resulting work was published in *Mathematische Zeitschrift* [32].

Yves Guiraud has written and defended his “Habilitation à diriger des recherches” manuscript, as a survey on rewriting methods in algebra based on Squier theory [26]. The defense was held in June 2019.

Yves Guiraud works with Dimitri Ara (Univ. Aix-Marseille), Albert Burroni, Philippe Malbos (Univ. Lyon 1), François Métayer (Univ. Nanterre) and Samuel Mimram (École Polytechnique) on a reference book on the theory of polygraphs and higher-dimensional categories, and their applications in rewriting theory and homotopical algebra.

Yves Guiraud works with Marcelo Fiore (Univ. Cambridge) on the theoretical foundations of higher-dimensional algebra, in order to develop a common setting to develop rewriting methods for various algebraic structures at the same time. Practically, they aim at a definition of polygraphic resolutions of monoids in monoidal categories, based on the recent notion of n -oid in an n -oidal category. This theory will subsume the known cases of monoids and associative algebras, and encompass a wide range of objects, such as Lawvere theories (for term rewriting), operads (for Gröbner bases) or higher-order theories (for the λ -calculus).

Building on [9], Yves Guiraud is currently finishing with Matthieu Picantin (Univ. Paris Diderot) a work that generalises already known constructions such as the bar resolution, several resolutions defined by Dehornoy and Lafont [73], and the main results of Gaussent, Guiraud and Malbos on coherent presentations of Artin monoids [11], to monoids with a Garside family. This allows an extension of the field of application of the rewriting methods to other geometrically interesting classes of monoids, such as the dual braid monoids.

Still with Matthieu Picantin, Yves Guiraud develops an improvement of the classical Knuth-Bendix completion procedure, called the KGB (for Knuth-Bendix-Garside) completion procedure. The original algorithm tries to compute, from an arbitrary terminating rewriting system, a finite convergent presentation, by adding relations to solve confluence issues. Unfortunately, this algorithm fails on standard examples, like most Artin monoids with their usual presentations. The KGB procedure uses the theory of Tietze transformations, together with Garside theory, to also add new generators to the presentation, trying to reach the convergent Garside

presentation identified in [9]. The KGB completion procedure is partially implemented in the prototype Rewr, developed by Yves Guiraud and Samuel Mimram.

Yves Guiraud has started a collaboration with Najib Idrissi (IMJ-PRG, Univ. Paris Diderot) whose aim is to understand the relation between several different methods known to compute small resolutions of algebras and operads: those based on rewriting methods (Anick, Squier) and those that stem from Koszul duality theory.

6.3.2. Normalisation of monoids

Alen Durić started his Phd thesis (supervised by Yves Guiraud and Pierre-Louis Curien) in October 2019. His work so far has been mostly bibliographical. The goal is to combine methods from rewriting theory (and in particular the method of homotopical completion and reduction developed by Guiraud-Malbos-Mimram) and methods developed by Dehornoy and his coauthors in the study of monoids with Garside families, and by Dehornoy-Guiraud in the study of normalisation for monoids. Alen Durić is currently experimenting with some examples taken from these latter works, with the goal of building coherent presentations for them using the former methods.

6.3.3. Topological aspects of polygraphs

Amar Hadzihasanović joined the team at the end of November 2019, as a one-year postdoc funded by FSMP. He has been working intensively on the study of shapes appropriate for the description of higher cells as needed in various approaches to higher categories and higher structures. Amar Hadzihasanović's project is to recast his ideas in the framework of polygraphs, with the aim of bringing topological insights into the study of higher-dimensional rewriting.

6.3.4. Opetopes

The work of Pierre-Louis Curien, Cédric Ho Thanh and Samuel Mimram on syntactic and type-theoretic presentations of opetopes and opetopic sets has been submitted to a journal, and a short version has been presented at the LICS conference in Vancouver this year [45].

Cédric Ho Thanh, in collaboration with Chaitanya Leena Subramaniam, has defined the notion of “opetopic algebras” that leverages the subtle combinatorics of opetopes. This framework encompasses categories, planar operads, and Loday's combinads over planar trees. They have defined an opetopic nerve functor that fully embeds each category of opetopic algebras into the category of opetopic sets. In particular, they obtain fully faithful opetopic nerve functors for categories and for planar coloured operads. These results have been written up in [51]. This work is the first in a series aimed at using opetopic spaces as models for higher algebraic structures. In particular, the aim is to provide new models for infinity-categories and infinity-operads.

6.3.5. Foundations and formalisation of higher algebra

Antoine Allieux (PhD started in February 2018), Eric Finster, Yves Guiraud and Matthieu Sozeau are exploring the development of higher algebra in type theory. To formalise higher algebra, one needs a new source of coherent structure in type theory. During the first year of Allieux's PhD, they studied an internalisation of polynomial monads (of which opetopes and ∞ -categories are instances) in type theory, which ought to provide such a coherent algebraic structure, inspired by the work of Kock et al [90]. They later realised that this internalisation is however incoherent as presented in pure type theory, essentially because of its reliance on equality types. Since then, they switched to a different view, describing opetopes as an external construction and relying on strict equalities in the metatheory to avoid the coherence problem. Opetopic type theory should then be, similarly to cubical type theory, a type theory indexed over these opetopic structures, where grafting and substitution are computational operations. They are now concentrating on showing that the modified inductive characterisation of opetopes and their algebras, still definable in type theory, gives rise to the standard notion of opetopes in mathematics, an original result in itself.

6.4. Incrementality

Participants: Thibaut Girka, Yann Régis-Gianas.

In collaboration with Paolo Giarrusso (EPFL, Switzerland), Philipp Shuster (Univ. of Tübingen, Germany), Yann Régis-Gianas developed a new method to incrementalise higher-order programs using formal derivatives and static caching. Yann Régis-Gianas has developed a mechanised proof for this transformation as well as a prototype language featuring efficient derivatives for functional programs. A paper has been presented at ESOP 2019 in Prague. Yann Régis-Gianas also presented this work at several places (Gallium seminar, Galinette seminar, and Choccola seminar).

In collaboration with Olivier Martinot (Paris Diderot), Yann Régis-Gianas studied a new technique to implement incrementalised operations on lists.

In collaboration with Faridath Akinotcho (Paris Diderot), Yann Régis-Gianas studied an incrementalisation of the Earley parsing algorithm.

6.5. Metatheory and development of Coq

Participants: Félix Castro, Emilio Jesús Gallego Arias, Gaëtan Gilbert, Hugo Herbelin, Pierre Letouzey, Cyprien Mangin, Thierry Martinez, Yann Régis-Gianas, Matthieu Sozeau, Théo Winterhalter, Théo Zimmermann.

6.5.1. Meta-programming and Metatheory of Coq

The MetaCoq project started last year, providing the means to program program transformations and general purpose plugins in Coq, using a quoting/unquoting mechanism. This year, they extended the framework to specify the theory, including the reduction, cumulativity and typing relations of the Polymorphic, Cumulative Calculus of Inductive Constructions at the basis of Coq. Matthieu Sozeau, together with Simon Boulier, Nicolas Tabareau and Théo Winterhalter at Galinette, Cyril Cohen at Marelle, Yannick Forster and Fabian Kunze at the University of Saarbrücken and Abhishek Anand and Gregory Malecha at BedRock Systems, Inc co-authored [54] a full description of the resulting theory (to appear in JAR). This allows for the verification of term manipulations with respect to typing: syntactic translations but also reflexive tactics glue code can hence be verified. The article also develops an alternative extraction mode to OCaml allowing the efficient compilation and execution of meta-programs written in the Template Monad. An example partial extraction of Coq programs to call-by-value pure lambda-calculus is developed this way.

Following up on this work, Matthieu Sozeau led a metatheoretical study of Coq in Coq, proving the basic metatheoretical properties of the typing relation, and developed together with Yannick Forster (Saarbrücken) and Simon Boulier, Nicolas Tabareau and Théo Winterhalter (Galinette) verified correct versions of type-checking and erasure for a large subset of Coq. This work involved the production of a fully-precise specification for the type theory implemented by Coq, cleaning up the previously untested typing specification, and variants of the algorithms used in its kernel amenable to proofs of correctness. The corresponding implementations can be extracted and provide an alternative, verified checker for Coq terms, that can run on medium-sized examples. This work will be presented [35] at POPL in New Orleans in January 2020.

6.5.2. Homotopy type theory

Hugo Moeneclaey started in September 2019 a PhD on the syntax of spheres in homotopy type theory, under the supervision of Hugo Herbelin.

Hugo Herbelin and Hugo Moeneclaey worked on the syntax of a variant of Cohen, Coquand, Huber and Mörtberg's Cubical Type Theory justified by an iterated parametricity model where equality on types is defined to be equivalence of types, thus satisfying univalence by construction.

6.5.3. Computational contents of the axiom of choice

Hugo Herbelin developed in collaboration with Nuria Brede (U. Potsdam) a unified logical structure for choice and bar induction principles.

6.5.4. Computational contents of Gödel's constructible universe

Félix Castro started his PhD under the supervision of Hugo Herbelin and Alexandre Miquel in September 2019. His PhD work will focus on the computational contents of Gödel's constructible universe. Previously, he worked on the formalisation of the ramified analytical hierarchy in classical second-order arithmetic.

6.5.5. Dependent pattern-matching and recursion

Together with Cyprien Mangin, Matthieu Sozeau refined the treatment of dependent pattern-matching in the Equations plugin. By carefully studying the type of equalities between indexed inductive types, he devised a new criterion for the elimination of equalities between inductive families based on the notion of forced arguments of constructors, resulting in a simplification of the setup of Cockx and Devriese [68] for simplification of dependent pattern-matching without K. This improved simplifier is part of the latest version of the Equations plugin, which also provides better support for the definition of mutual and well-founded recursive definitions on indexed inductive types. This work was presented at ICFP 2019 in Berlin [36]. A longer journal version is in preparation, along with a dedicated tutorial on Equations slated for inclusion in a new volume of the *Software Foundations* series dedicated to advanced tools.

Thierry Martinez continued part time the implementation of a dependent pattern-matching compilation algorithm in Coq based on the PhD thesis work of Pierre Boutillier and on the internship work of Meven Bertrand.

6.5.6. Software engineering aspects of the development of Coq

Théo Zimmermann has studied software engineering and open collaboration aspects of the development of Coq.

Following the migration of the Coq bug tracker from Bugzilla to GitHub which he conducted in 2017, he analysed data (extracted through the GitHub API), in collaboration with Annalí Casanueva Artís from the Paris School of Economics. The results show an increased number of bugs by core developers and an increased diversity of the people commenting bug reports. These quantitative results were completed with qualitative data coming from interviews with main Coq developers, which help interpret them. They validate *a posteriori* the usefulness of such a switch. A paper [43] has been published at ICSME 2019, which is the leading conference on the topic of Software Maintenance and Evolution.

Besides, Théo Zimmermann also studied and influenced the pull-based model that is now used for the development of Coq, he improved the release management process and tools, he studied package distribution and maintenance, in particular with the foundation of the coq-community organisation in 2018, which has taken off by attracting 19 maintainers, and hosting 25 projects. All of these topics are presented in the PhD thesis [28] that he defended in December 2019.

Emilio J. Gallego Arias and Théo Zimmerman took the roles of release managers for the Coq 8.12 and will oversee this release, planned for mid-2020.

Emilio J. Gallego Arias and Théo Zimmerman discussed on future plans for compositional proof checking using the Dune build system, which will include a new library format for Coq. The Dune team was informed, with Emilio J. Gallego Arias participating in the bi-weekly developer meetings. Emilio J. Gallego Arias also started discussion with the Debian OCaml maintainers (who are located at IRIF) as to see how to better integrate Dune with the Debian packaging workflow.

Emilio J. Gallego Arias designed the Coq instrumentation used in the [103] paper, which collects and analyses changes to proof scripts.

Emilio J. Gallego Arias and Karl Palmkog released a new version of the Coq SerAPI tool, which has been used in some recent proof engineering efforts, such as [65], the machine-learning environments CoqGYM and Proverbot9001 [108], [55], offering state of the art proof automation after training with proof data sets, and the educational user interface WaterProof [56]. SerAPI has also been used in some other works undergoing review and thus yet not public.

Emilio J. Gallego Arias and Shachar Itzhaky released a new version of the educational Coq frontend jsCoq [10], and assisted a few users who have been preparing courses using it.

Emilio J. Gallego Arias maintains an ongoing collaboration with the Deducteam group at Inria Saclay on the topic of interactive proof methods and standards; this has resulted in the release of an experimental LSP server for the Lambdapi theorem prover.

Emilio J. Gallego Arias, Hugo Herbelin, and Théo Zimmerman participate in the Logipedia project led by Gilles Dowek, which aims to develop a standard proof interchange format.

6.5.7. Software Infrastructure

Emilio J. Gallego Arias did significant work to refactor the Coq codebase in preparation for further work on incremental and multi-core aware type checking.

6.5.8. Dissemination activities

Emilio J. Gallego Arias and Théo Zimmerman organised the Coq meetup, an after-work event targeting industry and other communities outside academia.

6.5.9. Coordination of the development of Coq

Hugo Herbelin, Matthieu Sozeau, Emilio J. Gallego Arias and Théo Zimmermann, helped by members from Gallinette (Nantes) and Marelle (Sophia-Antipolis), devoted an important part of their time to coordinate the development, to review propositions of extensions of Coq from external and/or young contributors, and to propose themselves extensions.

6.6. Formalisation and verification

Participants: Pierre-Louis Curien, Lucien David, Emilio Jesús Gallego Arias, Kailiang Ji, Pierre Letouzey, Jean-Jacques Lévy, Cyprien Mangin, Daniel de Rauglaudre, Yann Régis-Gianas, Alexis Saurin, Matthieu Sozeau.

6.6.1. Proofs and surfaces

The joint work of Pierre-Louis Curien with Jovana Obradović (former PhD student of the team and now postdoc in Prague), Zoran Petrić and other Serbian colleagues on formalising proofs of incidence theorems (arising by repeated use of Menelaus theorem) by means of a cyclic sequent calculus, has been submitted to a journal, and has been presented at the conference Topology, Algebra, and Categories in Logic (TACL) 2019, Nice, in June 2019 [53].

6.6.2. A Coq formalisation of the first-order predicate calculus

In relation with a logic course for master students, Pierre Letouzey made a Coq formalisation of the first-order predicate calculus. The logical rules are expressed in a natural deduction style (with explicit contexts). Pierre Letouzey proposed two low-level representations of formulas : one based on quantifiers with names, the other using “locally nameless” techniques. The equivalence between the two settings has been proved correct. Using this deep embedding, Pierre Letouzey formalised in Coq the whole course notes (prepared some years ago by Alexandre Miquel), including the completeness theorem for this logic. This development is available at <https://gitlab.math.univ-paris-diderot.fr/letouzey/natded>.

6.6.3. A Coq formalisation of circular proofs and their validity condition

During the summer 2019, Alexis Saurin supervised Lucien David’s M1 internship on formalizing in Coq circular proofs and their meta-theory. This work built on Xavier Onfroy’s previous work as well as on Pierre Letouzey’s formalisation of the predicate calculus in natural deduction mentioned above. While the previous work by Xavier Onfroy was both contributing to the proof theory part and the ω -automata part (which is need for the decidability theorem), Lucien David completely focused on the the proof theory side. In particular, he was able to improve significantly on Xavier Onfroy’s formalisation by using ideas from Letouzey’s formalisation of natural deduction and by interacting with Pierre Letouzey and Alexis Saurin. This development is available at <https://github.com/LuluDavid/CircularProofsValidity>.

6.6.4. *Lexing and regular expressions in Coq*

Pierre Letouzey and Yann Régis-Gianas revisited in Coq classical techniques about lexing and regular expressions. In particular, regular expressions (with complement and conjunction) have been formalised, as well as their Brzozowski derivatives, and the finiteness theorem due to Brzozowski : a given regular expression admits only a finite number of derivatives (up to some equivalence). Both the general equivalence (based on language identity) and practical approximations (similarities) has been considered (and proved decidable). From that, the algorithms building recognizing automata (with derivatives as states) have been formalised and proved, leading to the minimal automata when using the general equivalence (but at a high cost), or to practical approximations of the minimal automata when using various similarities. This work is still ongoing. For instance, the correctness proof of a particular similarity used in an existing implementation (ml-ulex) is quite elusive for the moment. They also plan to extend this development up to a full-scale tool a la ocamllex in Coq.

6.6.5. *Real Numbers as sequences of digits in Coq*

Daniel de Rauglaudre has been continuing the formalisation of real numbers defined as sequences of digits in any radix with the LPO axiom/oracle (Limited Principle of Omniscience). Although the operations (additions and multiplications) work with this method, the proof of associativity of addition needs more work to be achieved. This development is available at https://github.com/roglo/coq_real/.

6.6.6. *Category theory in Coq*

Daniel de Rauglaudre started an implementation in Coq of Category theory in Coq, using in particular theorems coming from HOTT (HOMotopy Type theory) that he implemented some years ago. Several notions around Categories have been defined. For example, Yoneda Lemma, among others. This development is available at <https://github.com/roglo/mycoqhott/>.

6.6.7. *Number theory in Coq*

Daniel de Rauglaudre started and almost completed the formalisation in Coq of the proof of Euler's Product Formula, stating that the Riemann zeta function, which is a sum on all the natural numbers, is also a product on all the prime numbers. He also added several theorems about the prime numbers. This development is available at https://github.com/roglo/coq_euler_prod_form.

6.6.8. *Proofs of algorithms on graphs*

Jean-Jacques Lévy and Chen Ran (a PhD student at the Institute of Software, Beijing) pursued their work about formal proofs of graph algorithms. Their goal is to provide proofs of algorithms checked by computer and human readable. In 2019, they presented at ITP 2019 a joint paper with Cyril Cohen, Stephan Merz and Laurent Théry on this work [37]. This article compared formal proofs in three different systems (Why3, Coq, Isabelle/HOL) of Tarjan (1972) linear-time algorithm computing the strongly connected components in directed graphs.

The current work is to have a proof of the implementation of this algorithm with imperative programming and memory pointers. They also planed to produce formal proofs of other abstract algorithms such as the Hopcroft-Tarjan (1972) linear-time algorithm for planarity testing in undirected graphs.

6.6.9. *Certified compilation and meta-programming*

Matthieu Sozeau participates to the CertiCoq project led by Andrew Appel at Princeton (<https://www.cs.princeton.edu/~appel/certicoq>) whose aim is to verify a compiler from Coq's Gallina language down to CompCert C-light which provides itself a certified compilation path to assembly language. Together with Yannick Forster at the University of Saarbrücken and the MetaCoq team, Matthieu Sozeau focused the verification of type-checking and erasure which were previously trusted parts of the system. The new verified erasure function fills a gap in the proof of correctness of compilation from Gallina terms down to C-light. The whole compiler can be run on realistic examples (the erasure phase does take most of the compilation time and should be optimised further).

In collaboration with Xavier Denis (Paris Diderot), Yann Régis-Gianas formalised and built a compiler for Mtac2. A paper is in preparation.

POLSYS Project-Team

6. New Results

6.1. Fundamental algorithms and structured polynomial systems

The Berlekamp–Massey–Sakata algorithm and the Scalar-FGLM algorithm both compute the ideal of relations of a multidimensional linear recurrent sequence. Whenever querying a single sequence element is prohibitive, the bottleneck of these algorithms becomes the computation of all the needed sequence terms. As such, having adaptive variants of these algorithms, reducing the number of sequence queries, becomes mandatory. A native adaptive variant of the Scalar-FGLM algorithm was presented by its authors, the so-called Adaptive Scalar-FGLM algorithm. In [3], our first contribution is to make the Berlekamp–Massey–Sakata algorithm more efficient by making it adaptive to avoid some useless relation test-ings. This variant allows us to divide by four in dimension 2 and by seven in dimension 3 the number of basic operations performed on some sequence family. Then, we compare the two adaptive algorithms. We show that their behaviors differ in a way that it is not possible to tweak one of the algorithms in order to mimic exactly the behavior of the other. We detail precisely the differences and the similarities of both algorithms and conclude that in general the Adaptive Scalar-FGLM algorithm needs fewer queries and performs fewer basic operations than the Adaptive Berlekamp–Massey–Sakata algorithm. We also show that these variants are always more efficient than the original algorithms.

The problem of finding $m \times s$ matrices (with $m \geq s$) of rank r in a real affine subspace of dimension n has many applications in information and systems theory, where low rank is synonymous of structure and parsimony. In [8], we design computer algebra algorithms to solve this problem efficiently and exactly: the input are the rational coefficients of the matrices spanning the affine subspace as well as the expected maximum rank, and the output is a rational parametrization encoding a finite set of points that intersects each connected component of the low rank real algebraic set. The complexity of our algorithm is studied thoroughly. It is essentially polynomial in $n + m(s - r)$; it improves on the state-of-the-art in the field. Moreover, computer experiments show the practical efficiency of our approach.

Gröbner bases is one the most powerful tools in algorithmic non-linear algebra. Their computation is an intrinsically hard problem with a complexity at least single exponential in the number of variables. However, in most of the cases, the polynomial systems coming from applications have some kind of structure. For example, several problems in computer-aided design, robotics, vision, biology, kinematics, cryptography, and optimization involve sparse systems where the input polynomials have a few non-zero terms. In [16], our approach to exploit sparsity is to embed the systems in a semigroup algebra and to compute Gröbner bases over this algebra. Up to now, the algorithms that follow this approach benefit from the sparsity only in the case where all the polynomials have the same sparsity structure, that is the same Newton polytope. We introduce the first algorithm that overcomes this restriction. Under regularity assumptions, it performs no redundant computations. Further, we extend this algorithm to compute Gröbner basis in the standard algebra and solve sparse polynomial systems over the torus $(\mathbb{C}^{\star})^n$. The complexity of the algorithm depends on the Newton polytopes.

In [10], we consider the problem of approximating numerically the moments and the supports of measures which are invariant with respect to the dynamics of continuous- and discrete-time polynomial systems, under semialgebraic set constraints. First, we address the problem of approximating the density and hence the support of an invariant measure which is absolutely continuous with respect to the Lebesgue measure. Then, we focus on the approximation of the support of an invariant measure which is singular with respect to the Lebesgue measure. Each problem is handled through an appropriate reformulation into a linear optimization problem over measures, solved in practice with two hierarchies of finite-dimensional semidefinite moment-sum-of-square relaxations, also called Lasserre hierarchies. Under specific assumptions, the first Lasserre hierarchy allows to approximate the moments of an absolutely continuous invariant measure as close as desired and

to extract a sequence of polynomials converging weakly to the density of this measure. The second Lasserre hierarchy allows to approximate as close as desired in the Hausdorff metric the support of a singular invariant measure with the level sets of the Christoffel polynomials associated to the moment matrices of this measure. We also present some application examples together with numerical results for several dynamical systems admitting either absolutely continuous or singular invariant measures.

6.2. Solving systems over the reals and applications

It is well-known that every non-negative univariate real polynomial can be written as the sum of two polynomial squares with real coefficients. When one allows a weighted sum of finitely many squares instead of a sum of two squares, then one can choose all coefficients in the representation to lie in the field generated by the coefficients of the polynomial. In particular, this allows an effective treatment of polynomials with rational coefficients. In [11], we describe, analyze and compare both from the theoretical and practical points of view, two algorithms computing such a weighted sums of squares decomposition for univariate polynomials with rational coefficients. The first algorithm, due to the third author relies on real root isolation, quadratic approximations of positive polynomials and square-free decomposition but its complexity was not analyzed. We provide bit complexity estimates, both on the runtime and the output size of this algorithm. They are exponential in the degree of the input univariate polynomial and linear in the maximum bitsize of its complexity. This analysis is obtained using quantifier elimination and root isolation bounds. The second algorithm, due to Chevillard, Harrison, Joldes and Lauter, relies on complex root isolation and square-free decomposition and has been introduced for certifying positiveness of poly-nomials in the context of computer arithmetics. Again, its complexity was not analyzed. We provide bit complexity estimates, both on the runtime and the output size of this algorithm, which are polynomial in the degree of the input polynomial and linear in the maximum bitsize of its complexity. This analysis is obtained using Vieta's formula and root isolation bounds. Finally, we report on our implementations of both algorithms and compare them in practice on several application benchmarks. While the second algorithm is, as expected from the complexity result, more efficient on most of examples, we exhibit families of non-negative polynomials for which the first algorithm is better.

[9] describes our freely distributed Maple library SPECTRA, for Semidefinite Programming solved Exactly with Computational Tools of Real Algebra. It solves linear matrix inequalities with symbolic computation in exact arithmetic and it is targeted to small-size, possibly degenerate problems for which symbolic infeasibility or feasibility certificates are required.

Let $S \subset \mathbb{R}^n$ be a compact basic semi-algebraic set defined as the real solution set of multivariate polynomial inequalities with rational coefficients. In [19], we design an algorithm which takes as input a polynomial system defining S and an integer $p \geq 0$ and returns the n -dimensional volume of S at absolute precision 2^{-p} . Our algorithm relies on the relationship between volumes of semi-algebraic sets and periods of rational integrals. It makes use of algorithms computing the Picard-Fuchs differential equation of appropriate periods, properties of critical points, and high-precision numerical integration of differential equations. The algorithm runs in essentially linear time with respect to p . This improves upon the previous exponential bounds obtained by Monte-Carlo or moment-based methods. Assuming a conjecture of Dimca, the arithmetic cost of the algebraic subroutines for computing Picard-Fuchs equations and critical points is singly exponential in n and polynomial in the maximum degree of the input.

Let $\mathbf{f} = (f_1, \dots, f_s)$ be a sequence of polynomials in $\mathbb{Q}[X_1, \dots, X_n]$ of maximal degree D and $V \subset \mathbb{C}^n$ be the algebraic set defined by \mathbf{f} and r be its dimension. The real radical $\sqrt[\mathbb{R}]{\langle \mathbf{f} \rangle}$ associated to \mathbf{f} is the largest ideal which defines the real trace of V . When V is smooth, we show in [13], that $\sqrt[\mathbb{R}]{\langle \mathbf{f} \rangle}$, has a finite set of generators with degrees bounded by $\deg V$. Moreover, we present a probabilistic algorithm of complexity $(snD^n)^{O(1)}$ to compute the minimal primes of $\sqrt[\mathbb{R}]{\langle \mathbf{f} \rangle}$. When V is not smooth, we give a probabilistic algorithm of complexity $s^{O(1)}(nD)^{O(nr2^r)}$ to compute rational parametrizations for all irreducible components of the real algebraic set $V \cap \mathbb{R}^n$.

Let (g_1, \dots, g_p) in $\mathbb{Q}[X_1, \dots, X_n]$ and S be the basic closed semi-algebraic set defined by $g_1 \geq 0, \dots, g_p \geq 0$. The S -radical of $\langle \mathbf{f} \rangle$, which is denoted by $\sqrt[S]{\langle \mathbf{f} \rangle}$, is the ideal associated to the Zariski closure of $V \cap S$.

We give a probabilistic algorithm to compute rational parametrizations of all irreducible components of that Zariski closure, hence encoding $\sqrt[s]{\mathbf{f}}$. Assuming now that D is the maximum of the degrees of the f_i 's and the g_i 's, this algorithm runs in time $2^p(s+p)^{O(1)}(nD)^{O(rn2^r)}$.

Experiments are performed to illustrate and show the efficiency of our approaches on computing real radicals.

In [14], we consider the second-order discontinuous differential equation $y'' + \eta \operatorname{sgn}(y) = \theta y + \alpha \sin(\beta t)$ where the parameters $\eta, \theta, \alpha, \beta$ are real. The main goal is to discuss the existence of periodic solutions. Under explicit conditions, the number of such solutions is given. Furthermore, for each of these periodic solutions, an explicit formula is provided.

6.3. Solving Systems in Finite Fields, Applications in Cryptology and Algebraic Number Theory.

6.3.1. Algebraic Cryptanalysis of a Quantum Money Scheme – The Noisy Case.

At STOC 2012, Aaronson and Christiano proposed a noisy and a noiseless version of the first public-key quantum money scheme endowed with a security proof. [5] addresses the so-called noisy hidden subspaces problem, on which the noisy version of their scheme is based. The first contribution of this work is a non-quantum cryptanalysis of the above-mentioned noisy quantum money scheme extended to prime fields \mathbb{F} , with $|\mathbb{F}| \neq 2$, that runs in randomised polynomial time. This finding is supported with experimental results showing that, in practice, the algorithm presented is efficient and succeeds with overwhelming probability. The second contribution is a non-quantum randomised polynomial-time cryptanalysis of the noisy quantum money scheme over \mathbb{F}_2 succeeding with a certain probability for values of the noise lying within a certain range. This result disproves a conjecture made by Aaronson and Christiano about the non-existence of an algorithm that solves the noisy hidden subspaces problem over \mathbb{F}_2 and succeeds with such probability.

6.3.2. On the Complexity of MQ in the Quantum Setting.

In August 2015 the cryptographic world was shaken by a sudden and surprising announcement by the US National Security Agency NSA concerning plans to transition to post-quantum algorithms. Since this announcement post-quantum cryptography has become a topic of primary interest for several standardization bodies. The transition from the currently deployed public-key algorithms to post-quantum algorithms has been found to be challenging in many aspects. In particular the problem of evaluating the quantum-bit security of such post-quantum cryptosystems remains vastly open. Of course this question is of primary concern in the process of standardizing the post-quantum cryptosystems. In [21] we consider the quantum security of the problem of solving a system of m Boolean multivariate quadratic equations in n variables (MQb); a central problem in post-quantum cryptography. When $n = m$, under a natural algebraic assumption, we present a Las-Vegas quantum algorithm solving MQb that requires the evaluation of, on average, $O(2^{0.462n})$ quantum gates. To our knowledge this is the fastest algorithm for solving MQb.

6.3.3. MQsoft.

In 2017, NIST shook the cryptographic world by starting a process for standardizing post-quantum cryptography. Sixty-four submissions have been considered for the first round of the on-going NIST Post-Quantum Cryptography (PQC) process. Multivariate cryptography is a classical post-quantum candidate that turns to be the most represented in the signature category. At this stage of the process, it is of primary importance to investigate efficient implementations of the candidates. [17] presents MQsoft, an efficient library which permits to implement HFE-based multivariate schemes submitted to the NIST PQC process such as *GeMSS*, *Gui* and *DualModeMS*. The library is implemented in C targeting Intel 64-bit processors and using *avx2* set instructions. We present performance results for our library and its application to *GeMSS*, *Gui* and *DualModeMS*. In particular, we optimize several crucial parts for these schemes. These include root finding for HFE polynomials and evaluation of multivariate quadratic systems in \mathbb{F}_2 . We propose a new method which accelerates root finding for specific HFE polynomials by a factor of two. For *GeMSS* and *Gui*, we obtain a speed-up of a factor between 2 and 19 for the keypair generation, between 1.2 and 2.5 for the signature generation, and between

1.6 and 2 for the verifying process. We have also improved the arithmetic in F_{2^n} by a factor of 4 compared to the NTL library. Moreover, a large part of our implementation is protected against timing attacks.

PROSECCO Project-Team

7. New Results

7.1. Verification of security protocols

Participants: Bruno Blanchet, Karthikeyan Bhargavan, Benjamin Lipp.

Our verification of the **WireGuard** open source Virtual Private Network (VPN) with CryptoVerif appears at EuroS&P 2019 [22], [27].

We continued the development of our protocol verification tools ProVerif and CryptoVerif. The new features of this year are detailed in the section on software.

In the setting of the ANR AnaStaSec project, we worked on the verification of avionic security protocols. More specifically, in 2015, we had verified the protocol of the Secure Dialog Service using ProVerif and CryptoVerif and recommended many changes to the specification. The ICAO started to take into account our remarks, and this year we analyzed a new version of the specification. Our analysis showed that many recommendations still need to be taken into account. Additionally, we also commented on the recent choice of using DTLS over UDP to secure the future ATN/IPS (Aeronautical Telecommunication Network / Internet Protocol Suite) network, which seems very positive. The details of these results are still confidential; they have been provided to ANR.

7.2. Verified Software for Cryptographic Web Applications

Participants: Karthikeyan Bhargavan, Benjamin Beurdouche, Denis Merigoux, Jonathan Protzenko.

WebAssembly in a new language runtime that is now supported by all major web browsers and web application frameworks. We developed a compiler from the Low* subset of the F* programming language to WebAssembly and used this compiler to translate our HACL* verified cryptographic library to WebAssembly, hence obtaining the first verified cryptographic library for the Web. We also used this framework to develop and verify an implementation of the Signal protocol in WebAssembly, and demonstrated how this implementation can be used as a drop-in replacement for the libsignal-protocol library used in mainstream messaging applications like Signal, WhatsApp, and Skype.

Our work was published at the IEEE Security and Privacy conference [24]. Our WebAssembly version of HACL* and our verified Signal implementation were publicly released as open source on GitHub.

7.3. Journey beyond full abstraction

Participants: Carmine Abate, Roberto Blanco, Deepak Garg [MPI-SWS], Catalin Hritcu, Marco Patrignani [Stanford and CISP], J  r  my Thibault.

Even for safe languages, all guarantees are lost when interacting with low-level code, for instance when using low-level libraries. A compromised or malicious library that gets linked in can easily read and write data and code, jump to arbitrary instructions, or smash the stack, blatantly violating any source-level abstraction and breaking any guarantee obtained by source-level reasoning. Our goal is to build formally secure compartmentalizing compilation chains that defend against such attacks. We started by investigating what it means for a compilation chain to provide secure interoperability between a safe source language and linked target-level code that is adversarial. In this model, a secure compilation chain ensures that even linked adversarial target-level code cannot break the security properties of a compiled program any more than some linked source-level code could. However, the precise class of security properties one chooses to preserve crucially impacts not only the supported security goals and the strength of the attacker model, but also the kind of protections the compilation chain has to introduce and the kind of proof techniques one can use to make sure that the protections are watertight. We are the first to thoroughly explore a large space of secure compilation criteria based on the preservation against adversarial contexts of various classes of trace properties such as safety, of hyperproperties such as noninterference, and of relational hyperproperties such as trace equivalence [17], [10].

7.4. Principles of Program Verification for Arbitrary Monadic Effects

Participants: Kenji Maillard, Danel Ahman [University of Ljubljana], Robert Atkey [University of Strathclyde], Guido Martinez, Catalin Hritcu, Exequiel Rivas, Éric Tanter, Antoine Van Muylder, Cezar Andrici.

We devised a principled semantic framework for verifying programs with arbitrary monadic effects in a generic way with respect to expressive specifications. The starting point are Dijkstra monads, which are monad-like structures that classify effectful computations satisfying a specification drawn from a monad. Dijkstra monads have already proven valuable in practice for verifying effectful code, and in particular, they allow the F* program verifier to compute verification conditions.

We provide the first semantic investigation of the algebraic structure underlying Dijkstra monads [13], [11] and unveil a close relationship between Dijkstra monads and effect observations, i.e., mappings between a computational and a specification monad that respect their monadic structure. Effect observations are flexible enough to provide various interpretations of effects, for instance total vs partial correctness, or angelic vs demonic nondeterminism. Our semantic investigation relies on a general theory of specification monads and effect observations, using an enriched notion of relative monads and relative monad morphisms. We moreover show that a large variety of specification monads can be obtained by applying monad transformers to various base specification monads, including predicate transformers and Hoare-style pre- and postconditions. For defining correct monad transformers, we design a language inspired by the categorical analysis of the relationship between monad transformers and algebras for a monad.

We also adapt our framework to relational verification [14], [11], i.e., proving relational properties between multiple runs of one or more programs, such as noninterference or program equivalence. For this we extend specification monads and effect observations to the relational setting and use them to derive the semantics and core rules of a relational program logic generically for any monadic effect. Finally, we identify and overcome conceptual challenges that prevented previous relational program logics from properly dealing with effects such as exceptions, and are the first to provide a proper semantic foundation and a relational program logic for exceptions.

7.5. Meta-F*: Proof automation with SMT, Tactics, and Metaprograms

Participants: Guido Martinez, Danel Ahman, Victor Dumitrescu, Nick Giannarakis [Princeton University], Chris Hawblitzel [Microsoft Research], Catalin Hritcu, Monal Narasimhamurthy [University of Colorado Boulder], Zoe Paraskevopoulou [Princeton University], Clément Pit-Claudel [MIT], Jonathan Protzenko [Microsoft Research], Tahina Ramananandro [Microsoft Research], Aseem Rastogi [Microsoft Research], Nikhil Swamy [Microsoft Research].

We introduced Meta-F*[23], a tactics and metaprogramming framework for the F* program verifier. The main novelty of Meta-F* is allowing to use tactics and metaprogramming to discharge assertions not solvable by SMT, or to just simplify them into well-behaved SMT fragments. Plus, Meta-F* can be used to generate verified code automatically.

Meta-F* is implemented as an F* effect, which, given the powerful effect system of F*, heavily increases code reuse and even enables the lightweight verification of metaprograms. Metaprograms can be either interpreted, or compiled to efficient native code that can be dynamically loaded into the F* type-checker and can interoperate with interpreted code. Evaluation on realistic case studies shows that Meta-F* provides substantial gains in proof development, efficiency, and robustness.

SECRET Project-Team

7. New Results

7.1. Symmetric cryptology

Participants: Xavier Bonnetain, Christina Boura, Anne Canteaut, Daniel Coggia, Pascale Charpin, Daniel Coggia, Gaëtan Leurent, María Naya Plasencia, Léo Perrin, André Schrottenloher, Ferdinand Sibleyras.

7.1.1. Block ciphers

Our recent results mainly concern either the analysis or the design of lightweight block ciphers.

Recent results:

- Design of SATURNIN a new lightweight block cipher for authenticated encryption [74], which is resistant to quantum cryptanalysis. SATURNIN has been submitted to the NIST competition for lightweight cryptography, and has been selected for the 2nd round of the competition ⁰.
- Mixture-differential distinguishers on AES-like ciphers [18].
- Cryptanalysis of the Sbox of the Russian standards, Streebog and Kuznyechik [31], [56]. This work by L. Perrin received the best paper award at *FSE 2019*. Moreover, L. Perrin has been invited to present his results to AFNOR. He is involved in the international standardization processes in symmetric cryptography [50], [86] and has been invited to ISO meetings on this topic.
- The work on the Streebog Sbox has led to a more general study on tools for quantifying anomalies in Sboxes [44].
- Design of BISON, the first concrete block cipher following the whitened swap-or-not construction [46].

7.1.2. MACs and hash functions

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.

Recent results:

- Chosen-prefix collision attack on SHA-1 [52]: A chosen-prefix collision attack is a stronger variant of a collision attack, where an arbitrary pair of challenge prefixes are turned into a collision. Chosen-prefix collisions are usually significantly harder to produce than (identical-prefix) collisions, but the practical impact of such an attack is much larger. G. Leurent and T. Peyrin proposed new techniques to turn collision attacks into chosen-prefix collision attacks, and present such an attack against SHA-1 with complexity between $2^{66.9}$ and $2^{69.4}$ (depending on assumptions about the cost of finding near-collision blocks).
- Design of lightweight MACs from universal hash functions [51]. Many constructions of MACs used in practice (such as GMAC or Poly1305-AES) follow the Wegman-Carter-Shoup construction, which is only secure up to 2^{64} queries with a 128-bit state. S. Duval and G. Leurent proposed new constructions to reach security beyond the birthday bound, and proposed a concrete instantiation, with very good performances on ARM micro-controllers.

⁰<https://csrc.nist.gov/CSRC/media/Projects/lightweight-cryptography/documents/round-2/spec-doc-rnd2/saturnin-spec-round2.pdf>

7.1.3. 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:

- Differential Equivalence of Sboxes: C. Boura, A. Canteaut and their co-authors have studied two notions of differential equivalence of Sboxes corresponding to the case when the functions have the same difference table, or when their difference tables have the same support [19]. They proved that these two notions do not coincide, and that they are invariant under some classical equivalence relations like EA and CCZ equivalence. They also proposed an algorithm for determining the whole equivalence class of a given function.
- Boomerang Uniformity of Sboxes: The boomerang attack is a cryptanalysis technique against block ciphers which combines two differentials for the upper part and the lower part of the cipher. The Boomerang Connectivity Table (BCT) is a tool introduced by Cid *et al.* at Eurocrypt 2018 for analysing the dependency between these two differentials. C. Boura and A. Canteaut have provided an in-depth analysis of BCT, by studying more closely differentially 4-uniform Sboxes. More recently, C. Boura, L. Perrin and S. Tian have obtained new results on the boomerang uniformity of several constructions of Sboxes [57].
- CCZ equivalence of Sboxes: A. Canteaut and L. Perrin have characterized CCZ-equivalence as a property of the zeroes in the Walsh spectrum of an Sbox (or equivalently in their DDT). They used this framework to show how to efficiently upper bound the number of distinct EA-equivalence classes in a given CCZ-equivalence class. More importantly, they proved that CCZ-equivalence can be reduced to the association of EA-equivalence and an operation called twisting. They then revisited several results from the literature on CCZ-equivalence and showed how they can be interpreted in light of this new framework [21], [58].
- Links between linear and differential properties of Sboxes: P. Charpin together with J. Peng has established new links between the differential uniformity and the nonlinearity of some Sboxes in the case of two-valued functions and quadratic functions. More precisely, they have exhibited a lower bound on the nonlinearity of monomial permutations depending on their differential uniformity, as well as an upper bound in the case of differentially two-valued functions [27]
- Study of the properties of the error-correcting codes associated to differentially 4-uniform Sboxes [26]. Most notably, this work analyzes the relationship between the number of low-weight codewords and the nonlinearity of the corresponding Sbox.
- Study of crooked and weakly-crooked functions [35]: Crooked functions form a family of APN functions whose derivatives take their values in an (affine) hyperplane.
- APN functions with the butterfly construction [22], [34]: the butterfly construction, originally introduced by Perrin *et al.*, is a general construction which includes the only known example of APN permutation operating on an even number of variables. A. Canteaut, L. Perrin and S. Tian have proved that the most recent generalization of this construction does not include any other APN function when the number of variables exceeds six.

7.1.4. Modes of operation and generic attacks

In order to use a block cipher in practice, and to achieve a given security notion, a mode of operation must be used on top of the block cipher. Modes of operation are usually studied through provable security, and we know that their use is secure as long as the underlying primitive is secure, and we respect some limits on the amount of data processed. The analysis of generic attack helps us understand what happens when the hypotheses of the security proofs do not hold, or the corresponding limits are not respected. Comparing proofs and attacks also shows gaps where our analysis is incomplete, and when improved proof or attacks are required.

Recent results:

- Low-memory attacks against the 2-round Even-Mansour construction, using the 3-xor problem [41]: G. Leurent and F. Sibleyras proved that attacking the 2-round Even-Mansour construction with blocksize n is related to the 3-XOR problem with elements on size $2n$. Then, they exhibited the first generic attacks on this construction where both the data and the memory complexity are significantly lower than 2^n .
- Generic attacks against the tweakable FX-construction [55]: F. Sibleyras exhibited a generic attack on the general tweakable iterated FX-construction, which provides an upper-bound on its security. Most notably, for two rounds, this upper bound matches the proof of the particular case of XHX2 by Lee and Lee at Asiacrypt 2018, thus proving for the first time its tightness.
- Modes for authenticated encryption: Besides the design of new lightweight authenticated encryption schemes, we also analyzed some modes of operation in case of release of unverified plaintext (RUP). Indeed, in this setting, an adversary gets separated access to the decryption and verification functionality, and has more power in breaking the scheme. Our results include a forgery attack against the GCM-RUP mode of operation [54], and the design of a new lightweight deterministic scheme, named ANYDAE, which is particularly efficient for short messages, and achieves both conventional security and RUP security [24].
- Generic attacks on hash combiners [15]: G. Leurent and his co-authors analyzed the security of hash combiners, i.e. of procedures that combine two or more hash functions in a way that is hopefully more secure than each of the underlying hash functions, or at least remains secure as long as one of them is secure. They found generic attacks on the XOR combiner, on the concatenation of two Merkle-Damgård hash functions and on the Zipper hash and on the Hash-Twice combiners when they both use Merkle-Damgård hash constructions.

7.2. Code-based cryptography

Participants: Magali Bardet, Kevin Carrier, André Chailloux, Thomas Debris, Matthieu Lequesne, Rocco Mora, Nicolas Sendrier, Jean-Pierre Tillich, Valentin Vasseur.

In recent years, there has been a substantial amount of research on quantum computers. Such computers would be a major threat for all the public-key cryptosystems used in practice, since all these systems rely on the hardness of integer factoring or discrete logarithms, and these problems are easy on a quantum computer. This has prompted NIST to launch a standardization process in 2017 for quantum-safe alternatives to those cryptosystems. This concerns all three major asymmetric primitives, namely public-key encryption schemes, key-exchange protocols and digital signatures. There were 69 valid submissions to this call in November 2017, with numerous lattice-based, code-based and multivariate-cryptography submissions and some submission based either on hashing or on supersingular elliptic curve isogenies. NIST expects to perform multiple rounds of evaluation, over a period of three to five years. The goal of this process is to select a number of acceptable candidate cryptosystems for standardization. The second round of evaluation started in February 2019.

The research of the project-team in this field is focused on the design and cryptanalysis of cryptosystems making use of coding theory. 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). Our recent work on code-based cryptography has to be seen in the context of the recently launched NIST competition for quantum-safe primitives. We have proposed five code-based candidates to the NIST call for the first two primitives, namely public key encryption and key exchange protocols. Our contributions in this area are two-fold and consist in:

- designing and analysis new code-based solutions;
- cryptanalyzing code-based schemes, especially candidates to the NIST competition.

We have also been organizing since 2015 a working group held every month or every two months on code-based cryptography that structures the French efforts on this topic: every meeting is attended by most of the groups working in France on this topic (project-team GRACE, University of Bordeaux, University of Limoges, University of Rennes and University of Rouen).

7.2.1. Design of new code-based solutions

The members of the project-team have submitted several candidates to the NIST competition and have designed new code-based primitives.

Recent results:

- Design of a new code-based signature scheme [49]: T. Debris, N. Sendrier and JP Tillich recently proposed a "hash-and-sign" code-based signature scheme called WAVE, which uses a family of ternary generalized $(U, U + V)$ codes. WAVE achieves existential unforgeability under adaptive-chosen-message attacks in the random oracle model with a tight reduction to two assumptions from coding theory: one is a distinguishing problem that is related to the trapdoor inserted in the scheme, the other one is a multiple-target version of syndrome decoding. This scheme enjoys efficient signature and verification algorithms. For 128-bit security, signature are 8000-bit long and the public-key size is slightly smaller than one megabyte.
- Analysis of the ternary Syndrome Decoding problem [45]: R. Bricout, A. Chailloux, T. Debris and M. Lequesne have performed an algorithmic study of this decoding problem in large weight, which corresponds to the underlying problem in the WAVE signature scheme. Most notably, their study results in an update of the Wave parameters. It also shows that ternary Syndrome Decoding with large weight is a really harder problem than the binary Syndrome Decoding problem, and could have several applications for the design of code-based cryptosystems.

7.2.2. Cryptanalysis of code-based schemes

Recent results:

- Attack against RLCE [48]: M. Lequesne and JP Tillich, together with A. Couvreur, recently presented a key-recovery attack against the Random Linear Code Encryption (RLCE) scheme recently submitted by Y. Wang to the NIST competition. This attack recovers the secret-key for all the short key-parameters proposed by the author. It uses a polynomial-time algorithm based on a square code distinguisher.
- Analysis of an encryption scheme based on the rank syndrome decoding problem [61]: D. Coggia and A. Couvreur presented an attack against a cryptosystem proposed by Loidreau, which used an intermediary version between Gabidulin codes and LRPC codes. This attack has polynomial time for some parameters of the scheme.
- Decoding algorithm for codes with a non-trivial automorphism group [47]: R. Canto-Torres and JP Tillich presented an algorithm which is able to speed up the decoding of a code with a non-trivial automorphism group. For a certain range of parameters, this results in a decoding that is faster by an exponential factor in the code length when compared to the best algorithms for decoding generic linear codes. This algorithm was then used to break several proposals of public-key cryptosystems based on codes with a non-trivial automorphism group.

7.3. Quantum Information

Participants: Simon Apers, Ivan Bardet, Xavier Bonnetain, Rémi Bricout, André Chailloux, Simona Etinski, Antonio Florez Gutierrez, Shouvik Ghorai, Antoine Grospellier, Lucien Grouès, Anthony Leverrier, Vivien Londe, María Naya Plasencia, Andrea Olivo, Jean-Pierre Tillich, André Schrottenloher, Christophe Vuillot.

Our research in quantum information focusses on several axes: quantum codes with the goal of developing better error-correction strategies to build large quantum computers, quantum cryptography which exploits the laws of quantum mechanics to derive security guarantees, relativistic cryptography which exploits in addition the fact that no information can travel faster than the speed of light and finally quantum cryptanalysis which investigates how quantum computers could be harnessed to attack classical cryptosystems.

7.3.1. Quantum codes

Protecting quantum information from external noise is an issue of paramount importance for building a quantum computer. It is 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.

Two PhD theses have been defended this year within the project-team on this topic. First, Antoine Grospellier, co-advised by A. Leverrier and O. Fawzi (Ens Lyon), studied efficient decoding algorithms for quantum LDPC codes [13]. Beyond their intrinsic interest for channel-coding problems, such algorithms would be particularly relevant in the context of quantum fault-tolerance, since they would allow to considerably reduce the required overhead to obtain fault-tolerance in quantum computation. Vivien Londe, co-advised by A. Leverrier and G. Zémor (IMB), worked on the design of better quantum LDPC codes [14]: the main idea is to generalize the celebrated toric code of Kitaev by considering cellulations of manifolds in higher dimensions. A surprising result was that this approach leads to a much better behaviour than naively expected and a major challenge is to explore the mathematics behind this phenomenon in order to find even better constructions, or to uncover potential obstructions.

Lucien Grouès, who did an internship this summer in the project-team, has recently started a PhD with A. Leverrier and O. Fawzi on decoding quantum LDPC codes, and preliminary numerical results have already appeared in [62].

Ivan Bardet joined the project-team as a postdoc in March 2019, and will start a starting research position in 2020. His research focusses on the study of open-system dynamics as well as mixing times of Markovian dissipative evolutions with the goal of better understanding the lifetime of quantum memories.

Recent results:

- Decoding algorithms for Hypergraph Product Codes [62]: this work deals with numerical simulation of several variants of the SMALL-SET-FLIP decoder for hypergraph product codes. While this decoder had already been studied analytically in previous work in the regime of extremely low noise, we are focussing here on understanding its performance for a realistic noise model.
- Towards Low Overhead Magic State Distillation [30]: the major source of overhead in quantum fault-tolerance usually lies in the primitive called magic state distillation which takes a number of noisy versions of a specific quantum state and prepares a new state with less noise. An important question is to understand how efficient this procedure can be. In this work, we prove that magic state distillation can perform much more efficiently than expected when working with quantum systems of large dimension instead of qubits.

7.3.2. Quantum cryptography

Quantum cryptography exploits the laws of quantum physics to establish the security of certain cryptographic primitives. The most studied one is certainly quantum key distribution, which allows two distant parties to establish a secret using an untrusted quantum channel. Our activity in this field is particularly focussed on protocols with continuous variables, which are well-suited to implementations. The interest of continuous

variables for quantum cryptography was recently recognized by being awarded a 10 M€ funding from the Quantum Flagship and SECRET contributes to this project by studying the security of new key distribution protocols.

Recent results:

- Security proof for two-way continuous-variable quantum key distribution [28]: while many quantum key distribution protocols are one-way in the sense that quantum information is sent from one party to the other, it can be beneficial in terms of performance to consider two-way protocols where the quantum states perform a round-trip between the two parties. In this paper, we show how to exploit the symmetries of the protocols in phase-space to establish their security against the most general attacks allowed by quantum theory.
- Asymptotic security of continuous-variable quantum key distribution with a discrete modulation [29]: in this work, we establish a lower bound on the secret key rate of a practical quantum key distribution protocol that will be implemented in the context of the H2020 project CiViQ.

7.3.3. Quantum cryptanalysis of symmetric primitives and quantum algorithms

Symmetric cryptography seems at first sight much less affected in the post-quantum world than asymmetric cryptography: its main known threat seemed for a long time Grover's algorithm, which allows for an exhaustive key search in the square root of the normal complexity. For this reason, it was 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, as well as to find new dedicated quantum attacks. M. Naya Plasencia has been awarded an ERC Starting grant for her project named QUASYModo on this topic.

In parallel to this work, S. Apers is developing generic quantum algorithms solving combinatorial problems, notably in graphs. He also recently proposed a unified framework of quantum walk search, that will likely find applications in the context of quantum cryptanalysis.

Recent results:

- Quantum algorithm for the k -XOR problem and for list merging: The k -XOR (or generalized birthday) problem aims at finding k elements of n -bits, drawn at random, such that the XOR of all of them is 0. The algorithms proposed by Wagner more than 15 years ago remain the best known classical algorithms for solving it, when disregarding logarithmic factors. A. Chailloux, M. Naya-Plasencia and A. Schrottenloher, together with L. Grassi, studied this problem in the quantum setting and provided algorithms with the best known quantum time-complexities [38], [39].
- Quantum security of AES [17]: In order to determine the post-quantum security margin of AES-256, X. Bonnetain and M. Naya-Plasencia have proposed generalized and quantized versions of the best known cryptanalysis on reduced-round versions of AES-256, including a quantum Demirci-Selçuk meet-in-the-middle attack.
- Quantum attacks without superposition queries : In symmetric cryptanalysis, the model of superposition queries has led to surprising results, but the practical implications of these attacks remain blurry. In contrast, the results obtained so far for a quantum adversary making classical queries only were less impressive. For the first time, M. Naya-Plasencia and A. Schrottenloher, together with A. Hosoyamada and Y. Sasaki, managed to leverage the algebraic structure of some cryptosystems in the context of a quantum attacker limited to classical queries and offline quantum computations. Most notably, they are able to break the Even-Mansour construction in quantum time $\tilde{O}(2n/3)$ with $\mathcal{O}(2n/3)$ classical queries and $\mathcal{O}(n^2)$ qubits only.
- Quantum cryptanalysis of CSIDH and Ordinary Isogeny-based Schemes [16]: CSIDH is a recent proposal by Castryck et al. for post-quantum non-interactive key-exchange. It is similar in design to a scheme by Couveignes, Rostovtsev and Stolbunov, but it replaces ordinary elliptic curves by supersingular elliptic curves. Although CSIDH uses supersingular curves, it can be attacked by a quantum subexponential hidden shift algorithm due to Childs et al. While the designers of CSIDH

claimed that the parameters they suggested ensures security against this algorithm, X. Bonnetain and A. Schrottenloher showed that these security parameters were too optimistic: they improved the hidden shift algorithm and gave a precise complexity analysis in this context, which greatly reduced the complexity. For example, they showed that only 2^{35} quantum equivalents of a key-exchange are sufficient to break the 128-bit classical, 64-bit quantum security parameters proposed, instead of 2^{62} . They also extended their analysis to ordinary isogeny computations, and showed that an instance proposed by De Feo, Kieffer and Smith and expected to offer 56 bits of quantum security can be broken in 2^{38} quantum evaluations of a key exchange.

- New graph-related quantum algorithms. A first paper presents an approach to improve expansion testing using quantum Fast-Forwarding and growing seed sets [64]. A second paper introduces a graph sparsification algorithm [65], which when combined with existing classical algorithms yields the first quantum speedup for approximating the max cut, min cut, min st -cut, sparsest cut and balanced separator of a graph. Moreover, combining it with a classical Laplacian solver yields a similar speedup for Laplacian solving, for approximating effective resistances, cover times and eigenvalues of the Laplacian, and for spectral clustering.
- Quantum walks: in a first work, S. Apers describes a new quantum algorithm for quantum walk sampling using growing seed sets [42] with applications for st -connectivity and problems related to graph isomorphism. A second work [66] introduces a new quantum walk search framework that unifies and strengthens the existing ones.
- Quantum query lower bounds [59], [60]: Many computational problems, such as finding collisions in a function, are symmetric in their inputs. A. Chailloux showed that for this class of problems, any quantum algorithm can have at most a cubic advantage over the best classical algorithm in the query model, while the previously known bound gave up to 7th root advantage. This result enhances our understanding on the limitations of quantum algorithms.

CAGE Project-Team

6. New Results

6.1. Geometry of vision and sub-Riemannian geometry: new results

Let us list here our new results in the geometry of vision axis and, more generally, on hypoelliptic diffusion and sub-Riemannian geometry.

- In [12] we propose a variational model for joint image reconstruction and motion estimation applicable to spatiotemporal imaging. This model consists of two parts, one that conducts image reconstruction in a static setting and another that estimates the motion by solving a sequence of coupled indirect image registration problems, each formulated within the large deformation diffeomorphic metric mapping framework. The proposed model is compared against alternative approaches (optical flow based model and diffeomorphic motion models). Next, we derive efficient algorithms for a time-discretized setting and show that the optimal solution of the time-discretized formulation is consistent with that of the time-continuous one. The complexity of the algorithm is characterized and we conclude by giving some numerical examples in 2D space + time tomography with very sparse and/or highly noisy data.
- The article [16] presents a method to incorporate a deformation prior in image reconstruction via the formalism of deformation modules. The framework of deformation modules allows to build diffeomorphic deformations that satisfy a given structure. The idea is to register a template image against the indirectly observed data via a modular deformation, incorporating this way the deformation prior in the reconstruction method. We show that this is a well-defined regularization method (proving existence, stability and convergence) and present numerical examples of reconstruction from 2-D tomographic simulations and partially-observed images.
- The article [28] adapts the framework of metamorphosis to the resolution of inverse problems with shape prior. The metamorphosis framework allows to transform an image via a balance between geometrical deformations and changes in intensities (that can for instance correspond to the appearance of a new structure). The idea developed here is to reconstruct an image from noisy and indirect observations by registering, via metamorphosis, a template to the observed data. Unlike a registration with only geometrical changes, this framework gives good results when intensities of the template are poorly chosen. We show that this method is a well-defined regularization method (proving existence, stability and convergence) and present several numerical examples.
- In [8] we prove the C^1 regularity for a class of abnormal length-minimizers in rank 2 sub-Riemannian structures. As a consequence of our result, all length-minimizers for rank 2 sub-Riemannian structures of step up to 4 are of class C^1 .
- In [33] we show that, for a sub-Laplacian Δ on a 3-dimensional manifold M , no point interaction centered at a point $q_0 \in M$ exists.
- In [39] we consider a one-parameter family of Grushin-type singularities on surfaces, and discuss the possible diffusions that extend Brownian motion to the singularity. This gives a quick proof and clear intuition for the fact that heat can only cross the singularity for an intermediate range of the parameter. When crossing is possible and the singularity consists of one point, we give a complete description of these diffusions, and we describe a “best” extension, which respects the isometry group of the surface and also realizes the unique symmetric one-point extension of the Brownian motion, in the sense of Chen-Fukushima. This extension, however, does not correspond to the bridging extension, which was introduced by Boscain-Prandi, when they previously considered self-adjoint extensions of the Laplace-Beltrami operator on the Riemannian part for these surfaces. We clarify that several of the extensions they considered induce diffusions that are carried by the Marin compactification at the singularity, which is much larger than the (one-point) metric

completion. In the case when the singularity is more than one-point, a complete classification of diffusions extending Brownian motion would be unwieldy. Nonetheless, we again describe a “best” extension which respects the isometry group, and in this case, this diffusion corresponds to the bridging extension. A prominent role is played by Bessel processes (of every real dimension) and the classical theory of one-dimensional diffusions and their boundary conditions.

- In [50] we study the notion of geodesic curvature of smooth horizontal curves parametrized by arc length in the Heisenberg group, that is the simplest sub-Riemannian structure. Our goal is to give a metric interpretation of this notion of geodesic curvature as the first corrective term in the Taylor expansion of the distance between two close points of the curve.

We would also like to mention the monograph [30] and the PhD thesis of Mathieu Kohli [3].

6.2. Quantum control: new results

Let us list here our new results in quantum control theory.

- In [29], we discuss the compatibility between the rotating-wave and the adiabatic approximations for controlled quantum systems. Although the paper focuses on applications to two-level quantum systems, the main results apply in higher dimension. Under some suitable hypotheses on the time scales, the two approximations can be combined. As a natural consequence of this, it is possible to design control laws achieving transitions of states between two energy levels of the Hamiltonian that are robust with respect to inhomogeneities of the amplitude of the control input.
- In [34] we study one-parametric perturbations of finite dimensional real Hamiltonians depending on two controls, and we show that generically in the space of Hamiltonians, conical intersections of eigenvalues can degenerate into semi-conical intersections of eigenvalues. Then, through the use of normal forms, we study the problem of ensemble controllability between the eigenstates of a generic Hamiltonian.
- In [35] we discuss which controllability properties of classical Hamiltonian systems are preserved after quantization. We discuss some necessary and some sufficient conditions for small-time controllability of classical systems and quantum systems using the WKB method. In particular, we investigate the conjecture that if the classical system is not small-time controllable, then the corresponding quantum system is not small-time controllable either.
- In [40] we study the controllability problem for a symmetric-top molecule, both for its classical and quantum rotational dynamics. As controlled fields we consider three orthogonally polarized electric fields which interact with the electric dipole of the molecule. We characterize the controllability in terms of the dipole position: when it lies along the symmetry axis of the molecule nor the classical neither the quantum dynamics are controllable, due to the presence of a conserved quantity, the third component of the total angular momentum; when it lies in the orthogonal plane to the symmetry axis, a quantum symmetry arises, due to the superposition of symmetric states, which as no classical counterpart. If the dipole is neither along the symmetry axis nor orthogonal to it, controllability for the classical dynamics and approximate controllability for the quantum dynamics is proved to hold.

We would also like to mention the defense of the PhD thesis of Nicolas Augier (not yet on TEL) on the subject.

6.3. Stability and uncertain dynamics: new results

Let us list here our new results about stability and stabilization of control systems, on the properties of systems with uncertain dynamics.

- In an open channel, a hydraulic jump is an abrupt transition between a torrential (super-critical) flow and a fluvial (subcritical) flow. In [9] hydraulic jumps are represented by discontinuous shock solutions of hyperbolic Saint-Venant equations. Using a Lyapunov approach, we prove that we can stabilize the state of the system in H^2 -norm as well as the hydraulic jump location, with simple feedback boundary controls and an arbitrary decay rate, by appropriately choosing the gains of the feedback boundary controls.

- In [10], we study the exponential stabilization of a shock steady state for the inviscid Burgers equation on a bounded interval. Our analysis relies on the construction of an explicit strict control Lyapunov function. We prove that by appropriately choosing the feedback boundary conditions, we can stabilize the state as well as the shock location to the desired steady state in H^2 -norm, with an arbitrary decay rate.
- We develop in [19] a method ensuring robustness properties to bang-bang strategies, for general nonlinear control systems. Our main idea is to add bang arcs in the form of needle-like variations of the control. With such bang-bang controls having additional degrees of freedom, steering the control system to some given target amounts to solving an overdetermined nonlinear shooting problem, what we do by developing a least-square approach. In turn, we design a criterion to measure the quality of robustness of the bang-bang strategy, based on the singular values of the end-point mapping, and which we optimize. Our approach thus shows that redundancy implies robustness, and we show how to achieve some compromises in practice, by applying it to the attitude control of a 3d rigid body.
- Partial stability characterizes dynamical systems for which only a part of the state variables exhibits a stable behavior. In his book on partial stability, Vorotnikov proposed a sufficient condition to establish this property through a Lyapunov-like function whose total derivative is upper-bounded by a negative definite function involving only the sub-state of interest. In [20], we show with a simple two-dimensional system that this statement is wrong in general. More precisely, we show that the convergence rate of the relevant state variables may not be uniform in the initial state. We also discuss the impact of this lack of uniformity on the connected issue of robustness with respect to exogenous disturbances.
- The paper [21] elaborates control strategies to prevent clustering effects in opinion formation models. This is the exact opposite of numerous situations encountered in the literature where, on the contrary, one seeks controls promoting consensus. In order to promote declustering, instead of using the classical variance that does not capture well the phenomenon of dispersion, we introduce an entropy-type functional that is adapted to measuring pairwise distances between agents. We then focus on a Hegselmann-Krause-type system and design declustering sparse controls both in finite-dimensional and kinetic models. We provide general conditions characterizing whether clustering can be avoided as function of the initial data. Such results include the description of black holes (where complete collapse to consensus is not avoidable), safety zones (where the control can keep the system far from clustering), basins of attraction (attractive zones around the clustering set) and collapse prevention (when convergence to the clustering set can be avoided).
- The goal of [23] is to compute a boundary control of reaction-diffusion partial differential equation. The boundary control is subject to a constant delay, whereas the equation may be unstable without any control. For this system equivalent to a parabolic equation coupled with a transport equation, a prediction-based control is explicitly computed. To do that we decompose the infinite-dimensional system into two parts: one finite-dimensional unstable part, and one stable infinite-dimensional part. A finite-dimensional delay controller is computed for the unstable part, and it is shown that this controller succeeds in stabilizing the whole partial differential equation. The proof is based on an explicit form of the classical Artstein transformation, and an appropriate Lyapunov function. A numerical simulation illustrates the constructive design method.
- Given a linear control system in a Hilbert space with a bounded control operator, we establish in [26] a characterization of exponential stabilizability in terms of an observability inequality. Such dual characterizations are well known for exact (null) controllability. Our approach exploits classical Fenchel duality arguments and, in turn, leads to characterizations in terms of observability inequalities of approximately null controllability and of α -null controllability. We comment on the relationships between those various concepts, at the light of the observability inequalities that characterize them.
- In [37] we propose an extension of the theory of control sets to the case of inputs satisfying a dwell-time constraint. Although the class of such inputs is not closed under concatenation, we propose a

suitably modified definition of control sets that allows to recover some important properties known in the concatenable case. In particular we apply the control set construction to dwell-time linear switched systems, characterizing their maximal Lyapunov exponent looking only at trajectories whose angular component is periodic. We also use such a construction to characterize supports of invariant measures for random switched systems with dwell-time constraints.

- In [41] we study asymptotic stability of continuous-time systems with mode-dependent guaranteed dwell time. These systems are reformulated as special cases of a general class of mixed (discrete-continuous) linear switching systems on graphs, in which some modes correspond to discrete actions and some others correspond to continuous-time evolutions. Each discrete action has its own positive weight which accounts for its time-duration. We develop a theory of stability for the mixed systems; in particular, we prove the existence of an invariant Lyapunov norm for mixed systems on graphs and study its structure in various cases, including discrete-time systems for which discrete actions have inhomogeneous time durations. This allows us to adapt recent methods for the joint spectral radius computation (Gripenberg's algorithm and the Invariant Polytope Algorithm) to compute the Lyapunov exponent of mixed systems on graphs.
- Given a discrete-time linear switched system associated with a finite set of matrices, we consider the measures of its asymptotic behavior given by, on the one hand, its deterministic joint spectral radius and, on the other hand, its probabilistic joint spectral radius for Markov random switching signals with given transition matrix and corresponding invariant probability. In [42], we investigate the cases of equality between the two measures.
- In [45] we address the question of the exponential stability for the C^1 norm of general 1-D quasilinear systems with source terms under boundary conditions. To reach this aim, we introduce the notion of basic C^1 Lyapunov functions, a generic kind of exponentially decreasing function whose existence ensures the exponential stability of the system for the C^1 norm. We show that the existence of a basic C^1 Lyapunov function is subject to two conditions: an interior condition, intrinsic to the system, and a condition on the boundary controls. We give explicit sufficient interior and boundary conditions such that the system is exponentially stable for the C^1 norm and we show that the interior condition is also necessary to the existence of a basic C^1 Lyapunov function. Finally, we show that the results conducted in this article are also true under the same conditions for the exponential stability in the C^p norm, for any $p \geq 1$.
- In [46] we study the exponential stability for the C^1 norm of general 2×2 1-D quasilinear hyperbolic systems with source terms and boundary controls. When the eigenvalues of the system have the same sign, any nonuniform steady-state can be stabilized using boundary feedbacks that only depend on measurements at the boundaries and we give explicit conditions on the gain of the feedback. In other cases, we exhibit a simple numerical criterion for the existence of basic C^1 Lyapunov function, a natural candidate for a Lyapunov function to ensure exponential stability for the C^1 norm.
- In [47] we study the exponential stability in the H^2 norm of the nonlinear Saint-Venant (or shallow water) equations with arbitrary friction and slope using a single Proportional-Integral (PI) control at one end of the channel. Using a local dissipative entropy we find a simple and explicit condition on the gain the PI control to ensure the exponential stability of any steady-states. This condition is independent of the slope, the friction, the length of the river, the inflow disturbance and, more surprisingly, the steady-state considered. When the inflow disturbance is time-dependent and no steady-state exist, we still have the Input-to-State stability of the system, and we show that changing slightly the PI control enables to recover the exponential stability of slowly varying trajectories.
- In [48], we address the problem of the exponential stability of density-velocity systems with boundary conditions. Density-velocity systems are omnipresent in physics as they encompass all systems that consist in a flux conservation and a momentum equation. In this paper we show that any such system can be stabilized exponentially quickly in the H^2 norm using simple local feedbacks, provided a condition on the source term which holds for most physical systems, even when it is not

dissipative. Besides, the feedback laws obtained only depends on the target values at the boundaries, which implies that they do not depend on the expression of the source term or the force applied on the system and makes them very easy to implement in practice and robust to model errors. For instance, for a river modeled by Saint-Venant equations this means that the feedback laws do not require any information on the friction model, the slope or the shape of the channel considered. This feat is obtained by showing the existence of a basic H^2 Lyapunov functions and we apply it to numerous systems: the general Saint-Venant equations, the isentropic Euler equations, the motion of water in rigid-pipe, the osmosis phenomenon, etc.

- The general context of [56] is the feedback control of an infinite-dimensional system so that the closed-loop system satisfies a fading-memory property and achieves the setpoint tracking of a given reference signal. More specifically, this paper is concerned with the Proportional Integral (PI) regulation control of the left Neumann trace of a one-dimensional reaction-diffusion equation with a delayed right Dirichlet boundary control. In this setting, the studied reaction-diffusion equation might be either open-loop stable or unstable. The proposed control strategy goes as follows. First, a finite-dimensional truncated model that captures the unstable dynamics of the original infinite-dimensional system is obtained via spectral decomposition. The truncated model is then augmented by an integral component on the tracking error of the left Neumann trace. After resorting to the Artstein transformation to handle the control input delay, the PI controller is designed by pole shifting. Stability of the resulting closed-loop infinite-dimensional system, consisting of the original reaction-diffusion equation with the PI controller, is then established thanks to an adequate Lyapunov function. In the case of a time-varying reference input and a time-varying distributed disturbance, our stability result takes the form of an exponential Input-to-State Stability (ISS) estimate with fading memory. Finally, another exponential ISS estimate with fading memory is established for the tracking performance of the reference signal by the system output. In particular, these results assess the setpoint regulation of the left Neumann trace in the presence of distributed perturbations that converge to a steady-state value and with a time-derivative that converges to zero. Numerical simulations are carried out to illustrate the efficiency of our control strategy.
- There exist many ways to stabilize an infinite-dimensional linear autonomous control systems when it is possible. Anyway, finding an exponentially stabilizing feedback control that is as simple as possible may be a challenge. The Riccati theory provides a nice feedback control but may be computationally demanding when considering a discretization scheme. Proper Orthogonal Decomposition (POD) offers a popular way to reduce large-dimensional systems. In [59], we establish that, under appropriate spectral assumptions, an exponentially stabilizing feedback Riccati control designed from a POD finite-dimensional approximation of the system stabilizes as well the infinite-dimensional control system.
- In [60] we consider a 1-D linear transport equation on the interval $(0, L)$, with an internal scalar control. We prove that if the system is controllable in a periodic Sobolev space of order greater than 1, then the system can be stabilized in finite time, and we give an explicit feedback law.
- In [61] we use the backstepping method to study the stabilization of a 1-D linear transport equation on the interval $(0, L)$, by controlling the scalar amplitude of a piecewise regular function of the space variable in the source term. We prove that if the system is controllable in a periodic Sobolev space of order greater than 1, then the system can be stabilized exponentially in that space and, for any given decay rate, we give an explicit feedback law that achieves that decay rate.

6.4. Controllability: new results

Let us list here our new results on controllability beyond the quantum control framework.

- In [13], we study approximate and exact controllability of linear difference equations using as a basic tool a representation formula for its solution in terms of the initial condition, the control, and some suitable matrix coefficients. When the delays are commensurable, approximate and exact controllability are equivalent and can be characterized by a Kalman criterion. The paper focuses on providing

characterizations of approximate and exact controllability without the commensurability assumption. In the case of two-dimensional systems with two delays, we obtain an explicit characterization of approximate and exact controllability in terms of the parameters of the problem. In the general setting, we prove that approximate controllability from zero to constant states is equivalent to approximate controllability in L^2 . The corresponding result for exact controllability is true at least for two-dimensional systems with two delays.

- In [14] we consider the 2D incompressible Navier-Stokes equation in a rectangle with the usual no-slip boundary condition prescribed on the upper and lower boundaries. We prove that for any positive time, for any finite energy initial data, there exist controls on the left and right boundaries and a distributed force, which can be chosen arbitrarily small in any Sobolev norm in space, such that the corresponding solution is at rest at the given final time. Our work improves earlier results where the distributed force is small only in a negative Sobolev space. It is a further step towards an answer to Jacques-Louis Lions' question about the small-time global exact boundary controllability of the Navier-Stokes equation with the no-slip boundary condition, for which no distributed force is allowed. Our analysis relies on the well-prepared dissipation method already used for Burgers and for Navier-Stokes in the case of the Navier slip-with-friction boundary condition. In order to handle the larger boundary layers associated with the no-slip boundary condition, we perform a preliminary regularization into analytic functions with arbitrarily large analytic radius and prove a long-time nonlinear Cauchy-Kovalevskaya estimate relying only on horizontal analyticity.
- We consider the wave equation on a closed Riemannian manifold. We observe the restriction of the solutions to a measurable subset ω along a time interval $[0, T]$ with $T > 0$. It is well known that, if ω is open and if the pair (ω, T) satisfies the Geometric Control Condition then an observability inequality is satisfied, comparing the total energy of solutions to their energy localized in $\omega \times (0, T)$. The observability constant $C_T(\omega)$ is then defined as the infimum over the set of all nontrivial solutions of the wave equation of the ratio of localized energy of solutions over their total energy. In [17], we provide estimates of the observability constant based on a low/high frequency splitting procedure allowing us to derive general geometric conditions guaranteeing that the wave equation is observable on a measurable subset ω . We also establish that, as $T \rightarrow +\infty$, the ratio $C_T(\omega)/T$ converges to the minimum of two quantities: the first one is of a spectral nature and involves the Laplacian eigenfunctions, the second one is of a geometric nature and involves the average time spent in ω by Riemannian geodesics.
- In [22] we consider the problem of controlling parabolic semilinear equations arising in population dynamics, either in finite time or infinite time. These are the monostable and bistable equations on $(0, L)$ for a density of individuals $0 \leq y(t, x) \leq 1$, with Dirichlet controls taking their values in $[0, 1]$. We prove that the system can never be steered to extinction (steady state 0) or invasion (steady state 1) in finite time, but is asymptotically controllable to 1 independently of the size L , and to 0 if the length L of the interval domain is less than some threshold value L^* , which can be computed from transcendental integrals. In the bistable case, controlling to the other homogeneous steady state $0 < \theta < 1$ is much more intricate. We rely on a staircase control strategy to prove that θ can be reached in finite time if and only if $L < L^*$. The phase plane analysis of those equations is instrumental in the whole process. It allows us to read obstacles to controllability, compute the threshold value for domain size as well as design the path of steady states for the control strategy.
- The paper [27] deals with the controllability problem of a linearized Korteweg-de Vries equation on bounded interval. The system has a homogeneous Dirichlet boundary condition and a homogeneous Neumann boundary condition at the right end-points of the interval, a non homogeneous Dirichlet boundary condition at the left end-point which is the control. We prove the null controllability by using a backstepping approach, a method usually used to handle stabilization problems.
- The paper [44] is devoted to the controllability of a general linear hyperbolic system in one space dimension using boundary controls on one side. Under precise and generic assumptions on the boundary conditions on the other side, we previously established the optimal time for the null and

the exact controllability for this system for a generic source term. In this work, we prove the null-controllability for any time greater than the optimal time and for any source term. Similar results for the exact controllability are also discussed.

- Given any measurable subset ω of a closed Riemannian manifold and given any $T > 0$, we study in [49] the smallest average time over $[0, T]$ spent by all geodesic rays in ω . This quantity appears naturally when studying observability properties for the wave equation on M , with ω as an observation subset.
- Our goal is to study controllability and observability properties of the 1D heat equation with internal control (or observation) set an interval of size $\epsilon \rightarrow 0$. For any ϵ fixed, the heat equation is controllable in any time $T > 0$. It is known that depending on arithmetic properties of the center of the interval, there may exist a minimal time of pointwise control of the heat equation. We relate these two phenomena in [54].
- Our goal in [55] is to relate the observation (or control) of the wave equation on observation domains which evolve in time with some dynamical properties of the geodesic flow. In comparison to the case of static domains of observation, we show that the observability of the wave equation in any dimension of space can be improved by allowing the domain of observation to move.
- In [57] we consider the controllability problem for finite-dimensional linear autonomous control systems with nonnegative controls. Despite the Kalman condition, the unilateral nonnegativity control constraint may cause a positive minimal controllability time. When this happens, we prove that, if the matrix of the system has a real eigenvalue, then there is a minimal time control in the space of Radon measures, which consists of a finite sum of Dirac impulses. When all eigenvalues are real, this control is unique and the number of impulses is less than half the dimension of the space. We also focus on the control system corresponding to a finite-difference spatial discretization of the one-dimensional heat equation with Dirichlet boundary controls, and we provide numerical simulations.

Let us also mention the book chapter [31], which has been published this year.

6.5. Optimal control: new results

Let us list here our new results in optimal control theory beyond the sub-Riemannian framework.

- In order to determine the optimal strategy to run a race on a curved track according to the lane number, we introduce in [7] a model based on differential equations for the velocity, the propulsive force and the anaerobic energy which takes into account the centrifugal force. This allows us to analyze numerically the different strategies according to the types of track since different designs of tracks lead to straights of different lengths. In particular, we find that the tracks with shorter straights lead to better performances, while the double bend track with the longest straight leads to the worst performances and the biggest difference between lanes. Then for a race with two runners, we introduce a psychological interaction: there is an attraction to follow someone just ahead, but after being overtaken, there is a delay before any benefit from this interaction occurs. We provide numerical simulations in different cases. Overall, the results agree with the IAAF rules for lane draws in competition, where the highest ranked athletes get the center lanes, the next ones the outside lanes, while the lowest ranked athletes get the inside lanes.
- Consider a general nonlinear optimal control problem in finite dimension, with constant state and/or control delays. By the Pontryagin Maximum Principle, any optimal trajectory is the projection of a Pontryagin extremal. In [11] we establish that, under appropriate assumptions, Pontryagin extremals depend continuously on the parameter delays, for adequate topologies. The proof of the continuity of the trajectory and of the control is quite easy, however, for the adjoint vector, the proof requires a much finer analysis. The continuity property of the adjoint with respect to the parameter delay opens a new perspective for the numerical implementation of indirect methods, such as the shooting method. We also discuss the sharpness of our assumptions.

- In [15] we are concerned about the controllability of a general linear hyperbolic system in one space dimension using boundary controls on one side. More precisely, we establish the optimal time for the null and exact controllability of the hyperbolic system under some generic setting. We also present examples which yield that the generic requirement is necessary. Our approach is based on the backstepping method paying a special attention on the construction of the kernel and the selection of controls.
- A new approach to estimate traffic energy consumption via traffic data aggregation in (speed, acceleration) probability distributions is proposed in [18]. The aggregation is done on each segment composing the road network. In order to reduce data occupancy, clustering techniques are used to obtain meaningful classes of traffic conditions. Different times of the day with similar speed patterns and traffic behavior are thus grouped together in a single cluster. Different energy consumption models based on the aggregated data are proposed to estimate the energy consumption of the vehicles in the road network. For validation purposes, a microscopic traffic simulator is used to generate the data and compare the estimated energy consumption to the reference one. A thorough sensitivity analysis with respect to the parameters of the proposed method (i.e. number of clusters, size of the distributions support, etc.) is also conducted in simulation. Finally, a real-life scenario using floating car data is analyzed to evaluate the applicability and the robustness of the proposed method.
- In [24] we consider a spectral optimal design problem involving the Neumann traces of the Dirichlet-Laplacian eigenfunctions on a smooth bounded open subset Ω of \mathbf{R}^n . The cost functional measures the amount of energy that Dirichlet eigenfunctions concentrate on the boundary and that can be recovered with a bounded density function. We first prove that, assuming a L^1 constraint on densities, the so-called *Rellich functions* maximize this functional. Motivated by several issues in shape optimization or observation theory where it is relevant to deal with bounded densities, and noticing that the L^∞ -norm of *Rellich functions* may be large, depending on the shape of Ω , we analyze the effect of adding pointwise constraints when maximizing the same functional. We investigate the optimality of *bang-bang* functions and *Rellich densities* for this problem. We also deal with similar issues for a close problem, where the cost functional is replaced by a spectral approximation. Finally, this study is completed by the investigation of particular geometries and is illustrated by several numerical simulations.
- In [25] we consider the task of solving an aircraft trajectory optimization problem where the system dynamics have been estimated from recorded data. Additionally, we want to avoid optimized trajectories that go too far away from the domain occupied by the data, since the model validity is not guaranteed outside this region. This motivates the need for a proximity indicator between a given trajectory and a set of reference trajectories. In this presentation, we propose such an indicator based on a parametric estimator of the training set density. We then introduce it as a penalty term in the optimal control problem. Our approach is illustrated with an aircraft minimal consumption problem and recorded data from real flights. We observe in our numerical results the expected trade-off between the consumption and the penalty term.
- In [36] we study how bad can be the singularities of a time-optimal trajectory of a generic control affine system. In the case where the control is scalar and belongs to a closed interval it was recently shown that singularities cannot be, generically, worse than finite order accumulations of Fuller points, with order of accumulation lower than a bound depending only on the dimension of the manifold where the system is set. We extend here such a result to the case where the control has an even number of scalar components and belongs to a closed ball.
- In [38] we develop a geometric analysis and a numerical algorithm, based on indirect methods, to solve optimal guidance of endo-atmospheric launch vehicle systems under mixed control-state constraints. Two main difficulties are addressed. First, we tackle the presence of Euler singularities by introducing a representation of the configuration manifold in appropriate local charts. In these local coordinates, not only the problem is free from Euler singularities but also it can be recast as an optimal control problem with only pure control constraints. The second issue concerns the

initialization of the shooting method. We introduce a strategy which combines indirect methods with homotopies, thus providing high accuracy. We illustrate the efficiency of our approach by numerical simulations on missile interception problems under challenging scenarios.

- We introduce and study in [51] the turnpike property for time-varying shapes, within the viewpoint of optimal control. We focus here on second-order linear parabolic equations where the shape acts as a source term and we seek the optimal time-varying shape that minimizes a quadratic criterion. We first establish existence of optimal solutions under some appropriate sufficient conditions. We then provide necessary conditions for optimality in terms of adjoint equations and, using the concept of strict dissipativity, we prove that state and adjoint satisfy the measure-turnpike property, meaning that the extremal time-varying solution remains essentially close to the optimal solution of an associated static problem. We show that the optimal shape enjoys the exponential turnpike property in term of Hausdorff distance for a Mayer quadratic cost. We illustrate the turnpike phenomenon in optimal shape design with several numerical simulations.
- The work [52] proposes a new approach to optimize the consumption of a hybrid electric vehicle taking into account the traffic conditions. The method is based on a bi-level decomposition in order to make the implementation suitable for online use. The offline lower level computes cost maps thanks to a stochastic optimization that considers the influence of traffic, in terms of speed/acceleration probability distributions. At the online upper level, a deterministic optimization computes the ideal state of charge at the end of each road segment, using the computed cost maps. Since the high computational cost due to the uncertainty of traffic conditions has been managed at the lower level, the upper level is fast enough to be used online in the vehicle. Errors due to discretization and computation in the proposed algorithm have been studied. Finally, we present numerical simulations using actual traffic data, and compare the proposed bi-level method to a deterministic optimization with perfect information about traffic conditions. The solutions show a reasonable over-consumption compared with deterministic optimization, and manageable computational times for both the offline and online parts.
- An extension of the bi-level optimization for the energy management of hybrid electric vehicles (HEVs) proposed in [52] to the eco-routing problem is presented in [53]. Using the knowledge of traffic conditions over the entire road network, we search both the optimal path and state of charge trajectory. This problem results in finding the shortest path on a weighted graph whose nodes are (position, state of charge) pairs for the vehicle, the edge cost being evaluated thanks to the cost maps from optimization at the 'micro' level of a bi-level decomposition. The error due to the discretization of the state of charge is proven to be linear if the cost maps are Lipschitz. The classical A^* algorithm is used to solve the problem, with a heuristic based on a lower bound of the energy needed to complete the travel. The eco-routing method is validated by numerical simulations and compared to the fastest path on a synthetic road network.
- In [58] we study a driftless system on a three-dimensional manifold driven by two scalar controls. We assume that each scalar control has an independent bound on its modulus and we prove that, locally around every point where the controlled vector fields satisfy some suitable nondegeneracy Lie bracket condition, every time-optimal trajectory has at most five bang or singular arcs. The result is obtained using first- and second-order necessary conditions for optimality.

MATHERIALS Project-Team

6. New Results

6.1. Electronic structure calculations and related problems

Participants: Robert Benda, Éric Cancès, Virginie Ehrlicher, Luca Gorini, Gaspard Kemlin, Claude Le Bris, Antoine Levitt, Sami Siraj-Dine, Gabriel Stoltz.

6.1.1. Mathematical analysis

The members of the team have continued their systematic study of the properties of materials in the reduced Hartree-Fock (rHF) approximation, a model striking a good balance between mathematical tractability and the ability to reproduce qualitatively complex effects.

In collaboration with L. Cao, E. Cancès and G. Stoltz have studied the nuclear dynamics of infinite crystals with local defects within the Born-Oppenheimer approximation, using the reduced Hartree-Fock model to compute the electronic ground state. In this model, nuclei obey an autonomous classical Hamiltonian dynamics on a potential energy surface obtained by rHF electronic ground-state calculations. One of the main motivations for this work is to study the *nonlinear* collective excitations of nuclei in a crystal, in order to go beyond the simple harmonic approximation of non-interacting phonons. rHF ground states associated with generic nuclear displacements with respect to the periodic configuration are not mathematically well-defined at the time of writing. However, by relying on results by Cancès, Deleurence and Lewin for the rHF ground states of crystals with local defects, it is possible to study the fully nonlinear rHF Born-Oppenheimer dynamics of nuclei in the neighborhood of an equilibrium periodic configuration of a crystal. A Hilbert space of admissible nuclear displacements, and an infinite-dimensional Hamiltonian describing the dynamics of nuclei can then be defined. For small initial data, it is proved that the Cauchy problem associated with this Hamiltonian dynamics is well posed for short times (see the PhD thesis of Lingling Cao). The existence and uniqueness for arbitrary initial data, and/or long times requires a perturbation analysis of the rHF model when the Fermi level is occupied, which is work in progress.

6.1.2. Numerical analysis

E. Cancès has pursued his long-term collaboration with Y. Maday (Sorbonne Université) on the numerical analysis of linear and nonlinear eigenvalue problems. Together with G. Dusson (Besançon), B. Stamm (Aachen, Germany), and M. Vohralik (Inria SERENA), they have designed *a posteriori* error estimates for conforming numerical approximation of eigenvalue clusters of second-order self-adjoint operators on bounded domains [44]. Given a cluster of eigenvalues, they have estimated the error in the sum of the eigenvalues, as well as the error in the eigenvectors represented through the density matrix, i.e. the orthogonal projector on the associated eigenspace. This allows them to deal with degenerate (multiple) eigenvalues within this framework. The bounds are guaranteed and converge at the same rate as the exact error. They can be turned into fully computable bounds as soon as an estimate on the dual norm of the residual is available, which is notably the case (i) for the Laplace eigenvalue problem discretized with conforming finite elements, and (ii) for a Schrödinger operator with periodic boundary conditions discretized with plane waves.

R. Benda, E. Cancès and B. Leventhal (Ecole Polytechnique) have initiated the design and analysis of multiscale models for the electrical conductivity of networks of functionalized carbon nanotubes. Such devices are used as nanosensors, for instance to monitor the quality of water. In [11], they study by means of Monte-Carlo numerical simulations the resistance of two-dimensional random percolating networks of stick, widthless nanowires. They use the multi-nodal representation (MNR) to model a nanowire network as a graph. They derive numerically from this model the expression of the total resistance as a function of all meaningful parameters, geometrical and physical, over a wide range of variation for each. They justify their choice of non-dimensional variables applying Buckingham π -theorem. The effective resistance of 2D random percolating

networks of nanowires is found to have a nice expression in terms of the geometrical parameters (number of wires, aspect ratio of electrode separation over wire length) and the physical parameters (nanowire linear resistance per unit length, nanowire/nanowire contact resistance, metallic electrode/nanowire contact resistance). The dependence of the resistance on the geometry of the network, on the one hand, and on the physical parameters (values of the resistances), on the other hand, is thus clearly separated thanks to this expression, much simpler than the previously reported analytical expressions. In parallel, atomic scale models based on electronic structure theory are being developed to parameterize these mesoscale models (PhD thesis of R. Benda).

C. Le Bris has pursued his long term collaboration with Pierre Rouchon (Ecole des Mines de Paris and Inria QUANTIC) on the study of high dimensional Lindblad type equations at play in the modelling of open quantum systems. They have co-supervised the M2 internship of Luca Gorini, that was focused on the simulation of some simple quantum gates, and has investigated several discretization strategies based upon the choice of suitable basis sets.

V. Ehrlacher, L. Grigori (Inria ALPINES), D. Lombardi (Inria COMMEDIA) and H. Song (Inria ALPINES) have designed a new numerical method for the compression of high-order tensors [49]. The principle of the algorithm consists in constructing an optimal partition of the set of indices of the tensor, and construct an approximation of the tensor on each indices subdomain by means of an adapted High-Order Singular Value Decomposition. This method was used, among other examples, for the reduction of the solution of the Vlasov-Poisson system, and enabled to reach very significant compression factors. They also obtained very encouraging results on the compression of the Coulomb potential, which could be very interesting with a view to the resolution of the time-dependent Schrödinger equation in high dimension, which is currently work in progress.

A. Alfonsi, R. Coyaud (Ecole des Ponts), V. Ehrlacher and D. Lombardi (Inria COMMEDIA) studied a different approach for discretizing optimal transport problems, which relies in relaxing the marginal constraints in a finite number of marginal moment constraints, while keeping an infinite state space [40]. The advantage of such an approach is that the approximate solution of the multi-marginal optimal transport problem with Coulomb cost, which is the semi-classical limit of the so-called Lévy-Lieb functional, can be represented as a discrete measure charging a low number of points, thus avoiding the curse of dimensionality when the number of electrons is large.

M. Herbst and his collaborators have developed the `adcc` Python/C++ software package for performing excited state calculations based on algebraic-diagrammatic construction methods. It connects to four SCF packages (`pyscf`, `psifour`, `molsturm` and `veloxchem`), allows the inclusion of environmental effects through implicit or explicit solvent models, and implements methods up to third order in perturbation theory. Its features are summarized in [54].

6.2. Computational Statistical Physics

Participants: Manon Baudel, Qiming Du, Grégoire Ferré, Frédéric Legoll, Tony Lelièvre, Mouad Ramil, Geneviève Robin, Laura Silva Lopes, Gabriel Stoltz.

The objective of computational statistical physics is to compute macroscopic properties of materials starting from a microscopic description, using concepts of statistical physics (thermodynamic ensembles and molecular dynamics). The contributions of the team can be divided into four main topics: (i) the development of methods for sampling the configuration space; (ii) the efficient computation of dynamical properties which requires to sample metastable trajectories; (iii) the simulation of nonequilibrium systems and the computation of transport coefficients; (iv) coarse-graining techniques to reduce the computational cost of molecular dynamic simulations and gain some insights on the models.

6.2.1. Sampling of the configuration space: new algorithms and applications

The work [52] by G. Ferré and G. Stoltz considers fluctuations of empirical averages for stochastic differential equations. Such averages are commonly used to compute ergodic averages in statistical physics in order to

estimate macroscopic quantities, but they are subject to fluctuations. If small deviations are described by the central limit theorem, important fluctuations enter the large deviations framework. This theory is well understood when considering bounded observables of a stochastic differential equation, but quantities of interest are generally unbounded. The authors identify the class of unbounded functions which enter the "usual" regime of large deviations. The answer is not trivial, and suggests that many physical observables satisfy another type of large deviations, which leads to further works. Additionally, the influence of irreversibility on the fluctuations was studied by providing a mathematical illustration of the second law of thermodynamics, stating that irreversible dynamics generate more entropy, or more disorder, than reversible ones.

The team also pursued its endeavour to study and improve free energy biasing techniques, such as adaptive biasing force or metadynamics. The gist of these techniques is to bias the original metastable dynamics used to sample the target probability measure in the configuration space by an approximation of the free energy along well-chosen reaction coordinates. This approximation is built on the fly, using empirical measures over replicas, or occupations measures over the trajectories of the replicas. Two works have been performed on such methods

- First, in [50], V. Ehrlacher, T. Lelièvre and P. Monmarché (Sorbonne Université) have developed a new numerical method in order to compute the free energy and the biased potential given by the Adaptive Biasing Force method in the case where the number of reaction coordinates in the system is too large to apply standard grid-based approximation techniques. The algorithm uses a greedy algorithm and a tensor product approximation. Convergence proofs of both the underlying ABF technique (which uses an unbiased occupation measure) and the greedy tensor-product approximation are provided.
- Second, in [55], T. Lelièvre together with B. Jourdain (Ecole des Ponts) and P.-A. Zitt (Université Paris Est) have used a parallel between metadynamics and self interacting models for polymers to study the longtime convergence of the original metadynamics algorithm in the adiabatic setting, namely when the dynamics along the collective variables decouple from the dynamics along the other degrees of freedom. The bias which is introduced when the adiabatic assumption does not hold is also discussed.

The team has also considered new applications in terms of sampling, and the analysis of related sampling methods:

- For large scale Bayesian inference, B. Leimkuhler (Edinburgh, United Kingdom), M. Sachs (Duke, USA) and G. Stoltz have studied in [57] the convergence of Adaptive Langevin dynamics, which is a method for sampling the Boltzmann-Gibbs distribution at a prescribed temperature in cases where the potential gradient is subject to stochastic perturbation of unknown magnitude. The method replaces the friction in underdamped Langevin dynamics with a dynamical variable, updated according to a negative feedback loop control law as in the Nose-Hoover thermostat. Hypocoercive techniques allow to show that the law of Adaptive Langevin dynamics converges exponentially rapidly to the stationary distribution, with a rate that can be quantified in terms of the key parameters of the dynamics. This implies in particular that a central limit theorem holds for the time averages computed along a stochastic path.
- For the simulation of log-gases, G. Ferré and G. Stoltz have studied in [46] with D. Chafaï (Université Paris Dauphine) a follow up to a former project on the efficient simulation of Coulomb and logarithmic gases. A previous work has demonstrated the usefulness of Hybrid Monte Carlo techniques for sampling the invariant measure of such gases. Gases under constraint are now considered. First, the algorithm proposed in [31] was used to numerically explore the situation. Then, large deviations techniques were employed to study the limiting behaviour of the conditioned gas when the number of particles gets large. For a class of constraints, the equation solved by the limiting empirical density shows in particular cases a spectacular behaviour. This work suggests to further explore some research paths, such as the limiting distribution for large constraints.

6.2.2. Sampling of dynamical properties and rare events

In the preprint [48], T. Lelièvre uses the quasi-stationary distribution approach to study the first exit point distribution from a bounded domain of the overdamped Langevin dynamics, in collaboration with G. Di Gesù (TU Wien, Austria), B. Nectoux (Université Blaise Pascal) and D. Le Peutrec (Université Paris-Sud). The quasi-stationary distribution approach has been developed by T. Lelièvre and collaborators over the past years in order to rigorously model the exit event from a metastable state by a jump Markov process, and to study this exit event in the small temperature regime. In [48], the authors prove that in the small temperature regime and under rather general assumptions on the initial conditions and on the potential function, the support of the distribution of the first exit point concentrates on some points realizing the minimum of the potential on the boundary. The proof relies on tools to study tunnelling effects in semi-classical analysis. This preprint has been divided into two separate articles for publication: the first one [22] has been accepted for publication; the second one is currently under review.

6.2.3. Nonequilibrium systems and computation of transport coefficients

Stemming from the IHP trimester "Stochastic Dynamics Out of Equilibrium" held at Institut Henri Poincaré in April-July 2017, a collection of contributions has been grouped in a volume of proceedings [38], focusing on aspects of nonequilibrium dynamics and its ongoing developments. This volume has been edited by G. Giacomin (Université Paris Diderot), S. Olla (Université Paris Dauphine), E. Saada (CNRS and Université Paris Descartes), H. Söfn (TU Munich, Germany) and G. Stoltz. It includes contributions from various events relating to three domains: (i) transport in non-equilibrium statistical mechanics; (ii) the design of more efficient simulation methods; (iii) life sciences.

In addition, P. Plechac (University of Delaware, USA), T. Wang (Army Research Lab, USA) and G. Stoltz have considered in [61] numerical schemes for computing the linear response of steady-state averages of stochastic dynamics with respect to a perturbation of the drift part of the stochastic differential equation. The schemes are based on Girsanov's change-of-measure theory to reweight trajectories with factors derived from a linearization of the Girsanov weights. Both the discretization error and the finite time approximation error have been investigated. The designed numerical schemes have been shown to be of bounded variance with respect to the integration time, which is a desirable feature for long time simulation. The discretization error has been shown to be improved to second order accuracy in the time step by modifying the weight process in an appropriate way.

6.2.4. Coarse-graining

Two works have been done to explore new methods to define "good" reaction coordinates:

- The estimation of the Poincaré constant of a given probability measure allows to quantify the typical convergence rate of reversible diffusions to their equilibrium measure. Loucas Pillaud-Vivien, F. Bach and A. Rudi (Inria SIERRA), together with T. Lelièvre and G. Stoltz, have shown in [58], both theoretically and experimentally how to estimate the Poincaré constant given sufficiently many samples of the probability measure under consideration, using reproducing Hilbert kernel spaces. As a by-product of this estimation, they have also derived an algorithm that captures a low dimensional representation of the data by finding directions which are difficult to sample – reaction coordinates in the language of molecular dynamics. This amounts to finding the marginal of the high dimensional sampled measure for which the Poincaré constant is the largest possible.
- In [60], T. Lelièvre together with B. Leimkuhler and Z. Trstanova (University of Edinburgh, Scotland) has explored numerically the interest of using diffusion maps to define reaction coordinates or metastable states. Diffusion maps approximate the generator of Langevin dynamics from simulation data, and the idea is thus to use the eigenvalues and eigenvectors to build coarse-grained variables. They have also discussed the use of diffusion maps to define local reaction coordinates within the metastable sets, formalising the locality via the concept of quasi-stationary distribution and justifying the convergence of diffusion maps applied to samples within a metastable set.

Another coarse-graining procedure was considered to justify the approximation of an infinite system of ordinary differential equations (the Becker-Doring equations, describing coagulation/fragmentation processes of species of integer sizes) in terms of a partial differential equation. Formal Taylor expansions motivate that the dynamics at large sizes should be dictated by an advection-diffusion equation, called Fokker-Planck equation. P. Terrier and G. Stoltz rigorously proved in [59] the link between these two descriptions for evolutions on finite times rather than in some hydrodynamic limit, motivated by the results of numerical simulations and the construction of dedicated algorithms based on splitting strategies. In fact, the Becker-Doring equations and the Fokker-Planck equation are related through some pure diffusion with unbounded diffusion coefficient. The crucial point in the analysis is to obtain decay estimates for the solution of this pure diffusion and its derivatives to control remainders in the Taylor expansions. The small parameter in this analysis is the inverse of the minimal size of the species.

6.3. Homogenization

Participants: Xavier Blanc, Virginie Ehrlacher, Olga Gorynina, Rémi Goudey, Claude Le Bris, Frédéric Legoll, Adrien Lesage, Pierre-Loïc Rothé.

In homogenization theory, members of the project-team have pursued their ongoing systematic study of perturbations of periodic problems (by local and nonlocal defects). This has been done in several different directions.

6.3.1. Deterministic non-periodic systems

For linear elliptic equations with highly oscillating coefficients, X. Blanc and C. Le Bris have recently developed, in collaboration with P.-L. Lions (Collège de France), a theory in the case of periodic problems with local defects. In particular, the existence of a corrector function for such problems has been shown. More details on the quality of approximation achieved by their theory have been recently provided. The fact that a corrector exists with suitable properties indeed allows one to quantify the rate of convergence of the two-scale expansion (which uses that corrector) to the actual exact solution, as the small homogenization parameter ε vanishes. In that spirit, some of these works by X. Blanc and C. Le Bris, in collaboration with M. Josien (former PhD student in the team, now at MPI Leipzig, Germany), have been presented in [12], [13].

Also in the context of homogenization theory, O. Gorynina, C. Le Bris and F. Legoll have explored the question of how to determine the homogenized coefficient of heterogeneous media without explicitly performing an homogenization approach. This work is a follow-up on earlier works by C. Le Bris and F. Legoll in collaboration with K. Li and next S. Lemaire over the years. During the year, O. Gorynina, C. Le Bris and F. Legoll have mathematically studied a computational approach initially introduced by R. Cottreau (CNRS Marseille). This approach combines, in the Arlequin framework, the original fine-scale description of the medium (modelled by an oscillatory coefficient) with an effective description (modelled by a constant coefficient) and optimizes upon the coefficient of the effective medium to best fit the response of a purely homogeneous medium. In the limit of asymptotically infinitely fine structures, the approach yields the value of the homogenized coefficient. The aim is to mathematically study the problem and to investigate how to improve on the practical algorithm, in order to obtain a procedure as efficient as possible. Results will be presented in a couple of manuscripts in preparation.

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 are practically relevant and keep the computational workload limited.

Using 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, and the team has proposed, over the past years, many approaches to improve on the computation of the homogenized tensor.

Besides the averaged behavior of the oscillatory solution u_ε on large space scales (which is given by its homogenized limit), a question of interest is to describe how u_ε fluctuates. This question has been investigated in the PhD thesis of P.-L. Rothé, both from a theoretical and a numerical viewpoints. First, theoretical results have been obtained for a weakly stochastic setting (where the coefficient is the sum of a periodic coefficient and a small random perturbation). It has been shown that, at the first order and when ε is small, the localized fluctuations (characterized by a test function g) of u_ε are Gaussian. The corresponding variance depends on the localization function g , on the right-hand side f of the problem satisfied by u_ε , and on a fourth order tensor Q which is defined in terms of the corrector. Since the corrector function is challenging to compute, so is Q . A numerical approach (based on using the standard truncated corrector problem) has been designed to approximate Q and its convergence has been proven, again in a weakly stochastic setting. All these theoretical results critically depend on detailed properties of the Green function associated to the periodic operator. Second, numerical experiments in more general settings (i.e. full stochastic case) following the same approach have been performed, in order to investigate the generality of the obtained results. First, the convergence of the approximation of Q has been monitored. Second, it has been checked that the localized fluctuations of u_ε indeed become Gaussian when ε decreases, and that their variance can be related to Q . These promising numerical results, which are consistent with the theoretical results obtained in the weakly stochastic setting, are presented in a manuscript in preparation.

6.3.3. Multiscale Finite Element approaches

From a numerical perspective, 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 a sufficiently accurate approximation).

During the year, several research tracks have been pursued in this general direction.

The MsFEM approach uses a Galerkin approximation on a pre-computed basis, obtained by solving local problems mimicking the problem at hand at the scale of mesh elements, with carefully chosen right-hand sides and boundary conditions. The initially proposed version of MsFEM uses as basis functions the solutions to these local problems, posed on each mesh element, with null right-hand sides and with the coarse P1 elements as Dirichlet boundary conditions. Various improvements have next been proposed, such as the *oversampling* variant, which solves local problems on larger domains and restricts their solutions to the considered element. In collaboration with U. Hetmaniuk (University of Washington in Seattle, USA), C. Le Bris, F. Legoll and P.-L. Rothé have completed the study of a MsFEM method improved differently. They have considered a variant of the classical MsFEM approach with enrichments based on Legendre polynomials, both in the bulk of the mesh elements and on their interfaces. A convergence analysis of this new variant has been performed. In addition, residue type a posteriori error estimators have been proposed and certified, leading to a numerical strategy where the degree of enrichment is *locally* adapted in order to reach, at the smallest computational cost, a given error. The promising numerical results are currently being collected in a manuscript in preparation.

Many numerical analysis studies of the MsFEM are focused on obtaining a priori error bounds. In collaboration with L. Chamoin, who was on leave in the project-team a few years ago from ENS Cachan, members of the project-team have been working on a posteriori error analysis for MsFEM approaches, with the aim of developing 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 perform all expensive computations in an offline stage, independently of the right-hand side. The standard CRE approach has thus been adapted to that context. In the recent work [47], the approach has also been adapted towards the design of adaptive algorithms for specific quantities of interest (in the so-called “goal-oriented” setting). It provides an accurate estimation of the error, and leads to a discretization which is efficiently tailored to the specific quantity under consideration.

One of the perspectives of the team, through the PhD thesis of A. Lesage, is the development of Multiscale Finite Element Methods for thin heterogeneous plates. The fact that one of the dimension of the domain of interest scales as the typical size of the heterogeneities within the material induces theoretical and practical difficulties that have to be carefully taken into account (see [37]). The first steps of the work of V. Ehrlacher, F. Legoll and A. Lesage, in collaboration with A. Lebé (Ecole des Ponts) have consisted in studying the homogenized limit (and the two-scale expansion) of problems posed on thin heterogeneous plates. After having considered the case of a diffusion equation, the more challenging case of elasticity has been studied. In the so-called membrane case (that is, when the loading is in the in-plane directions), an approximation result for the two-scale expansion has been obtained. Several MsFEM variants have been proposed and compared numerically. The results will be presented in a forthcoming manuscript.

6.4. Various topics

Participants: Sébastien Boyaval, Virginie Ehrlacher.

A new mathematical framework has been identified for the modelling of complex fluids in [42]. It allows one to incorporate rheological features of a real (non-ideal, visous and compressible) fluid and, at the same time, to compute flows as solution to a *hyperbolic system of conservation laws* complemented by an initial value. In [42], the framework is specified for 2D hydrostatic flows of *Maxwell fluids*, with a numerical finite-volume scheme preserving the positivity of mass and a Clausius-Duem inequality. Formally, the (macroscopic) model has a (microscopic) molecular justification using a generalized Langevin equation for the distortion of the fluid texture.

On the other hand, recall that stochastic models are also used for the numerical simulation of hydrodynamical turbulence. In particular, a generalized Langevin equation can be used to model the "thermostated" velocity fluctuations in a "stationary" turbulent flow modelled as an invariant measure. But the interest for the effective numerical simulation of turbulent flows is not fully understood yet. In [43], S. Boyaval with S. Martel and J. Reygner (Ecole des Ponts) have studied the convergence of a discretization of a 1D stochastic scalar viscous conservation laws (a toy-model), for the numerical simulation of its invariant measure.

A. Benaceur (EDF), V. Ehrlacher and A. Ern (École des Ponts and Inria SERENA) developed a new EIM/reduced-basis method [10] for the reduction of parametrized variational inequalities with nonlinear constraints, and applied this method to the reduction of contact mechanics problems with non-coincident meshes.

V. Ehrlacher, D. Lombardi (Inria COMMEDIA), O. Mula (Université Paris-Dauphine) and F-X. Vialard (Université Paris-Est) developed new model-order reduction techniques based on the use of Wasserstein spaces for transport-dominated problems [51], which gives very encouraging results on several classes of conservative transport problems like the Burger's equation. Theoretical convergence rates are proved on some particular test cases.

J. Berendsen (Chemnitz, Germany), Martin Burger (Erlangen, Germany), V. Ehrlacher, J-F. Pietschmann (Chemnitz, Germany) proved the existence and uniqueness of strong solutions and weak-strong stability in a particular system of cross-diffusion equations [41]. It is in general very difficult to obtain such kind of results for general cross-diffusion systems. The proof for the particular system studied here relies on the fact that, when all the cross-diffusion coefficients of the system are equal to the same constant, the system boils down to a set of independent heat equations, for which uniqueness of strong solutions is trivial. The uniqueness of strong solutions was proved under the assumption that the cross-diffusion coefficients should not be close enough to one another.

MATHRISK Project-Team

7. New Results

7.1. Control of systemic risk in a dynamic framework

Agnès Sulem, Andreea Minca (Cornell University), Hamed Amini (J. Mack Robinson College of Business, Georgia State University) and Rui Chen have studied a Dynamic Contagion Risk Model With Recovery Features [27]. In this paper, they introduce threshold growth in the classical threshold contagion model, in which nodes have downward jumps when there is a failure of a neighboring node. Choosing the configuration model as underlying graph, they prove fluid limits for the baseline model, as well as extensions to the directed case, state-dependent inter-arrival times and the case of growth driven by upward jumps. They obtain explicit ruin probabilities for the nodes according to their characteristics: initial threshold and in- (and out-) degree. They then allow nodes to choose their connectivity by trading off link benefits and contagion risk. They define a rational equilibrium concept in which nodes choose their connectivity according to an expected failure probability of any given link, and then impose condition that the expected failure probability coincides with the actual failure probability under the optimal connectivity. Existence of an asymptotic equilibrium is shown as well as convergence of the sequence of equilibria on the finite networks. In particular, these results show that systems with higher overall growth may have higher failure probability in equilibrium.

The results have been presented in Lisbon at the COMPLEX NETWORKS 2019 conference. Rui Chen has defended his thesis in July 2019 on this topic [10].

7.2. Mean-field BSDEs and systemic risk measures

Agnès Sulem with her PhD student Rui Chen, Andreea Minca and Roxana Dumitrescu have studied mean-field BSDEs with a generalized mean-field operator that can capture the average intensity in an inhomogeneous random graph. Comparison and strict comparison results have been obtained. Based on these, they interpret the BSDE solution as a global dynamic risk measure that can account for the intensity of system interactions and therefore incorporate systemic risk. Using Fenchel-Legendre transforms, they establish a dual representation for the expectation of the risk measure, and exhibit its dependence on the mean-field operator [31].

7.3. Risk management in finance and insurance

7.3.1. Option pricing in a non-linear incomplete market model with default

Agnès Sulem has studied with Miryana Grigorova (University of Leeds) and Marie-Claire Quenez (Université Paris Denis Diderot) superhedging prices and the associated superhedging strategies for both European and American options (see [33] and [32] in a non-linear incomplete market model with default. The underlying market model consists of a risk-free asset and a risky asset driven by a Brownian motion and a compensated default martingale. The portfolio processes follow non-linear dynamics with a non-linear driver f .

7.3.2. Neural network regression for Bermudan option pricing

The pricing of Bermudan options amounts to solving a dynamic programming principle, in which the main difficulty, especially in high dimension, comes from the conditional expectation involved in the computation of the continuation value. These conditional expectations are classically computed by regression techniques on a finite dimensional vector space. In [36], Bernard Lapeyre and Jérôme Lelong study neural networks approximations of conditional expectations. They prove the convergence of the well-known Longstaff and Schwartz algorithm when the standard least-square regression is replaced by a neural network approximation. They illustrate the numerical efficiency of neural networks as an alternative to standard regression methods for approximating conditional expectations on several numerical examples.

7.3.3. Hybrid numerical method for option pricing

With Giulia Terenzi, Lucia Caramellino (Tor Vergata University), and Maya Briani (CNR Roma), Antonino Zanette develop and study stability properties of a hybrid approximation of functionals of the Bates jump model with stochastic interest rate that uses a tree method in the direction of the volatility and the interest rate and a finite-difference approach in order to handle the underlying asset price process. They also propose hybrid simulations for the model, following a binomial tree in the direction of both the volatility and the interest rate, and a space-continuous approximation for the underlying asset price process coming from a Euler–Maruyama type scheme. They test their numerical schemes by computing European and American option prices [17].

7.3.4. American options

With his PhD student Giulia Terenzi, Damien Lamberton has been working on American options in Heston's model [22]. He is currently preparing his contribution to a winter school on "Theory and practice of optimal stopping and free boundary problems" (cf. <https://conferences.leeds.ac.uk/osfbp/>).

7.3.5. Solvency Capital Requirement in Insurance

A. Alfonsi has obtained a grant from AXA Foundation on a Joint Research Initiative with a team of AXA France working on the strategic asset allocation. This team has to make recommendations on the investment over some assets classes as, for example, equity, real estate or bonds. In order to do that, each side of the balance sheet (assets and liabilities) is modeled in order to take into account their own dynamics but also their interactions. Given that the insurance products are long time contracts, the projections of the company's margins have to be done considering long maturities. When doing simulations to assess investment policies, it is necessary to take into account the SCR which is the amount of cash that has to be settled to manage the portfolio. Typically, the computation of the future values of the SCR involve expectations under conditional laws, which is greedy in computation time.

A. Alfonsi and his PhD student A. Cherchali have developed a model of the ALM management of insurance companies that takes into account the regulatory constraints on life-insurance [25]. We now focus on developing Multilevel Monte-Carlo methods to approximate the SCR (Solvency Capital Requirement).

7.3.6. Pricing and hedging variable annuities of GMWB type in advanced stochastic models

Antonino Zanette with Ludovic Goudenège (Ecole Centrale de Paris) and Andrea Molent (University of Udine) study the valuation of a particular type of variable annuity called GMWB when advanced stochastic models are considered. As remarked by Yang and Dai (Insur Math Econ 52(2):231–242, 2013), and Dai et al. (Insur Math Econ 64:364–379, 2015), the Black–Scholes framework seems to be inappropriate for such a long maturity products. Also Chen et al. (Insur Math Econ 43(1):165–173, 2008) show that the price of GMWB variable annuities is very sensitive to the interest rate and the volatility parameters. They propose here to use a stochastic volatility model (the Heston model) and a Black–Scholes model with stochastic interest rate (the Black–Scholes Hull–White model). For this purpose, they consider four numerical methods: a hybrid tree-finite difference method, a hybrid tree-Monte Carlo method, an ADI finite difference scheme and a Standard Monte Carlo method. These approaches are employed to determine the no-arbitrage fee for a popular version of the GMWB contract and to calculate the Greeks used in hedging. Both constant withdrawal and dynamic withdrawal strategies are considered. Numerical results are presented, which demonstrate the sensitivity of the no-arbitrage fee to economic and contractual assumptions as well as the different features of the proposed numerical methods [18].

7.4. Stochastic Analysis and probabilistic numerical methods

7.4.1. Particles approximation of mean-field SDEs

O. Bencheikh and Benjamin Jourdain analysed the rate of convergence of a system of N interacting particles with mean-field rank based interaction in the drift coefficient and constant diffusion coefficient [16], [30]. They first adapted arguments by Kolli and Shkolnikhov to check trajectorial propagation of chaos with optimal rate

$N^{-1/2}$ to the associated stochastic differential equations nonlinear in the sense of McKean. They next relaxed the assumption needed by Bossy to check convergence in $L^1(\mathbf{R})$ of the empirical cumulative distribution function of the Euler discretization with step h of the particle system to the solution of a one dimensional viscous scalar conservation law with rate $\mathcal{O}\left(\frac{1}{\sqrt{N}} + h\right)$. Last, they proved that the bias of this stochastic particle method behaves in $\mathcal{O}\left(\frac{1}{N} + h\right)$, which is confirmed by numerical experiments.

7.4.2. Abstract Malliavin calculus and convergence in total variation

In collaboration with L. Caramellino (University Tor Vergata) and with G. Poly (University of Rennes), V. Bally has settled a Malliavin type calculus for a general class of random variables, which are not supposed to be Gaussian (as it is the case in the standard Malliavin calculus). This is an alternative to the Γ calculus settled by Bakry, Gentile and Ledoux. The main application is the estimate in total variation distance of the error in general convergence theorems. This is done in [29].

7.4.3. Invariance principles

As an application of the methodology mentioned above, V. Bally and coauthors have studied several limit theorems of Central Limit type - (see [14] and [15]). In particular they have estimate the total variation distance between random polynomials on one hand, and proved an universality principle for the variance of the number of roots of trigonometric polynomials with random coefficients, on the other hand.

7.4.4. Regularity of the law of the solution of jump type equations

V. Bally, L. Caramellino and G. Poly obtained some new regularity results for the solution of the 2 dimensional Boltzmann equation (see [13]). Moreover, in collaboration with L. Caramellino and A. Kohatsu Higa, V. Bally has started a research program on the regularity of the solutions of jump type equations. A first result in this sense is contained in [28].

7.4.5. Approximation of ARCH models

Benjamin Jourdain and Gilles Pagès (LPSM) are interested in proposing approximations of a sequence of probability measures in the convex order by finitely supported probability measures still in the convex order [35]. They propose to alternate transitions according to a martingale Markov kernel mapping a probability measure in the sequence to the next and dual quantization steps. In the case of ARCH models and in particular of the Euler scheme of a driftless Brownian diffusion, the noise has to be truncated to enable the dual quantization step. They analyze the error between the original ARCH model and its approximation with truncated noise and exhibit conditions under which the latter is dominated by the former in the convex order at the level of sample-paths. Last, they analyse the error of the scheme combining the dual quantization steps with truncation of the noise according to primal quantization.

7.4.6. Convergence of metadynamics

By drawing a parallel between metadynamics and self interacting models for polymers, B. Jourdain, T. Lelièvre (Cermics / ENPC) and P.-A. Zitt (LAMA) study the longtime convergence of the original metadynamics algorithm in the adiabatic setting, namely when the dynamics along the collective variables decouples from the dynamics along the other degrees of freedom. They also discuss the bias which is introduced when the adiabatic assumption does not holds [34].

7.4.7. Optimal transport

With V. Ehrlacher, D. Lombardi and R. Coyaud, Aurelien Alfonsi is working on numerical approximations of the optimal transport between two (or more) probability measures [26].

7.4.8. Generic approximation schemes for Markov semigroups.

A. Alfonsi and V. Bally have produced a general approximation scheme for Markov semigroups, based on random grids. This is a new approach to approximation schemes which is an alternative to the multi level method and the Romberg method [24].

7.4.9. Approximation with rough paths

A. Alfonsi and A. Kebaier are working on the approximation of some processes with rough paths.

MOKAPLAN Project-Team

5. New Results

5.1. An augmented Lagrangian approach to Wasserstein gradient flows and applications

Benamou, Jean-David and Carlier, Guillaume and Laborde, Maxime. In [2] : Taking advantage of the Benamou-Brenier dynamic formulation of optimal transport, we propose 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.

5.2. An entropy minimization approach to second-order variational mean-field games

Benamou, Jean-David and Carlier, Guillaume and Marino, Simone Di and Nenna, Luca. In [18] : We propose an entropy minimization viewpoint on variational meanfield games with diffusion and quadratic Hamiltonian. We carefully analyze the time-discretization of such problems, establish Γ -convergence results as the time step vanishes and propose an efficient algorithm relying on this entropic interpretation as well as on the Sinkhorn scaling algorithm.

5.3. Minimal convex extensions and finite difference discretization of the quadratic Monge-Kantorovich problem

Benamou, Jean-David and Duval, Vincent. In [3] : We propose an adaptation of the MA-LBR scheme to the Monge-Ampère equation with second boundary value condition, provided the target is a convex set. This yields a fast adaptive method to numerically solve the Optimal Transport problem between two absolutely continuous measures, the second of which has convex support. The proposed numerical method actually captures a specific Brenier solution which is minimal in some sense. We prove the convergence of the method as the grid stepsize vanishes and we show with numerical experiments that it is able to reproduce subtle properties of the Optimal Transport problem

We propose an entropy minimization viewpoint on variational meanfield games with diffusion and quadratic Hamiltonian. We carefully analyze the time-discretization of such problems, establish Γ -convergence results as the time step vanishes and propose an efficient algorithm relying on this entropic interpretation as well as on the Sinkhorn scaling algorithm.

5.4. Sparse analysis methods for Mesoscale Convective Systems (MCS)

Jean-Baptiste Courbot, Vincent Duval and Bernard Legras. In [20] : Mesoscale Convective Systems (MCS) are organized cloud systems which constitute the major part of the high cloud cover in the tropical region, where they can reach sizes of up to 1000 km in diameter. Under some favorable circumstances, they can eventually organize as tropical cyclones. Due to the complexity and the huge amount of available information, there is a need for an automatic tool able to detect and keep track of MCS in time series of these images. In [20] we have proposed a new method for the tracking of those cloud systems in infrared satellite images.

5.5. Pareto optimality with transport costs

Bruno Nazaret, Xavier Bacon, Guillaume Carlier, Thomas Gallouët. Arrived in September 2019, B. Nazaret, with his Phd student X. Bacon, has started a collaboration with G. Carlier on a Pareto optimality model with transport costs. He also has initiated a study with T.O. Gallouët of a first order proximal scheme for densities on Euclidean half-spaces, with the additional feature that the boundary moves while the mass still remains constant along the time.

5.6. An epigraphical approach to the representer theorem

Duval, Vincent. In [22] : Describing the solutions of inverse problems arising in signal or image processing is an important issue both for theoretical and numerical purposes. In [22], we have proposed a principle which describes the solutions to convex variational problems involving a finite number of measurements. We discuss its optimality on various problems concerning the recovery of Radon measures.

5.7. Approximate Optimal Designs for Multivariate Polynomial Regression

De Castro, Yohann et al. In [8] : We introduce a new approach aiming at computing approximate optimal designs for multivariate polynomial regressions on compact (semi-algebraic) design spaces. We use the moment-sum-of-squares hierarchy of semidefinite programming problems to solve numerically the approximate optimal design problem. The geometry of the design is recovered via semidefinite programming duality theory. This article shows that the hierarchy converges to the approximate optimal design as the order of the hierarchy increases. Furthermore, we provide a dual certificate ensuring finite convergence of the hierarchy and showing that the approximate optimal design can be computed numerically with our method. As a byproduct, we revisit the equivalence theorem of the experimental design theory: it is linked to the Christoffel polynomial and it characterizes finite convergence of the moment-sum-of-square hierarchies. Describing the solutions of inverse problems arising

5.8. Sparse Recovery from Extreme Eigenvalues Deviation Inequalities

De Castro, Yohann et al. In [7] : This article provides a new toolbox to derive sparse recovery guarantees. It is usually referred to as “stable and robust sparse regression” (SRSR) ’ from deviations on extreme singular values or extreme eigenvalues obtained in Random Matrix Theory. This work is based on Restricted Isometry Constants (RICs) which are a pivotal notion in Compressed Sensing and High-Dimensional Statistics as these constants finely assess how a linear operator is conditioned on the set of sparse vectors and hence how it performs in SRSR. While it is an open problem to construct deterministic matrices with apposite RICs, one can prove that such matrices exist using random matrices models. In this paper, we show upper bounds on RICs for Gaussian and Rademacher matrices using state-of-the-art deviation estimates on their extreme eigenvalues. This allows us to derive a lower bound on the probability of getting SRSR. One benefit of this paper is a direct and explicit derivation of upper bounds on RICs and lower bounds on SRSR from deviations on the extreme eigenvalues given by Random Matrix theory.

5.9. Some new Stein operators for product distributions

Mijoule G. et al. In [12] : We provide a general result for finding Stein operators for the product of two independent random variables whose Stein operators satisfy a certain assumption, extending a recent result of Gaunt, R. E., Mijoule, G. and Swan. This framework applies to non-centered normal and non-centered gamma random variables, as well as a general sub-family of the variance-gamma distributions. Curiously, there is an increase in complexity in the Stein operators for products of independent normals as one moves, for example, from centered to non-centered normals. As applications, we give a simple derivation of the characteristic function of the product of independent normals, and provide insight into why the probability density function of this distribution is much more complicated in the non-centered case than the centered case.

5.10. Simulation of multiphase porous media flows with minimizing movement and finite volume schemes

C. Cancès, T. O. Gallouët, M. Laborde, L. Monsaingeon. In [6]: the Wasserstein gradient flow structure of the PDE system governing multiphase flows in porous media was recently highlighted in [68]. The model can thus be approximated by means of the minimizing movement (or JKO) scheme. We solve the JKO scheme using the ALG2-JKO scheme proposed in [2]. The numerical results are compared to a classical upstream mobility Finite Volume scheme, for which strong stability properties can be established.

5.11. An unbalanced optimal transport splitting scheme for general advection-reaction-diffusion problems

T. O. Gallouët, M. Laborde, L. Monsaingeon. In [10] the authors show that unbalanced optimal transport provides a convenient framework to handle reaction and diffusion processes in a unified metric framework. We use a constructive method, alternating minimizing movements for the Wasserstein distance and for the Fisher-Rao distance, and prove existence of weak solutions for general scalar reaction-diffusion-advection equations. We extend the approach to systems of multiple interacting species, and also consider an application to a very degenerate diffusion problem involving a Gamma-limit. Moreover, some numerical simulations are included.

5.12. An unbalanced optimal transport splitting scheme for general advection-reaction-diffusion problems

C. Cancès, T. O. Gallouët, G. Todeschi. We propose a variational finite volume scheme to approximate the solutions to Wasserstein gradient flows. The time discretization is based on an implicit linearization of the Wasserstein distance expressed thanks to Benamou-Brenier formula, whereas space discretization relies on upstream mobility two-point flux approximation finite volumes. Our scheme is based on a first discretize then optimize approach in order to preserve the variational structure of the continuous model at the discrete level. Our scheme can be applied to a wide range of energies, guarantees non-negativity of the discrete solutions as well as decay of the energy. We show that our scheme admits a unique solution whatever the convex energy involved in the continuous problem, and we prove its convergence in the case of the linear Fokker-Planck equation with positive initial density. Numerical illustrations show that it is first order accurate in both time and space, and robust with respect to both the energy and the initial profile.

5.13. Generalized compressible fluid flows and solutions of the Camassa-Holm variational model

T. O. Gallouët, A. Natale, F-X. Vialard. In [11]: The Camassa-Holm equation on a domain $M \in \mathbb{R}^d$, in one of its possible multi-dimensional generalizations, describes geodesics on the group of diffeomorphisms with respect to the $H(\text{div})$ metric. It has been recently reformulated as a geodesic equation for the L^2 metric on a subgroup of the diffeomorphism group of the cone over M . We use such an interpretation to construct an analogue of Brenier's generalized incompressible Euler flows for the Camassa-Holm equation. This involves describing the fluid motion using probability measures on the space of paths on the cone, so that particles are allowed to split and cross. Differently from Brenier's model, however, we are also able to account for compressibility by employing an explicit probabilistic representation of the Jacobian of the flow map. We formulate the boundary value problem associated to the Camassa-Holm equation using such generalized flows. We prove existence of solutions and that, for short times, smooth solutions of the Camassa-Holm equations are the unique solutions of our model. We propose a numerical scheme to construct generalized solutions on the cone and present some numerical results illustrating the relation between the generalized Camassa-Holm and incompressible Euler solutions.

5.14. Metric completion of $\text{Diff}([0, 1])$ with the right-invariant H^1 metric

S. di Marino, A. Natale, R. Tahraoui, F-X. Vialard. In [21]: We consider the group of smooth increasing diffeomorphisms Diff on the unit interval endowed with the right-invariant H^1 metric. We compute the metric completion of this space which appears to be the space of increasing maps of the unit interval with boundary conditions at 0 and 1. We compute the lower-semicontinuous envelope associated with the length minimizing geodesic variational problem. We discuss the Eulerian and Lagrangian formulation of this relaxation and we show that smooth solutions of the EPDiff equation are length minimizing for short times.

5.15. Embedding Camassa-Holm equations in incompressible Euler

A. Natale, F-X. Vialard. In [13]: In this article, we show how to embed the so-called CH2 equations into the geodesic flow of the $H(\text{div})$ metric in 2D, which, itself, can be embedded in the incompressible Euler equation of a non compact Riemannian manifold. The method consists in embedding the incompressible Euler equation with a potential term coming from classical mechanics into incompressible Euler of a manifold and seeing the CH2 equation as a particular case of such fluid dynamic equation.

5.16. Second order models for optimal transport and cubic splines on the Wasserstein space

J-D. Benamou, T. O. Gallouët, F-X. Vialard. In [4]: On the space of probability densities, we extend the Wasserstein geodesics to the case of higher-order interpolation such as cubic spline interpolation. After presenting the natural extension of cubic splines to the Wasserstein space, we propose a simpler approach based on the relaxation of the variational problem on the path space. We explore two different numerical approaches, one based on multi-marginal optimal transport and entropic regularization and the other based on semi-discrete optimal transport.

QUANTIC Project-Team

6. New Results

6.1. Highly coherent spin states in carbon nanotubes coupled to cavity photons

Participants: Zaki Leghtas

Spins confined in quantum dots are considered as a promising platform for quantum information processing. While many advanced quantum operations have been demonstrated, experimental as well as theoretical efforts are now focusing on the development of scalable spin quantum bit architectures. One particularly promising method relies on the coupling of spin quantum bits to microwave cavity photons. This would enable the coupling of distant spins via the exchange of virtual photons for two qubit gate applications, which still remains to be demonstrated with spin qubits. Here, we use a circuit QED spin-photon interface to drive a single electronic spin in a carbon nanotube based double quantum dot using cavity photons. The microwave spectroscopy allows us to identify an electrically controlled spin transition with a decoherence rate which can be tuned to be as low as 250kHz. We show that this value is consistent with the expected hyperfine coupling in carbon nanotubes. These coherence properties, which can be attributed to the use of pristine carbon nanotubes stapled inside the cavity, should enable coherent spin-spin interaction via cavity photons and compare favourably to the ones recently demonstrated in Si-based circuit QED experiments. This experimental result is a collaboration between Zaki Leghtas (QUANTIC) and the group of Takis Kontos at ENS and was published in [13].

6.2. Escape of a Driven Quantum Josephson Circuit into Unconfined States

Participants: Raphaël Lescanne, Zaki Leghtas, Mazyar Mirrahimi and Lucas Verney

Josephson circuits have been ideal systems to study complex nonlinear dynamics that can lead to chaotic behavior and instabilities. More recently, Josephson circuits in the quantum regime, particularly in the presence of microwave drives, have demonstrated their ability to emulate a variety of Hamiltonians that are useful for the processing of quantum information. In this work, we show that these drives lead to an instability that results in the escape of the circuit mode into states that are not confined by the Josephson cosine potential. We observe this escape in a ubiquitous circuit: a transmon embedded in a 3D cavity. When the transmon occupies these free-particle-like states, the circuit behaves as though the junction had been removed and all nonlinearities are lost. This work deepens our understanding of strongly driven Josephson circuits, which is important for fundamental and application perspectives, such as the engineering of Hamiltonians by parametric pumping. This experimental work published in [17] demonstrates elements of the theory derived by [22].

6.3. Structural Instability of Driven Josephson Circuits Prevented by an Inductive Shunt

Participants: Raphaël Lescanne, Zaki Leghtas, Mazyar Mirrahimi and Lucas Verney

Superconducting circuits are a versatile platform to implement a multitude of Hamiltonians that perform quantum computation, simulation, and sensing tasks. A key ingredient for realizing a desired Hamiltonian is the irradiation of the circuit by a strong drive. These strong drives provide an in situ control of couplings, which cannot be obtained by near-equilibrium Hamiltonians. However, as shown in this theoretical study, out-of-equilibrium systems are easily plagued by complex dynamics, leading to instabilities. The prediction and prevention of these instabilities is crucial, both from a fundamental and application perspective. We propose an inductively shunted transmon as the elementary circuit optimized for strong parametric drives. Developing a numerical approach that avoids the built-in limitations of perturbative analysis, we demonstrate that adding the inductive shunt significantly extends the range of pump powers over which the circuit behaves in a stable manner. This theoretical result was published in [22] and analyzes the experiment [17].

6.4. Fast and virtually exact quantum gate generation in $U(n)$ via Iterative Lyapunov Methods

Participants: Pierre Rouchon

This work presents an iterative algorithm published in [20] and named RIGA for Reference Input Generation Algorithm. This algorithm constructs smooth control pulses for quantum gate preparations of closed quantum systems. It combines right translation invariance and Lyapunov trajectory tracking. It exhibits exponential convergence when the system is controllable. It can be seen as a closed-loop version of the widely used GRAPE algorithm. Two numerical case-studies borrowed from the recent literature are addressed. The first one is relative to a system of 10 coupled qubits with local controls. The second one considers a C-NOT gate generation involving the lower levels of two coupled nonlinear cavities (transmon-qubits).

6.5. Benchmarking maximum-likelihood state estimation with an entangled two-cavity state

Participants: Pierre Rouchon

The efficient quantum state reconstruction algorithm described in the PhD of Pierre Six, a former student of the Quantic team, (see [89]) is experimentally implemented on the non-local state of two microwave cavities entangled by a circular Rydberg atom. In [19], we use information provided by long sequences of measurements performed by resonant and dispersive probe atoms over time scales involving the system decoherence. Moreover, we benefit from the consolidation, in the same reconstruction, of different measurement protocols providing complementary information. Finally, we obtain realistic error bars for the matrix elements of the reconstructed density operator. These results demonstrate the pertinence and precision of the method, directly applicable to any complex quantum system.

6.6. Towards tight impossibility and possibility results for string stability

Participants: Alain Sarlette

This is the last step of the PhD thesis of Arash Farnam under the direction of A. Sarlette at Ghent University. The aim was to study which elements are really essential in so-called “string instability” results, which exist in several variants and with several assumptions. In [15], we have significantly extended the often frequency-based linear approach, by showing how the assumptions lead to string instability also in any nonlinear systems with reasonable bandwidth. This should allow to clarify that how the problem setting must be adapted in order to obtain more positive results. In [14], we show how other elements do not help, and we clarify how the knowledge of individual vehicles’ absolute velocity is key to enable strong versions of string stability, in conjunction with PID control. Previous studies had only considered weaker versions, with PD type control. A more detailed study of string stability with absolute velocity control has also been published in [25].

6.7. Stabilization of quantum systems under continuous non-demolition measurements

Participants: Gerardo Cardona, Alain Sarlette and Pierre Rouchon

The stabilization of quantum states or quantum subspaces using feedback signals from quantum non-demolition measurements is a basic control task; in discrete-time, this is the fundamental control property shown in the first quantum feedback experiment by Serge Haroche. In continuous-time, the problem is harder. So-called Markovian feedback can stabilize some states, but in particular the quantum non-demolition eigenstates which would be marginally stable under measurements, cannot be stabilized asymptotically with this technique. Stochastic control techniques, based on feedback from a full state estimator, have been proposed and analyzed to stabilize such eigenstates, proving convergence but not much more. In [12], we prove how a relatively simple controller, feeding back Wiener noise with a gain that depends on eigenstate populations, allows to exponentially stabilize the target eigenstate. This generalizes our previous results about the qubit.

In [24], we provide a similar scheme and convergence proof for stabilizing an invariant subspace of the measurement, namely the codespace of a repetition code for quantum error correction. To the best of our knowledge there was no convergence proof so far for stabilizing such subspaces on the basis of continuous measurements.

6.8. Modified Integral Control Globally Counters Symmetry-Breaking Biases

Participants: Alain Sarlette

This is the end of the PhD thesis of Zhifei Zhang, under joint supervision of A.Sarlette at Ghent University and Zhihao Ling at ECUST Shanghai. The work [23] builds on our earlier proposal of formulating integral control on nonlinear groups as the integral of proportional correcting feedback actions: since these actions belong to a Lie algebra, they can be integrated in this vector space and applied at the current point. We here show how this controller can be modified in order to recover coordinated motion among steering-controlled vehicles, in a situation where biases would make the standard controller fail. We prove how the simple addition of the integral controller allows to recover global convergence towards the coordinated motion, restoring symmetry exactly.

6.9. Quantum Fast-Forwarding: Markov Chains and graph property testing

Participants: Alain Sarlette

This is the end of the PhD thesis of S.Apers, under supervision of A.Sarlette at Ghent University. In [11], we propose a quantum algorithmic routine, called Quantum Fast-Forwarding, which allows to simulate a Markov chain quadratically faster on a quantum computer than on a classical one. The key novelty, from an application point of view, is that we can achieve this acceleration not only for reaching the asymptotic distribution, but also for any intermediate time that one would be interested in. Such transient behaviors of Markov chains are important in algorithmic context, for instance to distinguish clusters in graphs. We explicitly work out those applications on graph properties.

6.10. Quantum Adiabatic Elimination: extension to rotating systems

Participants: Paolo Forni, Timothée Launay, Alain Sarlette and Pierre Rouchon

Adiabatic elimination is a technique to eliminate fast converging variables of a large system, while retaining their impact on slower dynamics of interest. Its most extreme form is a standard procedure when neglecting the dynamics of e.g. actuators or measurement devices in dynamical systems. In quantum systems it is particularly relevant to eliminate subsystems in tensor product structure. However, a major constraint is to obtain a reduced system in quantum form (Lindblad equations), preserving positivity and the unit trace. After having set up the framework for quantum adiabatic elimination to arbitrary order as a series expansion during the thesis of Rémi Azouit, we had worked out first- and second-order Lindblad equations only. With Paolo Forni, we have been pursuing the development of explicit formulas for higher-order cases. In [26], we present an extension of the technique for the case where the slowly decaying subsystem of interest, is subject to fast Hamiltonian dynamics. This appears e.g. in systems with significant detunings, where a description in rotating frame would lead to time-dependent equations if one does not want to neglect fast oscillating terms.

6.11. Minimizing decoherence on target in bipartite open quantum systems

Participants: Paolo Forni and Alain Sarlette

We consider a target quantum system, coupled to an auxiliary quantum system which dissipates rapidly at somewhat adjustable rates. The goal is to minimize the dissipation induced on the target system by this coupling. In [27], we use explicit model reduction formulas to express this as a quadratic optimization problem. We prove that maybe counterintuitively, when the auxiliary system dissipates along Hermitian (entropy-increasing) channels, the minimum induced dissipation is reached by maximizing the dissipation rate of the auxiliary system. This may be interpreted as a dynamical decoupling among the target system and the auxiliary one, induced not by standard Hamiltonian control acting on the target, but by noise acting on the environment. This link has been pursued with PhD student Michiel Burgelman and should lead to further results next year.

6.12. Repetition Cat Qubits for Fault-Tolerant Quantum Computation

Participants: Jérémie Guillaud and Mazyar Mirrahimi

We present a 1D repetition code based on the so-called cat qubits as a viable approach toward hardware-efficient universal and fault-tolerant quantum computation. The cat qubits that are stabilized by a two-photon driven-dissipative process exhibit a tunable noise bias where the effective bit-flip errors are exponentially suppressed with the average number of photons. We propose a realization of a set of gates on the cat qubits that preserve such a noise bias. Combining these base qubit operations, we build, at the level of the repetition cat qubit, a universal set of fully protected logical gates. This set includes single-qubit preparations and measurements, not, controlled-not, and controlled-controlled-not (Toffoli) gates. Remarkably, this construction avoids the costly magic state preparation, distillation, and injection. Finally, all required operations on the cat qubits could be performed with slight modifications of existing experimental setups.

This result was recently published in *Physical Review X* [16].

6.13. Experimental Implementation of a Raman-Assisted Eight-Wave Mixing Process

Participants: Mazyar Mirrahimi

Nonlinear processes in the quantum regime are essential for many applications, such as quantum-limited amplification, measurement, and control of quantum systems. In particular, the field of quantum error correction relies heavily on high-order nonlinear interactions between various modes of a quantum system. However, the required order of nonlinearity is often not directly available or weak compared to dissipation present in the system. Here, following our earlier theoretical proposal [72] we experimentally demonstrate a route to obtain higher-order nonlinearity by combining more easily available lower-order nonlinear processes, using a generalization of the Raman transition. In particular, we show a transformation of four photons of a high-Q superconducting resonator into two excitations of a superconducting transmon mode and two pump photons, and vice versa. The resulting eight-wave mixing process is obtained by cascading two fourth-order nonlinear processes through a virtual state. We expect this type of process to become a key component of hardware-efficient quantum error correction using continuous-variable error-correction codes. This work in collaboration with the group of Michel Devoret at Yale university was published in [18].

6.14. Stabilized Cat in a Driven Nonlinear Cavity: A Fault-Tolerant Error Syndrome Detector

Participants: Philippe Campagne-Ibarcq and Mazyar Mirrahimi

In quantum error correction, information is encoded in a high-dimensional system to protect it from the environment. A crucial step is to use natural, two-body operations with an ancilla to extract information about errors without causing backaction on the encoded information. Essentially, ancilla errors must not propagate to the encoded system and induce errors beyond those which can be corrected. The current schemes for achieving this fault tolerance to ancilla errors come at the cost of increased overhead requirements. An efficient way to extract error syndromes in a fault-tolerant manner is by using a single ancilla with a strongly biased noise channel. Typically, however, required elementary operations can become challenging when the noise is extremely biased. In this collaborative work with the groups of Steven Girvin and Michel Devoret at Yale University, we propose to overcome this shortcoming by using a bosonic-cat ancilla in a parametrically driven nonlinear oscillator. Such a cat qubit experiences only bit-flip noise, while the phase flips are exponentially suppressed. To highlight the flexibility of this approach, we illustrate the syndrome extraction process in a variety of codes such as qubit-based toric, bosonic-cat, and Gottesman-Kitaev-Preskill codes. Our results open a path for realizing hardware-efficient, fault-tolerant error syndrome extraction. This work was published in [21].

SIERRA Project-Team

7. New Results

7.1. Fast Gradient Methods for Symmetric Nonnegative Matrix Factorization.

We describe fast gradient methods for solving the symmetric nonnegative matrix factorization problem (SymNMF). We use recent results on non-Euclidean gradient methods and show that the SymNMF problem is smooth relative to a well-chosen Bregman divergence. This approach provides a simple hyperparameter-free method which comes with theoretical convergence guarantees. We also discuss accelerated variants. Numerical experiments on clustering problems show that our algorithm scales well and reaches both state of the art convergence speed and clustering accuracy for SymNMF methods.

7.2. Naive Feature Selection: Sparsity in Naive Bayes.

Due to its linear complexity, naive Bayes classification remains an attractive supervised learning method, especially in very large-scale settings. We propose a sparse version of naive Bayes, which can be used for feature selection. This leads to a combinatorial maximum-likelihood problem, for which we provide an exact solution in the case of binary data, or a bound in the multinomial case. We prove that our bound becomes tight as the marginal contribution of additional features decreases. Both binary and multinomial sparse models are solvable in time almost linear in problem size, representing a very small extra relative cost compared to the classical naive Bayes. Numerical experiments on text data show that the naive Bayes feature selection method is as statistically effective as state-of-the-art feature selection methods such as recursive feature elimination, ℓ_1 -penalized logistic regression and LASSO, while being orders of magnitude faster. For a large data set, having more than with 1.6 million training points and about 12 million features, and with a non-optimized CPU implementation, our sparse naive Bayes model can be trained in less than 15 seconds.

7.3. Regularity as Regularization: Smooth and Strongly Convex Brenier Potentials in Optimal Transport.

The problem of estimating Wasserstein distances in high-dimensional spaces suffers from the curse of dimensionality: Indeed, one needs an exponential (w.r.t. dimension) number of samples for the distance between the two samples to be comparable to that between the two measures. Therefore, regularizing the optimal transport (OT) problem is crucial when using Wasserstein distances in machine learning. One of the greatest achievements of the OT literature in recent years lies in regularity theory: one can prove under suitable hypothesis that the OT map between two measures is Lipschitz, or, equivalently when studying 2-Wasserstein distances, that the Brenier convex potential (whose gradient yields an optimal map) is a smooth function. We propose in this work to go backwards, to adopt instead regularity as a regularization tool. We propose algorithms working on discrete measures that can recover nearly optimal transport maps that have small distortion, or, equivalently, nearly optimal Brenier potential that are strongly convex and smooth. For univariate measures, we show that computing these potentials is equivalent to solving an isotonic regression problem under Lipschitz and strong monotonicity constraints. For multivariate measures the problem boils down to a non-convex QCQP problem. We show that this QCQP can be lifted a semidefinite program. Most importantly, these potentials and their gradient can be evaluated on the measures themselves, but can more generally be evaluated on any new point by solving each time a QP. Building on these two formulations we propose practical algorithms to estimate and evaluate transport maps, and illustrate their performance statistically as well as visually on a color transfer task.

7.4. Ranking and synchronization from pairwise measurements via SVD.

Given a measurement graph $G = (V, E)$ and an unknown signal $r \in R^n$, we investigate algorithms for recovering r from pairwise measurements of the form $r_i - r_j; i, j$ in E . This problem arises in a variety of applications, such as ranking teams in sports data and time synchronization of distributed networks. Framed in the context of ranking, the task is to recover the ranking of n teams (induced by r) given a small subset of noisy pairwise rank offsets. We propose a simple SVD-based algorithmic pipeline for both the problem of time synchronization and ranking. We provide a detailed theoretical analysis in terms of robustness against both sampling sparsity and noise perturbations with outliers, using results from matrix perturbation and random matrix theory. Our theoretical findings are complemented by a detailed set of numerical experiments on both synthetic and real data, showcasing the competitiveness of our proposed algorithms with other state-of-the-art methods.

7.5. Polyak Steps for Adaptive Fast Gradient Methods.

Accelerated algorithms for minimizing smooth strongly convex functions usually require knowledge of the strong convexity parameter μ . In the case of an unknown μ , current adaptive techniques are based on restart schemes. When the optimal value f^* is known, these strategies recover the accelerated linear convergence bound without additional grid search. In this paper we propose a new approach that has the same bound without any restart, using an online estimation of strong convexity parameter. We show the robustness of the Fast Gradient Method when using a sequence of upper bounds on μ . We also present a good candidate for this estimate sequence and detail consistent empirical results.

7.6. An Accelerated Decentralized Stochastic Proximal Algorithm for Finite Sums.

Modern large-scale finite-sum optimization relies on two key aspects: distribution and stochastic updates. For smooth and strongly convex problems, existing decentralized algorithms are slower than modern accelerated variance-reduced stochastic algorithms when run on a single machine, and are therefore not efficient. Centralized algorithms are fast, but their scaling is limited by global aggregation steps that result in communication bottlenecks. In this work, we propose an efficient Accelerated Decentralized stochastic algorithm for Finite Sums named ADFS, which uses local stochastic proximal updates and randomized pairwise communications between nodes. On n machines, ADFS learns from nm samples in the same time it takes optimal algorithms to learn from m samples on one machine. This scaling holds until a critical network size is reached, which depends on communication delays, on the number of samples m , and on the network topology. We provide a theoretical analysis based on a novel augmented graph approach combined with a precise evaluation of synchronization times and an extension of the accelerated proximal coordinate gradient algorithm to arbitrary sampling. We illustrate the improvement of ADFS over state-of-the-art decentralized approaches with experiments.

7.7. On Lazy Training in Differentiable Programming.

In a series of recent theoretical works, it was shown that strongly overparameterized neural networks trained with gradient-based methods could converge exponentially fast to zero training loss, with their parameters hardly varying. In this work, we show that this “lazy training” phenomenon is not specific to overparameterized neural networks, and is due to a choice of scaling, often implicit, that makes the model behave as its linearization around the initialization, thus yielding a model equivalent to learning with positive-definite kernels. Through a theoretical analysis, we exhibit various situations where this phenomenon arises in non-convex optimization and we provide bounds on the distance between the lazy and linearized optimization paths. Our numerical experiments bring a critical note, as we observe that the performance of commonly used non-linear deep convolutional neural networks in computer vision degrades when trained in the lazy regime. This makes it unlikely that “lazy training” is behind the many successes of neural networks in difficult high dimensional tasks.

7.8. Implicit Regularization of Discrete Gradient Dynamics in Linear Neural Networks.

When optimizing over-parameterized models, such as deep neural networks, a large set of parameters can achieve zero training error. In such cases, the choice of the optimization algorithm and its respective hyper-parameters introduces biases that will lead to convergence to specific minimizers of the objective. Consequently, this choice can be considered as an implicit regularization for the training of over-parameterized models. In this work, we push this idea further by studying the discrete gradient dynamics of the training of a two-layer linear network with the least-squares loss. Using a time rescaling, we show that, with a vanishing initialization and a small enough step size, this dynamics sequentially learns the solutions of a reduced-rank regression with a gradually increasing rank.

7.9. Efficient Primal-Dual Algorithms for Large-Scale Multiclass Classification.

We develop efficient algorithms to train ℓ_1 -regularized linear classifiers with large dimensionality d of the feature space, number of classes k , and sample size n . Our focus is on a special class of losses that includes, in particular, the multiclass hinge and logistic losses. Our approach combines several ideas: (i) passing to the equivalent saddle-point problem with a quasi-bilinear objective; (ii) applying stochastic mirror descent with a proper choice of geometry which guarantees a favorable accuracy bound; (iii) devising non-uniform sampling schemes to approximate the matrix products. In particular, for the multiclass hinge loss we propose a sublinear algorithm with iterations performed in $O(d + n + k)$ arithmetic operations.

7.10. Fast and Faster Convergence of SGD for Over-Parameterized Models (and an Accelerated Perceptron).

Modern machine learning focuses on highly expressive models that are able to fit or interpolate the data completely, resulting in zero training loss. For such models, we show that the stochastic gradients of common loss functions satisfy a strong growth condition. Under this condition, we prove that constant step-size stochastic gradient descent (SGD) with Nesterov acceleration matches the convergence rate of the deterministic accelerated method for both convex and strongly-convex functions. We also show that this condition implies that SGD can find a first-order stationary point as efficiently as full gradient descent in non-convex settings. Under interpolation, we further show that all smooth loss functions with a finite-sum structure satisfy a weaker growth condition. Given this weaker condition, we prove that SGD with a constant step-size attains the deterministic convergence rate in both the strongly-convex and convex settings. Under additional assumptions, the above results enable us to prove an $O(1/k^2)$ mistake bound for k iterations of a stochastic perceptron algorithm using the squared-hinge loss. Finally, we validate our theoretical findings with experiments on synthetic and real datasets.

7.11. Globally Convergent Newton Methods for Ill-conditioned Generalized Self-concordant Losses.

In this project, we study large-scale convex optimization algorithms based on the Newton method applied to regularized generalized self-concordant losses, which include logistic regression and softmax regression. We first prove that our new simple scheme based on a sequence of problems with decreasing regularization parameters is provably globally convergent, that this convergence is linear with a constant factor which scales only logarithmically with the condition number. In the parametric setting, we obtain an algorithm with the same scaling than regular first-order methods but with an improved behavior, in particular in ill-conditioned problems. Second, in the non parametric machine learning setting, we provide an explicit algorithm combining the previous scheme with Nyström projection techniques, and prove that it achieves optimal generalization bounds with a time complexity of order $O(nd_{eff}(\lambda))$, a memory complexity of order $O(d_{eff}(\lambda)^2)$ and no dependence on the condition number, generalizing the results known for least-squares regression. Here n is the

number of observations and λ is the associated degrees of freedom. In particular, this is the first large-scale algorithm to solve logistic and softmax regressions in the non-parametric setting with large condition numbers and theoretical guarantees.

7.12. Efficient online learning with kernels for adversarial large scale problems

We are interested in a framework of online learning with kernels for low-dimensional but large-scale and potentially adversarial datasets. We study the computational and theoretical performance of online variations of kernel Ridge regression. Despite its simplicity, the algorithm we study is the first to achieve the optimal regret for a wide range of kernels with a per-round complexity of order n^α with $\alpha < 2$. The algorithm we consider is based on approximating the kernel with the linear span of basis functions. Our contributions is two-fold: 1) For the Gaussian kernel, we propose to build the basis beforehand (independently of the data) through Taylor expansion. For d -dimensional inputs, we provide a (close to) optimal regret of order $O((\log n)^{d+1})$ with per-round time complexity and space complexity $O((\log n)^{2d})$. This makes the algorithm a suitable choice as soon as $n \geq e^d$ which is likely to happen in a scenario with small dimensional and large-scale dataset; 2) For general kernels with low effective dimension, the basis functions are updated sequentially in a data-adaptive fashion by sampling Nyström points. In this case, our algorithm improves the computational trade-off known for online kernel regression

7.13. Affine Invariant Covariance Estimation for Heavy-Tailed Distributions

In this work we provide an estimator for the covariance matrix of a heavy-tailed multivariate distribution. We prove that the proposed estimator \widehat{S} admits an *affine-invariant* bound of the form

$$(1 - \epsilon)S \leq \widehat{S}(1 + \epsilon)S$$

in high probability, where S is the unknown covariance matrix, and \leq is the positive semidefinite order on symmetric matrices. The result only requires the existence of fourth-order moments, and allows for $\epsilon = O(\sqrt{k^4 d \log(d/\delta)/n})$ where k^4 is a measure of kurtosis of the distribution, d is the dimensionality of the space, n is the sample size, and $1-\delta$ is the desired confidence level. More generally, we can allow for regularization with level λ , then d gets replaced with the degrees of freedom number. Denoting $\text{cond}(S)$ the condition number of S , the computational cost of the novel estimator is $O(d^2 n + d^3 \log(\text{cond}(S)))$, which is comparable to the cost of the sample covariance estimator in the statistically interesting regime $n \geq d$. We consider applications of our estimator to eigenvalue estimation with relative error, and to ridge regression with heavy-tailed random design.

7.14. Beyond Least-Squares: Fast Rates for Regularized Empirical Risk Minimization through Self-Concordance

We consider learning methods based on the regularization of a convex empirical risk by a squared Hilbertian norm, a setting that includes linear predictors and non-linear predictors through positive-definite kernels. In order to go beyond the generic analysis leading to convergence rates of the excess risk as $O(\sqrt{1/n})$ from n observations, we assume that the individual losses are self-concordant, that is, their third-order derivatives are bounded by their second-order derivatives. This setting includes least-squares, as well as all generalized linear models such as logistic and softmax regression. For this class of losses, we provide a bias-variance decomposition and show that the assumptions commonly made in least-squares regression, such as the source and capacity conditions, can be adapted to obtain fast non-asymptotic rates of convergence by improving the bias terms, the variance terms or both.

7.15. Statistical Estimation of the Poincaré constant and Application to Sampling Multimodal Distributions

Poincaré inequalities are ubiquitous in probability and analysis and have various applications in statistics (concentration of measure, rate of convergence of Markov chains). The Poincaré constant, for which the inequality is tight, is related to the typical convergence rate of diffusions to their equilibrium measure. In this paper, we show both theoretically and experimentally that, given sufficiently many samples of a measure, we can estimate its Poincaré constant. As a by-product of the estimation of the Poincaré constant, we derive an algorithm that captures a low dimensional representation of the data by finding directions which are difficult to sample. These directions are of crucial importance for sampling or in fields like molecular dynamics, where they are called reaction coordinates. Their knowledge can leverage, with a simple conditioning step, computational bottlenecks by using importance sampling techniques.

7.16. Stochastic first-order methods: non-asymptotic and computer-aided analyses via potential functions

We provide a novel computer-assisted technique for systematically analyzing first-order methods for optimization. In contrast with previous works, the approach is particularly suited for handling sublinear convergence rates and stochastic oracles. The technique relies on semidefinite programming and potential functions. It allows simultaneously obtaining worst-case guarantees on the behavior of those algorithms, and assisting in choosing appropriate parameters for tuning their worst-case performances. The technique also benefits from comfortable tightness guarantees, meaning that unsatisfactory results can be improved only by changing the setting. We use the approach for analyzing deterministic and stochastic first-order methods under different assumptions on the nature of the stochastic noise. Among others, we treat unstructured noise with bounded variance, different noise models arising in over-parametrized expectation minimization problems, and randomized block-coordinate descent schemes.

7.17. Optimal Complexity and Certification of Bregman First-Order Methods

We provide a lower bound showing that the $O(1/k)$ convergence rate of the NoLips method (a.k.a. Bregman Gradient) is optimal for the class of functions satisfying the h -smoothness assumption. This assumption, also known as relative smoothness, appeared in the recent developments around the Bregman Gradient method, where acceleration remained an open issue. On the way, we show how to constructively obtain the corresponding worst-case functions by extending the computer-assisted performance estimation framework of Drori and Teboulle (*Mathematical Programming*, 2014) to Bregman first-order methods, and to handle the classes of differentiable and strictly convex functions.

7.18. Efficient First-order Methods for Convex Minimization: a Constructive Approach

We describe a novel constructive technique for devising efficient first-order methods for a wide range of large-scale convex minimization settings, including smooth, non-smooth, and strongly convex minimization. The technique builds upon a certain variant of the conjugate gradient method to construct a family of methods such that a) all methods in the family share the same worst-case guarantee as the base conjugate gradient method, and b) the family includes a fixed-step first-order method. We demonstrate the effectiveness of the approach by deriving optimal methods for the smooth and non-smooth cases, including new methods that forego knowledge of the problem parameters at the cost of a one-dimensional line search per iteration, and a universal method for the union of these classes that requires a three-dimensional search per iteration. In the strongly convex case, we show how numerical tools can be used to perform the construction, and show that the resulting method offers an improved worst-case bound compared to Nesterov's celebrated fast gradient method.

ANGE Project-Team

7. New Results

7.1. Numerical methods for fluid flows

7.1.1. *PARAOPT: A parareal algorithm for optimality systems*

Member: J. Salomon,

Coll.: Martin Gander, Felix Kwok

The time parallel solution of optimality systems arising in PDE constraint optimization could be achieved by simply applying any time parallel algorithm, such as Parareal, to solve the forward and backward evolution problems arising in the optimization loop. We propose in [21] a different strategy by devising directly a new time parallel algorithm, which we call ParaOpt, for the coupled forward and backward non-linear partial differential equations. ParaOpt is inspired by the Parareal algorithm for evolution equations, and thus is automatically a two-level method. We provide a detailed convergence analysis for the case of linear parabolic PDE constraints. We illustrate the performance of ParaOpt with numerical experiments both for linear and nonlinear optimality systems.

7.1.2. *Dynamical Behavior of a Nondiffusive Scheme for the Advection Equation*

Member: N. Aguillon,

Coll.: Pierre-Antoine Guihéneuf

In [16], we study the long time behaviour of a dynamical system strongly linked to the anti-diffusive scheme of Després and Lagoutiere for the 1-dimensional transport equation. This scheme is overcompressive when the Courant–Friedrichs–Levy number is $1/2$: when the initial data is nondecreasing, the approximate solution becomes a Heaviside function. In a special case, we also understand how plateaus are formed in the solution and their stability, a distinctive feature of the Després and Lagoutiere scheme.

7.1.3. *Convergence of numerical schemes for a conservation equation with convection and degenerate diffusion*

Member: C. Guichard,

Coll.: Robert Eymard, Xavier Lhébrard

In [20], the approximation of problems with linear convection and degenerate nonlinear diffusion, which arise in the framework of the transport of energy in porous media with thermodynamic transitions, is done using a θ -scheme based on the centered gradient discretisation method. The convergence of the numerical scheme is proved, although the test functions which can be chosen are restricted by the weak regularity hypotheses on the convection field, owing to the application of a discrete Gronwall lemma and a general result for the time translate in the gradient discretisation setting. Some numerical examples, using both the Control Volume Finite Element method and the Vertex Approximate Gradient scheme, show the role of θ for stabilising the scheme.

7.1.4. *Gradient-based optimization of a rotating algal biofilm process*

Members: N. Aguillon, J. Sainte-Marie,

Coll.: Pierre-Olivier Lamare, Jérôme Grenier, Hubert Bonnefond, Olivier Bernard

Microalgae are microorganisms that have only very recently been used for bio-technological applications and more specifically for the production of bio-fuel. In the report [15] we focus on the shape optimization and optimal control of an innovative process where microalgae are fixed on a support. They are successively exposed to light and darkness. The resulting growth rate can be represented by a dynamic system describing the denaturation of key proteins due to excess light. A Partial Derivative Equation (PDE) model for the Rotary Algae Biofilm (RAB) is proposed. It represents the local growth of microalgae subjected to time-varying light. A gradient method based on the calculation of the model adjoint is proposed to identify the optimal (constant) folding of the process and the (time-varying) speed of the biofilm. Once this method is used in a realistic case, the optimization results in a configuration that significantly improves productivity compared to the case where the biofilm is fixed.

7.2. Modelling

7.2.1. *Accurate steam-water equation of state for two-phase flow LMNC model with phase transition*

Member: Y. Penel,

Coll.: *Stéphane Dellacherie, Bérénice Grec, Gloria Faccanoni*

The paper [9] is dedicated to the design of incomplete equations of state for a two-phase flow with phase transition that are specific to the low Mach number regime. It makes use of the fact that the thermodynamic pressure has small variations in this regime. These equations of state supplement the 2D LMNC model introduced in previous works. This innovative strategy relies on tabulated values and is proven to satisfy crucial thermodynamic requirements such as positivity, monotonicity, continuity. In particular, saturation values are exact. This procedure is assessed by means of analytical steady solutions and comparisons with standard analytical equations of state, and shows a great improvement in accuracy.

7.2.2. *Numerical simulations of Serre - Green-Naghdi type models for dispersive free surface flows*

Members: Y. Penel, J. Sainte-Marie

Coll.: *Enrique D. Fernandez-Nieto, Tomas Morales de Luna, Cipriano Escalante Sanchez*

The Serre - Green-Naghdi equations are simulated under their non-hydrostatic formulation by means of a projection-correction method. This is then extended to the layerwise discretisation of the Euler equations with a special care to the computational cost. An original alternating direction method is used and relies on the tools designed for the monolayer case.

7.2.3. *Entropy-satisfying scheme for a hierarchy of dispersive reduced models of free surface flow*

Member: M. Parisot

The work [12] is devoted to the numerical resolution in the multidimensional framework of a hierarchy of reduced models of the water wave equations, such as the Serre-Green-Naghdi model. A particular attention is paid to the dissipation of mechanical energy at the discrete level, that act as a stability argument of the scheme, even with source terms such space and time variation of the bathymetry. In addition, the analysis leads to a natural way to deal with dry areas without leakage of energy. To illustrate the accuracy and the robustness of the strategy, several numerical experiments are carried out. In particular, the strategy is capable of treating dry areas without special treatment.

7.2.4. *Congested shallow water model: on floating body*

Members: E. Godlewski, M. Parisot, J. Sainte-Marie, F. Wahl

In [22], we are interested in the numerical modeling of body floating freely on the water such as icebergs or wave energy converters. The fluid-solid interaction is formulated using a congested shallow water model for the fluid and Newton's second law of motion for the solid. We make a particular focus on the energy transfer between the solid and the water since it is of major interest for energy production. A numerical approximation based on the coupling of a finite volume scheme for the fluid and a Newmark scheme for the solid is presented. An entropy correction based on an adapted choice of discretization for the coupling terms is made in order to ensure a dissipation law at the discrete level. Simulations are presented to verify the method and to show the feasibility of extending it to more complex cases.

7.2.5. Pseudo-compressibility, dispersive model and acoustic waves in shallow water flows

Members: E. Godlewski, M-O. Bristeau, J. Sainte-Marie

In this paper we study a dispersive shallow water type model derived from the compressible Navier-Stokes system. The compressible effects allow to capture the acoustic waves propagation and can be seen as a relaxation of an underlying incompressible model. Hence, the pseudo-compressibility terms circumvent the resolution of an elliptic equation for the non-hydrostatic part of the pressure. For the numerical approximation of shallow water type models, the hyperbolic part, often approximated using explicit time schemes, is constrained by a CFL condition. Since the approximation of the dispersive terms – implicit in time – generally requires the numerical resolution of an elliptic equation, it is very costly. In this paper, we show that when considering the pseudo-compressibility terms a fully explicit in time scheme can be derived. This drastically reduces the cost of the numerical resolution of dispersive models especially in 2d and 3d.

7.2.6. Some quasi-analytical solutions for propagative waves in free surface Euler equations

Members: B. Di Martino, M-O. Bristeau, J. Sainte-Marie, A. Mangeney, F. Souillé

This note describes some quasi-analytical solutions for wave propagation in free surface Euler equations and linearized Euler equations. The obtained solutions vary from a sinusoidal form to a form with singularities. They allow a numerical validation of the free-surface Euler codes.

7.2.7. Challenges and prospects for dynamical cores of oceanic models across all scales

Members: E. Audusse, J. Sainte-Marie

Review paper, more than 30 co-authors

The paper [11] provides an overview of the recent evolution and future challenges of oceanic models dynamical cores used for applications ranging from global paleoclimate scales to short-term prediction in estuaries and shallow coastal areas. The dynamical core is responsible for the discrete approximation in space and time of the resolved processes, as opposed to the physical parameterizations which represent unresolved or under-resolved processes. The paper reviews the challenges and prospects outlined by the modeling groups that participated to the Community for the Numerical Modeling of the Global, Regional, and Coastal Ocean (COMMODORE) workshop. The topics discussed in the paper originate from the experience acquired during the development of 16 dynamical cores representative of the variety of numerical methods implemented in models used for realistic ocean simulations. The topics of interest include the choice of model grid and variables arrangement, vertical coordinate, temporal discretization, and more practical aspects about the evolution of code architecture and development practices.

7.2.8. The Navier-Stokes system with temperature and salinity for free surface flows Part I: Low-Mach approximation & layer-averaged formulation

Members: M-O. Bristeau, L. Boittin, A. Mangeney, J. Sainte-Marie, F. Bouchut

In this paper, we are interested in free surface flows where density variations coming e.g. from temperature or salinity differences play a significant role in the hydrodynamic regime. In water, acoustic waves travel much faster than gravity and internal waves, hence the study of models arising in compressible fluid mechanics often requires a decoupling between these waves. Starting from the compressible Navier-Stokes system, we derive the so-called Navier-Stokes-Fourier system in an incompressible context (the density does not depend on the fluid pressure) using the low-Mach scaling. Notice that a modified low-Mach scaling is necessary to obtain a model with a thermo-mechanical compatibility. The case where the density depends only on the temperature is studied first. Then the variations of the fluid density with respect to the temperature and the salinity are considered. We give a layer-averaged formulation of the obtained models in an hydrostatic context. Allowing to derive numerical schemes endowed with strong stability properties – that are presented in a companion paper – the layer-averaged formulation is very useful for the numerical analysis and the numerical simulations of the models. Several stability properties of the layer-averaged Navier-Stokes-Fourier system are proved.

7.2.9. The Navier-Stokes system with temperature and salinity for free surface flows - Part II: Numerical scheme and validation

Members: M-O. Bristeau, L. Boittin, A. Mangeney, J. Sainte-Marie, F. Bouchut

In this paper, we propose a numerical scheme for the layer-averaged Euler with variable density and the Navier-Stokes-Fourier systems presented in part I. These systems model hydrostatic free surface flows with density variations. We show that the finite volume scheme presented is well balanced with regards to the steady state of the lake at rest and preserves the positivity of the water height. A maximum principle on the density is also proved as well as a discrete entropy inequality in the case of the Euler system with variable density. Some numerical validations are finally shown with comparisons to 3D analytical solutions and experiments.

7.3. Functional analysis of PDE models in Fluid Mechanics

7.3.1. On the rigid-lid approximation of shallow water Bingham model

Member: J. Sainte-Marie

Coll.: Bilal Al Taki, Khawla Msheik

The paper [17] discusses the well posedness of an initial value problem describing the motion of a Bingham fluid in a basin with a degenerate bottom topography. A physical interpretation of such motion is discussed. The system governing such motion is obtained from the Shallow Water-Bingham models in the regime where the Froude number degenerates, i.e taking the limit of such equations as the Froude number tends to zero. Since we are considering equations with degenerate coefficients, then we shall work with weighted Sobolev spaces in order to establish the existence of a weak solution. In order to overcome the difficulty of the discontinuity in Bingham's constitutive law, we follow a similar approach to that introduced in [G. Duvaut and J.-L. Lions, Springer-Verlag, 1976]. We study also the behavior of this solution when the yield limit vanishes. Finally, a numerical scheme for the system in 1D is furnished.

7.3.2. Global $bmo-1(\mathbb{R}^N)$ radially symmetric solution for compressible Navier-Stokes equations with initial density in $\mathbb{L}^\infty(\mathbb{R}^N)$

Member: B. Haspot

In [24], we investigate the question of the existence of global weak solution for the compressible Navier Stokes equations provided that the initial momentum belongs to $bmo-1(\mathbb{R}^N)$ with $N = 2, 3$ and is radially symmetric. We prove then a equivalent of the so-called Koch-Tataru theorem for the compressible Navier-Stokes equations. In addition we assume that the initial density is only bounded in $\mathbb{L}^\infty(\mathbb{R}^N)$, it allows us in particular to consider initial density admitting shocks. Furthermore we show that if the coupling between the density and the velocity is sufficiently strong, then the initial density which admits initially shocks is instantaneously regularizing inasmuch as the density becomes Lipschitz. To finish we prove the global existence of strong solution for large initial data provided that the initial data are radially symmetric and sufficiently regular in dimension $N = 2, 3$ for γ -law pressure.

7.3.3. New effective pressure and existence of global strong solution for compressible Navier-Stokes equations with general viscosity coefficient in one dimension

Member: Boris Haspot

Coll.: Cosmin Burtea

In this paper we prove the existence of global strong solution for the Navier-Stokes equations with general degenerate viscosity coefficients. The cornerstone of the proof is the introduction of a new effective pressure which allows to obtain an Oleinik-type estimate for the so called effective velocity. In our proof we make use of additional regularizing effects on the velocity which requires to extend the techniques developed by Hoff for the constant viscosity case.

7.4. Assessments of models by means of experimental data and assimilation

7.4.1. Metamodeling corrected by observational data

Members: V. Mallet, J. Hammond

An air quality model at urban scale computes the air pollutant concentrations at street resolution based on various emissions, meteorology, imported pollution and city geometry. Because of the computational cost of such model, we previously designed a metamodel using dimension reduction and statistical emulation, and then corrected this metamodel with observational data. Novel work was dedicated to the error modeling for a more balanced integration of the observations. The work was also applied to air quality simulation over Paris using several months of data.

7.4.2. Metamodeling of a complete air quality simulation chain

Members: A. Lesieur, V. Mallet

Coll.: Ruiwei Chen

With the objective of uncertainty quantification, we worked on the generation of a metamodel for the simulation of urban air quality, using a complete simulation chain including dynamic traffic assignment, the computation of air pollutant emissions and the dispersion of the pollutant in a city. The traffic model and the dispersion model are computationally costly and operate in high dimension. We employed dimension reduction, and coupled it with Kriging in order to build a metamodel for the complete simulation chain.

7.4.3. Artificial neural networks for the modeling of air pollution

Member: V. Mallet

Air quality simulations at national, continental or global scales are subject to large uncertainties which are typically mitigated by data assimilation techniques. Another approach to improve the forecasts is to design an error model, learning from historical discrepancies between simulations and observations. Such a model was built using an artificial neural network trained with many meteorological and geographical data. Further studies showed that the technique could successfully generate not only an error model (to improve pre-existing simulations), but also a complete model (without the need for pre-existing simulations) whose forecasts are more accurate than those of traditional models.

7.4.4. Uncertainty quantification in atmospheric dispersion of radionuclides

Members: V. Mallet,

Coll.: Irène Korsakissok

In collaboration with IRSN (Institute of Radiation Protection and Nuclear Safety), we investigated the uncertainties of the atmospheric-dispersion forecasts that are used during an accidental release of radionuclides such as the Fukushima disaster. These forecasts are subject to considerable uncertainties which originate from inaccurate weather forecasts, poorly known source term and modeling shortcomings. In order to quantify the uncertainties, we designed a metamodel and carried out the calibration of the metamodel input distributions using Markov chain Monte Carlo.

7.4.5. Meta-modeling for urban noise mapping

Members: A. Lesieur, V. Mallet

Coll.: Pierre Aumond, Arnaud Can

Noise computing software can require several hours to produce a map over an urban center for a given set of input data. This computational cost makes the models unsuitable for applications like uncertainty quantification or data assimilation where thousands of simulations, or more, can be required. One solution is to replace the physical model with a meta-model which is very fast and yet fairly reproduces the results of the physical model. The strategy is first to reduce the dimension of both inputs and outputs of the physical model, which leads to a reduced model. This reduced model is then replaced by a statistical emulator. The emulator is trained with calls to the reduced model for a set of chosen inputs. The emulator relies on the interpolation between the training output values.

7.4.6. Data assimilation for urban noise maps generated with a meta-model

Members: A. Lesieur, V. Mallet

Coll.: Pierre Aumond, Arnaud Can

In an urban area, it is increasingly common to have access to both a simulated noise map and a sensor network. A data assimilation algorithm is developed to combine data from both a noise map simulator and a network of acoustic sensors. One-hour noise maps are generated with a meta-model fed with hourly traffic and weather data. The data assimilation algorithm merges the simulated map with the sound level measurements into an improved noise map. The performance of this method relies on the accuracy of the meta-model, the input parameters selection and the model of the error covariance that describes how the errors of the simulated sound levels are correlated in space. The performance of the data assimilation is obtained with a leave-one-out cross-validation method.

7.4.7. Uncertainty quantification in wildland fire propagation

Members: F. Allaire, V. Mallet

Coll.: Jean-Baptiste Filippi

We worked further on the Monte Carlo simulation of wildland fires. We calibrated the input distributions that represent the uncertainties in the inputs of our fire spread predictions by using the observations of the final contours for a number of fire cases. We used a new metric to measure the dissimilarity between two burned surfaces that relies on the Wasserstein distance. We designed a metamodel and carried out the calibration of the model input distributions using Markov chain Monte Carlo.

7.4.8. A non-intrusive reduced order data assimilation method applied to the monitoring of urban flows

Member: J. Hammond

Coll.: R. Chakir

In [13], we investigate a variational data assimilation method to rapidly estimate urban pollutant concentration around an area of interest using measurement data and CFD based models in a non-intrusive and computationally efficient manner. In case studies presented here, we used a sample of solutions from a dispersion model with varying meteorological conditions and pollution emissions to build a Reduced Basis approximation space and combine it with concentration observations. The method allows to correct for unmodeled physics, while significantly reducing online computational time.

7.5. Software Developments

Members: C., A. El Baz, J. Sainte-Marie

Several improvements of FreshKiss3D software have been made:

1. the code can now be used on RPM-based systems (Fedora, Redhat) thanks to a user contribution (Julien Jerphanion);
2. the conda-based installation steps are now automatically tested by means of a continuous integration process performed on Linux and Mac virtual machines provided by the Inria platform <https://ci.inria.fr/>;
3. third-party libraries have been updated so as to benefit from their very latest features;
4. software quality is now monitored by static analysis via the Inria SonarQube platform (<https://sonarqube.inria.fr/>);
5. code optimization including parallelization with MPI is currently under development.

ARAMIS Project-Team

7. New Results

7.1. Predicting PET-derived Demyelination from Multimodal MRI using Sketcher-Refiner Adversarial Training for Multiple Sclerosis

Participants: Wen Wei, Emilie Poirion, Benedetta Bodini, Stanley Durrleman, Nicholas Ayache, Bruno Stankoff, Olivier Colliot [Correspondant].

Multiple sclerosis (MS) is the most common demyelinating disease. In MS, demyelination occurs in the white matter of the brain and in the spinal cord. It is thus essential to measure the tissue myelin content to understand the physiopathology of MS, track progression and assess treatment efficacy. Positron emission tomography (PET) with [11C]PIB is a reliable method to measure myelin content in vivo. However, the availability of PET in clinical centers is limited. Moreover, it is expensive to acquire and invasive due to the injection of a radioactive tracer. By contrast, MR imaging is non-invasive, less expensive and widely available, but conventional MRI sequences cannot provide a direct and reliable measure of myelin. In this work, we therefore propose, to the best of our knowledge for the first time, a method to predict the PET-derived myelin content map from multimodal MRI. To that purpose, we introduce a new approach called Sketcher-Refiner generative adversarial networks (GANs) with specifically designed adversarial loss functions. The first network (Sketcher) generates global anatomical and physiological information. The second network (Refiner) refines and generates the tissue myelin content. A visual attention saliency map is also proposed to interpret the attention of neural networks. Our approach is shown to outperform the state-of-the-art methods in terms of image quality and myelin content prediction. Particularly, our prediction results show similar results to the PET-derived gold standard at both global and voxel-wise levels indicating the potential for clinical management of patients with MS.

More details in [25].

7.2. Reproducible evaluation of methods for predicting progression to Alzheimer's disease from clinical and neuroimaging data

Participants: Jorge Samper-González, Ninon Burgos, Simona Bottani, Marie-Odile Habert, Stéphane Epelbaum, Theodoros Evgeniou, Olivier Colliot [Correspondant].

Various machine learning methods have been proposed for predicting progression of patients with mild cognitive impairment (MCI) to Alzheimer's disease (AD) using neuroimaging data. Even though the vast majority of these works use the public dataset ADNI, reproducing their results is complicated because they often do not make available elements that are essential for reproducibility, such as selected participants and input data, image preprocessing and cross-validation procedures. Comparability is also an issue. Specially, the influence of different components like preprocessing, feature extraction or classification algorithms on the performance is difficult to evaluate. Finally, these studies rarely compare their results to models built from clinical data only, a critical aspect to demonstrate the utility of neuroimaging. In our previous work, 1, 2 we presented a framework for reproducible and objective classification experiments in AD, that included automatic conversion of ADNI database into the BIDS community standard, image preprocessing pipelines and machine learning evaluation. We applied this framework to perform unimodal classifications of T1 MRI and FDG-PET images. In the present paper, we extend this work to the combination of multimodal clinical and neuroimaging data. All experiments are based on standard approaches (namely SVM and random forests). In particular, we assess the added value of neuroimaging over using only clinical data. We first demonstrate that using only demographic and clinical data (gender, education level, MMSE, CDR sum of boxes, ADASCog) results in a balanced accuracy of 75% (AUC of 0.84). This performance is higher than that of standard

neuroimaging-based classifiers. We then propose a simple trick to improve the performance of neuroimaging-based classifiers: training from AD patients and controls (rather than from MCI patients) improves the performance of FDG-PET classification by 5 percent points, reaching the level of the clinical classifier. Finally, combining clinical and neuroimaging data, prediction results further improved to 80% balanced accuracy and an AUC of 0.88). These prediction accuracies, obtained in a reproducible way, provide a base to develop on top of it and, to compare against, more sophisticated methods. All the code of the framework and the experiments is publicly available at <https://github.com/aramis-lab/AD-ML>.

More details in [34].

7.3. Disrupted core-periphery structure of multimodal brain networks in Alzheimer's disease

Participants: Jeremy Guillon, Mario Chavez, Federico Battiston, Yohan Attal, Valentina Corte, Michel Thiebaut de Schotten, Bruno Dubois, Denis Schwartz, Olivier Colliot, Fabrizio de Vico Fallani [Correspondant].

In Alzheimer's disease (AD), the progressive atrophy leads to aberrant network reconfigurations both at structural and functional levels. In such network reorganization, the core and peripheral nodes appear to be crucial for the prediction of clinical outcome because of their ability to influence large-scale functional integration. However, the role of the different types of brain connectivity in such prediction still remains unclear. Using a multiplex network approach we integrated information from DWI, fMRI, and MEG brain connectivity to extract an enriched description of the core-periphery structure in a group of AD patients and age-matched controls. Globally, the regional coreness—that is, the probability of a region to be in the multiplex core—significantly decreased in AD patients as result of a random disconnection process initiated by the neurodegeneration. Locally, the most impacted areas were in the core of the network—including temporal, parietal, and occipital areas—while we reported compensatory increments for the peripheral regions in the sensorimotor system. Furthermore, these network changes significantly predicted the cognitive and memory impairment of patients. Taken together these results indicate that a more accurate description of neurodegenerative diseases can be obtained from the multimodal integration of neuroimaging-derived network data.

More details in [20]

7.4. Network neuroscience for optimizing brain-computer interfaces

Participants: Fabrizio de Vico Fallani [Correspondant], Danielle Bassett.

Human-machine interactions are being increasingly explored to create alternative ways of communication and to improve our daily life. Based on a classification of the user's intention from the user's underlying neural activity, brain-computer interfaces (BCIs) allow direct interactions with the external environment while bypassing the traditional effector of the musculoskeletal system. Despite the enormous potential of BCIs, there are still a number of challenges that limit their societal impact, ranging from the correct decoding of a human's thoughts, to the application of effective learning strategies. Despite several important engineering advances, the basic neuroscience behind these challenges remains poorly explored. Indeed, BCIs involve complex dynamic changes related to neural plasticity at a diverse range of spatiotemporal scales. One promising antidote to this complexity lies in network science, which provides a natural language in which to model the organizational principles of brain architecture and function as manifest in its interconnectivity. Here, we briefly review the main limitations currently affecting BCIs, and we offer our perspective on how they can be addressed by means of network theoretic approaches. We posit that the emerging field of network neuroscience will prove to be an effective tool to unlock human-machine interactions.

More details in [13]

7.5. Quality Assessment of Single-Channel EEG for Wearable Devices

Participants: Fanny Grosselin, Xavier Navarro-Sune, Alessia Vozzi, Katerina Pandremmenou, Fabrizio de Vico Fallani, Yohan Attal, Mario Chavez [Correspondant].

The recent embedding of electroencephalographic (EEG) electrodes in wearable devices raises the problem of the quality of the data recorded in such uncontrolled environments. These recordings are often obtained with dry single-channel EEG devices, and may be contaminated by many sources of noise which can compromise the detection and characterization of the brain state studied. In this paper, we propose a classification-based approach to effectively quantify artefact contamination in EEG segments, and discriminate muscular artefacts. The performance of our method were assessed on different databases containing either artificially contaminated or real artefacts recorded with different type of sensors, including wet and dry EEG electrodes. Furthermore, the quality of unlabelled databases was evaluated. For all the studied databases, the proposed method is able to rapidly assess the quality of the EEG signals with an accuracy higher than 90

More details in [19]

7.6. Reduction of recruitment costs in preclinical AD trials. Validation of automatic pre-screening algorithm for brain amyloidosis

Participants: Manon Ansart [correspondant], Stéphane Epelbaum, Geoffroy Gagliardi, Olivier Colliot, Didier Dormont, Bruno Dubois, Harald Hampel, Stanley Durrleman.

We propose a method for recruiting asymptomatic Amyloid positive individuals in clinical trials, using a two-step process. We first select during a pre-screening phase a subset of individuals which are more likely to be amyloid positive based on the automatic analysis of data acquired during routine clinical practice, before doing a confirmatory PET-scan to these selected individuals only. This method leads to an increased number of recruitments and to a reduced number of PET-scans, resulting in a decrease in overall recruitment costs. We validate our method on 3 different cohorts, and consider 5 different classification algorithms for the pre-screening phase. We show that the best results are obtained using solely cognitive, genetic and socio-demographic features, as the slight increased performance when using MRI or longitudinal data is balanced by the cost increase they induce. We show that the proposed method generalizes well when tested on an independent cohort, and that the characteristics of the selected set of individuals are identical to the characteristics of a population selected in a standard way. The proposed approach shows how Machine Learning can be used effectively in practice to optimize recruitment costs in clinical trials.

More details in[7]

7.7. Learning low-dimensional representations of shape data sets with diffeomorphic autoencoders

Participants: Alexandre Bône [Correspondant], Maxime Louis, Olivier Colliot, Stanley Durrleman.

Contemporary deformation-based morphometry offers parametric classes of diffeomorphisms that can be searched to compute the optimal transformation that warps a shape into another, thus defining a similarity metric for shape objects. Extending such classes to capture the geometrical variability in always more varied statistical situations represents an active research topic. This quest for genericity however leads to computationally-intensive estimation problems. Instead, we propose in this work to learn the best-adapted class of diffeomorphisms along with its parametrization, for a shape data set of interest. Optimization is carried out with an auto-encoding variational inference approach, offering in turn a coherent model-estimator pair that we name diffeomorphic auto-encoder. The main contributions are: (i) an original network-based method to construct diffeomorphisms, (ii) a current-splating layer that allows neural network architectures to process meshes, (iii) illustrations on simulated and real data sets that show differences in the learned statistical distributions of shapes when compared to a standard approach.

More details in[30]

7.8. Learning disease progression models with longitudinal data and missing values

Participants: Raphaël Couronné [correspondant], Marie Vidailhet, Jean-Christophe Corvol, Stéphane Lehéricy, Stanley Durrleman.

Statistical methods have been developed for the analysis of longitudinal data in neurodegenerative diseases. To cope with the lack of temporal markers- i.e. to account for subject-specific disease progression in regard to age- a common strategy consists in realigning the individual sequence data in time. Patient's specific trajectories can indeed be seen as spatiotemporal perturbations of the same normative disease trajectory. However, these models do not easily allow one to account for multimodal data, which more than often include missing values. Indeed, it is rare that imaging and clinical examinations for instance are performed at the same frequency in clinical protocols. Multimodal models also need to allow a different profile of progression for data with different structure and representation. We propose to use a generative mixed effect model that considers the progression trajectories as curves on a Riemannian Manifold. We use the concept of product manifold to handle multimodal data, and leverage the generative aspect of our model to handle missing values. We assess the robustness of our methods toward missing values frequency on both synthetic and real data. Finally we apply our model on a real-world dataset to model Parkinson's disease progression from data derived from clinical examination and imaging.

More details in[31]

7.9. Learning the clustering of longitudinal shape data sets into a mixture of independent or branching trajectories

Participants: Vianney Debavelaere [correspondant], Stéphanie Allasonnière, Stanley Durrleman.

Given repeated observations of several subjects over time, i.e. a longitudinal data set, this work introduces a new model to learn a classification of the shapes progression in an unsupervised setting: we automatically cluster a longitudinal data set in different classes without labels. Our method learns for each cluster an average shape trajectory (or representative curve) and its variance in space and time. Representative trajectories are built as the combination of pieces of curves. This mixture model is flexible enough to handle independent trajectories for each cluster as well as fork and merge scenarios. The estimation of such non linear mixture models in high dimension is known to be difficult because of the trapping states effect that hampers the optimisation of cluster assignments during training. We address this issue by using a tempered version of the stochastic EM algorithm. Finally, we apply our algorithm on different data sets. First, synthetic data are used to show that a tempered scheme achieves better convergence. We then apply our method to different real data sets: 1D RECIST score used to monitor tumors growth, 3D facial expressions and meshes of the hippocampus. In particular, we show how the method can be used to test different scenarios of hippocampus atrophy in ageing by using an heterogeneous population of normal ageing individuals and mild cognitive impaired subjects.

More details in[32]

7.10. Auto-encoding meshes of any topology with the current-splatting and exponentiation layers

Participants: Alexandre Bône [Correspondant], Olivier Colliot, Stanley Durrleman.

Deep learning has met key applications in image computing, but still lacks processing paradigms for meshes, i.e. collections of elementary geometrical parts such as points, segments or triangles. Meshes are both a powerful representation for geometrical objects, and a challenge for network architectures because of their inherent irregular structure. This work contributes to adapt classical deep learning paradigms to this particular type of data in three ways. First, we introduce the current-splatting layer which embeds meshes in a metric space, allowing the downstream network to process them without any assumption on their topology: they may be composed of varied numbers of elements or connected components, contain holes, or bear high

levels of geometrical noise. Second, we adapt to meshes the exponentiation layer which, from an upstream image array, generates shapes with a diffeomorphic control over their topology. Third, we take advantage of those layers to devise a variational auto-encoding architecture, which we interpret as a generative statistical model that learns adapted low-dimensional representations for mesh data sets. An explicit norm-control layer ensures the correspondence between the latent-space Euclidean metric and the shape-space log-Euclidean one. We illustrate this method on simulated and real data sets, and show the practical relevance of the learned representation for visualization, classification and mesh synthesis.

More details in[29]

7.11. Riemannian Geometry Learning for Disease Progression Modelling

Participants: Maxime Louis, Raphael Couronne, Igor Koval, Benjamin Charlier, Stanley Durrleman.

The analysis of longitudinal trajectories is a longstanding problem in medical imaging which is often tackled in the context of Riemannian geometry: the set of observations is assumed to lie on an a priori known Riemannian manifold. When dealing with high-dimensional or complex data, it is in general not possible to design a Riemannian geometry of relevance. In this work, we perform Riemannian manifold learning in association with the statistical task of longitudinal trajectory analysis. After inference, we obtain both a submanifold of observations and a Riemannian metric so that the observed progressions are geodesics. This is achieved using a deep generative network, which maps trajectories in a low-dimensional Euclidean space to the observation space.

More details in[33]

7.12. How many patients are eligible for disease-modifying treatment in Alzheimer's disease? A French national observational study over 5 years.

Participants: Stéphane Epelbaum [Correspondant], Claire Paquet, Jacques Hugon, Julien Dumurgier, David Wallon, Didier Hannequin, Thérèse Jonveaux, Annick Besozzi, Stéphane Pouponneau, Caroline Hommet, Frédéric Blanc, Laetitia Berly, Adrien Julian, Marc Paccalin, Florence Pasquier, Julie Bellet, Claire Boutoleau-Bretonniere, Tiphaine Charriau, Olivier Rouaud, Olivier Madec, Aurélie Mouton, Renaud David, Samir Bekadar, Roxanne Fabre, Emmanuelle Liegey, Walter Deberdt, Philippe Robert, Bruno Dubois.

We aimed to study the epidemiology of the prodromal and mild stages of Alzheimer's disease (AD) patients who are eligible for clinical trials with disease-modifying therapies. We analyzed two large complementary databases to study the incidence and characteristics of this population on a nationwide scope in France from 2014 to 2018. The National Alzheimer Database contains data from 357 memory centres and 90 private neurologists. Data from 2014 to 2018 have been analyzed. Patients, 50–85 years old, diagnosed with AD who had an Mini-Mental State Exam (MMSE) score greater or equal to 20 were included. We excluded patients with mixed and non-AD neurocognitive disorders. Descriptive statistics of the population of interest was the primary measure. Results In the National Alzheimer Database, 550,198 patients were assessed. Among them, 72,174 (13.1%) were diagnosed with AD and had an MMSE greater or equal to 20. Using corrections for specificity of clinical diagnosis of AD, we estimated that about 50,000 (9.1%) had a prodromal or mild AD. In the combined electronic clinical records database of 11 French expert memory centres, a diagnosis of prodromal or mild AD, certified by the use of cerebrospinal fluid AD biomarkers, could be established in 195 (1.3%) out of 14 596 patients. AD was not frequently diagnosed at a prodromal or mild dementia stage in France in 2014 to 2018. Diagnosis rarely relied on a pathophysiological marker even in expert memory centres. National databases will be valuable to monitor early stage AD diagnosis efficacy in memory centres when a disease-modifying treatment becomes available. More details in [15]

7.13. EEG evidence of compensatory mechanisms in preclinical Alzheimer's disease

Participants: Sinead Gaubert, Federico Raimondo, Marion Houot, Marie-Constance Corsi, Jacobo Diego Sitt, Bertrand Hermann, Delphine Oudiette, Geoffroy Gagliardi, Marie Odile Habert, Bruno Dubois, Fabrizio de Vico Fallani, Hovagim Bakardjian, Stéphane Epelbaum [Correspondant].

Early biomarkers are needed to identify individuals at high risk of preclinical Alzheimer's disease and to better understand the pathophysiological processes of disease progression. Preclinical Alzheimer's disease EEG changes would be non-invasive and cheap screening tools and could also help to predict future progression to clinical Alzheimer's disease. However, the impact of amyloid-beta deposition and neurodegeneration on EEG biomarkers needs to be elucidated. We included participants from the INSIGHT-preAD cohort, which is an ongoing single-centre multimodal observational study that was designed to identify risk factors and markers of progression to clinical Alzheimer's disease in 318 cognitively normal individuals aged 70-85 years with a subjective memory complaint. We divided the subjects into four groups, according to their amyloid status (based on 18F-florbetapir PET) and neurodegeneration status (evidenced by 18F-fluorodeoxyglucose PET brain metabolism in Alzheimer's disease signature regions). The first group was amyloid-positive and neurodegeneration-positive, which corresponds to stage 2 of preclinical Alzheimer's disease. The second group was amyloid-positive and neurodegeneration-negative, which corresponds to stage 1 of preclinical Alzheimer's disease. The third group was amyloid-negative and neurodegeneration-positive, which corresponds to 'suspected non-Alzheimer's pathophysiology'. The last group was the control group, defined by amyloid-negative and neurodegeneration-negative subjects. We analysed 314 baseline 256-channel high-density eyes closed 1-min resting state EEG recordings. EEG biomarkers included spectral measures, algorithmic complexity and functional connectivity assessed with a novel information-theoretic measure, weighted symbolic mutual information. The most prominent effects of neurodegeneration on EEG metrics were localized in frontocentral regions with an increase in high frequency oscillations (higher beta and gamma power) and a decrease in low frequency oscillations (lower delta power), higher spectral entropy, higher complexity and increased functional connectivity measured by weighted symbolic mutual information in theta band. Neurodegeneration was associated with a widespread increase of median spectral frequency. We found a non-linear relationship between amyloid burden and EEG metrics in neurodegeneration-positive subjects, either following a U-shape curve for delta power or an inverted U-shape curve for the other metrics, meaning that EEG patterns are modulated differently depending on the degree of amyloid burden. This finding suggests initial compensatory mechanisms that are overwhelmed for the highest amyloid load. Together, these results indicate that EEG metrics are useful biomarkers for the preclinical stage of Alzheimer's disease.

More details in [17]

7.14. Latent class analysis identifies functional decline with Amsterdam IADL in preclinical Alzheimer's disease

Participants: Sarah-Christine Villeneuve, Marion Houot, Federica Cacciamani, Merike Verrijp, Marie Odile Habert, Bruno Dubois, Sietske Sikkes, Stéphane Epelbaum [Correspondant].

Trials in Alzheimer's disease (AD) now include participants at the earliest stages to prevent further decline. However, the lack of tools sensitive to subtle functional changes in early-stage AD hinders the development of new therapies as it is difficult to prove their clinical relevance. We assessed functional changes over three years in 289 elderly memory complainers from the Investigation of Alzheimer's Predictors in subjective memory complainers cohort using the Amsterdam Instrumental-Activities-of-Daily-Living questionnaire (A-IADL-Q). No overall functional decline related to AD imaging markers was evidenced. However, five distinct classes of A-IADL-Q trajectories were identified. The largest class (212 [73.4%]) had stable A-IADL-Q scores over 3 years. A second group (23 [8.0%]) showed a persistent functional decline, higher amyloid load ($P < .0005$), and lower education ($P < .0392$). The A-IADL-Q identified a subtle functional decline in asymptomatic at-risk AD individuals. This could have important implications in the field of early intervention in AD

More details in [24]

COMMEDIA Project-Team

7. New Results

7.1. Cardiovascular hemodynamics

Participant: Miguel Ángel Fernández Varela.

Mitral regurgitation is one of the most prevalent valvular heart disease. Proper evaluation of its severity is necessary to choose appropriate treatment. The PISA method, based on Color Doppler echocardiography, is widely used in the clinical setting to estimate various relevant quantities related to the severity of the disease. In [19], the use of a pipeline to quickly generate image-based numerical simulation of intracardiac hemodynamics is investigated. The pipeline capabilities are evaluated on a database of twelve volunteers. Full pre-processing is achieved completely automatically in 55 minutes, on average, with small registration errors compared to the image spatial resolution. This pipeline is then used to study the intracardiac hemodynamics in the presence of diseased mitral valve. A strong variability among the simulated cases, mainly due to the valve geometry and regurgitation specifics, is found. The results from those numerical simulations is used to assess the potential limitations of the PISA method with respect to different MR types. While the PISA method provides reasonable estimates in the case of a simple circular regurgitation, it is shown that unsatisfying estimates are obtained in the case of non-circular leakage. Moreover, it is shown that the choice of high aliasing velocities can lead to difficulties in quantifying MR.

7.2. Respiratory flows

Participant: Céline Grandmont.

In [21], we propose a coupled fluid-kinetic model taking into account the radius growth of aerosol particles due to humidity in the respiratory system. We aim to numerically investigate the impact of hygroscopic effects on the particle behaviour. The air flow is described by the incompressible Navier-Stokes equations, and the aerosol by a Vlasov-type equation involving the airhumidity and temperature, both quantities satisfying a convection-diffusion equation with a source term. Conservations properties are checked and an explicit time-marching scheme is proposed. Two-dimensional numerical simulations in a branched structure show the influence of the particle size variations on the aerosol dynamics.

7.3. Fluid flow reconstruction from medical imaging

Participants: Muriel Boulakia, Miguel Ángel Fernández Varela, Felipe Galarce Marin, Damiano Lombardi, Olga Mula, Colette Voisembert.

In [22], we are interested in designing and analyzing a finite element data assimilation method for laminar steady flow described by the linearized incompressible Navier-Stokes equation. We propose a weakly consistent stabilized finite element method which reconstructs the whole fluid flow from velocity measurements in a subset of the computational domain. Using the stability of the continuous problem in the form of a three balls inequality, we derive quantitative local error estimates for the velocity. Numerical simulations illustrate these convergences properties and we finally apply our method to the flow reconstruction in a blood vessel.

In [29] a state estimation problem is investigated, that consists in reconstructing the blood flow from ultrasound Doppler images. The method proposed is based on a reduced-order technique. Semi-realistic 3D configurations are tested.

7.4. Safety pharmacology

Participants: Damiano Lombardi, Fabien Raphael.

In [31] a greedy method is used to classify molecules action on cardiac myocytes. The method is applied to a realistic dataset: the experiments were performed at Ncardia (Netherlands). The experimental dataset is complemented by a synthetic dataset, obtained by simulating experimental meaningful scenarios. The results obtained are encouraging.

7.5. Mathematical analysis of PDEs

Participants: Muriel Boulakia, Jean-Jerome Casanova, Céline Grandmont.

In [23], we consider a reaction-diffusion equation where the reaction term is given by a cubic function and we are interested in the numerical reconstruction of the time-independent part of the source term from measurements of the solution. For this identification problem, we present an iterative algorithm based on Carleman estimates which consists of minimizing at each iteration strongly convex cost functionals. Despite the nonlinear nature of the problem, we prove that our algorithm globally converges and the convergence speed evaluated in weighted norm is linear. In the last part of the paper, we illustrate the effectiveness of our algorithm with several numerical reconstructions in dimension one or two.

In [25] a coupled system of pdes modelling the interaction between a two-dimensional incompressible viscous fluid and a one-dimensional elastic beam located on the upper part of the fluid domain boundary is considered. A good functional framework to define weak solutions in case of contact between the elastic beam and the bottom of the fluid cavity is designed. It is then prove that such solutions exist globally in time regardless a possible contact by approximating the beam equation by a damped beam and letting this additional viscosity vanishes.

7.6. Numerical methods for multi-physics problems

Participants: Miguel Ángel Fernández Varela, Fannie Gerosa.

In [24] we introduce a mixed dimensional Stokes-Darcy coupling where a d -dimensional Stokes' flow is coupled to a Darcy model on the $d - 1$ dimensional boundary of the domain. The porous layer introduces tangential creeping flow along the boundary and allows for the modelling of boundary flow due to surface roughness. This leads to a new model of flow in fracture networks with reservoirs in an impenetrable bulk matrix. Exploiting this modelling capability, we then formulate a fluid-structure interaction method with contact, where the porous layer allows for mechanically consistent contact and release. Physical seepage in the contact zone due to rough surfaces is modelled by the porous layer. Some numerical examples are reported, both on the Stokes-Darcy coupling alone and on the fluid-structure interaction with contact in the porous boundary layer.

Unfitted mesh finite element approximations of immersed incompressible fluid-structure interaction problems which efficiently avoid strong coupling without compromising stability and accuracy are rare in the literature. Moreover, most of the existing approaches introduce additional unknowns or are limited by penalty terms which yield ill conditioning issues. In [28], we introduce a new unfitted mesh semi-implicit coupling scheme which avoids these issues. To this purpose, we provide a consistent generalization of the projection based semi-implicit coupling paradigm of [Int. J. Num. Meth. Engrg.,69(4):794-821, 2007] to the unfitted mesh Nitsche-XFEM framework.

7.7. Statistical learning and mathematical modeling interactions

Participants: Damiano Lombardi, Fabien Raphael.

In [30] a greedy dimension reduction method is proposed to deal with classification problems. The method proposed can be seen as a goal oriented dimension reduction method. Elements of a Stiefel manifold (whose dimension is not fixed a priori) are computed in such a way that the classification score is maximised. Several examples are proposed to illustrate the method features and to highlight its differences with classical reduction methods used in classification.

7.8. Tensor approximation and HPC

Participant: Damiano Lombardi.

In [26] a hierarchical adaptive piece-wise tensor decomposition is proposed to approximate high-dimensional functions. Neither the subtensor partitioning nor the rank of the approximation in each of the partitions are fixed a priori. Instead, they are computed to fulfill a prescribed accuracy. Two main contributions are proposed. A greedy error distribution scheme, that allows to adaptively construct the approximation in each of the partitions and a hierarchical tree algorithm that optimise the subtensor partitioning to minimise the storage. Several example on challenging functions are proposed.

7.9. Miscellaneous

Participants: Damiano Lombardi, Olga Mula.

In [20] an approximated formulation of the multi-marginal optimal transport problem (Kantorovich formulation) is proposed. In the formulation, called MCOT, the constraints on the marginal densities are replaced by moments of the densities. This formulation allows to deal simply with a wide spectrum of high-dimensional multi-marginal problems, with non-standard (martingale) constraints.

In [27] a reduced-order modeling framework is proposed, in which a set of model instances is part of a metric space. The introduction of the exponential and logarithmic maps (Riemannian geometry) makes it possible to reduce in an effective way solutions that are classically challenging for standard model reduction methods. Some examples on 1D hyperbolic equations and Wasserstein distance are proposed.

MAMBA Project-Team

7. New Results

7.1. Direct and inverse Problems in Structured-population equations

7.1.1. Modelling Polymerization Processes

In 2017, we evidenced the presence of several polymeric species by using data assimilation methods to fit experimental data from H. Rezaei's lab [64]; new experimental evidence reinforced these findings [19], [35]. The challenges are now to propose mathematical models capable of tracking such diversity while keeping sufficient simplicity to be tractable to analysis.

In collaboration with Klemens Fellner from the university of Graz, we propose a new model, variant of the Becker-Döring system but containing two monomeric species, capable of displaying sustained though damped oscillations as is experimentally observed. We also proposed a statistical test to validate or invalidate the presence of oscillations in experimental highly nonstationary signals [55].

7.1.2. Asymptotic behaviour of structured-population equations

Pierre Gabriel and Hugo Martin studied the mathematical properties of a model of cell division structured by two variables – the size and the size increment – in the case of a linear growth rate and a self-similar fragmentation kernel [16]. They first show that one can construct a solution to the related two dimensional eigenproblem associated to the eigenvalue 1 from a solution of a certain one dimensional fixed point problem. Then they prove the existence and uniqueness of this fixed point in the appropriate L^1 weighted space under general hypotheses on the division rate. Knowing such an eigenfunction proves useful as a first step in studying the long time asymptotic behaviour of the Cauchy problem.

Etienne Bernard, Marie Doumic and Pierre Gabriel proved in [9] that for the growth-fragmentation equation with fission into two equal parts and linear growth rate, under fairly general assumptions on the division rate, the solution converges towards an oscillatory function, explicitly given by the projection of the initial state on the space generated by the countable set of the dominant eigenvectors of the operator. Despite the lack of hypo-coercivity of the operator, the proof relies on a general relative entropy argument in a convenient weighted L^2 space, where well-posedness is obtained via semigroup analysis. They also propose a non-dissipative numerical scheme, able to capture the oscillations.

Pierre Gabriel and Hugo Martin then extended this asymptotic result in the framework of measure solutions [50]. To do so they adopt a duality approach, which is also well suited for proving the well-posedness when the division rate is unbounded. The main difficulty for characterizing the asymptotic behavior is to define the projection onto the subspace of periodic (rescaled) solutions. They achieve this by using the generalized relative entropy structure of the dual problem.

7.1.3. Estimating the division rate from indirect measurements of single cells

7.1.3.1. Marie Doumic and Adélaïde Olivier

Is it possible to estimate the dependence of a growing and dividing population on a given trait in the case where this trait is not directly accessible by experimental measurements, but making use of measurements of another variable? The article [46] addresses this general question for a very recent and popular model describing bacterial growth, the so-called incremental or adder model - the model studied by Hugo Martin and Pierre Gabriel in [16]. In this model, the division rate depends on the increment of size between birth and division, whereas the most accessible trait is the size itself. We prove that estimating the division rate from size measurements is possible, we state a reconstruction formula in a deterministic and then in a statistical setting, and solve numerically the problem on simulated and experimental data. Though this represents a severely ill-posed inverse problem, our numerical results prove to be satisfactory.

7.2. Stochastic Models of Biological Systems

7.2.1. Stochastic models for spike-timing dependent plasticity

7.2.1.1. *Ph. Robert and G. Vignoud*

Synaptic plasticity is a common mechanism used to model learning in stochastic neural networks, STDP is a great example of such mechanisms. We develop a simple framework composed by two neurons and one synaptic weight, seen as stochastic processes and study the existence and stability of such distributions, for a wide range of classical synaptic plasticity models. Using two simple examples of STDP, the calcium-based rule and the all-to-all pair-based rule, we apply stochastic averaging principles and obtain differential equations for the limit processes, based on the invariant distributions of the fast system when the slow variables are considered fixed. We study a general stochastic queue to approximate the calcium-based rule and are able to have an analytical solution for the invariant distribution of the fast synaptic processes. We also detail some simpler systems, either through some approximations or simulations to put into light the influences of different biologically-linked parameters on the dynamics of the synaptic weight.

7.2.2. Online Sequence Learning In The Striatum With Anti-Hebbian Spike-Timing-Dependent Plasticity

7.2.2.1. *G. Vignoud. Collaboration with J. Touboul (Brandeis University)*

Spike-Timing Dependent Plasticity (STDP) in the striatum is viewed as a substrate for procedural learning. Striatal projecting neurons (SPNs) express anti-Hebbian plasticity at corticostriatal synapses, (a presynaptic cortical spike followed by a postsynaptic striatal spike leads to the weakening of the connection, whereas the reverse pairing leads to potentiation). SPNs need to integrate many inputs to spike, and as such, their main role is to integrate context elements to choose between different sensorimotor associations. In this work, we develop a simple numerical model of the striatum, integrating cortical spiking inputs to study the role of anti-Hebbian STDP in pattern recognition and sequence learning. Cortical neurons are seen as binary input neurons and one striatal SPN is modeled as a leaky-integrate-and-fire neuron. Combined informations from the output, reward and timing between the different spikes modify the intensity of each connection, through two mechanisms: anti-Hebbian STDP and dopaminergic signaling, using three-factor learning rules. We have added a second output neuron with collateral inhibition which leads to an improvement of the global accuracy. In another project, we studied the dynamics of learning, by shutting off/on the dopaminergic plasticity, and compare it to DMS/DLS experimental and behavioral experiments. We show that anti-Hebbian STDP favors the learning of complete sequence of spikes, such as is needed in the striatum, whereas, even if Hebbian STDP helps to correlate the spiking of two connected neurons, it is not sufficient to integrate of long sequences of correlated inputs spikes.

7.2.3. D1/D2 detection from action-potential properties using machine learning approach in the dorsal striatum

7.2.3.1. *G. Vignoud. Collaboration, with Team Venance (CIRB/Collège de France)*

Striatal medium spiny neurons (MSNs) are segregated into two subpopulations, the D1 receptor-expressing MSNs (the direct striatonigral pathway) and the D2 receptor-expressing MSNs (the indirect striatopallidal pathway). The fundamental role of MSNs as output neurons of the striatum, and the necessary distinction between D1- and D2-expressing neurons accentuate the need to clearly distinguish both subpopulations in electrophysiological recordings in vitro and in vivo. Currently, fluorescent labelling of the dopaminergic receptors in mice enables a clear differentiation. However, multiplying in vivo the number of genetic markers (optogenetics, fluorescence) hinders possibilities for other genetic manipulations. Moreover, electrophysiological properties of fluorescent neurons can slightly differ from “native” cells and false-positive can be observed. The lack of a proper way to separate D1- and D2-MSNs based on electrophysiological properties led us to devise a detection algorithm based on action potential profile. We used more than 450 D1/D2 labelled MSNs from in vitro patch-clamp recordings (different experimentalists, different setups and protocols), to characterize and identify properties that facilitate the MSN discrimination. After analyzing passive and active MSN membrane

properties, we built an extensive dataset and fed it into classical machine learning classification methods. The training of the different algorithms (k-nearest neighbors, random forest, deep neural networks, ...) was performed with the scikit-learn Python library, and the optimized classifier was able to correctly discriminate neurons in the dorsolateral striatum at 76% (and up to 83% if we allow the classifier to reject some MSNs). This study developed an efficient classification algorithm for D1/D2-MSNs, facilitating cell discrimination without specific genetic fluorescent labelling, leaving some room for other genetic markers and optogenetic labeling.

7.2.4. The Stability of Non-Linear Hawkes Processes

7.2.4.1. Ph. Robert and G. Vignoud

We have investigated the asymptotic properties of self-interacting point processes introduced by Kerstan (1964) and Hawkes and Oakes (1974). These point processes have the property that the intensity at some point $t \in (-\infty, +\infty)$ is a functional of all points of the point process before t . Such a process is said to be stable if it has a version whose distribution is invariant by translation. By using techniques of coupling and Markovian methods, we have been able to obtain some existence and uniqueness results with weaker conditions than in the current literature.

7.2.5. Mathematical Models of Gene Expression

7.2.5.1. Ph. Robert

In Robert [30] we analyze the equilibrium properties of a large class of stochastic processes describing the fundamental biological process within bacterial cells, *the production process of proteins*. Stochastic models classically used in this context to describe the time evolution of the numbers of mRNAs and proteins are presented and discussed. An extension of these models, which includes elongation phases of mRNAs and proteins, is introduced. A convergence result to equilibrium for the process associated to the number of proteins and mRNAs is proved and a representation of this equilibrium as a functional of a Poisson process in an extended state space is obtained. Explicit expressions for the first two moments of the number of mRNAs and proteins at equilibrium are derived, generalizing some classical formulas. Approximations used in the biological literature for the equilibrium distribution of the number of proteins are discussed and investigated in the light of these results. Several convergence results for the distribution of the number of proteins at equilibrium are in particular obtained under different scaling assumptions.

7.2.6. Stochastic modelling of molecular motors

7.2.6.1. Marie Doumic, Dietmar Oelz, Alex Mogilner

It is often assumed in biophysical studies that when multiple identical molecular motors interact with two parallel microtubules, the microtubules will be crosslinked and locked together. The aim of the article [4] is to examine this assumption mathematically. We model the forces and movements generated by motors with a time-continuous Markov process and find that, counter-intuitively, a tug-of-war results from opposing actions of identical motors bound to different microtubules. The model shows that many motors bound to the same microtubule generate a great force applied to a smaller number of motors bound to another microtubule, which increases detachment rate for the motors in minority, stabilizing the directional sliding. However, stochastic effects cause occasional changes of the sliding direction, which has a profound effect on the character of the long-term microtubule motility, making it effectively diffusion-like. Here, we estimate the time between the rare events of switching direction and use them to estimate the effective diffusion coefficient for the microtubule pair. Our main result is that parallel microtubules interacting with multiple identical motors are not locked together, but rather slide bidirectionally. We find explicit formulae for the time between directional switching for various motor numbers.

7.3. Analysis and control of mosquito populations

7.3.1. Control Strategies for Sterile Insect Techniques

We proposed different models to serve as a basis for the design of control strategies relying on releases of sterile male mosquitoes (*Aedes spp*) and aiming at elimination of wild vector population. Different types of releases were considered (constant, periodic or impulsive) and sufficient conditions to reach elimination were provided in each case [152]. We also estimated sufficient and minimal treatment times. A feedback approach was introduced, in which the impulse amplitude is chosen as a function of the actual wild population [152].

7.3.2. Optimal replacement strategies, application to *Wolbachia*

We modelled and designed optimal release control strategy with the help of a least square problem. In a nutshell, one wants to minimize the number of uninfected mosquitoes at a given time horizon, under relevant biological constraints. We derived properties of optimal controls and studied a limit problem providing useful asymptotic properties of optimal controls [8], [42].

7.3.3. Oscillatory regimes in population models

Understanding mosquitoes life cycle is of great interest presently because of the increasing impact of vector borne diseases. Observations yields evidence of oscillations in these populations independent of seasonality, still unexplained. We proposed [33] a simple mathematical model of egg hatching enhancement by larvae which produces such oscillations that conveys a possible explanation.

On the other hand, population oscillations may be induced by seasonal changes. We considered a biological population whose environment varies periodically in time, exhibiting two very different “seasons”, favorable and unfavorable. We addressed the following question: the system’s period being fixed, under what conditions does there exist a critical duration above which the population cannot sustain and extincts, and below which the system converges to a unique periodic and positive solution? We obtained [153], [154] sufficient conditions for such a property to occur for monotone differential models with concave nonlinearities, and applied the obtained criterion to a two-dimensional model featuring juvenile and adult insect populations.

7.3.4. Feedback control principles for population replacement by *Wolbachia*

The issue of effective scheduling of the releases of *Wolbachia*-infected mosquitoes is an interesting problem for Control theory. Having in mind the important uncertainties present in the dynamics of the two populations in interaction, we attempted to identify general ideas for building release strategies, which should apply to several models and situations [39]. These principles were exemplified by two interval observer-based feedback control laws whose stabilizing properties were demonstrated when applied to a model retrieved from [76].

7.4. Bacterial motion by Run and tumble

Collective motion of chemotactic bacteria such as *Escherichia coli* relies, at the individual level, on a continuous reorientation by runs and tumbles. It has been established that the length of run is decided by a stiff response to a temporal sensing of chemical cues along the pathway. We describe a novel mechanism for pattern formation stemming from the stiffness of chemotactic response relying on a kinetic chemotaxis model which includes a recently discovered formalism for the bacterial chemotaxis [142]. We prove instability both for a microscopic description in the space-velocity space and for the macroscopic equation, a flux-limited Keller-Segel equation, which has attracted much attention recently. A remarkable property is that the unstable frequencies remain bounded, as it is the case in Turing instability. Numerical illustrations based on a powerful Monte Carlo method show that the stationary homogeneous state of population density is destabilized and periodic patterns are generated in realistic ranges of parameters. These theoretical developments are in accordance with several biological observations.

This motivates also our study of traveling wave and aggregation in population dynamics of chemotactic cells based on the FLKS model with a population growth term [86]. Our study includes both numerical and theoretical contributions. In the numerical part, we uncover a variety of solution types in the one-dimensional FLKS model additionally to standard Fisher/KPP type traveling wave. The remarkable result is a counter-intuitive backward traveling wave, where the population density initially saturated in a stable state transits toward an un-stable state in the local population dynamics. Unexpectedly, we also find that the backward traveling wave solution transits to a localized spiky solution as increasing the stiffness of chemotactic response. In the theoretical part, we obtain a novel analytic formula for the minimum traveling speed which includes the counter-balancing effect of chemotactic drift vs. reproduction/diffusion in the propagating front. The front propagation speeds of numerical results only slightly deviate from the minimum traveling speeds, except for the localized spiky solutions, even for the backward traveling waves. We also discover an analytic solution of unimodal traveling wave in the large-stiffness limit, which is certainly unstable but exists in a certain range of parameters.

Another activity concerns the relation between the tumbling rate and the internal state of bacteria. The study [58] aims at deriving models at the macroscopic scale from assumptions on the microscopic scales. In particular we are interested in comparisons between the stiffness of the response and the adaptation time. Depending on the asymptotics chosen both the standard Keller-Segel equation and the flux-limited Keller-Segel (FLKS) equation can appear. An interesting mathematical issue arises with a new type of equilibrium equation leading to solution with singularities.

7.5. Numerical methods for cell aggregation by chemotaxis

Three-dimensional cultures of cells are gaining popularity as an in vitro improvement over 2D Petri dishes. In many such experiments, cells have been found to organize in aggregates. We present new results of three-dimensional in vitro cultures of breast cancer cells exhibiting patterns. Understanding their formation is of particular interest in the context of cancer since metastases have been shown to be created by cells moving in clusters. In the paper [82], we propose that the main mechanism which leads to the emergence of patterns is chemotaxis, i.e., oriented movement of cells towards high concentration zones of a signal emitted by the cells themselves. Studying a Keller-Segel PDE system to model chemotactical auto-organization of cells, we prove that it is subject to Turing instability if a time-dependent condition holds. This result is illustrated by two-dimensional simulations of the model showing spheroidal patterns. They are qualitatively compared to the biological results and their variability is discussed both theoretically and numerically.

This motivates to study parabolic-elliptic Keller-Segel equation with sensitivity saturation, because of its pattern formation ability, is a challenge for numerical simulations. We provide two finite-volume schemes that are shown to preserve, at the discrete level, the fundamental properties of the solutions, namely energy dissipation, steady states, positivity and conservation of total mass [131]. These requirements happen to be critical when it comes to distinguishing between discrete steady states, Turing unstable transient states, numerical artifacts or approximate steady states as obtained by a simple upwind approach. These schemes are obtained either by following closely the gradient flow structure or by a proper exponential rewriting inspired by the Scharfetter-Gummel discretization. An interesting fact is that upwind is also necessary for all the expected properties to be preserved at the semi-discrete level. These schemes are extended to the fully discrete level and this leads us to tune precisely the terms according to explicit or implicit discretizations. Using some appropriate monotonicity properties (reminiscent of the maximum principle), we prove well-posedness for the scheme as well as all the other requirements. Numerical implementations and simulations illustrate the respective advantages of the three methods we compare.

7.6. Focus on cancer

Modelling Acute Myeloid Leukaemia (AML) and its control by anticancer drugs by PDEs and Delay Differential equations

This theme has continued to be developed in collaboration with Catherine Bonnet, Inria DISCO (Saclay) [93], [94]. Without control by drugs, but with representation of mutualistic interactions between tumor cells and their surrounding support stromal cells, it has also, in collaboration with Delphine Salort and Thierry Jaffredo (LCQB-IBPS) given rise to a recent work by Thanh Nam Nguyen, hired as HTE and ERC postdoctoral fellow at LCQB, submitted as full article [24].

Adaptive dynamics setting to model and circumvent evolution towards drug resistance in cancer by optimal control

The research topic “Evolution and cancer”, designed in the framework of adaptive dynamics to represent and overcome acquired drug resistance in cancer, initiated in [128], [127] and later continued in [91], [92], [126], has been recently summarised in [60] and has been the object of the PhD thesis work of Camille Pouchol, see above “Cell population dynamics and its control”. It is now oriented, thanks to work underway by Cécile Carrère, Jean Clairambault, Tommaso Lorenzi and Grégoire Nadin, in particular towards the mathematical representation of *bet hedging* in cancer, namely a supposed optimal strategy consisting for cancer cell populations under life-threatening cell stress in diversifying their phenotypes according to several resistance mechanisms, such as overexpression of ABC transporters (P-glycoprotein and many others), of DNA repair enzymes or of intracellular detoxication processes. According to different deadly insults the cancer cell population is exposed to, some phenotypes may be selected, any such successful subpopulation being able to store the cell population genome (or subclones of it if the cell population is already genetically heterogeneous) and make it amenable to survival and renewed replication.

Philosophy of cancer biology

This new research topic in Mamba, dedicated to explore possibly underinvestigated, from the mathematical modelling point of view, parts of the field of cancer growth, evolution and therapy, has been the object of a presentation by Jean Clairambault at the recent workshop “Philosophy of cancer biology”

<https://www.philinbiomed.org/event/philosophy-of-cancer-biology-workshop/>.

This workshop gathered most members worldwide of this small, but very active in publishing, community of philosophers of science whose field of research is “philosophy of cancer”, as they call it themselves. This topic offers a clear point of convergence between mathematics, biology and social and human sciences.

7.7. Deformable Cell Modeling: biomechanics and Liver regeneration

- Biomechanically mediated growth control of cancer cells The key intriguing novelty was that the same agent-based model after a single parameter has been calibrated with growth data for multicellular spheroids without application of external mechanical stress by adapting a single parameter, permitted to correctly predict the growth speed of multicellular spheroids of 5 different cell lines subject of external mechanical stress. Hereby the same mechanical growth control stress function was used without any modification [123]. The prediction turned out to be correct independent of the experimental method used to exert the stress, whereby once a mechanical capsule has been used, once dextran has been used in the experiments.
- Regeneration of liver with the Deformable Cell Model. The key novelty was the implementation of the model itself, but an interesting novel result is that the DCM permits closure of a pericentral liver lobule lesion generated by drug-induced damage with about 5 times smaller active migration force due to the ability of the cell to strongly deform and squeeze into narrow spaces between the capillaries. This finding stresses that a precise mechanical description is important in view of quantitatively correct modeling results [155]. The deformable cell model however could be used to calibrate the interaction forces of the computationally much cheaper center-based model to arrive at almost the same results.

REO Team

6. New Results

6.1. Numerical methods for fluid mechanics and application to blood flows

Participants: Irene Vignon-Clementel

If abdominal aortic aneurysms (AAA) are known to be associated with altered morphology and blood flow, intraluminal thrombus deposit and clinical symptoms, the growth mechanisms are yet to be fully understood. In this retrospective longitudinal study of 138 scans, morphological analysis and blood flow simulations for 32 patients with clinically diagnosed AAAs and several follow-up CT-scans, are performed and compared to 9 control subjects [21]. Local correlations between hemodynamic metrics and AAA growth are also explored. Finally, high-risk predictors trained with successively clinical, morphological, hemodynamic and all data, and their link to the AAA evolution are built from supervise learning.

In this paper [19], we perform a verification study of the Coupled-Momentum Method (CMM), a 3D fluid-structure interaction (FSI) model which uses a thin linear elastic membrane and linear kinematics to describe the mechanical behavior of the vessel wall. The verification of this model is done using Womersley's deformable wall analytical solution for pulsatile flow in a semi-infinite cylindrical vessel. This solution is, under certain premises, the analytical solution of the CMM and can thus be used for model verification. For the numerical solution, we employ an impedance boundary condition to define a reflection-free outflow boundary condition and thus mimic the physics of the analytical solution, which is defined on a semi-infinite domain. We first provide a rigorous derivation of Womersley's deformable wall theory via scale analysis. We then illustrate different characteristics of the analytical solution and verification tests comparing the CMM with Womersley's theory.

Superior cavopulmonary circulation can be achieved by either the Hemi-Fontan or Bidirectional Glenn connection. Debate remains as to which results in best hemodynamic results. In [22], adopting patient-specific multiscale computational modeling, we examined both the local dynamics and global physiology to determine if surgical choice can lead to different hemodynamic outcomes.

6.2. Liver biomedical research

Participants: Irene Vignon-Clementel, Nicolas Golse

Nicolas Golse, as part of his medical activity has published 7 articles in 2019 that are not reported here.

The hepatic volume gain following resection is essential for clinical recovery. Previous studies have focused on cellular regeneration. In [13], the study aims to explore the rate of hepatic regeneration of the porcine liver following major resection, highlighting estimates of the early microarchitectural changes that occur during the cellular regeneration. Nineteen large white pigs had 75% resection with serial measurements of the hepatic volume, density, blood flow, and architectural changes that are analyzed at different days to highlight differences pre-resection and in the days following resection.

SERENA Project-Team

7. New Results

7.1. Hybrid high-order methods for nonlinear mechanics

Participants: Alexandre Ern, Nicolas Pignet.




Figure 1. Torsion of a square-section bar: Equivalent plastic strain p (values between 0 (blue) to 0.49 (red)) for HHO with polynomial degree $k = 1$ at the quadrature points for different rotation angles Θ . From left to right: $\Theta = 0^\circ$, $\Theta = 90^\circ$, $\Theta = 180^\circ$, $\Theta = 270^\circ$, and $\Theta = 360^\circ$

Our team contributes actively to the development of hybrid high-order (HHO) methods for nonlinear solid mechanics. Within the PhD of Nicolas Pignet in collaboration with EDF we have addressed several nonlinearities, including plasticity, large deformations, contact, and (Tresca) friction [15], [14], [49]. The advantage with respect to conforming finite elements is the robustness with respect to volumetric locking. The advantage with respect to mixed approaches is computational efficiency avoiding saddle-point formulations and additional unknowns. The advantage with respect to discontinuous Galerkin methods is avoiding the integration of the nonlinear behavior law at face quadrature nodes and the use of symmetric tangent matrices within Newton's method. The torsion of a square-section elastoplastic bar is presented in Figure 1. The color filling reports the equivalent plastic strain. The solution is obtained with the HHO method using the polynomial degree $k = 1$ for the face and the cell unknowns.

7.2. A hybrid high-order method for flow simulations in discrete fracture networks

Participants: Florent Hédin, Géraldine Pichot, Alexandre Ern.




Figure 2. Examples of DFN: (left) B1: 19,007 fractures; (center) B2: 152,399 fractures ; (right) B3: 508,338 fractures.

In [36], we are interested in solving flow in large trimensional Discrete Fracture Networks (DFN) (cf Figure 2) with the hybrid high-order (HHO) method. The objectives of this paper are: (1) to demonstrate the benefit of using a high-order method for computing macroscopic quantities, like the equivalent permeability of fracture rocks; (2) to present the computational efficiency of our C++ software, NEF++, which implements the solving of flow in fractures based on the HHO method.

7.3. Analytic expressions of the solutions of advection-diffusion problems in 1D with discontinuous coefficients

Participants: Antoine Lejay, Lionel Lenôtre, Géraldine Pichot.

In [30], we provide a method to compute analytic expressions of the resolvent kernel of differential operators of the diffusion type with discontinuous coefficients in one dimension. Then we apply it when the coefficients are piecewise constant. We also perform the Laplace inversion of the resolvent kernel to obtain expressions of the transition density functions or fundamental solutions. We show how these explicit formula are useful to simulate advection-diffusion problems using particle tracking techniques.

7.4. An exponential timestepping algorithm for diffusion with discontinuous coefficients

Participants: Antoine Lejay, Lionel Lenôtre, Géraldine Pichot.

In [29], we present a new Monte Carlo algorithm to simulate diffusion processes in presence of discontinuous convective and diffusive terms. The algorithm is based on the knowledge of close form analytic expressions of the resolvents of the diffusion processes which are usually easier to obtain than close form analytic expressions of the density. In the particular case of diffusion processes with piecewise constant coefficients, known as Skew Diffusions, such close form expressions for the resolvent are available. Then we apply our algorithm to this particular case and we show that the approximate densities of the particles given by the algorithm replicate well the particularities of the true densities (discontinuities, bimodality, ...) Besides, numerical experiments show a quick convergence.

7.5. Polynomial-degree-robust multilevel algebraic error estimator & solver

Participants: Ani Miraci, Jan Papež, Martin Vohralfk.

In [58], we devise a novel multilevel a posteriori estimator of the algebraic error. It delivers a fully computable, guaranteed lower bound on the error between an unknown exact solution of a system of linear algebraic equations and its approximation by an algebraic solver. The bound is also proved to be efficient, i.e., it also gives an upper bound on the algebraic error. Remarkably, the quality of these bounds is independent of the approximation polynomial degree. The derived estimates give immediately rise to a multilevel iterative algebraic solver whose contraction factor is independent of the polynomial degree of the approximation. We actually prove an equivalence between efficiency of the estimator and contraction of the solver. The estimator/solver are based on a global coarsest-level solve of lowest-order ($p = 1$), followed by local patchwise p -degree problems solved on the other levels. It corresponds to a V-cycle geometric multigrid solver with zero pre- and one post-smoothing step via block-Jacobi. A salient feature is the choice of the optimal step size for the descent direction.

7.6. Local- and global-best equivalence, simple projector, and optimal hp approximation in $\mathbf{H}(\text{div})$

Participants: Alexandre Ern, Thirupathi Gudi, Iain Smears, Martin Vohralík.

$$\begin{aligned} & \min_{\substack{\mathbf{v}_{\mathcal{T}} \in \mathbf{RTN}_p(\mathcal{T}) \cap \mathbf{H}_{0,\Gamma_N}(\text{div}, \Omega) \\ \nabla \cdot \mathbf{v}_{\mathcal{T}} = \Pi_{\mathcal{T}}^p(\nabla \cdot \mathbf{v})}} \|\mathbf{v} - \mathbf{v}_{\mathcal{T}}\|_{\Omega}^2 \\ & \approx \sum_{K \in \mathcal{T}} \min_{\mathbf{v}_K \in \mathbf{RTN}_p(K)} \|\mathbf{v} - \mathbf{v}_K\|_K^2 \end{aligned}$$

Figure 3. Equivalence between global-best and local-best approximation for any $\mathbf{H}(\text{div})$ function \mathbf{v} with zero normal flux over part of the boundary and piecewise polynomial divergence

In [53], we prove that a global-best approximation in $\mathbf{H}(\text{div})$, with constraints on normal component continuity and divergence, is equivalent to the sum of independent local-best approximations, without any constraints, as illustrated in Figure 3. This may seem surprising on a first sight since the right term in Figure 3 is seemingly much smaller (since the minimization set is unconstrained and thus much bigger). This result leads to optimal a priori hp -error estimates for mixed and least-squares finite element methods, which were missing in the literature until 2019. Additionally, the construction we devise gives rise to a simple stable local commuting projector in $\mathbf{H}(\text{div})$, which delivers approximation error equivalent to the local-best approximation and applies under the minimal necessary Sobolev $\mathbf{H}(\text{div})$ regularity, which is another result that has been sought for a very long time.

ALPINES Project-Team

7. New Results

7.1. Adaptive Domain Decomposition Method for Saddle Point Problem

In [37], we introduce an adaptive domain decomposition (DD) method for solving saddle point problems defined as a block two by two matrix. The algorithm does not require any knowledge of the constrained space. We assume that all sub matrices are sparse and that the diagonal blocks are the sum of positive semi definite matrices. The latter assumption enables the design of adaptive coarse space for DD methods.

7.2. A Class of Efficient Locally Constructed Preconditioners Based on Coarse Spaces

In [24] we present a class of robust and fully algebraic two-level preconditioners for SPD matrices. We introduce the notion of algebraic local SPSD splitting of an SPD matrix and we give a characterization of this splitting. It helps construct *algebraically and locally* a class of efficient coarse subspaces which bound the spectral condition number of the preconditioned system by a number defined a priori. Some PDEs-dependant preconditioners correspond to a special case of the splitting. The examples of the algebraic coarse subspaces in this paper are not practical due to expensive construction. We propose an heuristic approximation that is not costly. Numerical experiments illustrate the efficiency of the proposed method.

7.3. A Multilevel Schwarz Preconditioner Based on a Hierarchy of Robust Coarse Spaces

In [32] we present a multilevel preconditioner for SPD matrices. Robust two-level additive Schwarz preconditioners guarantee a fast convergence of the Krylov method. To maintain the robustness each subdomain contributes a small number of vectors to construct a basis for the second level (the coarse space). As long as the dimension of the coarse space is reasonable i.e., direct solvers can be used efficiently, the two-level method scales well. However, the bottleneck arises when factoring the coarse space matrix becomes costly. Using an iterative Krylov method on the second level might be the right choice. Nevertheless, the condition number of the coarse space matrix is typically larger than the one of the first level. One of the difficulties of using two-level methods to solve the coarse problem is that the matrix does not arise from a PDE anymore. We introduce in this paper a practical method of applying a multilevel additive Schwarz preconditioner efficiently. This multilevel preconditioner is implemented in HPDDM and the code for reproducing the results from the paper is available [here](#).

7.4. Inverse scattering problems without knowing the source term

The solution of inverse scattering problems always presupposed knowledge of the incident wave-field and require repeated computations of the forward problem, for which knowing the source term is crucial. In [26], we present a three-step strategy to solve inverse scattering problems when the time signature of the source is unknown. The proposed strategy combines three recent techniques: (i) wave splitting to retrieve the incident and the scattered wavefields, (ii) time-reversed absorbing conditions (TRAC) for redatuming the data inside the computational domain, (iii) adaptive eigenspace inversion (AEI) to solve the inverse problem. Numerical results illustrate step-by-step the feasibility of the proposed strategy.

7.5. Envelope following methods

One difficulty when solving problems in plasma physics is the behaviour at several scales in time and space of the solutions of equations. For example, central equations in this domain of application are highly oscillatory in time. The multiscale aspect makes the models difficult to tackle when we aim at avoiding a high computational cost. A solution to this problem is to solve the models by designing adapted numerical methods with a low computational cost and which are able to deal efficiently with rapid and slow scales in time. In this direction, we worked on envelope following methods, which have been efficiently applied in the community of oscillators in RF circuits. The method has (at least) two variants: in a first place, it is based on the concept of using extra variables to represent the changing rapid period and the cumulative effect of changing periods and then, use of Newton iterations allows to find these unknowns. In a second place, we adopt a similar strategy except that the rapid period is not an extra variable but a direct outcome of the numerical integration by the use of the Poincaré map. We implemented and tested both approaches for equations of interest in plasma physics and we observed that these methods didn't perform accurate results.

7.6. Domain decomposition preconditioning for high frequency wave propagation problems

The work about domain decomposition preconditioning for Maxwell equations has been published in [21]. It studies two-level preconditioners where the coarse space is based on the discretisation of the PDE on a coarse mesh. The PDE is discretised using finite-element methods of fixed, arbitrary order. The theoretical part of this work is the Maxwell analogue of a previous work for Helmholtz equation, and shows that for Maxwell problems with absorption, if the absorption is large enough and if the subdomain and coarse mesh diameters are chosen appropriately, then classical two-level overlapping Additive Schwarz Domain Decomposition preconditioning performs optimally – in the sense that GMRES converges in a wavenumber-independent number of iterations. The theory is also illustrated by various numerical experiments.

Ongoing studies are being conducted on recursive one-level optimized Schwarz methods for the high frequency Helmholtz and Maxwell equations. The method consists in solving the subdomain problems in a one-level optimized Schwarz preconditioner only approximately, using inner GMRES iterations preconditioned again by a one-level method, with smaller subdomains. The asymptotic behaviour and parallel scalability of the method are being investigated. Exhaustive numerical experiments are being conducted to compare the efficiency of this method with two-level preconditioners, including cavity problems and benchmarks in seismic imaging.

7.7. The boundary element method in FreeFEM

The BemTool and HTOOL libraries developed by the team, implementing respectively the Boundary Element Method and Hierarchical Matrices, have been interfaced with FreeFEM to allow FreeFEM users to use the Boundary Element Method (BEM) in their FreeFEM scripts. New additions to the Domain Specific Language (DSL) of FreeFEM allows the user to define and manipulate curved (1D) and surface (2D) meshes, as well as define and solve BEM variational problems in a high-level manner, similarly to FEM problems. The parallelization of the HTOOL library allows the user to assemble and solve their BEM problems in parallel in a transparent way.

Ongoing work consists in finalizing the BEM DSL to propose complete and documented features to the FreeFEM user in the next release, as well as investigating FEM-BEM coupling.

7.8. New Optimised Schwarz Method for dealing with cross-points

We consider a scalar wave propagation in harmonic regime modelled by Helmholtz equation with heterogeneous coefficients. Using the Multi-Trace Formalism (MTF), we propose a new variant of the Optimized Schwarz Method (OSM) that can accommodate the presence of cross-points in the subdomain partition. This leads to the derivation of a strongly coercive formulation of our Helmholtz problem posed on the union of all interfaces. The corresponding operator takes the form "identity + contraction".

7.9. Two-level preconditioning for h-version boundary element approximation of hypersingular operator with GenEO

We consider symmetric positive definite operators stemming from boundary integral equation (BIE), and we analysed a two-level preconditioner where the coarse space is built using local generalized eigenproblems in the overlap. We will refer to this coarse space as the GenEO coarse space. We obtained bounds on the condition number of the preconditioned system. In this work package, we also performed large scale numerical experiments for testing the scalability of our approach. We relied on parallel implementation of our algorithm.

7.10. Adaptive resolution of linear systems based on a posteriori error estimators

In [18] we discuss a new adaptive approach for iterative solution of sparse linear systems arising from partial differential equations (PDEs) with self-adjoint operators. The idea is to use the a posteriori estimated local distribution of the algebraic error in order to steer and guide the solve process in such way that the algebraic error is reduced more efficiently in the consecutive iterations. We first explain the motivation behind the proposed procedure and show that it can be equivalently formulated as constructing a special combination of preconditioner and initial guess for the original system. We present several numerical experiments in order to identify when the adaptive procedure can be of practical use.

7.11. Adaptive hierarchical subtensor partitioning for tensor compression

In [33] a numerical method is proposed to compress a tensor by constructing a piece-wise tensor approximation. This is defined by partitioning a tensor into sub-tensors and by computing a low-rank tensor approximation (in a given format) in each sub-tensor. Neither the partition nor the ranks are fixed a priori, but, instead, are obtained in order to fulfill a prescribed accuracy and optimize, to some extent, the storage. The different steps of the method are detailed and some numerical experiments are proposed to assess its performances.

7.12. Frictionless contact problem for hyper-elastic materials with interior point optimizer

In [35] we present a method to solve the mechanical problems undergoing finite deformations and the unilateral contact problems without friction for hyperelastic materials. We apply it to an industrial application: contact between a mechanical gasket and an obstacle. The main idea is to formulate the contact problem into an optimization one, in order to use the Interior Point OPTimizer (IPOPT) to solve it. Finally, the FreeFEM software is used to compute and solve the contact problem. Our method is validated against several benchmarks and used on an industrial application example.

7.13. A posteriori error estimates for Darcy's problem coupled with the heat equation

In [25] we derive a posteriori error estimates, in two and three dimensions, for the heat equation coupled with Darcy's law by a nonlinear viscosity depending on the temperature. We introduce two variational formulations and discretize them by finite element methods. We prove optimal a posteriori errors with two types of computable error indicators. The first one is linked to the linearization and the second one to the discretization. Then we prove upper and lower error bounds under regularity assumptions on the solutions. Finally, numerical computations are performed to show the effectiveness of the error indicators.

DELYS Project-Team

5. New Results

5.1. Distributed Algorithms for Dynamic Networks and Fault Tolerance

Participants: Luciana Bezerra Arantes [correspondent], Sébastien Bouchard, Marjorie Bournat, João Paulo de Araujo, Swan Dubois, Laurent Feuilloley, Denis Jeanneau, Jonathan Lejeune, Franck Petit, 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.

The fact that computing units may leave, join, or move may result of an intentional behavior or not. In the latter case, the system may be subject to disruptions due to component faults that can be permanent, transient, exogenous, evil-minded, etc. It is therefore crucial to come up with solutions tolerating some types of faults.

In 2019, we obtained the following results.

5.1.1. Failure detectors

Mutual exclusion is one of the fundamental problems in distributed computing but existing mutual exclusion algorithms are unadapted to the dynamics and lack of membership knowledge of current distributed systems (e.g., mobile ad-hoc networks, peer-to-peer systems, etc.). Additionally, in order to circumvent the impossibility of solving mutual exclusion in asynchronous message passing systems where processes can crash, some solutions include the use of $(\mathcal{T}+\Sigma^l)$, which is the weakest failure detector to solve mutual exclusion in known static distributed systems. In [28], we define a new failure detector $\mathcal{T}\Sigma^{lr}$ which is equivalent to $(\mathcal{T}+\Sigma^l)$ in known static systems, and prove that $\mathcal{T}\Sigma^{lr}$ is the weakest failure detector to solve mutual exclusion in unknown dynamic systems with partial memory losses. We consider that crashed processes may recover.

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 [19] 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.

5.1.2. Scheduler Tolerant to Temporal Failures in Clouds

Cloud platforms offer different types of virtual machines which ensure different guarantees in terms of availability and volatility, provisioning the same resource through multiple pricing models. For instance, in Amazon EC2 cloud, the user pays per hour for on-demand instances while spot instances are unused resources available for a lower price. Despite the monetary advantages, a spot instance can be terminated or hibernated by EC2 at any moment. Using both hibernation prone spot instances (for cost sake) and on-demand instances, we propose in [31] a static scheduling for applications which are composed of independent tasks (bag-of-task) with deadline constraints. However, if a spot instance hibernates and it does not resume within a time which guarantees the application's deadline, a temporal failure takes place. Our scheduling, thus, aims at minimizing monetary costs of bag-of-tasks applications in EC2 cloud, respecting its deadline and avoiding temporal failures. Performance results with task execution traces, configuration of Amazon EC2 virtual machines, and EC2 market history confirms the effectiveness of our scheduling and that it tolerates temporal failures. In [30], we extend our approach for dynamic scheduling.

5.1.3. Gathering of Mobile Agents

Gathering a group of mobile agents is a fundamental task in the field of distributed and mobile systems. It consists of bringing agents that initially start from different positions to meet all together in finite time. In the case when there are only two agents, the gathering problem is often referred to as the rendezvous problem.

In [14] we show that rendezvous under the strong scenario is possible for agents with asynchrony restricted in the following way: agents have the same measure of time but the adversary can impose, for each agent and each edge, the speed of traversing this edge by this agent. The speeds may be different for different edges and different agents but all traversals of a given edge by a given agent have to be at the same imposed speed. We construct a deterministic rendezvous algorithm for such agents, working in time polynomial in the size of the graph, in the length of the smaller label, and in the largest edge traversal time.

5.1.4. Perpetual self-stabilizing exploration of dynamic environments

In [15], we deal with the classical problem of exploring a ring by a cohort of synchronous robots. We focus on the perpetual version of this problem in which it is required that each node of the ring is visited by a robot infinitely often. We assume that the robots evolve in ring-shape TVGs, *i.e.*, the static graph made of the same set of nodes and that includes all edges that are present at least once over time forms a ring of arbitrary size. We also assume that each node is infinitely often reachable from any other node. In this context, we aim at providing a self-stabilizing algorithm to the robots (*i.e.*, the algorithm must guarantee an eventual correct behavior regardless of the initial state and positions of the robots). We show that this problem is deterministically solvable in this harsh environment by providing a self-stabilizing algorithm for three robots.

5.1.5. Torus exploration by oblivious robots

In [17], we deal with a team of autonomous robots that are endowed with motion actuators and visibility sensors. Those robots are weak and evolve in a discrete environment. By weak, we mean that they are anonymous, uniform, unable to explicitly communicate, and oblivious. We first show that it is impossible to solve the terminating exploration of a simple torus of arbitrary size with less than 4 or 5 such robots, respectively depending on whether the algorithm is probabilistic or deterministic. Next, we propose in the SSYNC model a probabilistic solution for the terminating exploration of torus-shaped networks of size $\ell \times L$, where $7 \leq \ell \leq L$, by a team of 4 such weak robots. So, this algorithm is optimal *w.r.t.* the number of robots.

5.1.6. Explicit communication among stigmergic robots

In [18], we investigate avenues for the exchange of information (explicit communication) among deaf and mute mobile robots scattered in the plane. We introduce the use of movement-signals (analogously to flight signals and bees waggle) as a mean to transfer messages, enabling the use of distributed algorithms among robots. We propose one-to-one deterministic movement protocols that implement explicit communication among semi-synchronous robots. Our protocols enable the use of distributed algorithms based on message exchanges among swarms of stigmergic robots. They also allow robots to be equipped with the means of communication to tolerate faults in their communication devices.

5.1.7. Gradual stabilization

In [13], we introduce the notion of *gradual stabilization under (τ, ρ) -dynamics* (gradual stabilization, for short). A gradually stabilizing algorithm is a self-stabilizing algorithm with the following additional feature: after up to τ dynamic steps of a given type ρ occur starting from a legitimate configuration, it first quickly recovers to a configuration from which a specification offering a minimum quality of service is satisfied.

It then gradually converges to specifications offering stronger and stronger safety guarantees until reaching a configuration (1) from which its initial (strong) specification is satisfied again, and (2) where it is ready to achieve gradual convergence again in case of up to τ new dynamic steps of type ρ . A gradually stabilizing algorithm being also self-stabilizing, it still recovers within finite time (yet more slowly) after any other finite number of transient faults, including for example more than τ arbitrary dynamic steps or other failure patterns such as memory corruptions. We illustrate this new property by considering three variants of a synchronization problem respectively called *strong*, *weak*, and *partial* unison. We propose a self-stabilizing unison algorithm

which achieves gradual stabilization in the sense that after one dynamic step of a certain type *BULCC* (such a step may include several topological changes) occurs starting from a configuration which is legitimate for the strong unison, it maintains clocks almost synchronized during the convergence to strong unison: it satisfies partial unison immediately after the dynamic step, then converges in at most one round to weak unison, and finally re-stabilizes to strong unison.

5.2. Distributed systems and Large-scale data distribution

Participants: Guillaume Fraysse, Saalik Hatia, Mesaac Makpangou, Sreeja Nair, Jonathan Sid-Otmame, Pierre Sens, Marc Shapiro, Ilyas Toumlilt, Dimitrios Vasilas.

5.2.1. Proving the safety of highly-available distributed objects

To provide high availability in distributed systems, object replicas allow concurrent updates. Although replicas eventually converge, they may diverge temporarily, for instance when the network fails. This makes it difficult for the developer to reason about the object's properties, and in particular, to prove invariants over its state. For the sub-class of state-based distributed systems, we propose a proof methodology for establishing that a given object maintains a given invariant, taking into account any concurrency control. Our approach allows reasoning about individual operations separately. We demonstrate that our rules are sound, and we illustrate their use with some representative examples. We automate the rule using Boogie, an SMT-based tool.

This work is accepted for publication at the 29th European Symposium on Programming (ESOP), April 2020, Dublin, Ireland [34]. Preliminary results were presented at the Workshop on Principles and Practice of Consistency for Distributed Data (PaPoC), March 2019, Dresden, Germany [29].

5.3. Resource management in system software

Participants: Jonathan Lejeune, Marc Shapiro, Julien Sopena, Francis Laniel.

5.3.1. MemOpLight: Leveraging applicative feedback to improve container memory consolidation

The container mechanism supports consolidating several servers on the same machine, thus amortizing cost. To ensure performance isolation between containers, Linux relies on memory limits. However these limits are static, but application needs are dynamic; this results in poor performance. To solve this issue, MemOpLight reallocates memory to containers based on dynamic applicative feedback. MemOpLight rebalances physical memory allocation, in favor of under-performing ones, with the aim of improving overall performance. Our research explores the issues, addresses the design of MemOpLight, and validates it experimentally. Our approach increases total satisfaction by 13% compared to the default.

It is standard practice in Infrastructure as a Service to *consolidate* several logical servers on the same physical machine, thus amortizing cost. However, the execution of one logical server should not disturb the others: the logical servers should remain *isolated* from one another.

To ensure both consolidation and isolation, a recent approach is “containers,” a group of processes with sharing and isolation properties. To ensure *memory performance isolation*, *i.e.*, guaranteeing to each container enough memory for it to perform well, the administrator limits the total amount of physical memory that a container may use at the expense of others. In previous work, we showed that these limits impede memory consolidation [26]. Furthermore, the metrics available to the kernel to evaluate its policies (*e.g.*, frequency of page faults, I/O requests, use of CPU cycles, *etc.*), are not directly relevant to performance as experienced from the application perspective, which is better characterized by, for instance, response time or throughput measured at application level.

To solve these problems, we propose a new approach, called the Memory Optimization Light (MemOpLight). It is based on application-level feedback from containers. Our mechanism aims to rebalance memory allocation in favor of unsatisfied containers, while not penalizing the satisfied ones. By doing so, we guarantee application satisfaction, while consolidating memory; this also improves overall resource consumption.

Our main contributions are the following:

- An experimental demonstration of the limitations of the existing Linux mechanisms.
- The design of a simple feedback mechanism from application to the kernel.
- An algorithm for adapting container memory allocation.
- And implementation in Linux and experimental confirmation.

This work is currently under submission at a major conference. Some preliminary results are published at NCA 2019 [26].

DYOGENE Project-Team

7. New Results

7.1. Distributed network control and smart-grids

1. Distributed Control of Thermostatically Controlled Loads: Kullback-Leibler Optimal Control in Continuous Time [20] The paper develops distributed control techniques to obtain grid services from flexible loads. The Individual Perspective Design (IPD) for local (load level) control is extended to piecewise deterministic and diffusion models for thermostatically controlled load models. The IPD design is formulated as an infinite horizon average reward optimal control problem, in which the reward function contains a term that uses relative entropy rate to model deviation from nominal dynamics. In the piecewise deterministic model, the optimal solution is obtained via the solution to an eigenfunction problem, similar to what is obtained in prior work. For a jump diffusion model this simple structure is absent. The structure for the optimal solution is obtained, which suggests an ODE technique for computation that is likely far more efficient than policy-or value-iteration.

2. Optimal Control of Dynamic Bipartite Matching Models [23] A dynamic bipartite matching model is given by a bipartite matching graph which determines the possible matchings between the various types of supply and demand items. Both supply and demand items arrive to the system according to a stochastic process. Matched pairs leave the system and the others wait in the queues, which induces a holding cost. We model this problem as a Markov Decision Process and study the discounted cost and the average cost case. We first consider a model with two types of supply and two types of demand items with an N-shaped matching graph. For linear cost function, we prove that an optimal matching policy gives priority to the end edges of the matching graph and is of threshold type for the diagonal edge. In addition, for the average cost problem, we compute the optimal threshold value. According to our numerical experiments, threshold-type policies perform also very well for more general bipartite graphs.

3. Kullback-Leibler-Quadratic Optimal Control of Flexible Power Demand [24] A new stochastic control methodology is introduced for distributed control, motivated by the goal of creating virtual energy storage from flexible electric loads, i.e. Demand Dispatch. In recent work, the authors have introduced Kullback-Leibler-Quadratic (KLQ) optimal control as a stochastic control methodology for Markovian models. This paper develops KLQ theory and demonstrates its applicability to demand dispatch. In one formulation of the design, the grid balancing authority simply broadcasts the desired tracking signal, and the heterogeneous population of loads ramps power consumption up and down to accurately track the signal. Analysis of the Lagrangian dual of the KLQ optimization problem leads to a menu of solution options, and expressions of the gradient and Hessian suitable for Monte-Carlo-based optimization. Numerical results illustrate these theoretical results.

4. Bike sharing systems: a new incentive rebalancing method based on spatial outliers detection [8] Since its launch, Velib' (the Bike Sharing System-BSS-in Paris) has emerged in the Parisian landscape and has been a model for similar systems in many cities. A major problem with BSS is the stations' heterogeneity caused by the attractivity of some stations located in particular areas. In this paper, we focus on spatial outliers defined as stations having a behavior significantly different from their neighboring stations. First, we propose an improved version of Moran scatterplot to exploit the similarity between neighbors, and we test it on a real dataset issued from Velib' system to identify outliers. Then, we design a new method that globally improves the resources' availability in bike stations by adapting the users' trips to the resources' availability. Results show that with a partial collaboration of the users or a limitation to the rush hours, the proposed method enhances significantly the resources' availability in Velib' system.

5. Stochastic Battery Operations using Deep Neural Networks [25] In this paper, we introduce a scenario-based optimal control framework to account for the forecast uncertainty in battery arbitrage problems. Due to the uncertainty of prices and variations of forecast errors, it is challenging for battery operators to design profitable strategies in electricity markets. Without any explicit assumption or model for electricity price

forecasts' uncertainties, we generate future price scenarios via a data-driven, learning-based approach. By aiding the predictive control with such scenarios representing possible realizations of future markets, our proposed real-time controller seeks the optimal charge/discharge levels to maximize profits. Simulation results on a case-study of California-based batteries and prices show that our proposed method can bring higher profits for different battery parameters.

6. Aggregate capacity for TCLs providing virtual energy storage with cycling constraints [26] The coordination of thermostatically controlled loads (TCLs) is challenging due to the need to meet individual loads quality of service (QoS), such as indoor temperature constraints. Since these loads are usually on/off type, cycling rate is one of their QoS metrics; frequent cycling between on and off states is detrimental to them. While significant prior work has been done on the coordination of air conditioning TCLs, the question of cycling QoS has not been investigated in a principled manner. In this work we propose a method to characterize aggregate capacity of a collection of air conditioning TCLs that respects the loads cycling rate constraints (maximum number of cycles in a given time period). The development is done within the framework of randomized local control in which a load makes on/off decisions probabilistically. This characterization allows us to propose a reference planning problem to generate feasible reference trajectories for the ensemble that respect cycling constraints. The reference planning problem manifests itself in the form a Nonlinear Programming problem (NLP), that can be efficiently solved. Our proposed method is compared to previous methods in the literature that do not enforce aggregate cycling. Enforcing individual cycling constraint without taking that into account in reference generation leads to poor reference tracking.

7. Optimal Storage Arbitrage under Net Metering using Linear Programming [29] We formulate the optimal energy arbitrage problem for a piecewise linear cost function for energy storage devices using linear programming (LP). The LP formulation is based on the equivalent minimization of the epigraph. This formulation considers ramping and capacity constraints, charging and discharging efficiency losses of the storage, inelastic consumer load and local renewable generation in presence of net-metering which facilitates selling of energy to the grid and incentivizes consumers to install renewable generation and energy storage. We consider the case where the consumer loads, electricity prices, and renewable generations at different instances are uncertain. These uncertain quantities are predicted using an Auto-Regressive Moving Average (ARMA) model and used in a model predictive control (MPC) framework to obtain the arbitrage decision at each instance. In numerical results we present the sensitivity analysis of storage performing arbitrage with varying ramping batteries and different ratio of selling and buying price of electricity.

8. Energy Storage in Madeira, Portugal: Co-optimizing for Arbitrage, Self-Sufficiency, Peak Shaving and Energy Backup [30] Energy storage applications are explored from a prosumer (consumers with generation) perspective for the island of Madeira in Portugal. These applications could also be relevant to other power networks. We formulate a convex co-optimization problem for performing arbitrage under zero feed-in tariff, increasing self-sufficiency by increasing self-consumption of locally generated renewable energy, provide peak shaving and act as a backup power source during anticipated and scheduled power outages. Using real data from Madeira we perform short and long timescale simulations in order to select end-user contract which maximizes their gains considering storage degradation based on operational cycles. We observe energy storage ramping capability decides peak shaving potential, fast ramping batteries can significantly reduce peak demand charge. The numerical experiment indicates that storage providing backup does not significantly reduce gains performing arbitrage and peak demand shaving. Furthermore, we also use AutoRegressive Moving Average (ARMA) forecasting along with Model Predictive Control (MPC) for real-time implementation of the proposed optimization problem in the presence of uncertainty.

9. Sensitivity to forecast errors in energy storage arbitrage for residential consumers [34] With the massive deployment of distributed energy resources, there has been an increase in the number of end consumers that own photovoltaic panels and storage systems. The optimal use of such storage when facing Time of Use (ToU) prices is directly related to the quality of the load and generation forecasts as well as the algorithm that controls the battery. The sensitivity of such control to different forecasts techniques is studied in this paper. It is shown that good and bad forecasts can result in losses in particularly bad days. Nevertheless, it is observed that performing Model Predictive Control with a simple forecast that is representative of the pasts

can be profitable under different price and battery scenarios. We use real data from Pecan Street and ToU price levels with different buying and selling price for the numerical experiments.

10. Sizing and Profitability of Energy Storage for Prosumers in Madeira, Portugal [47] This paper proposes a framework to select the best-suited battery for co-optimizing for peak demand shaving, energy arbitrage and increase self-sufficiency in the context of power network in Madeira, Portugal. Feed-in-tariff for electricity network in Madeira is zero, which implies consumers with excess production should locally consume the excess generation rather than wasting it. Further, the power network operator applies a peak power contract for consumers which imposes an upper bound on the peak power seen by the power grid interfaced by energy meter. We investigate the value of storage in Madeira, using four different types of prosumers, categorized based on the relationship between their inelastic load and renewable generation. We observe that the marginal increase in the value of storage deteriorates with increase in size and ramping capabilities. We propose the use of profit per cycle per unit of battery capacity and expected payback period as indices for selecting the best-suited storage parameters to ensure profitability. This mechanism takes into account the consumption and generation patterns, profit, storage degradation, and cycle and calendar life of the battery. We also propose the inclusion of a friction coefficient in the original co-optimization formulation to increase the value of storage by reducing the operational cycles and eliminate low returning transactions.

11. Arbitrage with Power Factor Correction using Energy Storage [48] The importance of reactive power compensation for power factor (PF) correction will significantly increase with the large-scale integration of distributed generation interfaced via inverters producing only active power. In this work, we focus on co-optimizing energy storage for performing energy arbitrage as well as local power factor corrections. The joint optimization problem is non-convex, but can be solved efficiently using a McCormick relaxation along with penalty-based schemes. Using numerical simulations on real data and realistic storage profiles, we show that energy storage can correct PF locally without reducing arbitrage gains. It is observed that active and reactive power control is largely decoupled in nature for performing arbitrage and PF correction (PFC). Furthermore, we consider a stochastic online formulation of the problem with uncertain load, renewable and pricing profiles. We develop a model predictive control based storage control policy using ARMA forecast for the uncertainty. Using numerical simulations we observe that PFC is primarily governed by the size of the converter and therefore, look-ahead in time in the online setting does not affect PFC noticeably. However, arbitrage gains are more sensitive to uncertainty for batteries with faster ramp rates compared to slow ramping batteries.

12. A Utility Optimization Approach to Network Cache Design [11] In any caching system, the admission and eviction policies determine which contents are added and removed from a cache when a miss occurs. Usually, these policies are devised so as to mitigate staleness and increase the hit probability. Nonetheless, the utility of having a high hit probability can vary across contents. This occurs, for instance, when service level agreements must be met, or if certain contents are more difficult to obtain than others. In this paper, we propose utility-driven caching, where we associate with each content a utility, which is a function of the corresponding content hit probability. We formulate optimization problems where the objectives are to maximize the sum of utilities over all contents. These problems differ according to the stringency of the cache capacity constraint. Our framework enables us to reverse engineer classical replacement policies such as LRU and FIFO, by computing the utility functions that they maximize. We also develop online algorithms that can be used by service providers to implement various caching policies based on arbitrary utility functions.

13. Rapid Mixing of Dynamic Graphs with Local Evolution Rules [15] Dynamic graphs arise naturally in many contexts. In peer-to-peer networks, for instance, a participating peer may replace an existing connection with one neighbor by a new connection with a neighbor of that neighbor. Several such local rewiring rules have been proposed to ensure that peer-to-peer networks achieve good connectivity properties (e.g. high expansion) at equilibrium. However, the question of whether there exists such a rule that converges rapidly to equilibrium has remained open. In this work, we provide an affirmative answer: we exhibit a local rewiring rule that converges to equilibrium after each participating node has undergone only a number of changes that is at most poly-logarithmic in the system size. As a byproduct, we derive new results for random walks on graphs, bounding the spread of their law throughout the transient phase, i.e. prior to mixing. These rely on an extension of Cheeger's inequality, based on generalized isoperimetric constants, and may be of independent interest.

7.2. Reinforcement learning

14. On Matrix Momentum Stochastic Approximation and Applications to Q-learning [27] Stochastic approximation (SA) algorithms are recursive techniques used to obtain the roots of functions that can be expressed as expectations of a noisy parameterized family of functions. In this paper two new SA algorithms are introduced: 1) PolSA, an extension of Polyak's momentum technique with a specially designed matrix momentum, and 2) NeSA, which can either be regarded as a variant of Nesterov's acceleration method, or a simplification of PolSA. The rates of convergence of SA algorithms is well understood. Under special conditions, the mean square error of the parameter estimates is bounded by $\sigma^2/n + o(1/n)$, where $\sigma^2 \geq 0$ is an identifiable constant. If these conditions fail, the rate is typically sub-linear. There are two well known SA algorithms that ensure a linear rate, with minimal value of variance, σ^2 : the Ruppert-Polyak averaging technique, and the stochastic Newton-Raphson (SNR) algorithm. It is demonstrated here that under mild technical assumptions, the PolSA algorithm also achieves this optimality criteria. This result is established via novel coupling arguments: It is shown that the parameter estimates obtained from the PolSA algorithm couple with those of the optimal variance (but computationally more expensive) SNR algorithm, at a rate $O(1/n^2)$. The newly proposed algorithms are extended to a reinforcement learning setting to obtain new Q-learning algorithms, and numerical results confirm the coupling of PolSA and SNR.

15. Zap Q-Learning - A User's Guide [28] There are two well known Stochastic Approximation techniques that are known to have optimal rate of convergence (measured in terms of asymptotic variance): the Stochastic Newton-Raphson (SNR) algorithm (a matrix gain algorithm that resembles the deterministic Newton-Raphson method), and the Ruppert-Polyak averaging technique. This paper surveys new applications of these concepts for Q-learning: (i) The Zap Q-Learning algorithm was introduced by the authors in a NIPS 2017 paper. It is based on a variant of SNR, designed to more closely mimic its deterministic cousin. The algorithm has optimal rate of convergence under general assumptions, and showed astonishingly quick convergence in numerical examples. These algorithms are surveyed and illustrated with numerical examples. A potential difficulty in implementation of the Zap-Q-Learning algorithm is the matrix inversion required in each iteration. (ii) Remedies are proposed based on stochastic approximation variants of two general deterministic techniques: Polyak's momentum algorithms and Nesterov's acceleration technique. Provided the hyper-parameters are chosen with care, the performance of these algorithms can be comparable to the Zap algorithm, while computational complexity per iteration is far lower.

16. Zap Q-Learning With Nonlinear Function Approximation [44] The Zap stochastic approximation (SA) algorithm was introduced recently as a means to accelerate convergence in reinforcement learning algorithms. While numerical results were impressive, stability (in the sense of boundedness of parameter estimates) was established in only a few special cases. This class of algorithms is generalized in this paper, and stability is established under very general conditions. This general result can be applied to a wide range of algorithms found in reinforcement learning. Two classes are considered in this paper: (i) The natural generalization of Watkins' algorithm is not always stable in function approximation settings. Parameter estimates may diverge to infinity even in the *linear* function approximation setting with a simple finite state-action MDP. Under mild conditions, the Zap SA algorithm provides a stable algorithm, even in the case of *nonlinear* function approximation. (ii) The GQ algorithm of Maei et. al. 2010 is designed to address the stability challenge. Analysis is provided to explain why the algorithm may be very slow to converge in practice. The new Zap GQ algorithm is stable even for nonlinear function approximation.

17. Zap Q-Learning for Optimal Stopping Time Problems [43] The objective in this paper is to obtain fast converging reinforcement learning algorithms to approximate solutions to the problem of discounted cost optimal stopping in an irreducible, uniformly ergodic Markov chain, evolving on a compact subset of IR^n . We build on the dynamic programming approach taken by Tsitsikilis and Van Roy, wherein they propose a Q-learning algorithm to estimate the optimal state-action value function, which then defines an optimal stopping rule. We provide insights as to why the convergence rate of this algorithm can be slow, and propose a fast-converging alternative, the "Zap-Q-learning" algorithm, designed to achieve optimal rate of convergence. For the first time, we prove the convergence of the Zap-Q-learning algorithm under the assumption of linear

function approximation setting. We use ODE analysis for the proof, and the optimal asymptotic variance property of the algorithm is reflected via fast convergence in a finance example.

7.3. Mathematics of wireless cellular networks

18. Performance analysis of cellular networks with opportunistic scheduling using queueing theory and stochastic geometry [6] Combining stochastic geometric approach with some classical results from queueing theory, in this paper we propose a comprehensive framework for the performance study of large cellular networks featuring opportunistic scheduling. Rapid and verifiable with respect to real data, our approach is particularly useful for network dimensioning and long term economic planning. It is based on a detailed network model combining an information-theoretic representation of the link layer, a queueing-theoretic representation of the users' scheduler, and a stochastic-geometric representation of the signal propagation and the network cells. It allows one to evaluate principal characteristics of the individual cells, such as loads (defined as the fraction of time the cell is not empty), the mean number of served users in the steady state, and the user throughput. A simplified Gaussian approximate model is also proposed to facilitate study of the spatial distribution of these metrics across the network. The analysis of both models requires only simulations of the point process of base stations and the shadowing field to estimate the expectations of some stochastic-geometric functionals not admitting explicit expressions. A key observation of our approach, bridging spatial and temporal analysis, relates the SINR distribution of the typical user to the load of the typical cell of the network. The former is a static characteristic of the network related to its spectral efficiency while the latter characterizes the performance of the (generalized) processor sharing queue serving the dynamic population of users of this cell.

19. Two-tier cellular networks for throughput maximization of static and mobile users [10] In small cell networks, high mobility of users results in frequent handoff and thus severely restricts the data rate for mobile users. To alleviate this problem, we propose to use heterogeneous, two-tier network structure where static users are served by both macro and micro base stations, whereas the mobile (i.e., moving) users are served only by macro base stations having larger cells; the idea is to prevent frequent data outage for mobile users due to handoff. We use the classical two-tier Poisson network model with different transmit powers, assume independent Poisson process of static users and doubly stochastic Poisson process of mobile users moving at a constant speed along infinite straight lines generated by a Poisson line process. Using stochastic geometry, we calculate the average downlink data rate of the typical static and mobile (i.e., moving) users, the latter accounted for handoff outage periods. We consider also the average throughput of these two types of users defined as their average data rates divided by the mean total number of users co-served by the same base station. We find that if the density of a homogeneous network and/or the speed of mobile users is high, it is advantageous to let the mobile users connect only to some optimal fraction of BSs to reduce the frequency of handoffs during which the connection is not assured. If a heterogeneous structure of the network is allowed, one can further jointly optimize the mean throughput of mobile and static users by appropriately tuning the powers of micro and macro base stations subject to some aggregate power constraint ensuring unchanged mean data rates of static users via the network equivalence property.

20. Location Aware Opportunistic Bandwidth Sharing between Static and Mobile Users with Stochastic Learning in Cellular Networks [9] We consider location-dependent opportunistic bandwidth sharing between static and mobile downlink users in a cellular network. Each cell has some fixed number of static users. Mobile users enter the cell, move inside the cell for some time and then leave the cell. In order to provide higher data rate to mobile users, we propose to provide higher bandwidth to the mobile users at favourable times and locations, and provide higher bandwidth to the static users in other times. We formulate the problem as a long run average reward Markov decision process (MDP) where the per-step reward is a linear combination of instantaneous data volumes received by static and mobile users, and find the optimal policy. The transition structure of this MDP is not known in general. To alleviate this issue, we propose a learning algorithm based on single timescale stochastic approximation. Also, noting that the unconstrained MDP can be used to solve a constrained problem, we provide a learning algorithm based on multi-timescale stochastic approximation. The results are extended to address the issue of fair bandwidth sharing between the two classes of users. Numerical

results demonstrate performance improvement by our scheme, and also the trade-off between performance gain and fairness.

21. Per-Link Reliability and Rate Control: Two Facets of the SIR Meta Distribution [13] The meta distribution (MD) of the signal-to-interference ratio (SIR) provides fine-grained reliability performance in wireless networks modeled by point processes. In particular, for an ergodic point process, the SIR MD yields the distribution of the per-link reliability for a target SIR. Here we reveal that the SIR MD has a second important application, which is rate control. Specifically, we calculate the distribution of the SIR threshold (equivalently, the distribution of the transmission rate) that guarantees each link a target reliability and show its connection to the distribution of the per-link reliability. This connection also permits an approximate calculation of the SIR MD when only partial (local) information about the underlying point process is available.

22. Simple Approximations of the SIR Meta Distribution in General Cellular Networks [14] Compared to the standard success (coverage) probability, the meta distribution of the signal-to-interference ratio (SIR) provides much more fine-grained information about the network performance. We consider general heterogeneous cellular networks (HCNs) with base station tiers modeled by arbitrary stationary and ergodic non-Poisson point processes. The exact analysis of non-Poisson network models is notoriously difficult, even in terms of the standard success probability, let alone the meta distribution. Hence we propose a simple approach to approximate the SIR meta distribution for non-Poisson networks based on the ASAPPP ("approximate SIR analysis based on the Poisson point process") method. We prove that the asymptotic horizontal gap G_0 between its standard success probability and that for the Poisson point process exactly characterizes the gap between the b th moment of the conditional success probability, as the SIR threshold goes to 0. The gap G_0 allows two simple approximations of the meta distribution for general HCNs: 1) the per-tier approximation by applying the shift G_0 to each tier and 2) the effective gain approximation by directly shifting the meta distribution for the homogeneous independent Poisson network. Given the generality of the model considered and the fine-grained nature of the meta distribution, these approximations work surprisingly well.

23. Interference Queueing Networks [16] This work features networks of coupled processor sharing queues in the Euclidean space, where customers arrive according to independent Poisson point processes at every queue, are served, and then leave the network. The coupling is through service rates. In any given queue, this rate is inversely proportional to the interference seen by this queue, which is determined by the load in neighboring queues, attenuated by some distance-based path-loss function. The main focus is on the infinite grid network and translation invariant path-loss case. The model is a discrete version of a spatial birth and death process where customers arrive to the Euclidean space according to Poisson rain and leave it when they have transferred an exponential file, assuming that the instantaneous rate of each transfer is determined through information theory by the signal to interference and noise ratio experienced by the user. The stability condition is identified. The minimal stationary regime is built using coupling from the past techniques. The mean queue size of this minimal stationary regime is determined in closed form using the rate conservation principle of Palm calculus. When the stability condition holds, for all bounded initial conditions, there is weak convergence to this minimal stationary regime; however, there exist translation invariant initial conditions for which all queue sizes converge to infinity.

24. Statistical learning of geometric characteristics of wireless networks [19] Motivated by the prediction of cell loads in cellular networks, we formulate the following new, fundamental problem of statistical learning of geometric marks of point processes: An unknown marking function, depending on the geometry of point patterns, produces characteristics (marks) of the points. One aims at learning this function from the examples of marked point patterns in order to predict the marks of new point patterns. To approximate (interpolate) the marking function, in our baseline approach, we build a statistical regression model of the marks with respect to some local point distance representation. In a more advanced approach, we use a global data representation via the scattering moments of random measures, which build informative and stable to deformations data representation, already proven useful in image analysis and related application domains. In this case, the regression of the scattering moments of the marked point patterns with respect to the non-marked ones

is combined with the numerical solution of the inverse problem, where the marks are recovered from the estimated scattering moments. Considering some simple, generic marks, often appearing in the modeling of wireless networks, such as the shot-noise values, nearest neighbour distance, and some characteristics of the Voronoi cells, we show that the scattering moments can capture similar geometry information as the baseline approach, and can reach even better performance, especially for non-local marking functions. Our results motivate further development of statistical learning tools for stochastic geometry and analysis of wireless networks, in particular to predict cell loads in cellular networks from the locations of base stations and traffic demand.

25. Determinantal thinning of point processes with network learning applications [21] A new type of dependent thinning for point processes in continuous space is proposed, which leverages the advantages of determinantal point processes defined on finite spaces and, as such, is particularly amenable to statistical, numerical, and simulation techniques. It gives a new point process that can serve as a network model exhibiting repulsion. The properties and functions of the new point process, such as moment measures, the Laplace functional, the void probabilities, as well as conditional (Palm) characteristics can be estimated accurately by simulating the underlying (non-thinned) point process, which can be taken, for example, to be Poisson. This is in contrast (and preference to) finite Gibbs point processes, which, instead of thinning, require weighting the Poisson realizations, involving usually intractable normalizing constants. Models based on determinantal point processes are also well suited for statistical (supervised) learning techniques, allowing the models to be fitted to observed network patterns with some particular geometric properties. We illustrate this approach by imitating with determinantal thinning the well-known Matérn II hard-core thinning, as well as a soft-core thinning depending on nearest-neighbour triangles. These two examples demonstrate how the proposed approach can lead to new, statistically optimized, probabilistic transmission scheduling schemes.

26. Analyzing LoRa long-range, low-power, wide-area networks using stochastic geometry [22] In this paper we present a simple, stochastic-geometric model of a wireless access network exploiting the LoRa (Long Range) protocol, which is a non-expensive technology allowing for long-range, single-hop connectivity for the Internet of Things. We assume a space-time Poisson model of packets transmitted by LoRa nodes to a fixed base station. Following previous studies of the impact of interference, we assume that a given packet is successfully received when no interfering packet arrives with similar power before the given packet payload phase. This is as a consequence of LoRa using different transmission rates for different link budgets (transmissions with smaller received powers use larger spreading factors) and LoRa intra-technology interference treatment. Using our model, we study the scaling of the packet reception probabilities per link budget as a function of the spatial density of nodes and their rate of transmissions. We consider both the parameter values recommended by the LoRa provider, as well as proposing LoRa tuning to improve the equality of performance for all link budgets. We also consider spatially non-homogeneous distributions of LoRa nodes. We show also how a fair comparison to non-slotted Aloha can be made within the same framework.

27. Reliability and Local Delay in Wireless Networks: Does Bandwidth Partitioning Help? [33] In a series of papers initiated through a collaboration with Nokia Bell Labs, we study the effect of bandwidth partitioning (BWP) on the reliability and delay performance in infrastructureless wireless networks. The reliability performance is characterized by the density of concurrent transmissions that satisfy a certain reliability (outage) constraint and the delay performance by so-called local delay, defined as the average number of time slots required to successfully transmit a packet. We concentrate on the ultrareliable regime where the target outage probability is close to 0. BWP has two conflicting effects: while the interference is reduced as the concurrent transmissions are divided over multiple frequency bands, the signal-to-interference ratio (SIR) requirement is increased due to smaller allocated bandwidth if the data rate is to be kept constant. Instead, if the SIR requirement is to be kept the same, BWP reduces the data rate and in turn increases the local delay. For these two approaches with adaptive and fixed SIR requirements, we derive closed-form expressions of the local delay and the maximum density of reliable transmissions in the ultrareliable regime. Our analysis shows that, in the ultrareliable regime, BWP leads to the reliability-delay tradeoff.

28. The Influence of Canyon Shadowing on Device-to-Device Connectivity in Urban Scenario [35] In this work, we use percolation theory to study the feasibility of large-scale connectivity of relay-augmented device-to-device (D2D) networks in an urban scenario, featuring a haphazard system of streets and canyon shadowing allowing only for line-of-sight (LOS) communications in a limited finite range. We use a homogeneous Poisson-Voronoi tessellation (PVT) model of streets with homogeneous Poisson users (devices) on its edges and independent Bernoulli relays on the vertices. Using this model, we demonstrated the existence of a minimal threshold for relays below which large-scale connectivity of the network is not possible, regardless of all other network parameters. Through simulations, we estimated this threshold to 71.3%. Moreover, if the mean street length is not larger than some threshold (predicted to 74.3% of the communication range; which might be the case in a typical urban scenario) then any (whatever small) density of users can be compensated by equipping more crossroads with relays. Above this latter threshold, good connectivity requires some minimal density of users, compensated by the relays in a way we make explicit. The existence of the above regimes brings interesting qualitative arguments to the discussion on the possible D2D deployment scenarios.

29. Relay-assisted Device-to-Device Networks: Connectivity and Uberization Opportunities [46] It has been shown that deploying device-to-device (D2D) networks in urban environments requires equipping a considerable proportion of crossroads with relays. This represents a necessary economic investment for an operator. In this work, we tackle the problem of the economic feasibility of such relay-assisted D2D networks. First, we propose a stochastic model taking into account a positive surface for streets and crossroads, thus allowing for a more realistic estimation of the minimal number of needed relays. Secondly, we introduce a cost model for the deployment of relays, allowing one to study operators' D2D deployment strategies. We investigate the example of an uberizing neo-operator willing to set up a network entirely relying on D2D and show that a return on the initial investment in relays is possible in a realistic period of time, even if the network is funded by a very low revenue per D2D user. Our results bring quantitative arguments to the discussion on possible uberization scenarios of telecommunications networks.

30. Continuum Line-of-Sight Percolation on Poisson-Voronoi Tessellations [45] In this work, we study a new model for continuum line-of-sight percolation in a random environment given by a Poisson-Voronoi tessellation. The edges of this tessellation are the support of a Cox point process, while the vertices are the support of a Bernoulli point process. Taking the superposition of these two processes, two points are linked by an edge if and only if they are sufficiently close and located on the same edge of the supporting tessellation. We study the percolation of the random graph arising from this construction and prove that a subcritical phase as well as a supercritical phase exist under general assumptions. Our proofs are based on a renormalization argument with some notion of stabilization and asymptotic essential connectedness to investigate continuum percolation for Cox point processes. We also give numerical estimates of the critical parameters of the model. Our model can be seen as a good candidate for modelling telecommunications networks in a random environment with obstructive conditions for signal propagation.

7.4. High-dimensional statistical inference

31. Discrete Mean Field Games: Existence of Equilibria and Convergence [12] We consider mean field games with discrete state spaces (called discrete mean field games in the following) and we analyze these games in continuous and discrete time, over finite as well as infinite time horizons. We prove the existence of a mean field equilibrium assuming continuity of the cost and of the drift. These conditions are more general than the existing papers studying finite state space mean field games. Besides, we also study the convergence of the equilibria of N -player games to mean field equilibria in our four settings. On the one hand, we define a class of strategies in which any sequence of equilibria of the finite games converges weakly to a mean field equilibrium when the number of players goes to infinity. On the other hand, we exhibit equilibria outside this class that do not converge to mean field equilibria and for which the value of the game does not converge. In discrete time this non-convergence phenomenon implies that the Folk theorem does not scale to the mean field limit.

32. Modularity-based Sparse Soft Graph Clustering [32] Clustering is a central problem in machine learning for which graph-based approaches have proven their efficiency. In this paper, we study a relaxation

of the modularity maximization problem, well-known in the graph partitioning literature. A solution of this relaxation gives to each element of the dataset a probability to belong to a given cluster, whereas a solution of the standard modularity problem is a partition. We introduce an efficient optimization algorithm to solve this relaxation, that is both memory efficient and local. Furthermore, we prove that our method includes, as a special case, the Louvain optimization scheme, a state-of-the-art technique to solve the traditional modularity problem. Experiments on both synthetic and real-world data illustrate that our approach provides meaningful information on various types of data.

33. Phase Transitions, Optimal Errors and Optimality of Message-Passing in Generalized Linear Models [41] We consider generalized linear models where an unknown n -dimensional signal vector is observed through the successive application of a random matrix and a non-linear (possibly probabilistic) componentwise function. We consider the models in the high-dimensional limit, where the observation consists of m points, and $m/n \rightarrow \alpha$ where α stays finite in the limit $m, n \rightarrow \infty$. This situation is ubiquitous in applications ranging from supervised machine learning to signal processing. A substantial amount of work suggests that both the inference and learning tasks in these problems have sharp intrinsic limitations when the available data become too scarce or too noisy. Here, we provide rigorous asymptotic predictions for these thresholds through the proof of a simple expression for the mutual information between the observations and the signal. Thanks to this expression we also obtain as a consequence the optimal value of the generalization error in many statistical learning models of interest, such as the teacher-student binary perceptron, and introduce several new models with remarkable properties. We compute these thresholds (or "phase transitions") using ideas from statistical physics that are turned into rigorous methods thanks to a new powerful smart-path interpolation technique called the stochastic interpolation method, which has recently been introduced by two of the authors. Moreover we show that a polynomial-time algorithm referred to as generalized approximate message-passing reaches the optimal generalization performance for a large set of parameters in these problems. Our results clarify the difficulties and challenges one has to face when solving complex high-dimensional statistical problems.

34. Efficient inference in stochastic block models with vertex labels [18] We study the stochastic block model with two communities where vertices contain side information in the form of a vertex label. These vertex labels may have arbitrary label distributions, depending on the community memberships. We analyze a version of the popular belief propagation algorithm. We show that this algorithm achieves the highest accuracy possible whenever a certain function of the network parameters has a unique fixed point. When this function has multiple fixed points, the belief propagation algorithm may not perform optimally, where we conjecture that a non-polynomial time algorithm may perform better than BP. We show that increasing the information in the vertex labels may reduce the number of fixed points and hence lead to optimality of belief propagation.

35. Planting trees in graphs, and finding them back [36] In this paper we study detection and reconstruction of planted structures in Erdős-Rényi random graphs. Motivated by a problem of communication security, we focus on planted structures that consist in a tree graph. For planted line graphs, we establish the following phase diagram. In a low density region where the average degree λ of the initial graph is below some critical value $\lambda_c = 1$, detection and reconstruction go from impossible to easy as the line length K crosses some critical value $f(\lambda) \ln(n)$, where n is the number of nodes in the graph. In the high density region $\lambda > \lambda_c$, detection goes from impossible to easy as K goes from $o(\sqrt{n})$ to $\omega(\sqrt{n})$, and reconstruction remains impossible so long as $K = o(n)$. For D -ary trees of varying depth h and $2 \leq D \leq O(1)$, we identify a low-density region $\lambda < \lambda_D$, such that the following holds. There is a threshold $h^* = g(D) \ln(\ln(n))$ with the following properties. Detection goes from feasible to impossible as h crosses h^* . We also show that only partial reconstruction is feasible at best for $h \geq h^*$. We conjecture a similar picture to hold for D -ary trees as for lines in the high-density region $\lambda > \lambda_D$, but confirm only the following part of this picture: Detection is easy for D -ary trees of size $\omega(\sqrt{n})$, while at best only partial reconstruction is feasible for D -ary trees of any size $o(n)$. These results are in contrast with the corresponding picture for detection and reconstruction of *low rank* planted structures, such as dense subgraphs and block communities: We observe a discrepancy between detection and reconstruction, the latter being impossible for a wide range of parameters where detection is easy. This property does not hold for previously studied low rank planted structures.

36. Robustness of spectral methods for community detection [37] This work is concerned with community detection. Specifically, we consider a random graph drawn according to the stochastic block model: its vertex set is partitioned into blocks, or communities, and edges are placed randomly and independently of each other with probability depending only on the communities of their two endpoints. In this context, our aim is to recover the community labels better than by random guess, based only on the observation of the graph.

In the sparse case, where edge probabilities are in $O(1/n)$, we introduce a new spectral method based on the distance matrix D , where $D_{ij} = 1$ iff the graph distance between i and j , noted $d(i, j)$ is equal to ℓ . We show that when $\ell \sim c \log(n)$ for carefully chosen c , the eigenvectors associated to the largest eigenvalues of D provide enough information to perform non-trivial community recovery with high probability, provided we are above the so-called Kesten-Stigum threshold. This yields an efficient algorithm for community detection, since computation of the matrix D can be done in $O(n^{1+\kappa})$ operations for a small constant κ .

We then study the sensitivity of the eigendecomposition of D when we allow an adversarial perturbation of the edges of G . We show that when the considered perturbation does not affect more than $O(n^\varepsilon)$ vertices for some small $\varepsilon > 0$, the highest eigenvalues and their corresponding eigenvectors incur negligible perturbations, which allows us to still perform efficient recovery.

Our proposed spectral method therefore: i) is robust to larger perturbations than prior spectral methods, while semi-definite programming (or SDP) methods can tolerate yet larger perturbations; ii) achieves non-trivial detection down to the KS threshold, which is conjectured to be optimal and is beyond reach of existing SDP approaches; iii) is faster than SDP approaches.

7.5. Distributed optimization for machine learning

37. Optimal Convergence Rates for Convex Distributed Optimization in Networks [17] This work proposes a theoretical analysis of distributed optimization of convex functions using a network of computing units. We investigate this problem under two communication schemes (centralized and decentralized) and four classical regularity assumptions: Lipschitz continuity, strong convexity, smoothness, and a combination of strong convexity and smoothness. Under the decentralized communication scheme, we provide matching upper and lower bounds of complexity along with algorithms achieving this rate up to logarithmic constants. For non-smooth objective functions, while the dominant term of the error is in $O(1/\sqrt{t})$, the structure of the communication network only impacts a second-order term in $O(1/t)$, where t is time. In other words, the error due to limits in communication resources decreases at a fast rate even in the case of non-strongly convex objective functions. Such a convergence rate is achieved by the novel multi-step primal-dual (MSPD) algorithm. Under the centralized communication scheme, we show that the naive distribution of standard optimization algorithms is optimal for smooth objective functions, and provide a simple yet efficient algorithm called distributed randomized smoothing (DRS) based on a local smoothing of the objective function for non-smooth functions. We then show that DRS is within a $d^{1/4}$ multiplicative factor of the optimal convergence rate, where d is the underlying dimension.

38. Accelerated Decentralized Optimization with Local Updates for Smooth and Strongly Convex Objectives [31] In this paper, we study the problem of minimizing a sum of smooth and strongly convex functions split over the nodes of a network in a decentralized fashion. We propose the algorithm *ESDACD*, a decentralized accelerated algorithm that only requires local synchrony. Its rate depends on the condition number κ of the local functions as well as the network topology and delays. Under mild assumptions on the topology of the graph, *ESDACD* takes a time $O((\tau_{\max} + \Delta_{\max})\sqrt{\kappa/\gamma} \ln(\epsilon^{-1}))$ to reach a precision ϵ where γ is the spectral gap of the graph, τ_{\max} the maximum communication delay and Δ_{\max} the maximum computation time. Therefore, it matches the rate of *SSDA*, which is optimal when $\tau_{\max} = \Omega(\Delta_{\max})$. Applying *ESDACD* to quadratic local functions leads to an accelerated randomized gossip algorithm of rate $O(\sqrt{\theta_{\text{gossip}}/n})$ where θ_{gossip} is the rate of the standard randomized gossip. To the best of our knowledge, it is the first asynchronous gossip algorithm with a provably improved rate of convergence of the second moment of the error. We illustrate these results with experiments in idealized settings.

39. An Accelerated Decentralized Stochastic Proximal Algorithm for Finite Sums [49] Modern large-scale finite-sum optimization relies on two key aspects: distribution and stochastic updates. For smooth and strongly convex problems, existing decentralized algorithms are slower than modern accelerated variance-reduced stochastic algorithms when run on a single machine, and are therefore not efficient. Centralized algorithms are fast, but their scaling is limited by global aggregation steps that result in communication bottlenecks. In this work, we propose an efficient Accelerated, Decentralized stochastic algorithm for FiniteSums named ADFS, which uses local stochastic proximal updates and randomized pairwise communications between nodes. On machines, ADFS learns from samples in the same time it takes optimal algorithms to learn from samples on one machine. This scaling holds until a critical network size is reached, which depends on communication delays, on the number of samples, and on the network topology. We provide a theoretical analysis based on a novel augmented graph approach combined with a precise evaluation of synchronization times and an extension of the accelerated proximal coordinate gradient algorithm to arbitrary sampling. We illustrate the improvement of ADFS over state-of-the-art decentralized approaches with experiments.

7.6. Stochastic Geometry

40. On the Dimension of Unimodular Discrete Spaces, Part I: Definitions and Basic Properties [39] This work introduces two new notions of dimension, namely the *unimodular Minkowski and Hausdorff dimensions*, which are inspired from the classical analogous notions. These dimensions are defined for *unimodular discrete spaces*, introduced in this work, which provide a common generalization to stationary point processes under their Palm version and unimodular random rooted graphs. The use of unimodularity in the definitions of dimension is novel. Also, a toolbox of results is presented for the analysis of these dimensions. In particular, analogues of Billingsley's lemma and Frostman's lemma are presented. These lemmas are instrumental in deriving upper bounds on dimensions, whereas lower bounds are obtained from specific coverings. The notions of unimodular Hausdorff measure and unimodular dimension function are also introduced. This toolbox is used to connect the unimodular dimensions to various other notions such as growth rate, scaling limits, discrete dimension and amenability. It is also used to analyze the dimensions of a set of examples pertaining to point processes, branching processes, random graphs, random walks, and self-similar discrete random spaces.

41. On the Dimension of Unimodular Discrete Spaces, Part II: Relations with Growth Rate [40] The notions of unimodular Minkowski and Hausdorff dimensions are defined in [39] for unimodular random discrete metric spaces. This work is focused on the connections between these notions and the polynomial growth rate of the underlying space. It is shown that bounding the dimension is closely related to finding suitable equivariant weight functions (i.e., measures) on the underlying discrete space. The main results are unimodular versions of the mass distribution principle and Billingsley's lemma, which allow one to derive upper bounds on the unimodular Hausdorff dimension from the growth rate of suitable equivariant weight functions. Also, a unimodular version of Frostman's lemma is provided, which shows that the upper bound given by the unimodular Billingsley lemma is sharp. These results allow one to compute or bound both types of unimodular dimensions in a large set of examples in the theory of point processes, unimodular random graphs, and self-similarity. Further results of independent interest are also presented, like a version of the max-flow min-cut theorem for unimodular one-ended trees.

42. Doebelin trees [4] This work is centered on the random graph generated by a Doebelin-type coupling of discrete time processes on a countable state space whereby when two paths meet, they merge. This random graph is studied through a novel subgraph, called a bridge graph, generated by paths started in a fixed state at any time. The bridge graph is made into a unimodular network by marking it and selecting a root in a specified fashion. The unimodularity of this network is leveraged to discern global properties of the larger Doebelin graph. Bi-recurrence, i.e., recurrence both forwards and backwards in time, is introduced and shown to be a key property in uniquely distinguishing paths in the Doebelin graph, and also a decisive property for Markov chains indexed by \mathbb{Z} . Properties related to simulating the bridge graph are also studied.

43. The Stochastic Geometry of Unconstrained One-Bit Compression [5] A stationary stochastic geometric model is proposed for analyzing the data compression method used in one-bit compressed sensing. The data set is an unconstrained stationary set, for instance all of \mathbb{R}^n or a stationary Poisson point process in \mathbb{R}^n . It

is compressed using a stationary and isotropic Poisson hyperplane tessellation, assumed independent of the data. That is, each data point is compressed using one bit with respect to each hyperplane, which is the side of the hyperplane it lies on. This model allows one to determine how the intensity of the hyperplanes must scale with the dimension n to ensure sufficient separation of different data by the hyperplanes as well as sufficient proximity of the data compressed together. The results have direct implications in compressed sensing and in source coding.

44. Limit theory for geometric statistics of point processes having fast decay of correlations [7] We develop a limit theory (Laws of Large Numbers and Central Limit Theorems) for functionals of spatially correlated point processes. The “strength” of data correlation is captured and controlled by the speed of decay of the additive error in the asymptotic factorization the correlation functions, when the separation distance increases. In this way, the classical theory of Poisson and Bernoulli processes is extended to a larger class of data inputs, such as determinantal point processes with fast decreasing kernels, including the α -Ginibre ensembles, permanental point processes as well as the zero set of Gaussian entire functions. Both linear (U-statistics) and non-linear geometric statistics (such as clique counts, the number of Morse critical points, intrinsic volumes of the Boolean model, and total edge length of the k -nearest neighbor graph) are considered.

7.7. Information theory

45. Error Exponents for MAC Channelss [3] This work analyzes a class of Multiple Access Channels (MAC) where the sum of the dimensions of the transmitted signals matches that of the received signal. This channel is a classical object of information theory in the power constrained case. We first focus on the Poltyrev regime, namely the case without power constraint. Using point process techniques, we derive the capacity under general stationarity and ergodicity noise assumptions as well as a representation of the error probability. We use this to derive bounds on the error exponent in the Gaussian case. This also leads to new results on the power constrained error exponents.

EVA Project-Team

7. New Results

7.1. Falco startup launched!

Participants: Elsa Nicol, Keoma Brun-Laguna, Thomas Watteyne.

The Falco startup (<https://wefalco.com/>) was launched on 14 January 2019. During 2019, it developed a complete technical solution (PCB, assembly, networking, back-end) and completed a large market development campaign. Falco was selected to join the Parisian Incubator Agoranov. It was then awarded the prestigious Netva “Deeptech North America” program, and won the Favorite Startup Pitch battle at MassChallenge, Boston, as well as the Amplify Pitch battle. It had a booth at the Cap d’Agde and Paris Nautic shows. On 14 December 2019, Falco wins the Innovation Competition at the Paris Nautic Show.

7.2. 6TiSCH Standardization

Participants: Malisa Vucinic, Jonathan Muñoz, Tengfei Chang, Yasuyuki Tanaka, Thomas Watteyne.

The standardization work at 6TiSCH remains a strong federator of the work done in the team. In 2019, the working group finalized the work on the draft-ietf-6tisch-minimal-security and draft-ietf-6tisch-architecture specification, which are both in the editor queue. The draft-ietf-6tisch-msf has also passed the working group last call. This standardization work has resulted in several papers on 6TiSCH, including a tutorial [9], [11], [12] fragmentation in 6TiSCH [8], implementation details [17], simulating 6TiSCH [24], experimental approaches [21], [26], localization [25], multi-PHY extensions [10]. The HDR of Thomas Watteyne [2] reports on the work on 6TiSCH over the past years. Some work has started on implementing 6TiSCH on single-chip micro-motes [19], [23], [7], [18].

7.3. 6TiSCH Security

Participants: Malisa Vucinic, Thomas Watteyne.

The security work of Inria-EVA is a continuation of the efforts started during the H2020 ARMOUR project. The work focused on stabilizing the “Minimal Security” solution that has now been approved to be published as an RFC [13]. The solution that is standardized enables secure network access and configuration of 6TiSCH devices under the assumption that they have been provisioned with a secret key. Ongoing work extends this solution to support true zero-configuration network setup, under the assumption that the devices have been provisioned with certificates at manufacturing time.

7.4. 6TiSCH Benchmarking

Participants: Malisa Vucinic, Tengfei Chang, Yasuyuki Tanaka, Thomas Watteyne.

With the pure 6TiSCH standardizes coming to an end, the focus of the group is moving towards benchmarking how well it works. This has resulted in the following action. Although seemingly different, they all contribute to the overall goal of better understanding (the performance of) 6TiSCH.

We have built and put online the OpenTestbed, a collection of 80 OpenMote B boards deployed in 20 “pods”. These allow us to test the performance of the OpenWSN firmware in a realistic setting. You can access its management interface at <http://testbed.openwsn.org/>.

A tool complementary to the testbed is the 6TiSCH simulator (<https://bitbucket.org/6tisch/simulator>) which Yatsuyiki Tanaka is leading. The simulator now represents exactly the behavior of the 6TiSCH protocol stack, and has been a catalyst for benchmarking activities around 6TiSCH.

Beyond Inria, the benchmarking activity around 6TiSCH is a hot topic, with projects such as the 6TiSCH Open Data Action [26] (SODA, <http://www.soda.ucg.ac.me/>), the IoT Benchmarks Initiative (<https://www.iotbench.ethz.ch/>), and the Computer and Networking Experimental Research using Testbeds (CNERT) workshop at INFOCOM, all of which Inria-EVA is very involved in.

7.5. LAKE Standardization

Participants: Malisa Vucinic, Timothy Claeys, Thomas Watteyne.

In October 2019, a new working group was formed in the IETF with the goal of standardizing a lightweight authenticated key exchange protocol for IoT use cases. The group is co-chaired by Malisa Vucinic of Inria-EVA. Through our work in 6TiSCH and the requirements for the follow up work of the “Minimal Security Framework for 6TiSCH”, we directly contributed to the creation of this working group whose expected output is the key exchange protocol for the IoT. The document we lead in the LAKE working group [37] compiles the requirements for a lightweight authenticated key exchange protocol for OSCORE. OSCORE (RFC8613) is a lightweight communication security protocol providing end-to-end security on application layer for constrained IoT settings. It is expected to be deployed with standards and frameworks using CoAP such as 6TiSCH, LPWAN, OMA Specworks LwM2M, Fairhair Alliance and Open Connectivity Foundation.

7.6. IoT and Low-Power Wireless Meshed Networks

More than 50 billion devices will be connected in 2020. This huge infrastructure of devices, which is managed by highly developed technologies, is called the Internet of Things (IoT). The IoT provides advanced services, and brings economic and societal benefits. This is the reason why engineers and researchers in both industry and scientific communities are interested in this area. The Internet of Things enables the interconnection of smart physical and virtual objects, managed by highly developed technologies. Low-Power Wireless Meshed Network is an essential part of this paradigm. It uses smart, autonomous and usually limited capacity devices in order to sense and monitor their environment.

7.6.1. Centralized or Distributed Scheduling for IEEE 802.15.4e TSCH networks

Participants: Yasuyuki Tanaka, Pascale Minet, Thomas Watteyne, Malisa Vucinic, Tengfei Chang, Keoma Brun-Laguna.

The wireless TSCH (Time Slotted Channel Hopping) network specified in the e amendment of the IEEE 802.15.4 standard has many appealing properties. Its schedule of multichannel slotted data transmissions ensures the absence of collisions. Because there is no retransmission due to collisions, communication is faster. Since the devices save energy each time they do not take part in a transmission, the power autonomy of nodes is prolonged. Furthermore, channel hopping mitigates multipath fading and interferences.

All communication in a TSCH network is orchestrated by the communication schedule it is using. The scheduling algorithm used hence drives the latency and capacity of the network, and the power consumption of the nodes. To increase the flexibility and the self-organizing capacities required by IoT, the networks have to be able to adapt to changes. These changes may concern the application itself, the network topology by adding or removing devices, the traffic generated by increasing or decreasing the device sampling frequency, for instance. That is why flexibility of the schedule ruling all network communications is needed. We have designed a number of scheduling algorithms for TSCH networks, answering different needs. For instance, the centralized Load-based scheduler that assigns cells per flow, starting with the flow originating from the most loaded node has proved optimal for many configurations. Simulations with the 6TiSCH simulator showed that it gets latencies close to the optimal. They also highlighted that end-to-end latencies are positively impacted by message prioritization (i.e. each node transmits the oldest message first) at high loads, and negatively impacted by unreliable links, as presented at GlobeCom 2019 [30].

Among the distributed scheduling algorithms proposed in the literature, many rely on assumptions that may be violated by real deployments. This violation usually leads to conflicting transmissions of application data, decreasing the reliability and increasing the latency of data delivery. Others require a processing complexity that cannot be provided by sensor nodes of limited capabilities. Still others are unable to adapt quickly to traffic or topology changes, or are valid only for small traffic loads. We have designed MSF and YSF, two distributed scheduling algorithms that are adaptive and compliant with the standardized protocols used in the 6TiSCH working group at IETF. The Minimal Scheduling Function (MSF) is a distributed scheduling algorithm in which neighbor nodes locally negotiate adding and removing cells. MSF was evaluated by simulation and experimentation, before becoming the default scheduling algorithm of the IETF 6TiSCH working group, and now an official standard. We also designed LLSF, a scheduling algorithm focused on low latency communication. We proposed a full-featured 6TiSCH scheduling function called YSF, that autonomously takes into account all the aspects of network dynamics, including the network formation phase and parent switches. YSF aims at minimizing latency and maximizing reliability for data gathering applications. Simulation results obtained with the 6TiSCH simulator show that YSF yields lower end-to-end latency and higher end-to-end reliability than MSF, regardless of the network topology. Unlike other top-down scheduling functions, YSF does not rely on any assumption regarding network topology or traffic load, and is therefore more robust in real network deployments. An intensive simulation campaign made with the 6TiSCH simulator has provided comparative performance results. Our proposal outperforms MSF, the 6TiSCH Minimal Scheduling Function, in terms of end-to-end latency and end-to-end packet delivery ratio.

Furthermore we published additional research on computing the upper bounds on the end-to-end latency, finding the best trade-off between latency and network lifetime.

7.6.2. *Modeling and Improving Named Data Networking over IEEE 802.15.4*

Participants: Amar Abane, Samia Bouzebrane (Cnam), Paul Muhlethaler.

Enabling Named Data Networking (NDN) in real world Internet of Things (IoT) deployments becomes essential to benefit from Information Centric Networking (ICN) features in current IoT systems. One objective of the model is to show that caching can attenuate the number of transmissions generated by broadcast to achieve a reasonable overhead while keeping the data dissemination power of NDN. To design realistic NDN-based communication solutions for IoT, revisiting mainstream technologies such as low-power wireless standards may be the key. We explore the NDN forwarding over IEEE 802.15.4 by modeling a broadcast-based forwarding [27]. Based on the observations, we adapt the Carrier-Sense Multiple Access (CSMA) algorithm of 802.15.4 to improve NDN wireless forwarding while reducing broadcast effects in terms of packet redundancy, round-trip time and energy consumption. As future work, we aim to explore more complex CSMA adaptations for lightweight forwarding to make the most of NDN and design a general-purpose Named-Data CSMA.

7.6.3. *Evaluation of LORA with stochastic geometry*

Participants: Bartek Blaszczyzyn (Dyogene), Paul Muhlethaler.

We present a simple, stochastic-geometric model of a wireless access network exploiting the LoRA (Long Range) protocol, which is a non-expensive technology allowing for long-range, single-hop connectivity for the Internet of Things. We assume a space-time Poisson model of packets transmitted by LoRA nodes to a fixed base station. Following previous studies of the impact of interference, we assume that a given packet is successfully received when no interfering packet arrives with similar power before the given packet payload phase, see [16]. This is as a consequence of LoRa using different transmission rates for different link budgets (transmissions with smaller received powers use larger spreading factors) and LoRa intra-technology interference treatment. Using our model, we study the scaling of the packet reception probabilities per link budget as a function of the spatial density of nodes and their rate of transmissions. We consider both the parameter values recommended by the LoRa provider, as well as proposing LoRa tuning to improve the equality of performance for all link budgets. We also consider spatially non-homogeneous distributions of LoRa nodes. We show how a fair comparison to non-slotted Aloha can be made within the same framework.

7.6.4. *Position Certainty Propagation: A location service for MANETs*

Participants: Abdallah Sobehy, Paul Muhlethaler, Eric Renault (Telecom Sud-Paris).

A location method based on triangulation (via Channel State Information (CSI) based localization method is proposed [6]. A known method of triangulation is adopted to deduce the location of a node from 3 reference nodes (anchor nodes). We propose an optimized energy-aware and low computational solution, requiring 3-GPS equipped nodes (anchor nodes) in the network. Moreover, the computations are lightweight and can be implemented distributively among nodes. Knowing the maximum range of communication for all nodes and distances between 1-hop neighbors, each node localizes itself and shares its location with the network in an efficient manner. We simulate our proposed algorithm on a NS-3 simulator, and compare our solution with state-of-the-art methods. Our method is capable of localizing more nodes i.e. $\simeq 90\%$ of nodes in a network with an average degree $\simeq 10$.

7.7. Industry 4.0 and Low-Power Wireless Meshed Networks

The Internet of Things (IoT) connects tiny electronic devices able to measure a physical value (temperature, humidity, etc.) and/or to actuate on the physical world (pump, valve, etc). Due to their cost and ease of deployment, battery-powered wireless IoT networks are rapidly being adopted.

The promise of wireless communication is to offer wire-like connectivity. Major improvements have been made in that direction, but many challenges remain as industrial applications have strong operational requirements. This section of the IoT application is called Industrial IoT (IIoT).

By the year 2020, it is expected that the number of connected objects will exceed several billion devices. These objects will be present in everyday life for a smarter home and city as well as in future smart factories that will revolutionize the industry organization. This is actually the expected fourth industrial revolution, better known as Industry 4.0. In which, the Internet of Things (IoT) is considered as a key enabler for this major transformation. The IoT will allow more intelligent monitoring and self-organizing capabilities than traditional factories. As a consequence, the production process will be more efficient and flexible with products of higher quality.

To produce better quality products and improve monitoring in Industry 4.0, strong requirements in terms of latency, robustness and power autonomy have to be met by the networks supporting the Industry 4.0 applications.

7.7.1. Reliability for the Industrial Internet of Things (IIoT) and Industry 4.0

Participants: Yasuyuki Tanaka, Pascale Minet, Keoma Brun-Laguna, Thomas Watteyne.

The main IIoT requirement is reliability. Every bit of information that is transmitted in the network must not be lost. Current off-the-shelf solutions offer over 99.999% reliability.

To provide the end-to-end reliability targeted by industrial applications, we investigate an approach based on message retransmissions (on the same path). We propose two methods to compute the maximum number of transmissions per message and per link required to achieve the targeted end-to-end reliability. The MFair method is very easy to compute and provides the same reliability over each link composing the path, by means of different maximum numbers of transmissions, whereas the MOpt method minimizes the total number of transmissions necessary for a message to reach the sink. MOpt provides a better reliability and a longer lifetime than MFair, which provides a shorter average end-to-end latency. This study [5] was published in the Sensors journal in 2019.

7.8. Machine Learning applied to Networking

7.8.1. Machine Learning for energy-efficient and QoS-aware Data Centers

Participants: Ruben Milocco (Comahue University, Argentina, Invited Professor), Pascale Minet, Eric Renault (Telecom Sud-Paris), Selma Boumerdassi (Cnam).

To limit global warming, all industrial sectors must make effort to reduce their carbon footprint. Information and Communication Technologies (ICTs) alone generate 2% of global CO₂ emissions every year. Due to the rapid growth in Internet services, data centers have the largest carbon footprint of all ICTs. According to ARCEP (the French telecommunications regulator), Internet data traffic multiplied by 4.5 between 2011 and 2016. In order to support such a growth and maintain this traffic, data centers' energy consumption needs to be optimized.

We determine whether resource allocation in DCs can satisfy the three following requirements: 1) meet user requirements (e.g. short response times), 2) keep the data center efficient, and 3) reduce the carbon footprint.

An efficient way to reduce the energy consumption in a DC is to turn off servers that are not used for a minimum duration. The high dynamicity of the jobs submitted to the DC requires periodically adjusting the number of active servers to meet job requests. This is called Dynamic Capacity Provisioning. This provisioning can be based on prediction. In such a case, a proactive management of the DC is performed. The goal of this study is to provide a methodology to evaluate the energy cost reduction brought by proactive management, while keeping a high level of user satisfaction.

The state-of-the art shows that appropriate proactive management improves the cost, either by improving QoS or saving energy. As a consequence, there is great interest in studying different proactive strategies based on predictions of either the energy or the resources needed to serve CPU and memory requests. The cost depends on 1) the proactive strategy used, 2) the workload requested by jobs and 3) the prediction used. The problem complexity explains why, despite its importance, the maximum cost savings have not been evaluated in theoretical studies.

We propose a method to compute the upper bound of the relative cost savings obtained by proactive management compared to a purely reactive management based on the Last Value. With this method, it becomes possible to quantitatively compare the efficiency of two predictors.

We also show how to apply this method to a real DC and how to select the value of the DC parameters to get the maximum cost savings. Two types of predictors are studied: linear predictors, represented by the ARMA model, and nonlinear predictors obtained by maximizing the conditional probability of the next sample, given the past. They are both applied to the publicly available Google dataset collected over a period of 29 days. We evaluate the largest benefit that can be obtained with those two predictors. Some of these results have been presented at HPCS 2019 [20].

7.8.2. Machine Learning applied to IoT networks

Participants: Miguel Landry Foko Sindjoung (Phd Student, Dschang University, Cameroon, Inria Internship), Pascale Minet.

Knowledge of link quality in IoT networks allows a more accurate selection of wireless links to build the routes used for data gathering. The number of re-transmissions is decreased, leading to shorter end-to-end latency, better end-to-end reliability and a longer network lifetime.

We propose to predict link quality by means of machine learning techniques applied on two metrics: the Received Signal Strength Indicator (RSSI) and the Packet Delivery Ratio (PDR). These two metrics were selected because RSSI is a hardware metric that is easily obtained and PDR takes into account packets that are not successfully received, unlike RSSI.

The data set used in this study was collected from a TSCH network deployed in the Grenoble testbed consisting of 50 nodes operating on 16 channels. Data collected by Mercator include 108659 measurements of PDR and average RSSI. We train the model over the training set and predict the link quality on the channel considered for the samples in the validation set. By comparing the predicted values with the real values, the confusion matrix is computed by evaluating the number of true-positive, true-negative, false-positive and false-negative for the link and channel considered.

Whatever the link quality estimator used, RSSI, PDR or both, the Random Forest (RF) classifier model outperforms the other models studied: Linear Regression, Linear Support Vector Machine, Support Vector Machine.

Since using Bad links that have been predicted Good strongly penalizes network performance in terms of end-to-end latency, end-to-end reliability and network lifetime, the joint use of PDR and RSSI improves the accuracy of link quality prediction. Hence, we recommend using the Random Forest classifier applied on both PDR and RSSI metrics. This work has been presented at the PEMWN 2019 conference [33].

7.9. Machine Learning applied to Smart Farming

Participants: Jamal Ammouri (Internship Cnam), Malika Boudiaf (Ummto, Tizi-Ouzou, Algeria), Samia Bouzeffrane (Cnam), Pascale Minet, Meziane Yacoub (Cnam).

Intelligent Farming System (IFS) is made possible by the use of 4 elements: sensors and actuators, the Internet of Things (IoT), edge/cloud processing, and machine learning.

Soil degradation and a hot climate explain the poor yield of olive groves in North Algeria. Edaphic, climatic and geographical data were collected from 10 olive groves over several years and analyzed by means of Self-Organizing Maps (SOMs). SOM is a non-supervised neural network that projects high-dimensional data onto a low-dimension discrete space, called a topological map, such that close data are mapped onto nearby locations on the map. In the paper [28] presented at the PEMWN 2019 conference, we have shown how to use self-organizing maps to determine olive grove clusters with similar features, characterize each cluster and show the temporal evolution of each olive grove. With the SOM, it becomes possible to alert the farmer when some specific action needs to be done in the case of hydric stress, NPK stress, pest/disease attack. As a result, the nutritional quality of the oil produced is improved. SOM can be integrated in the Intelligent Farming System (IFS) to boost conservation agriculture.

This work requires a strong collaboration with agronomists. Malika Boudiaf (Laboratoire Ressources Naturelles, UMMTO, Tizi-Ouzou, Algeria) provided the data set and gave us many explanations about soil conservation. Meziane Yacoub (Cnam) is an expert in SOMs. Jamal Ammouri (Cnam) was co-advised by Samia Bouzeffrane, Pascale Minet and Meziane Yacoub.

7.10. Protocols and Models for Wireless Networks - Application to VANETs

7.10.1. Connection-less IoT - Protocol and models

Participants: Iman Hemdoush, Cédric Adjih, Paul Mühlethaler.

The goal is to construct some next-generation access protocols, for the IoT (or alternately for vehicular networks). One starting point are methods from the family of Non-Orthogonal Multiple Access (NOMA), where multiple transmissions can "collide" but can still be recovered - with sophisticated multiple access protocols (MAC) that take the physical layer/channel into account. One such example is the family of the Coded Slotted Aloha methods. Another direction is represented by some vehicular communications where vehicles communicate directly with each other without necessarily going through the infrastructure. This is also true more generally in any wireless network where the control is relaxed (such as in unlicensed IoT networks like LoRa). One observation is that in such distributed scenarios, explicit or implicit forms of signaling (with sensing, messaging, etc.), can be used for designing sophisticated protocols - including using machine learning techniques.

During this study, some of the following tools should be used: protocol/algorithm design (ensuring properties by construction), simulations (ns-2, ns-3, matlab, ...) on detailed or simplified network models, mathematical modeling (stochastic geometry, etc...); machine-learning techniques or modeling as code-on-graphs.

The first result we have obtained concerns Irregular Repetition Slotted Aloha (IRSA) which is a modern method of random access for packet networks that is based on repeating transmitted packets, and on successive interference cancellation at the receiver. In classical idealized settings of slotted random access protocols (where slotted ALOHA achieves $1/e$), it has been shown that IRSA could asymptotically achieve the maximal throughput of 1 packet per slot. Additionally, IRSA had previously been studied for many different variants and settings, including the case where the receiver is equipped with "multiple-packet reception" (MPR) capability. We extensively revisit the case of IRSA with MPR. We present a method to compute optimal

IRSA degree distributions with a given maximum degree n . A tighter bound for the load threshold (G/K) was proven, showing that plain K-IRSA cannot reach the asymptotic known bound $G/K = 1$ for $K > 1$, and we prove a new, lower bound for its performance. Numerical results illustrate that optimal degree distributions can approach this bound. Second, we analyze the error floor behavior of K-IRSA and provide an insightful approximation of the packet loss rate at low loads, and show its excellent performance. Third, we show how to formulate the search for the appropriate parameters of IRSA as an optimization problem, and how to solve it efficiently. By doing that for a comprehensive set of parameters, and by providing this work with simulations, we give numerical results that shed light on the performance of IRSA with MPR. A final open question is: what is the impact of introducing more structure in the slot selection (like Spatially Coupled Coded Slotted Aloha) and how best to do so?

7.10.2. Indoor positioning using Channel State Information (CSI) from a MIMO antenna

Participants: Abdallah Sobehy, Paul Muhlethaler, Eric Renault (Telecom Sud-Paris).

The channel status information is used for locating a node by applying machine learning [35] techniques. We propose a novel lightweight deep learning solution to the indoor positioning problem based on noise and dimensionality reduction of MIMO Channel State Information (CSI): real and imaginary parts of the signal received. Based on preliminary data analysis, the magnitude of the CSI is selected as the input feature for a Multilayer Perceptron (MLP) neural network. Polynomial regression is then applied to batches of data points to filter noise and reduce input dimensionality by a factor of 14. The MLP's hyper-parameters are empirically tuned to achieve the highest accuracy. The method is applied to a CSI dataset estimated at an 8×2 MIMO antenna that is published by the organizers of the Communication Theory Workshop Indoor Positioning Competition. The proposed solution is compared with a state-of-the-art method presented by the authors who designed the MIMO antenna that is used to generate the data-set. Our method yields a mean error which is 8 times less than that of its counterpart. We conclude that the arithmetic mean and standard deviation misrepresent the results since the errors follow a log-normal distribution. The mean of the log error distribution of our method translates to a mean error as low as 1.5 cm. We have shown that, using a K-nearest neighbor learning method an even better, indoor positioning is achieved. The input feature is the magnitude component of CSI which is pre-processed to reduce noise and allow for a quicker search. The Euclidean distance between CSI is the criterion chosen for measuring the closeness between samples. The proposed method is compared with three other methods, all based on deep learning approaches and tested with the same data-set. The K-nearest neighbor method presented in this paper achieves a Mean Square Error (MSE) of 2.4 cm, which outperforms its counterparts.

7.10.3. Predicting Vehicles Positions using Roadside Units: a Machine-Learning Approach

Participants: Samia Bouzeffrane (Cnam), Soumya Banerjee (Birla Institute Of Technology, Mesra), Paul Muhlethaler, Mamoudou Sangare.

We study positioning systems using Vehicular Ad Hoc Networks (VANETs) to predict the position of vehicles. We use the reception power of the packets received by the Road Side Units (RSUs) and sent by the vehicles on the roads. In fact, the reception power is strongly influenced by the distance between a vehicle and a RSU. We have already used and compared three widely recognized techniques : K Nearest Neighbors (KNN), Support Vector Machine (SVM) and Random Forest. We have studied these techniques in various configurations and discuss their respective advantages and drawbacks. We revisit the positioning problem VANETs but we also consider Neural Networks (NN) to predict the position [22]. The neural scheme we have tested in this paper consists of one hidden layer with three neurons. To boost this technique we use an ensemble neural network with 50 elements built with a bagging algorithm. The numerical experiments presented in this contribution confirm that a precise prediction can only be obtained when there is a main direct path of propagation. The prediction is altered when the training is incomplete or less precise but the precision remains acceptable. In contrast, with Rayleigh fading, the accuracy obtained is much less striking. We observe that the Neural Network is nearly always the best approach. With a direct path the ranking is: Neural Network, Random Forest, KNN and SVM except in the case when we have no measurement in [30m; 105m] where the ranking is Neural Network, Random Forest, SVM and KNN. When there is no direct path, the ranking is SVM, NN, RF and KNN but the difference in performance between SVM and NN is small.

7.10.4. Combining random access TDMA scheduling strategies for vehicular ad hoc networks

Participants: Fouzi Boukhalifa, Mohamed Hadded (Vedecom), Paul Mühlethaler, Oyunchimeg Shagdar (Vedecom).

This work is based on Fouzi Boukhalifa's PhD which started in October 2018, [29],[15]. The idea is to combine TDMA protocols with random access techniques to benefit from the advantages of both techniques. Fouzi Boukhalifa proposes to combine the DTMAC protocol introduced by Mohamed Hadded with a generalization of CSMA. This generalized CSMA uses active signaling; the idea is to send signaling bursts in order to select a unique transmitter. The protocol that Fouzi Boukhalifa obtains reduces the access and merging collisions of DTMAC but can also propose access with low latency for emergency traffic. The idea is that vehicles access their slots reserved with DTMAC but the transmission slots encompass a special section at the beginning with active signaling. The transmission of the signaling burst, during a mini-slot, is organized according to a random binary key. A '1' in the key means that a signaling burst will be transmitted, while a '0' means that the vehicle senses the channel on this mini-slot to potentially find the transmission of a signaling burst by another vehicle. Fouzi Boukhalifa shows that if we use a random key to transmit the signaling burst it very significantly decreases the collision rate (both merging and access collisions) and that emergency traffic can have a very small access delay. Fouzi Boukhalifa builds an analytical model which thoroughly confirms the simulation result. This model can encompass detection error in the selection process of the signaling bursts. It is shown that with a reasonable error rate the performance is only marginally affected.

7.10.5. Forecasting traffic accidents in VANETs

Participants: Samia Bouzefrane (Cnam), Soumya Banerjee (Birla Institute Of Technology, Mesra), Paul Mühlethaler, Mamoudou Sangare.

Road traffic accidents have become a major cause of death. With increasing urbanization and populations, the volume of vehicles has increased exponentially. As a result, traffic accident forecasting and the identification of the accident prone areas can help reduce the risk of traffic accidents and improve the overall life expectancy.

Conventional traffic forecasting techniques use either a Gaussian Mixture Model (GMM) or a Support Vector Classifier (SVC) to model accident features. A GMM on the one hand requires large amount of data and is computationally inexpensive, SVC on the other hand performs well with less data but is computationally expensive. We present a prediction model that combines the two approaches for the purpose of forecasting traffic accidents. A hybrid approach is proposed, which incorporates the advantages of both the generative (GMM) and the discriminant model (SVC). Raw feature samples are divided into three categories: those representing accidents with no injuries, accidents with non incapacitating injuries and those with incapacitating injuries. The output or the accident severity class was divided into three major categories namely: no injury in the accident, non-incapacitating injury in the accident and an incapacitating injury in the accident. A hybrid classifier is proposed which combines the descriptive strength of the baseline Gaussian mixture model (GMM) with the high performance classification capabilities of the support vector classifier (SVC). A new approach is introduced using the mean vectors obtained from the GMM model as input to the SVC. The model was supported with data pre-processing and re-sampling to convert the data points into suitable form and avoid any kind of biasing in the results. Feature importance ranking was also performed to choose relevant attributes with respect to accident severity. This hybrid model successfully takes advantage of both models and obtained a better accuracy than the baseline GMM model. The radial basis kernel outperforms the linear kernel by achieving an accuracy of 85.53%. Data analytics performed including the area under the receiver operating characteristics curve (AUC-ROC) and area under the precision/recall curve(AUC-PR) indicate the successful application of this model in traffic accident forecasting. Experimental results show that the proposed model can significantly improve the performance of accident prediction. Improvements of up to 24% are reported in the accuracy as compared to the baseline statistical model (GMM). The data about circumstances of personal injury in road accidents, the types of vehicles involved and the consequential casualties were obtained from data.govt.uk.

Although a significant improvement in accuracy has been observed, this study has several limitations. The first concerns the dataset used. This research is based on a road traffic accident dataset from the year of 2017 which contains very few data samples for the no injury and non-incapacitating injury types of accident. The data was unbalanced not just with respect to the output class but also with respect to the sub features of various attributes. Moreover, aggregating the accident severity into just three categories limits the scope of the study and the results obtained. The greater the number of severity classes, the less is the amount of extra training data required to feed in the SVC to avoid overfitting. Thus, datasets with sufficient records corresponding to each class are desirable and must be used for further study.

The second limitation concerns the dependence of the SVC model on parameters and attribute selection. In this study, the performance of SVC relies heavily on the feature selection results and the mean vectors obtained from the GMM. In order to improve the accuracy of the support vector classifier, other approaches like particle swarm optimization (PSO), ant colony optimization, genetic algorithms etc. could be used for effective parameter selection. In addition to this, more kernels like the polynomial kernel and the sigmoid kernel could be tested to improve future model performances.

GANG Project-Team

7. New Results

7.1. Graph and Combinatorial Algorithms

7.1.1. Fast Diameter Computation within Split Graphs

When can we compute the diameter of a graph in quasi linear time? In [22], we address this question for the class of *split graphs*, that we observe to be the hardest instances for deciding whether the diameter is at most two. We stress that although the diameter of a non-complete split graph can only be either 2 or 3, under the Strong Exponential-Time Hypothesis (SETH) we cannot compute the diameter of a split graph in less than quadratic time. Therefore it is worth to study the complexity of diameter computation on *subclasses* of split graphs, in order to better understand the complexity border. Specifically, we consider the split graphs with bounded *clique-interval number* and their complements, with the former being a natural variation of the concept of interval number for split graphs that we introduce in this paper. We first discuss the relations between the clique-interval number and other graph invariants such as the classic interval number of graphs, the treewidth, the *VC-dimension* and the *stabbing number* of a related hypergraph. Then, in part based on these above relations, we almost completely settle the complexity of diameter computation on these subclasses of split graphs:

- For the k -clique-interval split graphs, we can compute their diameter in truly subquadratic time if $k = \mathcal{O}(1)$, and even in quasi linear time if $k = o(\log n)$ and in addition a corresponding ordering is given. However, under SETH this cannot be done in truly subquadratic time for any $k = \omega(\log n)$.
- For the *complements* of k -clique-interval split graphs, we can compute their diameter in truly subquadratic time if $k = \mathcal{O}(1)$, and even in time $\mathcal{O}(km)$ if a corresponding ordering is given. Again this latter result is optimal under SETH up to polylogarithmic factors.

Our findings raise the question whether a k -clique interval ordering can always be computed in quasi linear time. We prove that it is the case for $k = 1$ and for some subclasses such as bounded-treewidth split graphs, threshold graphs and comparability split graphs. Finally, we prove that some important subclasses of split graphs – including the ones mentioned above – have a bounded clique-interval number.

7.1.2. Diameter computation on H -minor free graphs and graphs of bounded (distance)

VC-dimension

Under the Strong Exponential-Time Hypothesis, the diameter of general unweighted graphs cannot be computed in truly subquadratic time. Nevertheless there are several graph classes for which this can be done such as bounded-treewidth graphs, interval graphs and planar graphs, to name a few. We propose to study unweighted graphs of constant *distance VC-dimension* as a broad generalization of many such classes – where the distance VC-dimension of a graph G is defined as the VC-dimension of its ball hypergraph: whose hyperedges are the balls of all possible radii and centers in G . In particular for any fixed H , the class of H -minor free graphs has distance VC-dimension at most $|V(H)| - 1$. In [23], we show the following.

- Our first main result is a Monte Carlo algorithm that on graphs of distance VC-dimension at most d , for any fixed k , either computes the diameter or concludes that it is larger than k in time $\tilde{\mathcal{O}}(k \cdot mn^{1-\varepsilon_d})$, where $\varepsilon_d \in (0; 1)$ only depends on d . We thus obtain a *truly subquadratic-time parameterized* algorithm for computing the diameter on such graphs.
- Then as a byproduct of our approach, we get the first truly subquadratic-time randomized algorithm for *constant* diameter computation on all the *nowhere dense* graph classes. The latter classes include all proper minor-closed graph classes, bounded-degree graphs and graphs of bounded expansion.
- Finally, we show how to remove the dependency on k for *any* graph class that excludes a fixed graph H as a minor. More generally, our techniques apply to any graph with constant distance VC-dimension and *polynomial expansion* (or equivalently having strongly sublinear balanced separators). As a result for all such graphs one obtains a truly subquadratic-time randomized algorithm for computing their diameter.

We note that all our results also hold for *radius* computation. Our approach is based on the work of Chazelle and Welzl who proved the existence of spanning paths with strongly sublinear *stabbing number* for every hypergraph of constant VC-dimension. We show how to compute such paths efficiently by combining known algorithms for the stabbing number problem with a clever use of ε -nets, region decomposition and other partition techniques.

7.1.3. Approximation of eccentricities and distance using δ -hyperbolicity

In [9], we show that the eccentricities of all vertices of a δ -hyperbolic graph $G = (V, E)$ can be computed in linear time with an additive one-sided error of at most $c \cdot \delta$, i.e., after a linear time preprocessing, for every vertex v of G one can compute in $O(1)$ time an estimate $\overline{ecc}_G(v)$ of its eccentricity $ecc_G(v) := \max\{d_G(u, v) : u \in V\}$ such that $ecc_G(v) \leq \overline{ecc}_G(v) \leq ecc_G(v) + c \cdot \delta$ for a small constant c . We prove that every δ -hyperbolic graph G has a shortest path tree T , constructible in linear time, such that for every vertex v of G , $ecc_G(v) \leq ecc_T(v) \leq ecc_G(v) + c \cdot \delta$, where $ecc_T(v) := \max\{d_T(u, v) : u \in V\}$. These results are based on an interesting monotonicity property of the eccentricity function of hyperbolic graphs: the closer a vertex is to the center of G , the smaller its eccentricity is. We also show that the distance matrix of G with an additive one-sided error of at most $c' \cdot \delta$ can be computed in $O(|V|^2 \log^2 |V|)$ time, where $c' < c$ is a small constant. Recent empirical studies show that many real world graphs (including Internet application networks, web networks, collaboration networks, social networks, biological networks, and others) have small hyperbolicity. So, we analyze the performance of our algorithms for approximating eccentricities and distance matrix on a number of real-world networks. Our experimental results show that the obtained estimates are even better than the theoretical bounds.

7.1.4. Graph and Hypergraph Decompositions

In [26], we study modular decomposition of hypergraphs and propose some polynomial algorithms to this aim. We also study several notions of approximation of modular decomposition of graphs, by relaxing the definition of modules introducing a tolerance (ϵ edges can miss) this will be presented at CALDAM 2020, Hyderabad. Both topics can be seen as the search for new models of regularity in discrete structures, as in particular bipartite graphs. In both references our polynomial algorithms have to be improved before being applied on real-world data.

7.2. Distributed Computing

7.2.1. Distributed Interactive Proofs

In a distributed locally-checkable proof, we are interested in checking the legality of a given network configuration with respect to some Boolean predicate. To do so, the network enlists the help of a *prover* — a computationally-unbounded oracle that aims at convincing the network that its state is legal, by providing the nodes with certificates that form a distributed proof of legality. The nodes then verify the proof by examining their certificate, their local neighborhood and the certificates of their neighbors.

In [24], we examine the power of a *randomized* form of locally-checkable proof, called *distributed Merlin-Arthur protocols*, or *dMA* for short. In a *dMA* protocol, the prover assigns each node a short certificate, and the nodes then exchange *random messages* with their neighbors. We show that while there exist problems for which *dMA* protocols are more efficient than protocols that do not use randomness, for several natural problems, including Leader Election, Diameter, Symmetry, and Counting Distinct Elements, *dMA* protocols are no more efficient than standard nondeterministic protocols. This is in contrast with Arthur-Merlin (*dAM*) protocols and Randomized Proof Labeling Schemes (RPLS), which are known to provide improvements in certificate size, at least for some of the aforementioned properties.

The study of interactive proofs in the context of distributed network computing is a novel topic, recently introduced by Kol, Oshman, and Saxena [PODC 2018]. In the spirit of sequential interactive proofs theory, we study in [20] the power of distributed interactive proofs. This is achieved via a series of results establishing trade-offs between various parameters impacting the power of interactive proofs, including the number of interactions, the certificate size, the communication complexity, and the form of randomness used. Our results

also connect distributed interactive proofs with the established field of distributed verification. In general, our results contribute to providing structure to the landscape of distributed interactive proofs.

7.2.2. Topological Approach of Network Computing

More than two decades ago, combinatorial topology was shown to be useful for analyzing distributed fault-tolerant algorithms in shared memory systems and in message passing systems. In [18], we show that combinatorial topology can also be useful for analyzing distributed algorithms in networks of arbitrary structure. To illustrate this, we analyze consensus, set-agreement, and approximate agreement in networks, and derive lower bounds for these problems under classical computational settings, such as the LOCAL model and dynamic networks.

In [19], we study the number of rounds needed to solve consensus in a synchronous network G where at most t nodes may fail by crashing. This problem has been thoroughly studied when G is a complete graph, but very little is known when G is arbitrary. We define a notion of radius that considers all ways in which t nodes may crash, and present an algorithm that solves consensus in radius rounds. Then we derive a lower bound showing that our algorithm is optimal for vertex-transitive graphs, among oblivious algorithms.

7.2.3. Making Local Algorithms Wait-Free

When considering distributed computing, reliable message-passing synchronous systems on the one side, and asynchronous failure-prone shared-memory systems on the other side, remain two quite independently studied ends of the reliability/asynchrony spectrum. The concept of locality of a computation is central to the first one, while the concept of wait-freedom is central to the second one. In [8], we propose a new DECOUPLED model in an attempt to reconcile these two worlds. It consists of a synchronous and reliable communication graph of nodes, and on top a set of asynchronous crash-prone processes, each attached to a communication node. To illustrate the DECOUPLED model, the paper presents an asynchronous 3-coloring algorithm for the processes of a ring. From the processes point of view, the algorithm is wait-free. From a locality point of view, each process uses information only from processes at distance $O(\log^* n)$ from it. This local wait-free algorithm is based on an extension of the classical Cole and Vishkin's vertex coloring algorithm in which the processes are not required to start simultaneously.

In [31], we show that, for any task T associated to a locally checkable labeling (lcl), if T is solvable in t rounds by a deterministic algorithm in the local model, then T remains solvable by a deterministic algorithm in $O(t)$ rounds in an asynchronous variant of the local model whenever $t = O(\text{polylog} n)$.

7.2.4. Towards Synthesis of Distributed Algorithms with SMT Solvers

In [32], we consider the problem of synthesizing distributed algorithms working on a specific execution context. We show it is possible to use the linear time temporal logic in order to both specify the correctness of algorithms and their execution contexts. We then provide a method allowing to reduce the synthesis problem of finite state algorithms to some model-checking problems. We finally apply our technique to automatically generate algorithms for consensus and epsilon-agreement in the case of two processes using the SMT solver Z3.

7.2.5. On Weakest Failure Detector

Failure detectors are devices (objects) that provide the processes with information on failures. They were introduced to enrich asynchronous systems so that it becomes possible to solve problems (or implement concurrent objects) that are otherwise impossible to solve in pure asynchronous systems where processes are prone to crash failures. The most famous failure detector (which is called "eventual leader" and denoted Ω) is the weakest failure detector which allows consensus to be solved in n -process asynchronous systems where up to $t = n - 1$ processes may crash in the read/write communication model, and up to $t < n/2$ processes may crash in the message-passing communication model.

When looking at the mutual exclusion problem (or equivalently the construction of a lock object), while the weakest failure detectors are known for both asynchronous message-passing systems and read/write systems in which up to $t < n$ processes may crash, for the starvation-freedom progress condition, it is not yet known for weaker deadlock-freedom progress condition in read/write systems. In [34], we extend the previous results, namely, it presents the weakest failure detector that allows mutual exclusion to be solved in asynchronous n -process read/write systems where any number of processes may crash, whatever the progress condition (deadlock-freedom or starvation-freedom).

In read/read/write communication model, and in the message-passing communication model, all correct processes are supposed to participate in a consensus instance and in particular the eventual leader.

In [33], we considers the case where some subset of processes that do not crash (not predefined in advance) are allowed not to participate in a consensus instance. In this context Ω cannot be used to solve consensus as it could elect as eventual leader a non-participating process. This paper presents the weakest failure detector that allows correct processes not to participate in a consensus instance. This failure detector, denoted Ω^* , is a variant of Ω . The paper presents also an Ω^* -based consensus algorithm for the asynchronous read/write model, in which any number of processes may crash, and not all the correct processes are required to participate.

7.2.6. Multi-Round Cooperative Search Games with Multiple Players

We study search in the context of competing agents. The setting we consider combines game-theoretic concepts with notions related to parallel computing. Assume that a treasure is placed in one of M boxes according to a known distribution and that k searchers are searching for it in parallel during T rounds. In [27], we study the question of how to incentivize selfish players so that group performance would be maximized. Here, this is measured by the *success probability*, namely, the probability that at least one player finds the treasure. We focus on *congestion policies* $C(\ell)$ that specify the reward that a player receives if it is one of ℓ players that (simultaneously) find the treasure for the first time. Our main technical contribution is proving that the *exclusive policy*, in which $C(1) = 1$ and $C(\ell) = 0$ for $\ell > 1$, yields a *price of anarchy* of $(1 - (1 - 1/k)^k)^{-1}$, and that this is the best possible price among all symmetric reward mechanisms. For this policy we also have an explicit description of a symmetric equilibrium, which is in some sense unique, and moreover enjoys the best success probability among all symmetric profiles. For general congestion policies, we show how to polynomially find, for any $\theta > 0$, a symmetric multiplicative $(1 + \theta)(1 + C(k))$ -equilibrium.

Together with an appropriate reward policy, a central entity can suggest players to play a particular profile at equilibrium. As our main conceptual contribution, we advocate the use of symmetric equilibria for such purposes. Besides being fair, we argue that symmetric equilibria can also become highly robust to crashes of players. Indeed, in many cases, despite the fact that some small fraction of players crash (or refuse to participate), symmetric equilibria remain efficient in terms of their group performances and, at the same time, serve as approximate equilibria. We show that this principle holds for a class of games, which we call *monotonously scalable* games. This applies in particular to our search game, assuming the natural *sharing policy*, in which $C(\ell) = 1/\ell$. For the exclusive policy, this general result does not hold, but we show that the symmetric equilibrium is nevertheless robust under mild assumptions.

7.3. Models and Algorithms for Networks

7.3.1. Exploiting Hopsets: Improved Distance Oracles for Graphs of Constant Highway Dimension and Beyond

For fixed $h \geq 2$, we consider in [25] the task of adding to a graph G a set of weighted shortcut edges on the same vertex set, such that the length of a shortest h -hop path between any pair of vertices in the augmented graph is exactly the same as the original distance between these vertices in G . A set of shortcut edges with this property is called an *exact h -hopset* and may be applied in processing distance queries on graph G . In particular, a 2-hopset directly corresponds to a distributed distance oracle known as a *hub labeling*. In this work, we explore centralized distance oracles based on 3-hopsets and display their advantages in several practical scenarios. In particular, for graphs of constant highway dimension, and more generally for graphs

of constant skeleton dimension, we show that 3-hopsets require *exponentially* fewer shortcuts per node than any previously described distance oracle, and also offer a speedup in query time when compared to simple oracles based on a direct application of 2-hopsets. Finally, we consider the problem of computing minimum-size h -hopset (for any $h \geq 2$) for a given graph G , showing a polylogarithmic-factor approximation for the case of unique shortest path graphs. When $h = 3$, for a given bound on the space used by the distance oracle, we provide a construction of hopset achieving polylog approximation both for space and query time compared to the optimal 3-hopset oracle given the space bound.

7.3.2. Hardness of exact distance queries in sparse graphs through hub labeling

A *distance labeling scheme* is an assignment of bit-labels to the vertices of an undirected, unweighted graph such that the distance between any pair of vertices can be decoded solely from their labels. An important class of distance labeling schemes is that of *hub labelings*, where a node $v \in G$ stores its distance to the so-called hubs $S_v \subseteq V$, chosen so that for any $u, v \in V$ there is $w \in S_u \cap S_v$ belonging to some shortest uv path. Notice that for most existing graph classes, the best distance labelling constructions existing use at some point a hub labeling scheme at least as a key building block.

In [28], our interest lies in hub labelings of sparse graphs, i.e., those with $|E(G)| = O(n)$, for which we show a lowerbound of $\frac{n}{2^{O(\sqrt{\log n})}}$ for the average size of the hubsets. Additionally, we show a hub-labeling construction for sparse graphs of average size $O(\frac{n}{RS(n)^c})$ for some $0 < c < 1$, where $RS(n)$ is the so-called Ruzsa-Szemerédi function, linked to structure of induced matchings in dense graphs. This implies that further improving the lower bound on hub labeling size to $\frac{n}{2^{(\log n)^{o(1)}}$ would require a breakthrough in the study of lower bounds on $RS(n)$, which have resisted substantial improvement in the last 70 years.

For general distance labeling of sparse graphs, we show a lowerbound of $\frac{1}{2^{\Theta(\sqrt{\log n})}} \text{SumIndex}(n)$, where $\text{SumIndex}(n)$ is the communication complexity of the Sum-Index problem over Z_n . Our results suggest that the best achievable hub-label size and distance-label size in sparse graphs may be $\Theta(\frac{n}{2^{(\log n)^c}})$ for some $0 < c < 1$.

7.3.3. Fast Public Transit Routing with Unrestricted Walking through Hub Labeling

In [30], we propose a novel technique for answering routing queries in public transportation networks that allows unrestricted walking. We consider several types of queries: earliest arrival time, Pareto-optimal journeys regarding arrival time, number of transfers and walking time, and profile, i.e. finding all Pareto-optimal journeys regarding travel time and arrival time in a given time interval. Our techniques uses hub labeling to represent unlimited foot transfers and can be adapted to both classical algorithms RAPTOR and CSA. We obtain significant speedup compared to the state-of-the-art approach based on contraction hierarchies. A research report version is deposited on HAL with number hal-02161283.

7.3.4. Independent Lazy Better-Response Dynamics on Network Games

In [29], we study an *independent* best-response dynamics on network games in which the nodes (players) decide to revise their strategies independently with some probability. We are interested in the *convergence time* to the equilibrium as a function of this probability, the degree of the network, and the potential of the underlying games.

7.3.5. A Comparative Study of Neural Network Compression

There has recently been an increasing desire to evaluate neural networks locally on computationally-limited devices in order to exploit their recent effectiveness for several applications; such effectiveness has nevertheless come together with a considerable increase in the size of modern neural networks, which constitute a major downside in several of the aforementioned computationally-limited settings. There has thus been a demand of compression techniques for neural networks. Several proposal in this direction have been made, which famously include hashing-based methods and pruning-based ones. However, the evaluation of the efficacy of these techniques has so far been heterogeneous, with no clear evidence in favor of any of them over the others. In [36], we address this latter issue by providing a comparative study. While most previous studies test the capability of a technique in reducing the number of parameters of state-of-the-art networks, we follow [CWT

+ 15] in evaluating their performance on basic architectures on the MNIST dataset and variants of it, which allows for a clearer analysis of some aspects of their behavior. To the best of our knowledge, we are the first to directly compare famous approaches such as HashedNet, Optimal Brain Damage (OBD), and magnitude-based pruning with L1 and L2 regularization among them and against equivalent-size feed-forward neural networks with simple (fully-connected) and structural (convolutional) neural networks. Rather surprisingly, our experiments show that (iterative) pruning-based methods are substantially better than the HashedNet architecture, whose compression doesn't appear advantageous to a carefully chosen convolutional network. We also show that, when the compression level is high, the famous OBD pruning heuristics deteriorates to the point of being less efficient than simple magnitude-based techniques.

GANG Project-Team

7. New Results

7.1. Graph and Combinatorial Algorithms

7.1.1. Fast Diameter Computation within Split Graphs

When can we compute the diameter of a graph in quasi linear time? In [22], we address this question for the class of *split graphs*, that we observe to be the hardest instances for deciding whether the diameter is at most two. We stress that although the diameter of a non-complete split graph can only be either 2 or 3, under the Strong Exponential-Time Hypothesis (SETH) we cannot compute the diameter of a split graph in less than quadratic time. Therefore it is worth to study the complexity of diameter computation on *subclasses* of split graphs, in order to better understand the complexity border. Specifically, we consider the split graphs with bounded *clique-interval number* and their complements, with the former being a natural variation of the concept of interval number for split graphs that we introduce in this paper. We first discuss the relations between the clique-interval number and other graph invariants such as the classic interval number of graphs, the treewidth, the *VC-dimension* and the *stabbing number* of a related hypergraph. Then, in part based on these above relations, we almost completely settle the complexity of diameter computation on these subclasses of split graphs:

- For the k -clique-interval split graphs, we can compute their diameter in truly subquadratic time if $k = \mathcal{O}(1)$, and even in quasi linear time if $k = o(\log n)$ and in addition a corresponding ordering is given. However, under SETH this cannot be done in truly subquadratic time for any $k = \omega(\log n)$.
- For the *complements* of k -clique-interval split graphs, we can compute their diameter in truly subquadratic time if $k = \mathcal{O}(1)$, and even in time $\mathcal{O}(km)$ if a corresponding ordering is given. Again this latter result is optimal under SETH up to polylogarithmic factors.

Our findings raise the question whether a k -clique interval ordering can always be computed in quasi linear time. We prove that it is the case for $k = 1$ and for some subclasses such as bounded-treewidth split graphs, threshold graphs and comparability split graphs. Finally, we prove that some important subclasses of split graphs – including the ones mentioned above – have a bounded clique-interval number.

7.1.2. Diameter computation on H -minor free graphs and graphs of bounded (distance)

VC-dimension

Under the Strong Exponential-Time Hypothesis, the diameter of general unweighted graphs cannot be computed in truly subquadratic time. Nevertheless there are several graph classes for which this can be done such as bounded-treewidth graphs, interval graphs and planar graphs, to name a few. We propose to study unweighted graphs of constant *distance VC-dimension* as a broad generalization of many such classes – where the distance VC-dimension of a graph G is defined as the VC-dimension of its ball hypergraph: whose hyperedges are the balls of all possible radii and centers in G . In particular for any fixed H , the class of H -minor free graphs has distance VC-dimension at most $|V(H)| - 1$. In [23], we show the following.

- Our first main result is a Monte Carlo algorithm that on graphs of distance VC-dimension at most d , for any fixed k , either computes the diameter or concludes that it is larger than k in time $\tilde{\mathcal{O}}(k \cdot mn^{1-\varepsilon_d})$, where $\varepsilon_d \in (0; 1)$ only depends on d . We thus obtain a *truly subquadratic-time parameterized* algorithm for computing the diameter on such graphs.
- Then as a byproduct of our approach, we get the first truly subquadratic-time randomized algorithm for *constant* diameter computation on all the *nowhere dense* graph classes. The latter classes include all proper minor-closed graph classes, bounded-degree graphs and graphs of bounded expansion.
- Finally, we show how to remove the dependency on k for *any* graph class that excludes a fixed graph H as a minor. More generally, our techniques apply to any graph with constant distance VC-dimension and *polynomial expansion* (or equivalently having strongly sublinear balanced separators). As a result for all such graphs one obtains a truly subquadratic-time randomized algorithm for computing their diameter.

We note that all our results also hold for *radius* computation. Our approach is based on the work of Chazelle and Welzl who proved the existence of spanning paths with strongly sublinear *stabbing number* for every hypergraph of constant VC-dimension. We show how to compute such paths efficiently by combining known algorithms for the stabbing number problem with a clever use of ε -nets, region decomposition and other partition techniques.

7.1.3. Approximation of eccentricities and distance using δ -hyperbolicity

In [9], we show that the eccentricities of all vertices of a δ -hyperbolic graph $G = (V, E)$ can be computed in linear time with an additive one-sided error of at most $c \cdot \delta$, i.e., after a linear time preprocessing, for every vertex v of G one can compute in $O(1)$ time an estimate $\overline{ecc}_G(v)$ of its eccentricity $ecc_G(v) := \max\{d_G(u, v) : u \in V\}$ such that $ecc_G(v) \leq \overline{ecc}_G(v) \leq ecc_G(v) + c \cdot \delta$ for a small constant c . We prove that every δ -hyperbolic graph G has a shortest path tree T , constructible in linear time, such that for every vertex v of G , $ecc_G(v) \leq ecc_T(v) \leq ecc_G(v) + c \cdot \delta$, where $ecc_T(v) := \max\{d_T(u, v) : u \in V\}$. These results are based on an interesting monotonicity property of the eccentricity function of hyperbolic graphs: the closer a vertex is to the center of G , the smaller its eccentricity is. We also show that the distance matrix of G with an additive one-sided error of at most $c' \cdot \delta$ can be computed in $O(|V|^2 \log^2 |V|)$ time, where $c' < c$ is a small constant. Recent empirical studies show that many real world graphs (including Internet application networks, web networks, collaboration networks, social networks, biological networks, and others) have small hyperbolicity. So, we analyze the performance of our algorithms for approximating eccentricities and distance matrix on a number of real-world networks. Our experimental results show that the obtained estimates are even better than the theoretical bounds.

7.1.4. Graph and Hypergraph Decompositions

In [26], we study modular decomposition of hypergraphs and propose some polynomial algorithms to this aim. We also study several notions of approximation of modular decomposition of graphs, by relaxing the definition of modules introducing a tolerance (ϵ edges can miss) this will be presented at CALDAM 2020, Hyderabad. Both topics can be seen as the search for new models of regularity in discrete structures, as in particular bipartite graphs. In both references our polynomial algorithms have to be improved before being applied on real-world data.

7.2. Distributed Computing

7.2.1. Distributed Interactive Proofs

In a distributed locally-checkable proof, we are interested in checking the legality of a given network configuration with respect to some Boolean predicate. To do so, the network enlists the help of a *prover* — a computationally-unbounded oracle that aims at convincing the network that its state is legal, by providing the nodes with certificates that form a distributed proof of legality. The nodes then verify the proof by examining their certificate, their local neighborhood and the certificates of their neighbors.

In [24], we examine the power of a *randomized* form of locally-checkable proof, called *distributed Merlin-Arthur protocols*, or *dMA* for short. In a *dMA* protocol, the prover assigns each node a short certificate, and the nodes then exchange *random messages* with their neighbors. We show that while there exist problems for which *dMA* protocols are more efficient than protocols that do not use randomness, for several natural problems, including Leader Election, Diameter, Symmetry, and Counting Distinct Elements, *dMA* protocols are no more efficient than standard nondeterministic protocols. This is in contrast with Arthur-Merlin (*dAM*) protocols and Randomized Proof Labeling Schemes (RPLS), which are known to provide improvements in certificate size, at least for some of the aforementioned properties.

The study of interactive proofs in the context of distributed network computing is a novel topic, recently introduced by Kol, Oshman, and Saxena [PODC 2018]. In the spirit of sequential interactive proofs theory, we study in [20] the power of distributed interactive proofs. This is achieved via a series of results establishing trade-offs between various parameters impacting the power of interactive proofs, including the number of interactions, the certificate size, the communication complexity, and the form of randomness used. Our results

also connect distributed interactive proofs with the established field of distributed verification. In general, our results contribute to providing structure to the landscape of distributed interactive proofs.

7.2.2. Topological Approach of Network Computing

More than two decades ago, combinatorial topology was shown to be useful for analyzing distributed fault-tolerant algorithms in shared memory systems and in message passing systems. In [18], we show that combinatorial topology can also be useful for analyzing distributed algorithms in networks of arbitrary structure. To illustrate this, we analyze consensus, set-agreement, and approximate agreement in networks, and derive lower bounds for these problems under classical computational settings, such as the LOCAL model and dynamic networks.

In [19], we study the number of rounds needed to solve consensus in a synchronous network G where at most t nodes may fail by crashing. This problem has been thoroughly studied when G is a complete graph, but very little is known when G is arbitrary. We define a notion of radius that considers all ways in which t nodes may crash, and present an algorithm that solves consensus in radius rounds. Then we derive a lower bound showing that our algorithm is optimal for vertex-transitive graphs, among oblivious algorithms.

7.2.3. Making Local Algorithms Wait-Free

When considering distributed computing, reliable message-passing synchronous systems on the one side, and asynchronous failure-prone shared-memory systems on the other side, remain two quite independently studied ends of the reliability/asynchrony spectrum. The concept of locality of a computation is central to the first one, while the concept of wait-freedom is central to the second one. In [8], we propose a new DECOUPLED model in an attempt to reconcile these two worlds. It consists of a synchronous and reliable communication graph of nodes, and on top a set of asynchronous crash-prone processes, each attached to a communication node. To illustrate the DECOUPLED model, the paper presents an asynchronous 3-coloring algorithm for the processes of a ring. From the processes point of view, the algorithm is wait-free. From a locality point of view, each process uses information only from processes at distance $O(\log^* n)$ from it. This local wait-free algorithm is based on an extension of the classical Cole and Vishkin's vertex coloring algorithm in which the processes are not required to start simultaneously.

In [31], we show that, for any task T associated to a locally checkable labeling (lcl), if T is solvable in t rounds by a deterministic algorithm in the local model, then T remains solvable by a deterministic algorithm in $O(t)$ rounds in an asynchronous variant of the local model whenever $t = O(\text{polylog} n)$.

7.2.4. Towards Synthesis of Distributed Algorithms with SMT Solvers

In [32], we consider the problem of synthesizing distributed algorithms working on a specific execution context. We show it is possible to use the linear time temporal logic in order to both specify the correctness of algorithms and their execution contexts. We then provide a method allowing to reduce the synthesis problem of finite state algorithms to some model-checking problems. We finally apply our technique to automatically generate algorithms for consensus and epsilon-agreement in the case of two processes using the SMT solver Z3.

7.2.5. On Weakest Failure Detector

Failure detectors are devices (objects) that provide the processes with information on failures. They were introduced to enrich asynchronous systems so that it becomes possible to solve problems (or implement concurrent objects) that are otherwise impossible to solve in pure asynchronous systems where processes are prone to crash failures. The most famous failure detector (which is called "eventual leader" and denoted Ω) is the weakest failure detector which allows consensus to be solved in n -process asynchronous systems where up to $t = n - 1$ processes may crash in the read/write communication model, and up to $t < n/2$ processes may crash in the message-passing communication model.

When looking at the mutual exclusion problem (or equivalently the construction of a lock object), while the weakest failure detectors are known for both asynchronous message-passing systems and read/write systems in which up to $t < n$ processes may crash, for the starvation-freedom progress condition, it is not yet known for weaker deadlock-freedom progress condition in read/write systems. In [34], we extend the previous results, namely, it presents the weakest failure detector that allows mutual exclusion to be solved in asynchronous n -process read/write systems where any number of processes may crash, whatever the progress condition (deadlock-freedom or starvation-freedom).

In read/read/write communication model, and in the message-passing communication model, all correct processes are supposed to participate in a consensus instance and in particular the eventual leader.

In [33], we considers the case where some subset of processes that do not crash (not predefined in advance) are allowed not to participate in a consensus instance. In this context Ω cannot be used to solve consensus as it could elect as eventual leader a non-participating process. This paper presents the weakest failure detector that allows correct processes not to participate in a consensus instance. This failure detector, denoted Ω^* , is a variant of Ω . The paper presents also an Ω^* -based consensus algorithm for the asynchronous read/write model, in which any number of processes may crash, and not all the correct processes are required to participate.

7.2.6. Multi-Round Cooperative Search Games with Multiple Players

We study search in the context of competing agents. The setting we consider combines game-theoretic concepts with notions related to parallel computing. Assume that a treasure is placed in one of M boxes according to a known distribution and that k searchers are searching for it in parallel during T rounds. In [27], we study the question of how to incentivize selfish players so that group performance would be maximized. Here, this is measured by the *success probability*, namely, the probability that at least one player finds the treasure. We focus on *congestion policies* $C(\ell)$ that specify the reward that a player receives if it is one of ℓ players that (simultaneously) find the treasure for the first time. Our main technical contribution is proving that the *exclusive policy*, in which $C(1) = 1$ and $C(\ell) = 0$ for $\ell > 1$, yields a *price of anarchy* of $(1 - (1 - 1/k)^k)^{-1}$, and that this is the best possible price among all symmetric reward mechanisms. For this policy we also have an explicit description of a symmetric equilibrium, which is in some sense unique, and moreover enjoys the best success probability among all symmetric profiles. For general congestion policies, we show how to polynomially find, for any $\theta > 0$, a symmetric multiplicative $(1 + \theta)(1 + C(k))$ -equilibrium.

Together with an appropriate reward policy, a central entity can suggest players to play a particular profile at equilibrium. As our main conceptual contribution, we advocate the use of symmetric equilibria for such purposes. Besides being fair, we argue that symmetric equilibria can also become highly robust to crashes of players. Indeed, in many cases, despite the fact that some small fraction of players crash (or refuse to participate), symmetric equilibria remain efficient in terms of their group performances and, at the same time, serve as approximate equilibria. We show that this principle holds for a class of games, which we call *monotonously scalable* games. This applies in particular to our search game, assuming the natural *sharing policy*, in which $C(\ell) = 1/\ell$. For the exclusive policy, this general result does not hold, but we show that the symmetric equilibrium is nevertheless robust under mild assumptions.

7.3. Models and Algorithms for Networks

7.3.1. Exploiting Hopsets: Improved Distance Oracles for Graphs of Constant Highway Dimension and Beyond

For fixed $h \geq 2$, we consider in [25] the task of adding to a graph G a set of weighted shortcut edges on the same vertex set, such that the length of a shortest h -hop path between any pair of vertices in the augmented graph is exactly the same as the original distance between these vertices in G . A set of shortcut edges with this property is called an *exact h -hopset* and may be applied in processing distance queries on graph G . In particular, a 2-hopset directly corresponds to a distributed distance oracle known as a *hub labeling*. In this work, we explore centralized distance oracles based on 3-hopsets and display their advantages in several practical scenarios. In particular, for graphs of constant highway dimension, and more generally for graphs

of constant skeleton dimension, we show that 3-hopsets require *exponentially* fewer shortcuts per node than any previously described distance oracle, and also offer a speedup in query time when compared to simple oracles based on a direct application of 2-hopsets. Finally, we consider the problem of computing minimum-size h -hopset (for any $h \geq 2$) for a given graph G , showing a polylogarithmic-factor approximation for the case of unique shortest path graphs. When $h = 3$, for a given bound on the space used by the distance oracle, we provide a construction of hopset achieving polylog approximation both for space and query time compared to the optimal 3-hopset oracle given the space bound.

7.3.2. Hardness of exact distance queries in sparse graphs through hub labeling

A *distance labeling scheme* is an assignment of bit-labels to the vertices of an undirected, unweighted graph such that the distance between any pair of vertices can be decoded solely from their labels. An important class of distance labeling schemes is that of *hub labelings*, where a node $v \in G$ stores its distance to the so-called hubs $S_v \subseteq V$, chosen so that for any $u, v \in V$ there is $w \in S_u \cap S_v$ belonging to some shortest uv path. Notice that for most existing graph classes, the best distance labelling constructions existing use at some point a hub labeling scheme at least as a key building block.

In [28], our interest lies in hub labelings of sparse graphs, i.e., those with $|E(G)| = O(n)$, for which we show a lowerbound of $\frac{n}{2^{O(\sqrt{\log n})}}$ for the average size of the hubsets. Additionally, we show a hub-labeling construction for sparse graphs of average size $O(\frac{n}{RS(n)^c})$ for some $0 < c < 1$, where $RS(n)$ is the so-called Ruzsa-Szemerédi function, linked to structure of induced matchings in dense graphs. This implies that further improving the lower bound on hub labeling size to $\frac{n}{2^{(\log n)^{o(1)}}$ would require a breakthrough in the study of lower bounds on $RS(n)$, which have resisted substantial improvement in the last 70 years.

For general distance labeling of sparse graphs, we show a lowerbound of $\frac{1}{2^{\Theta(\sqrt{\log n})}} \text{SumIndex}(n)$, where $\text{SumIndex}(n)$ is the communication complexity of the Sum-Index problem over Z_n . Our results suggest that the best achievable hub-label size and distance-label size in sparse graphs may be $\Theta(\frac{n}{2^{(\log n)^c}})$ for some $0 < c < 1$.

7.3.3. Fast Public Transit Routing with Unrestricted Walking through Hub Labeling

In [30], we propose a novel technique for answering routing queries in public transportation networks that allows unrestricted walking. We consider several types of queries: earliest arrival time, Pareto-optimal journeys regarding arrival time, number of transfers and walking time, and profile, i.e. finding all Pareto-optimal journeys regarding travel time and arrival time in a given time interval. Our techniques uses hub labeling to represent unlimited foot transfers and can be adapted to both classical algorithms RAPTOR and CSA. We obtain significant speedup compared to the state-of-the-art approach based on contraction hierarchies. A research report version is deposited on HAL with number hal-02161283.

7.3.4. Independent Lazy Better-Response Dynamics on Network Games

In [29], we study an *independent* best-response dynamics on network games in which the nodes (players) decide to revise their strategies independently with some probability. We are interested in the *convergence time* to the equilibrium as a function of this probability, the degree of the network, and the potential of the underlying games.

7.3.5. A Comparative Study of Neural Network Compression

There has recently been an increasing desire to evaluate neural networks locally on computationally-limited devices in order to exploit their recent effectiveness for several applications; such effectiveness has nevertheless come together with a considerable increase in the size of modern neural networks, which constitute a major downside in several of the aforementioned computationally-limited settings. There has thus been a demand of compression techniques for neural networks. Several proposal in this direction have been made, which famously include hashing-based methods and pruning-based ones. However, the evaluation of the efficacy of these techniques has so far been heterogeneous, with no clear evidence in favor of any of them over the others. In [36], we address this latter issue by providing a comparative study. While most previous studies test the capability of a technique in reducing the number of parameters of state-of-the-art networks, we follow [CWT

+ 15] in evaluating their performance on basic architectures on the MNIST dataset and variants of it, which allows for a clearer analysis of some aspects of their behavior. To the best of our knowledge, we are the first to directly compare famous approaches such as HashedNet, Optimal Brain Damage (OBD), and magnitude-based pruning with L1 and L2 regularization among them and against equivalent-size feed-forward neural networks with simple (fully-connected) and structural (convolutional) neural networks. Rather surprisingly, our experiments show that (iterative) pruning-based methods are substantially better than the HashedNet architecture, whose compression doesn't appear advantageous to a carefully chosen convolutional network. We also show that, when the compression level is high, the famous OBD pruning heuristics deteriorates to the point of being less efficient than simple magnitude-based techniques.

MIMOVE Project-Team

7. New Results

7.1. Automated Synthesis of Mediators for Middleware-Layer Protocol Interoperability in the IoT

Participants: Georgios Bouloukakis, Nikolaos Georgantas, Patient Ntumba, Valérie Issarny (MiMove)

To enable direct Internet connectivity of Things, complete protocol stacks need to be deployed on resource-constrained devices. Such protocol stacks typically build on lightweight IPv6 adaptations and may even include a middleware layer supporting high-level application development. However, the profusion of IoT middleware-layer interaction protocols has introduced technology diversity and high fragmentation in the IoT systems landscape with siloed vertical solutions. To enable the interconnection of heterogeneous Things across these barriers, advanced interoperability solutions at the middleware layer are required. In this paper, we introduce a solution for the automated synthesis of protocol mediators that support the interconnection of heterogeneous Things. Our systematic approach relies on the Data eXchange (DeX) connector model, which comprehensively abstracts and represents existing and potentially future IoT middleware protocols. Thanks to DeX, Things seamlessly interconnect through lightweight mediators. We validate our solution with respect to: (i) the support to developers when developing heterogeneous IoT applications; (ii) the runtime performance of the synthesized mediators.

7.2. Probabilistic Event Dropping for Intermittently Connected Subscribers over Pub/Sub Systems

Participants: Georgios Bouloukakis, Nikolaos Georgantas (MiMove), Ioannis Moscholios (Univ of Peloponnese)

Internet of Things (IoT) aim to leverage data from multiple sensors, actuators and devices for improving peoples' daily life and safety. Multiple data sources must be integrated, analyzed from the corresponding application and notify interested stakeholders. To support the data exchange between data sources and stakeholders, the publish/subscribe (pub/sub) middleware is often employed. Pub/sub provides additional mechanisms such as reliable messaging, event dropping, prioritization, etc. The event dropping mechanism is often used to satisfy Quality of Service (QoS) requirements and ensure system stability. To enable event dropping, basic approaches apply finite buffers or data validity periods and more sophisticated ones are information-aware. In this paper, we introduce a pub/sub mechanism for probabilistic event dropping by considering the stakeholders' intermittent connectivity and QoS requirements. We model the pub/sub middleware as a network of queues which includes a novel ON/OFF queueing model that enables the definition of join probabilities. We validate our analytical model via simulation and compare our mechanism with existing ones. Experimental results can be used as insights for developing hybrid dropping mechanisms.

7.3. Adaptive Mediation for Data Exchange in IoT Systems

Participants: Georgios Bouloukakis (MiMove & Univ of California, Irvine), Andrew Chio, Sharad Mehrotra, Nalini Venkatasubramanian (Univ of California, Irvine), Cheng-Hsin Hsu (National Tsing Hua Univ)

Messaging and communication is a critical aspect of next generation Internet-of-Things (IoT) systems where interactions among devices, software systems/services and end-users is the expected mode of operation. Given the diverse and changing communication needs of entities, the data exchange interactions may assume different protocols (MQTT, CoAP, HTTP) and interaction paradigms (point to point, multicast, unicast). In this paper, we address the issue of supporting adaptive communications in IoT systems through a mediation-based architecture for data exchange. Here, components called mediators support protocol translation to bridge the heterogeneity gap. Aiming to provide a placement of mediators to nodes, we introduce an integer linear programming solution that takes as input: a set of Edge nodes, IoT devices, and networking semantics. Our proposed solution achieves adaptive placement resulting in timely interactions between IoT devices for larger topologies of IoT spaces.

7.4. Universal Social Network Bus: Toward the Federation of Heterogeneous Online Social Network Services

Participants: Valérie Issarny, Nikolaos Georgantas, Ehsan Ahvar, Bruno Lefèvre, Shohreh Ahvar (MiMove), Rafael Angarita (ISEP Paris)

Online Social Network Services (OSNSs) are changing the fabric of our society, impacting almost every aspect of it. Over the past few decades, an aggressive market rivalry has led to the emergence of multiple competing, “closed” OSNSs. As a result, users are trapped in the walled gardens of their OSNS, encountering restrictions about what they can do with their personal data, the people they can interact with, and the information they get access to. As an alternative to the platform lock-in, “open” OSNSs promote the adoption of open, standardized APIs. However, users still massively adopt closed OSNSs to benefit from the services’ advanced functionalities and/or follow their “friends,” although the users’ virtual social sphere is ultimately limited by the OSNSs they join. Our work aims at overcoming such a limitation by enabling users to meet and interact beyond the boundary of their OSNSs, including reaching out to “friends” of distinct closed OSNSs. We specifically introduce *Universal Social Network Bus* (USNB), which revisits the “service bus” paradigm that enables interoperability across computing systems to address the requirements of “social interoperability.” USNB features synthetic profiles and personae for interaction across the boundaries of closed and open and profile- and non-profile-based OSNSs through a reference social interaction service. We ran a 1-day workshop with a panel of users who experimented with the USNB prototype to assess the potential benefits of social interoperability for social network users. Results show the positive evaluation of users for USNB, especially as an enabler of applications for civic participation. This further opens up new perspectives for future work, among which includes enforcing security and privacy guarantees.

7.5. Social Middleware for Civic Engagement

Participants: Valérie Issarny, Nikolaos Georgantas, Grigoris Piperagkas (MiMove), Rafael Angarita (ISEP Paris)

Civic engagement refers to any collective action towards the identification and solving of public issues. Current civic technologies are traditional Web- or mobile-based platforms that make difficult, or just impossible, the participation of citizens via different communication technologies. Moreover, connected objects sensing physical-world data can nourish participatory processes by providing physical evidence to citizens; however, leveraging these data is not direct and still a time-consuming process for civic technologies developers. We introduce the concept of *social middleware* for civic engagement. Social middleware allows citizens to engage in participatory processes -supported by civic technologies- via their favorite communication tools, and to interact not only with other citizens but also with relevant connected objects and software platforms. The mission of social middleware goes beyond the connection of all these heterogeneous entities. It aims at easing the implementation of distributed applications oriented toward civic engagement by featuring dedicated built-in services.

7.6. Mobile Crowd-Sensing as a Resource for Contextualized Urban Public Policies

Participants: Valérie Issarny, Bruno Lefèvre, Rachit Agarwal (MiMove), Vivien Mallet (Inria Ange)

Environmental noise is a major pollutant in contemporary cities and calls for the active monitoring of noise levels to spot the locations where it most affects the people’s health and well-being. However, due to the complex relationship between environmental noise and its perception by the citizens, it is not sufficient to quantitatively measure environmental noise. We need to collect and aggregate contextualized –both quantitative and qualitative– data about the urban environmental noise so as to be able to study the objective and subjective relationships between sound and living beings. This complex knowledge is a prerequisite for making efficient territorial public policies for soundscapes that are inclined towards living beings welfare. In this paper, we investigate how Mobile Phone Sensing (MPS) –*aka* crowdsensing– enables the gathering of such knowledge, provided the implementation of sensing protocols that are customized according to the context of

use and the intended exploitation of the data. Through three case studies that we carried out in France and Finland, we show that MPS is not solely a tool that contributes to sensitizing citizens and decision-makers about noise pollution; it also contributes to increasing our knowledge about the impact of the environmental noise on people's health and well-being in relation to its physical and subjective perception.

7.7. Multi-Sensor Calibration Planning in IoT-Enabled Smart Spaces

Participants: Valérie Issarny (MiMove), Françoise Sailhan (CNAM), Qiuxi Zhu, Md Yusuf Sarwar Uddin, Nalini Venkatasubramanian (University of California, Irvine)

Emerging applications in smart cities and communities require massive IoT deployments using sensors/actuators (things) that can enhance citizens' quality of life and public safety. However, budget constraints often lead to limited instrumentation and/or the use of low-cost sensors that are subject to drift and bias. This raises concerns of robustness and accuracy of the decisions made on uncertain data. To enable effective decision making while fully exploiting the potential of low-cost sensors, we propose to send mobile units (e.g., trained personnel) equipped with high-quality (more expensive) and freshly-calibrated reference sensors so as to carry out calibration in the field. We design and implement an efficient cooperative approach to solve the calibration planning problem, which aims at minimizing the cost of the recurring calibration of multiple sensor types in the long-term operation. We propose a two-phase solution that consists of a sensor selection phase that minimizes the average cost of recurring calibration, and a path planning phase that minimizes the travel cost of multiple calibrators which have load constraints. We provide fast and effective heuristics for both phases. We further build a prototype that facilitates the mapping of the deployment field and provides navigation guidance to mobile calibrators. Extensive use-case-driven simulations show that our proposed approach significantly reduces the average cost compared to naive approaches: up to 30% in a moderate-sized indoor case, and higher in outdoor cases depending on scale

7.8. User-Centric Context Inference for Mobile Crowdsensing

Participants: Yifan Du, Valérie Issarny (MiMove), Françoise Sailhan (CNAM)

Mobile crowdsensing is a powerful mechanism to aggregate hyperlocal knowledge about the environment. Indeed, users may contribute valuable observations across time and space using the sensors embedded in their smartphones. However, the relevance of the provided measurements depends on the adequacy of the sensing context with respect to the phenomena that are analyzed. Our research concentrates more specifically on assessing the sensing context when gathering observations about the physical environment beyond its geographical position in the Euclidean space, i.e., whether the phone is in-/out-pocket, in-/out-door and on-/under-ground. We introduce an online learning approach to the local inference of the sensing context so as to overcome the disparity of the classification performance due to the heterogeneity of the sensing devices as well as the diversity of user behavior and novel usage scenarios. Our approach specifically features a hierarchical algorithm for inference that requires few opportunistic feedbacks from the user, while increasing the accuracy of the context inference per user.

7.9. Let Opportunistic Crowdsensors Work Together for Resource-efficient, Quality-aware Observations

Participants: Yifan Du, Valérie Issarny (MiMove), Françoise Sailhan (CNAM)

Opportunistic crowdsensing empowers citizens carrying hand-held devices to sense physical phenomena of common interest at a large and fine-grained scale without requiring the citizens' active involvement. However, the resulting uncontrolled collection and upload of the massive amount of contributed raw data incur significant resource consumption, from the end device to the server, as well as challenge the quality of the collected observations. Our research tackles both challenges raised by opportunistic crowdsensing, that is, enabling the resource-efficient gathering of relevant observations. To achieve so, we introduce the *BeTogether* middleware fostering context-aware, collaborative crowdsensing at the edge so that co-located crowdsensors operating in the same context, group together to share the work load in a cost- and quality-effective way. Our

implementation-driven evaluation of the proposed solution, which leverages a dataset embedding nearly one million entries contributed by 550 crowdsensors over a year, shows that *BeTogether* increases the quality of the collected data while reducing the overall resource cost compared to the cloud-centric approach.

7.10. Detecting Mobile Crowdsensing Context in the Wild

Participants: Rachit Agarwal, Shaan Chopra, Vassilis Christophides, Nikolaos Georgantas, Valérie Issarny (MiMove)

Understanding the sensing context of raw data is crucial for assessing the quality of large crowdsourced spatio-temporal datasets and supporting context-augmented personal trajectories. Detecting sensing contexts in the wild is a challenging task and requires features from smartphone sensors that are not always available. In this paper, we propose three heuristic algorithms for detecting sensing contexts such as in/out-pocket, under/over-ground, and in/out-door for crowdsourced spatio-temporal datasets. These are unsupervised binary classifiers with a small memory footprint and execution time. Using a segment of the Ambiciti real dataset-a feature-limited crowdsourced dataset-we report that our algorithms perform equally well in terms of balanced accuracy (within 4.3%) when compared to machine learning (ML) models reported by an AutoML tool.

7.11. Inferring Streaming Video Quality from Encrypted Traffic: Practical Models and Deployment Experience

Participants: Francesco Bronzino, Sara Ayoubi, Renata Teixeira (MiMove), Paul Schmitt (Princeton), Guilherme Martins, Nick Feamster (University of Chicago)

Inferring the quality of streaming video applications is important for Internet service providers, but the fact that most video streams are encrypted makes it difficult to do so. We develop models that infer quality metrics (i.e., startup delay and resolution) for encrypted streaming video services. Our paper builds on previous work, but extends it in several ways. First, the models work in deployment settings where the video sessions and segments must be identified from a mix of traffic and the time precision of the collected traffic statistics is more coarse (e.g., due to aggregation). Second, we develop a single composite model that works for a range of different services (i.e., Netflix, YouTube, Amazon, and Twitch), as opposed to just a single service. Third, unlike many previous models, our models perform predictions at finer granularity (e.g., the precise startup delay instead of just detecting short versus long delays) allowing to draw better conclusions on the ongoing streaming quality. Fourth, we demonstrate the models are practical through a 16-month deployment in 66 homes and provide new insights about the relationships between Internet "speed" and the quality of the corresponding video streams, for a variety of services; we find that higher speeds provide only minimal improvements to startup delay and resolution. This work was accepted for publication at the ACM SIGMETRICS conference. The models we developed in this work and the findings were the basis for a first-page story published on The Wall Street Journal ("The Truth About Faster Internet: It's Not Worth It").⁰

7.12. Implications of User Perceived Page Load Time Multi-Modality on Web QoE Measurement

Participants: Renata Teixeira, Vassilis Christophides (MiMove), Flavia Salutari, Diego Da Hora (Telecom Paris Tech), Matteo Varvello (Brave Software), Dario Rossi (Huawei)

⁰The article is available online at: https://www.wsj.com/graphics/faster-internet-not-worth-it/?mod=article_inline&mod=hp_lead_pos5.

Web browsing is one of the most popular applications for both desktop and mobile users. A lot of effort has been devoted to speedup the Web, as well as in designing metrics that can accurately tell whether a webpage loaded fast or not. An often implicit assumption made by industrial and academic research communities is that a *single* metric is sufficient to assess whether a webpage loaded fast. In this work we collect and make publicly available a unique dataset which contains webpage features (e.g., number and type of embedded objects) along with both *objective* and *subjective* Web quality metrics. This dataset was collected by crawling over 100 websites—representative of the top 1 M websites in the Web—while crowdsourcing 6,000 user opinions on *user perceived page load time* (uPLT). In contrast to related work, we show that the uPLT distribution is often multimodal and that, in practice, no more than three modes are present. The main conclusion drawn from our analysis is that, for complex webpages, each of the different objective QoE metrics proposed in the literature (such as AFT, TTI, PLT, etc.) is suited to approximate one of the different uPLT modes.

7.13. The News We Like Are Not the News We Visit: News Categories Popularity in Usage Data

Participants: Renata Teixeira (MiMove), Giuseppe Scavo (MiMove, Nokia Bell Labs), Zied Ben-Houidi (Nokia Bell-Labs), Stefano Traverso, Marco Mellia (Politecnico di Torino)

Most of our knowledge about online news consumption comes from survey-based news market reports, partial usage data from a single editor, or what people publicly share on social networks. Our work published on the 13th International AAAI Conference on Web and Social Media (ICWSM-2019) complements these sources by presenting the first holistic study of visits across online news outlets that a population uses to read news. We monitored the entire network traffic generated by Internet users in four locations in Italy. Together these users generated 80 million visits to 5.4 million news articles in about one year and a half. This unique view allowed us to evaluate how usage data complements existing data sources. We find for instance that only 16% of news visits in our datasets came from online social networks. In addition, the popularity of news categories when considering all visits is quite different from the one when considering only news discovered on social media, or visits to a single major news outlet. Interestingly, a substantial mismatch emerges between self-reported news-category preferences (as measured by Reuters Institute in the same year and same country) and their actual popularity in terms of visits in our datasets. In particular, unlike self-reported preferences expressed by users in surveys that put “Politics”, “Science” and “International” as the most appreciated categories, “Tragedies and Weird news” and “Sport” are by far the most visited. Our paper discusses two possible causes of this mismatch and conjecture that the most plausible reason is the disassociation that may occur between individuals’ cognitive values and their cue-triggered attraction.

7.14. Classification of Load Balancing in the Internet

Participants: Renata Teixeira (MiMove), Rafael Almeida, Ítalo Cunha (Universidade Federal de Minas Gerais), Darryl Veitch (University of Technology Sydney), Christophe Diot (Google)

Recent advances in programmable data planes, software-defined networking, and the adoption of IPv6, support novel, more complex load balancing strategies. We introduce the Multipath Classification Algorithm (MCA), a probing algorithm that extends traceroute to identify and classify load balancing in Internet routes. MCA extends existing formalism and techniques to consider that load balancers may use arbitrary combinations of bits in the packet header for load balancing. We propose optimizations to reduce probing cost that are applicable to MCA and existing load balancing measurement techniques. Through large-scale measurement campaigns, we characterize and study the evolution of load balancing on the IPv4 and IPv6 Internet with multiple transport protocols. Our results that will appear in the IEEE INFOCOM 2020 conference show that load balancing is more prevalent and that load balancing strategies are more mature than previous characterizations have found.

7.15. MinoanER: Schema-Agnostic, Non-Iterative, Massively Parallel Resolution of Web Entities

Participants: Vassilis Christophides (MiMove), Vasilis Efthymiou (IBM Almaden Research Center), George Papadakis (Univ of Athens), Kostas Stefanidis (Univ of Tampere)

Entity Resolution (ER) aims to identify different descriptions in various Knowledge Bases (KBs) that refer to the same entity. ER is challenged by the Variety, Volume and Veracity of entity descriptions published in the Web of Data. To address them, we propose the MinoanER framework that simultaneously fulfills full automation, support of highly heterogeneous entities, and massive parallelization of the ER process. MinoanER leverages a token-based similarity of entities to define a new metric that derives the similarity of neighboring entities from the most important relations, as they are indicated only by statistics. A composite blocking method is employed to capture different sources of matching evidence from the content, neighbors, or names of entities. The search space of candidate pairs for comparison is compactly abstracted by a novel disjunctive blocking graph and processed by a non-iterative, massively parallel matching algorithm that consists of four generic, schema-agnostic matching rules that are quite robust with respect to their internal configuration. We demonstrate that the effectiveness of MinoanER is comparable to existing ER tools over real KBs exhibiting low Variety, but it outperforms them significantly when matching KBs with high Variety.

WHISPER Project-Team

7. New Results

7.1. Software engineering for infrastructure software

Data races are often hard to detect in device drivers, due to the non-determinism of concurrent execution. With colleagues from Tsinghua University, we have addressed this issue using dynamic analysis. According to our study of Linux driver patches that fix data races, more than 38% of patches involve a pattern that we call inconsistent lock protection. Specifically, if a variable is accessed within two concurrently executed functions, the sets of locks held around each access are disjoint, at least one of the locksets is non-empty, and at least one of the involved accesses is a write, then a data race may occur. In a paper published at SANER 2019 [17], we present a runtime analysis approach, named DILP, to detect data races caused by inconsistent lock protection in device drivers. By monitoring driver execution, DILP collects the information about runtime variable accesses and executed functions. Then after driver execution, DILP analyzes the collected information to detect and report data races caused by inconsistent lock protection. We evaluate DILP on 12 device drivers in Linux 4.16.9, and find 25 real data races.

For waiting, the Linux kernel offers both sleep-able and non-sleep operations. However, only non-sleep operations can be used in atomic context. Detecting the possibility of execution in atomic context requires a complete inter-procedural flow analysis, often involving function pointers. Developers may thus conservatively use non-sleep operations even outside of atomic context, which may damage system performance, as such operations unproductively monopolize the CPU. Until now, no systematic approach has been proposed to detect such conservative non-sleep (CNS) defects. In a paper published at ASPLOS 2019 [14] with colleagues from Tsinghua University, we propose a practical static approach, named DCNS, to automatically detect conservative non-sleep defects in the Linux kernel. DCNS uses a summary-based analysis to effectively identify the code in atomic context and a novel file-connection-based alias analysis to correctly identify the set of functions referenced by a function pointer. We evaluate DCNS on Linux 4.16, and in total find 1629 defects. We manually check 943 defects whose call paths are not so difficult to follow, and find that 890 are real. We have randomly selected 300 of the real defects and sent them to kernel developers, and 251 have been confirmed.

In Linux device drivers, use-after-free (UAF) bugs can cause system crashes and serious security problems. We have addressed this issue in work with colleagues at Tsinghua University. According to our study of Linux kernel commits, 42% of the driver commits fixing use-after-free bugs involve driver concurrency. We refer to these use-after-free bugs as concurrency use-after-free bugs. Due to the non-determinism of concurrent execution, concurrency use-after-free bugs are often more difficult to reproduce and detect than sequential use-after-free bugs. In a paper published at USENIX ATC 2019 [13], we propose a practical static analysis approach named DCUAF, to effectively detect concurrency use-after-free bugs in Linux device drivers. DCUAF combines a local analysis analyzing the source code of each driver with a global analysis statistically analyzing the local results of all drivers, forming a local-global analysis, to extract the pairs of driver interface functions that may be concurrently executed. Then, with these pairs, DCUAF performs a summary-based lockset analysis to detect concurrency use-after-free bugs. We have evaluated DCUAF on the driver code of Linux 4.19, and found 640 real concurrency use-after-free bugs. We have randomly selected 130 of the real bugs and reported them to Linux kernel developers, and 95 have been confirmed.

Linux kernel stable versions serve the needs of users who value stability of the kernel over new features. The quality of such stable versions depends on the initiative of kernel developers and maintainers to propagate bug fixing patches to the stable versions. Thus, it is desirable to consider to what extent this process can be automated. A previous approach relies on words from commit messages and a small set of manually constructed code features. This approach, however, shows only moderate accuracy. In a tool paper published ICSE 2019 [11], in the context of the ANR-NRF ITrans project with colleagues from Singapore Management

University, paper, we investigate whether deep learning can provide a more accurate solution. We propose PatchNet, a hierarchical deep learning-based approach capable of automatically extracting features from commit messages and commit code and using them to identify stable patches. PatchNet contains a deep hierarchical structure that mirrors the hierarchical and sequential structure of commit code, making it distinctive from the existing deep learning models on source code. Experiments on 82,403 recent Linux patches confirm the superiority of PatchNet against various state-of-the-art baselines, including the one recently-adopted by Linux kernel maintainers.

Developing software often requires code changes that are widespread and applied to multiple locations. Previously, the Whisper team has addressed this problem with the tool Coccinelle. In a recent experience paper, published at ECOOP 2019 [21], in the context of the ANR-NRF ITrans project with colleagues from Singapore Management University, we have considered the benefits of extending Coccinelle to Java code. There are tools for Java that allow developers to specify patterns for program matching and source-to-source transformation. However, to our knowledge, none allows for transforming code based on its control-flow context. We prototype Coccinelle4J, an extension to Coccinelle, which is a program transformation tool designed for widespread changes in C code, in order to work on Java source code. We adapt Coccinelle to be able to apply scripts written in the Semantic Patch Language (SmPL), a language provided by Coccinelle, to Java source files. As a case study, we demonstrate the utility of Coccinelle4J with the task of API migration. We show 6 semantic patches to migrate from deprecated Android API methods on several open source Android projects. We describe how SmPL can be used to express several API migrations and justify several of our design decisions. This paper was accompanied by a tool demo.

A challenge in designing cooperative distributed systems is to develop feasible and cost-effective mechanisms to foster cooperation among selfish nodes, i.e., nodes that strategically deviate from the intended specification to increase their individual utility. Finding a satisfactory solution to this challenge may be complicated by the intrinsic characteristics of each system, as well as by the particular objectives set by the system designer. In a previous work we addressed this challenge by proposing RACOON, a general and semi-automatic framework for designing selfishness-resilient cooperative systems. RACOON relies on classical game theory and a custom built simulator to predict the impact of a fixed set of selfish behaviours on the designer's objectives. In a paper published in IEEE Transactions on Dependable and Secure Computing [12], we present RACOON++, which extends the previous framework with a declarative model for defining the utility function and the static behaviour of selfish nodes, along with a new model for reasoning on the dynamic interactions of nodes, based on evolutionary game theory. We illustrate the benefits of using RACOON++ by designing three cooperative systems: a peer-to-peer live streaming system, a load balancing protocol, and an anonymous communication system. Extensive experimental results using the state-of-the-art PeerSim simulator verify that the systems designed using RACOON++ achieve both selfishness-resilience and high performance.

7.2. Programming after the end of Moore's law

The end of Moore's law is a wake-up call that resonates across Computer Science at large. We are now firmly in an era of custom hardware design, as witnessed by the diversity of system-on-chip (SoC) and specialized processing units – such as graphics processing units (GPUs), tensor processing unit (TPUs) or programmable network adapters, to name but a few. This trend is justified by the existence of niche application domains (graphic processing, linear algebra, packet processing, etc.) that greatly benefit from specialized hardware. Faced with the imminent explosion of the number of niche applications and niche architectures, we are still grasping for a programming model that would accommodate this diversity.

The Usuba project is an exploratory effort in that direction. We chose a niche application domain (symmetric cryptographic algorithms), a specialized execution platform (Single Instruction Multiple Data, SIMD) processors and we set out to design a programming language faithfully describing our application domain as well as an optimizing compiler efficiently exploiting our target execution platform.

Indeed, cryptographic primitives are subject to diverging imperatives. Functional correctness and auditability pushes for the use of a high-level programming language. Performance and the threat of timing attacks push for directly programming in assembler to exploit (or avoid!) the micro-architectural features of a given machine.

In a paper published at PLDI 2019 [23], we have demonstrated that a suitable programming language could reconcile both views and actually improve on the state of the art of both.

USUBA is a dataflow programming language in which block ciphers become so simple as to be “obviously correct” and whose types document and enforce valid parallelization strategies at the granularity of individual bits. Its optimizing compiler, USUBAC, produces high-throughput, constant-time implementations performing on par with hand-tuned reference implementations. The cornerstone of our approach is a systematization and generalization of *bitslicing*, an implementation trick frequently used by cryptographers. We have shown that USUBA can produce code that executes between 5% slower to 22% faster than hand-tuned reference implementations while gracefully scaling across a wide range of architectures and automatically exploiting Single Instruction Multiple Data (SIMD) instructions whenever the cipher’s structure allows it.

7.3. Support for multicore machines

The complexity of computer architectures has risen since the early years of the Linux kernel: Simultaneous Multi-Threading (SMT), multicore processing, and frequency scaling with complex algorithms such as Intel Turbo Boost have all become omnipresent. In order to keep up with hardware innovations, the Linux scheduler has been rewritten several times, and many hardware-related heuristics have been added. Despite this, we have shown in a PLOS paper [16] that a fundamental problem was never identified: the POSIX process creation model, i.e., fork/wait, can behave inefficiently on current multicore architectures due to frequency scaling. We investigate this issue through a simple case study: the compilation of the Linux kernel source tree. To do this, we have developed SchedLog, a low-overhead scheduler tracing tool, and SchedDisplay, a scriptable tool to graphically analyze SchedLog’s traces efficiently. We implement two solutions to the problem at the scheduler level which improve the speed of compiling part of the Linux kernel by up to 26%, and the whole kernel by up to 10%.

In an Eurosys paper [15], we address the problem of efficiently virtualizing NUMA architectures. The major challenge comes from the fact that the hypervisor regularly reconfigures the placement of a virtual machine (VM) over the NUMA topology. However, neither guest operating systems (OSes) nor system runtime libraries (e.g., Hotspot) are designed to consider NUMA topology changes at runtime, leading end user applications to unpredictable performance. We present eXtended Para-Virtualization (XPV), a new principle to efficiently virtualize a NUMA architecture. XPV consists in revisiting the interface between the hypervisor and the guest OS, and between the guest OS and system runtime libraries (SRL) so that they can dynamically take into account NUMA topology changes. We introduce a methodology for systematically adapting legacy hypervisors, OSes, and SRLs. We have applied our approach with less than 2k line of codes in two legacy hypervisors (Xen and KVM), two legacy guest OSes (Linux and FreeBSD), and three legacy SRLs (Hotspot, TCMalloc, and jemalloc). The evaluation results showed that XPV outperforms all existing solutions by up to 304%.

Memory interferences may introduce important slowdowns in applications running on COTS multi-core processors. They are caused by concurrent accesses to shared hardware resources of the memory system. The induced delays are difficult to predict, making memory interferences a major obstacle to the adoption of COTS multi-core processors in real-time systems. In an RTSS paper [18], we propose an experimental characterization of applications’ memory consumption to determine their sensitivity to memory interferences. Thanks to a new set of microbenchmarks, we show the lack of precision of a purely quantitative characterization. To improve accuracy, we define new metrics quantifying qualitative aspects of memory consumption and implement a profiling tool using the VALGRIND framework. In addition, our profiling tool produces high resolution profiles allowing us to clearly distinguish the various phases in applications’ behavior. Using our microbenchmarks and our new characterization, we train a state-of-the-art regressor. The validation on applications from the MIBENCH and the PARSEC suites indicates significant gain in prediction accuracy compared to a purely quantitative characterization.

ALMANACH Project-Team

7. New Results

7.1. New results on text simplification

Participants: Benoît Sagot, Éric Villemonte de La Clergerie, Louis Martin.

Text simplification (TS) aims at making a text easier to read and understand by simplifying grammar and structure while keeping the underlying meaning and information identical. It is therefore an instance of language variation, based on language complexity. It can benefit numerous audiences, such as people with disabilities, language learners or even everyone, for instance when dealing with intrinsically complex texts such as legal documents.

We have initiated in 2017 a collaboration with the Facebook Artificial Intelligence Research (FAIR) lab in Paris and with the UNAPEI, the federation of French associations helping people with mental disabilities and their families. The objective of this collaboration is to develop tools for helping the simplification of texts aimed at mentally disabled people. More precisely, the is to develop a computer-assisted text simplification platform (as opposed to an automatic TS system). In this context, a CIFRE PhD thesis has started in collaboration with the FAIR on the TS task. We have first dedicated important efforts to the problem of the evaluation of TS systems, which remains an open challenge. As the task has common points with machine translation (MT), TS is often evaluated using MT metrics such as BLEU. However, such metrics require high quality reference data, which is rarely available for TS. TS has the advantage over MT of being a monolingual task, which allows for direct comparisons to be made between the simplified text and its original version. We compared multiple approaches to reference-less quality estimation of sentence-level TS systems, based on the dataset used for the QATS 2016 shared task. We distinguished three different dimensions: grammaticality, meaning preservation and simplicity. We have shown that n -gram-based MT metrics such as BLEU and METEOR correlate the most with human judgement of grammaticality and meaning preservation, whereas simplicity is best evaluated by basic length-based metrics [87]. Our implementations of several metrics have been made this year easily accessible and described in a demo paper in collaboration with the University of Sheffield [16].

In 2019, we have also investigated an important issue inherent to the TS task. Although it is often considered an all-purpose generic task where the same simplification is suitable for all, multiple audiences can benefit from simplified text in different ways. We have therefore introduced a discrete parametrisation mechanism that provides explicit control on TS systems based on Seq2Seq neural models. As a result, users can condition the simplifications returned by a model on parameters such as length and lexical complexity. We also show that carefully chosen values of these parameters allow out-of-the-box Seq2Seq neural models to outperform their standard counterparts on simplification benchmarks. Our best parametrised model improves over the previous state of the art performance [61].

Finally, we are involved in the development of a new text simplification corpus. In order to simplify a sentence, human editors perform multiple rewriting transformations: splitting it into several shorter sentences, paraphrasing (i.e. replacing complex words or phrases by simpler synonyms), reordering components, and/or deleting information deemed unnecessary. Despite the vast range of possible text alterations, current models for automatic sentence simplification are evaluated using datasets that are focused on single transformations, such as paraphrasing or splitting. This makes it impossible to understand the ability of simplification models in more abstractive and realistic settings. This is what motivated the development of ASSET, a new dataset for assessing sentence simplification in English, in collaboration with the University of Sheffield (United Kingdom). ASSET is a crowdsourced multi-reference corpus where each simplification was produced by executing several rewriting transformations. Through quantitative and qualitative experiments, we have shown that simplifications in ASSET are better at capturing characteristics of simplicity when compared to other standard evaluation datasets for the task. Furthermore, we have motivated the need for developing better methods for automatic evaluation using ASSET, since we show that current popular metrics may not be suitable for assessment when multiple simplification transformations were performed.

7.2. NLP and computational neurolinguistics

Participants: Éric Villemonte de La Clergerie, Murielle Fabre.

In the context of the CRCNS international network, the ANR-NSF NCM-ML project (dubbed “*Petit Prince* project”) aims to discover and explore correlations between features (or predictors) provided by NLP tools such as parsers, and brain imagery (fMRI) data resulting from listening of the novel *Le Petit Prince*. Following the availability of an increasing amount of fMRI datasets in French and English, the project has investigated the correlations between fMRI observations and an increasing number of parser-based features based on several parsers representing a number of architecture types (LSTM, RNN, Dyalog-SR [statistical], FRMG [hybrid symbolic/statistical]) [20].

While pursuing the purely computation goal of developing a method of variable beam size inference for Recurrent Neural Network Grammar (RNNG) the project investigated how different beam search methods can show different goodness of fit with fMRI signal recorded during naturalistic story listening [58]. This approach is part of a new trend that is now emerging under the name of cognitively inspired NLP, where the effort to leverage from what we know of human cognition to increase machine processing of language data. Drawing inspiration from sequential Monte-Carlo methods such as particle filtering, we illustrated the relevance of our new method for speeding up the computations of direct generative parsing for RNNG, and revealing the potential cognitive interpretation of the underlying representations built by the search method and its beam activity through the analysis of neuro-imaging signal.

A second focus of the project is on compositionality, memory retrieval and syntactic composition during language comprehension. By using quantifications of these hypothesised processes as obtained from computational linguistics we seek to highlight their neural substrates and better understand or model human cognition.

While linguistic expressions have been binarised as compositional and non-compositional given the lack of compositional linguistic analysis, the so-called Multi-word Expressions (MWEs) demonstrate finer-grained degrees of conventionalisation and predictability in psycho-linguistics, which can be quantified through computational Association Measures, like Point-wise Mutual Information and Dice’s Coefficient [57]. An fMRI analysis was conducted to investigate to what extent these computational measures and the underlying cognitive processes they reflect are observable during on-line naturalistic sentence processing. Our results show that predictability, as quantified through Dice’s Coefficient, is a better predictor of neural activation for processing MWEs and the more cognitively plausible computational metric. Computational results (1348) were obtained on MWE identification in French based on new method searching for frequent dependency-patterns [13]. These identifications in the *Little Prince* are contrasted with the ones published for English [69] and will yield an fMRI analysis comparing the two languages and the possible typological differences that the two languages may reflect in terms of morphological strategies to achieve lexical conventionalisation.

7.3. Large-scale raw corpus development

Participants: Benoît Sagot, Éric Villemonte de La Clergerie, Laurent Romary, Pedro Ortiz Suárez, Murielle Fabre, Louis Martin, Benjamin Muller, Yoann Dupont.

In order to be in phase (and comparable) with the US partners of the “*Petit-Prince*” ANR project, Murielle Fabre assembled two French corpora:

- a small corpus for domain adaptation to children’s books: it will permit the fine tuning of the different parsers to a great amount of dialogues and Q&A present in *Le Petit Prince*.
- a large corpus of Contemporary French oral transcriptions and texts to calculate lexical association measures (AM) like PMI (Point-wise Mutual information) or Dice scores on the MWEs found in *Le Petit Prince*. This corpus of approx. 600 millions words, called CaBerNET, represents a balanced counterpart to the American COCA corpus.⁰

⁰<https://corpus.byu.edu/coca/>

We have also developed a general, highly parallel, multi-threaded pipeline to clean and classify Common Crawl by language. Common Crawl is a huge (over 20TB), heterogeneous multilingual corpus comprised of documents crawled from the internet, not sorted per language. We designed our pipeline, called *goclassy*, so that it runs efficiently on medium to low resource infrastructures where I/O speeds are the main constraint. We have created and we distribute a 6.3TB version of Common Crawl, called OSCAR, which is filtered, classified by language, shuffled at line level in order to avoid copyright issues, and ready to be used for NLP applications [29]. OSCAR corpora served as input data to train a variety of neural language models, including the French BERT model CamemBERT (see relevant module for more information). Bridging corpus development, NLP and computational neurolinguistics on of our next step is to train BERT model with the above cited French balanced corpus CaBerNet to create CaBERTnet and extract from it parsing metrics that will be correlated with brain activity as measured by French fMRI recording while listening *Le Petit Prince* in French.

7.4. Neural language modelling

Participants: Benoît Sagot, Djamé Seddah, Éric Villemonte de La Clergerie, Laurent Romary, Louis Martin, Benjamin Muller, Pedro Ortiz Suárez, Yoann Dupont, Ganesh Jawahar.

Pretrained language models are now ubiquitous in Natural Language Processing. Despite their success, most available models have either been trained on English data or on the concatenation of data in multiple languages. This makes practical use of such models—in all languages except English—very limited. In 2019, one of the most visible achievements of the ALMANACH team was the training and release of CamemBERT, a BERT-like [75] (rather, RoBERTa-like) neural language model for French trained on the French section of our large-scale web-based OSCAR corpus, together with CamemBERT variants [60]. Our goal was to investigate the feasibility of training monolingual Transformer-based language models for other languages, taking French as an example and evaluating our language models on part-of-speech tagging, dependency parsing, named entity recognition and natural language inference tasks. We have shown that the use of web-crawled data such as found in OSCAR to train such language models is preferable to the use of Wikipedia data, because of the homogeneity of Wikipedia data. More surprisingly, we have also shown that a relatively small web crawled dataset (4GB randomly extracted from the French section of OSCAR) leads to results that are as good as those obtained using larger datasets (130+GB, i.e. the whole French section of OSCAR). CamemBERT allowed us to reach or improve the state of the art in all four downstream tasks.

Beyond training neural language models, we have reinforced the exploration of an active question, that of their interpretability. With the emergence of contextual vector representations of words, such as the ELMo [89] and BERT language models and word embeddings, the interpretability of neural models becomes a key research topic. It is a way to understand what such neural networks actually learn in an unsupervised way from (huge amounts of) textual data, and in which circumstances they manage to do so. The work carried out in the team this year to identify where morphological vs. syntactic vs. semantic information is stored in a BERT language model [26] was part of a more general trend (see for example [78]). And our work on training ELMo models for five mid-resourced languages has shown that such LSTM-based models, when trained on large scale although non edited dataset such as our web-based corpora OSCAR, can lead to outperforming state-of-the-art performance on a number of downstream tasks such as part-of-speech tagging and parsing. Finally, we have carried out comparative evaluations of the performance of CamemBERT and of ELMo models trained on the same French section of OSCAR on a number of downstream task, with an emphasis on named-entity recognition—a work that led us to publish a new version of the named-entity-annotated version of the French TreeBank [67] that we published in 2012 [99].

We have also investigated how word embeddings can capture the evolution of word usage and meaning over time, at a fine-grained scale. As part of the ANR SoSweet and the PHC Maimonide projects (in collaboration with Bar Ilan University for the latter), ALMANACH has invested a lot of efforts since 2018 into studying language variation within user-generated content (UGC), taking into account two main interrelated dimensions: how language variation is related to socio-demographic and dynamic network variables, and how UGC language evolves over time. Taking advantage of the SoSweet corpus (600 millions tweet) and of the Bar Ilan Hebrew Tweets (180M tweets) both collected over the last 5 years, we have been addressing the problem

of studying semantic changes via the use of dynamic word embeddings, that is embeddings evolving over time. We devised a novel attention model, based on Bernoulli word embeddings, that are conditioned on contextual extra-linguistic features such as network, spatial and socio-economic variables, which can be inferred from Twitter users metadata, as well as topic-based features. We posit that these social features provide an inductive bias that is susceptible to helping our model to overcome the narrow time-span regime problem. Our extensive experiments reveal that, as a result of being less biased towards frequency cues, our proposed model was able to capture subtle semantic shifts and therefore benefits from the inclusion of a reduced set of contextual features. Our model thus fit the data better than current state-of-the-art dynamic word embedding models and therefore is a promising tool to study diachronic semantic changes over small time periods. We published these ideas and results in [41].

A deep understanding of what is learned, and, beyond that, of how it is learned by neural language models, both synchronic and diachronic, will be a crucial step towards the improvement of such architectures (e.g. targeting low-resource languages or scenarios) and the design and deployment of new generations of neural networks for NLP. Particularly important is to assess the role of the training corpus size and heterogeneity, as well as the impact of the properties of the language at hand (e.g. morphological richness, token-type ratio, etc.). This line of research will also have an impact on our understanding of language variation and on our ability to improve the robustness of neural-network-based NLP tools to such variation.

7.5. Processing non-standard language: user-generated content and code-mixed language

Participants: Djamé Seddah, Benoît Sagot, Éric Villemonte de La Clergerie, Benjamin Muller, Ganesh Jawahar, Abhishek Srivastava, Jose Rosales Nuñez, Hafida Le Cloirec, Farah Essaidi, Matthieu Futral.

In 2019, we have resumed our long-lasting efforts towards increasing the robustness of our language analysis tools to the variation found in user-generated content (UGC). We have done this in two directions, in the context of the SoSweet and Parsiti projects.

Firstly, we have investigated how our state-of-the-art hybrid (symbolic and statistical) parsing architecture for French, based on SxPipe, FRMG and the Lefff, behaves on French UGC data, namely on 20 millions tweets from the SoSweet corpus. A first observation was that the current level of pre-parsing normalization was not sufficient to ensure a good parsing coverage with FRMG (around 67%, to be compared with around 93% on journalistic texts such as the French TreeBank), also leading to high parsing times because of correction strategies. However, we applied our error mining strategy [6] to identify a first set of easy errors. Clustering and word embedding were also tried for lemmas relying on the dependency parse trees, again leading to semi-successful results due to the poor quality of the pre-parsing phases.

Secondly, we have investigated the normalisation task, whose goal is to transform possibly noisy UGC into less noisy inputs that are more adapted to our standard neural analysis models (e.g. taggers and parsers). More precisely, we have investigated how useful a language model such as BERT [75], trained on standard data, can be in handling non-canonical text. We study the ability of BERT to perform lexical normalisation in a realistic, and therefore low-resource, English UGC scenario [28]. By framing lexical normalisation as a token prediction task, by enhancing its architecture and by carefully fine-tuning it, we have shown that BERT can be a competitive lexical normalisation model without the need of any UGC resources aside from 3,000 training sentences. To the best of our knowledge, it is the first work done in adapting and analysing the ability of this model to handle noisy UGC data.

Thirdly, we have compared the performances achieved by Phrase-Based Statistical Machine Translation systems (PBSMT) and attention-based Neural Machine Translation systems (NMT) when translating UGC from French to English [44]. We have shown that, contrary to what could have been expected, PBSMT outperforms NMT when translating non-canonical inputs. Our error analysis uncovers the specificities of UGC that are problematic for sequential NMT architectures and suggests new avenue for improving NMT models.

Finally, building natural language processing systems for highly variable and low resource languages is a hard challenge. The recent success of large-scale multilingual pretrained neural language models (including our CamemBERT language model for French) provides us with new modeling tools to tackle it. We have studied the ability of the multilingual version of BERT to model an unseen dialect, namely the Latin-script user-generated North African Arabic dialect called Arabizi. We have shown in different scenarios that multilingual language models are able to transfer to such an unseen dialect, specifically in two extreme cases: across script (Arabic to Latin) and from Maltese, a related language written in the Arabic script, unseen during pretraining. Preliminary results have already been published [66].

7.6. Long-range diachronic variation

Participants: Benoît Sagot, Laurent Romary, Éric Villemonte de La Clergerie, Clémentine Fourier, Gaël Guibon, Mathilde Regnault, Kim Gerdes.

ALMANaCH members have resumed their work on longer-range diachronic variation, in two distinct directions:

- Firstly, we have been working on resources and tools for Old French, using contemporary French as a starting point for which resources and tools are available. This work is carried out within the ANR project “Profiterole”, whose goal is to automatically annotate a large corpus of medieval French (9th-15th centuries) in dependency syntax and to provide a methodology for dealing with heterogeneous data as found in such a corpus. Indeed, Old French does not only involve diachronic variation when contrasted with contemporary French. It also involve large internal variation, notably because of diachronic (within Old French), dialectal, geographic, stylistic and genre-based variation. We have carried out experiments on morphosyntactic tagging by trying to determine which parameters and which training sets are the best ones to use when annotating a new text. We explored two approaches for parsing. On the one hand, an ongoing thesis aims at adapting the FRMG metagrammar to medieval French, notably by changing the constraints on certain syntactic phenomena and relaxing the order of words [31], [30]. This work relies on the new morphological and syntactic lexicon for Old French, OFrLex, developed at ALMANaCH [34]. On the other hand, we conducted parsing experiments with neural models (DyALog’s SRNN models).
- Secondly, we have started experiments to investigate whether and under which conditions neural networks can be used for learning sound correspondences between two related languages, i.e. for predicting cognates of source language words in a related target language. In order to obtain suitably large homogeneously phonetised data, we extracted bilingual lexicons and cognate sets from available resources, including our EtymDB etymological database, of which a new, extended version was created in 2019. This data was then used to train and evaluate several neural architectures (seq2seq, Siamese). Preliminary results are promising, but further investigation is required.

These two research directions will find a common ground now that we have begun to investigate, in the context of the Profiterole ANR project, how we can model the diachronic evolution of the lexicon from Old French to contemporary French. Moreover, our work on Basnage’s 1701 *Dictionnaire Universel*, in the context of the BASNUM ANR project might draw some inspiration from the Profiterole project. But since 1700’s French is much closer from contemporary French than Old French, another source of inspiration for BASNAGE might come from our work on sociolinguistic variation in contemporary French and more generally on our work on User-Generated Content (UCG).

7.7. Syntax and treebanking

Participants: Djamé Seddah, Benoît Sagot, Kim Gerdes, Benjamin Muller, Pedro Ortiz Suárez, Marine Courtin.

In 2019 we have introduced the first treebank for a romanized user-generated content of Algerian, a North-African Arabic dialect called Arabizi. It contains 1500 sentences, fully annotated in morpho-syntax and universal dependencies, and is freely available. We complement it with 50k unlabeled sentences that were collected using intensive data-mining techniques from Common Crawl and web-crawled data. Preliminary results show its usefulness for POS tagging and dependency parsing.

We have also developed the first syntactic treebank for spoken Naija, an English pidgin creole, which is rapidly spreading across Nigeria. The syntactic annotation is developed in the Surface-Syntactic Universal Dependency annotation scheme (SUD) [77] and automatically converted into Universal Dependencies (UD). A crucial step in the syntactic analysis of a spoken language consists in manually adding a markup onto the transcription, indicating the segmentation into major syntactic units and their internal structure. We have shown that this so-called “macrosyntactic” markup improves parsing results. We have also studied some iconic syntactic phenomena that clearly distinguish Naija from English. This work is published in [36].

We have carried out two pilot studies in empirical syntax based on UD treebanks. In a first study [38], we investigate the relationship between dependency distance and frequency based on the analysis of an English dependency treebank. The preliminary result shows that there is a non-linear relation between dependency distance and frequency. This relation between them can be further formalised as a power law function which can be used to predict the distribution of dependency distance in a treebank. In a second study [40], we discussed an empirical refoundation of selected Greenbergian word order universals based on a data analysis of the Universal Dependencies project. The nature of the data we worked on allows us to extract rich details for testing well-known typological universals and constitutes therefore a valuable basis for validating Greenberg’s universals. Our results show that we can refine some Greenbergian universals in a more empirical and accurate way by means of a data-driven typological analysis.

Finally, we have introduced a new schema to annotate Chinese Treebanks on the character level. The original UD and SUD projects provide token-level resources with rich morphosyntactic language details. However, without any commonly accepted word definition for Chinese, the dependency parsing always faces the dilemma of word segmentation. Therefore we have presented a character-level annotation schema integrated into the existing Universal Dependencies schema as an extension [39]. The different SUD projects were also presented at the Journées scientifiques “Linguistique informatique, formelle et de terrain” (LIFT 2019), Nov 28-29, 2019 at the University of Orléans.

7.8. Analysing and enriching legacy dictionaries

Participants: Laurent Romary, Benoît Sagot, Mohamed Khemakhem, Pedro Ortiz Suárez, Achraf Azhar.

2019 has been a year of deployment and large-scale experiment of the work initiated in 2016 on the analysis and enrichment of legacy dictionaries and implemented in the GROBID-dictionary framework [84]. GROBID-dictionary is an extension of the generic GROBID Suite [95] and implements an architecture of cascading CRF models with the purpose to parse and categorize components of a pdf documents, whether born-digital or resulting from an OCR. It is developed as part of the doctoral work of Mohamed Khemakhem. GROBID dictionaries produces an output that is conformant to the Text Encoding Initiative guideline and thus easy to distribute and further process in an open science context. We have had the opportunity to show the performances and robustness of the architecture on a variety of dictionaries and contexts resulting both from internal and external collaborations:

- In the context of the language documentation project of Jack Bowers dealing with Mixtepec-Mixtec (ISO 639-3: mix, [72]), we have been successful in completely parsing a new edition of an historical lexical resource of Colonial Mixtec ‘Voces del Dzaha Dzahui’ published by the Dominican fray Francisco Alvarado in the year 1593, published by Jansen and Perez Jiménez (2009). The result is now integrated into the reference lexical description maintained by Jack). See [18];
- Within the Nénufar project, a collaboration with the Praxiling laboratory in Montpellier, we have been contributing to the analyses and encoding of several editions of the Petit Larousse Illustré, a central legacy publication for the French language. [17], [27];
- For the ANR funded project BASNUM, we are deeply involved in understanding how a complex, semi-structured dictionary, for which we do not necessarily have a high quality digitized primary source, can be properly segmented in lexical entries and subfields from which we expect being able to extract fine-grained linguistic content (e.g. named entities for literary sources). In [42], we have shown for instant how the GROBID-dictionary framework could be robust to variations in scanning and thus OCR quality;

- In the same context of the BASNUM project, we have also started to explore the possibility of deploying deep learning components. As shown in [43], the main challenges is the lack of available annotated data in order to train machine learning models, decreased accuracy when using modern pre-trained models due to the differences between present-day and 18th century French, and even unreliable or low quality OCRisation;
- These various experiments have been accompanied by an intense training and hand-on activity in the context in particular of the Lexical Data Master Class and collaboration within the ELEXIS project, which has opted for using the system for building a dictionary matrix from legacy dictionaries ⁰.Further alignments with the ongoing standardisation activities around TEI Lex0 and ISO 24613 (LMF) has been carried out to ensure a proper standards compliance of the generated output.
- Finally, and as a nice example of the kind of DH collaborations that our researches can lead to, we should mention here the targeted experiments that we carried out on extending the GROBID-dictionary framework to deal with objects which, although analogous with dictionary entries from a distance, appear to have a highly specific structure. This is the case Manuscript Sales Catalogues, which are highly important for authenticating documents and studying the reception of authors. Their regular publication throughout Europe since the beginning of the 19th c. has raised the interest around scaling up the means for automatically structuring their contents. [33] presents the results of advanced tests of the system's capacity to handle a large corpus with MSC of different dealers, and therefore multiple layouts.

7.9. Coreference resolution

Participants: Loïc Grobol, Éric Villemonte de La Clergerie.

In 2019 we have resumed our work on coreference resolution for French with the release in [25] of the first end-to-end automatic coreference resolution system for spoken French by adapting state-of-the art neural network system to the case of noisy non-standard inputs.

This first release uses no external knowledge beyond pretrained non-contextual word embeddings, making it suitable for applications to languages with less pre-existing resources. We also investigated the integration of further knowledge, both in the form of contextual embedding techniques such as CamemBERT and syntactic parsers developed at ALMANACH (works to be published in 2020).

⁰<https://grobid.elex.is>

COML Team

6. New Results

6.1. Unsupervised learning

Humans learn to speak and to perceive the world in a largely self-supervised fashion. Yet, most of machine learning is still devoted to supervised algorithms that rely on abundant quantities of human labelled data. We have used humans as sources of inspiration for developing 3 novel machine learning benchmarks in order to push the field towards self-supervised learning.

- In the Zero Resource Speech Challenge 2019 [19], presented as a special session at Interspeech 2019, we propose to build a speech synthesizer without any text or phonetic labels: hence, TTS without T (text-to-speech without text). We provide raw audio for a target voice in an unknown language (the Voice dataset), but no alignment, text or labels. Participants must discover subword units in an unsupervised way (using the Unit Discovery dataset) and align them to the voice recordings in a way that works best for the purpose of synthesizing novel utterances from novel speakers, similar to the target speaker's voice. We describe the metrics used for evaluation, a baseline system consisting of unsupervised subword unit discovery plus a standard TTS system, and a topline TTS using gold phoneme transcriptions. We present an overview of the 19 submitted systems from 11 teams and discuss the main results.
- In [27], we introduce a new collection of spoken English audio suitable for training speech recognition systems under limited or no supervision. It is derived from open-source audio books from the LibriVox project. It contains over 60K hours of audio, which is, to our knowledge, the largest freely-available corpus of speech. The audio has been segmented using voice activity detection and is tagged with SNR, speaker ID and genre descriptions. Additionally, we provide baseline systems and evaluation metrics working under three settings: (1) the zero resource/unsupervised setting (ABX), (2) the semi-supervised setting (PER, CER) and (3) the distant supervision setting (WER). Settings (2) and (3) use limited textual resources (10 minutes to 10 hours) aligned with the speech. Setting (3) uses large amounts of unaligned text. They are evaluated on the standard LibriSpeech dev and test sets for comparison with the supervised state-of-the-art.
- In order to reach human performance on complex visual tasks, artificial systems need to incorporate a significant amount of understanding of the world in terms of macroscopic objects, movements, forces, etc. Inspired by work on intuitive physics in infants, we propose in [28] an evaluation framework which diagnoses how much a given system understands about physics by testing whether it can tell apart well matched videos of possible versus impossible events. The test requires systems to compute a physical plausibility score over an entire video. It is free of bias and can test a range of specific physical reasoning skills. We then describe the first release of a benchmark dataset aimed at learning intuitive physics in an unsupervised way, using videos constructed with a game engine. We describe two Deep Neural Network baseline systems trained with a future frame prediction objective and tested on the possible versus impossible discrimination task. The analysis of their results compared to human data gives novel insights in the potentials and limitations of next frame prediction architectures. This benchmark is currently being used in the DARPA project Machine Common Sense.

6.2. Language emergence in communicative agents

In this relatively new research topic, which is currently the focus of Rahma Chaabouni's PhD thesis, we study the inductive biases of neural systems by presenting them with few or no data.

- In [18], we study LSTMs' biases with respect to "natural" word-order constraints. To this end, we train them to communicate about trajectories in a grid world, using an artificial language that reflect or violate various natural language trends, such as the tendency to avoid redundancy or to minimize long-distance dependencies. We measure the speed of individual learning and the generational stability of language patterns in an iterative learning setting. Our results show a mixed picture. If LSTMs are affected by some "natural" word-order constraints, such as a preference for iconic orders and short-distance constructions, they have a preference toward redundant languages.
- In [25], we ask whether LSTMs have least-effort constraints and how this can affect their language. We let the neural systems develop their own language, to study a fundamental characteristic of natural language; Zipf's Law of Abbreviation (ZLA). In other words, we investigate if, even with the lack of the least-effort, LSTMs would produce a ZLA-like distribution like what we observe in natural language. Surprisingly, we find that networks develop an anti-efficient encoding scheme, in which the most frequent inputs are associated to the longest messages, and messages in general are skewed towards the maximum length threshold. This anti-efficient code appears easier to discriminate for the listener, and, unlike in human communication, the speaker does not impose a contrasting least-effort pressure towards brevity, as observed in [18]. Indeed, when the cost function includes a penalty for longer messages, the resulting message distribution starts respecting (ZLA). Our analysis stresses the importance of studying the basic features of emergent communication in a highly controlled setup, to ensure the latter will not strand too far from human language. Moreover, we present a concrete illustration of how different functional pressures can lead to successful communication codes that lack basic properties of human language, thus highlighting the role such pressures play in the latter.
- There is renewed interest in simulating language emergence among deep neural agents that communicate to jointly solve a task, spurred by the practical aim to develop language-enabled interactive AIs, and by theoretical questions about the evolution of human language. However, optimizing deep architectures connected by a discrete communication channel (such as that in which language emerges) is technically challenging. In [21], we introduce EGG, a toolkit that greatly simplifies the implementation of emergent-language communication experiments. EGG's modular design provides a set of building blocks that the user can combine to create new communication games, easily navigating the optimization and architecture space. We hope that the tool will lower the technical barrier, and encourage researchers from various backgrounds to do original work in this exciting area/

6.3. Evaluation of AI algorithms

Machine learning algorithms are typically evaluated in terms of end-to-end tasks, but it is very often difficult to get a grasp of how they achieve these tasks, what could be their break point, and more generally, how they would compare to the algorithms used by humans to do the same tasks. This is especially true of Deep Learning systems which are particularly opaque. The team develops evaluation methods based on psycholinguistic/linguistic criteria, and deploy them for systematic comparison of systems.

- Recurrent neural networks (RNNs) can learn continuous vector representations of symbolic structures such as sequences and sentences; these representations often exhibit linear regularities (analogies). Such regularities motivate our hypothesis that RNNs that show such regularities implicitly compile symbolic structures into tensor product representations (TPRs; Smolensky, 1990), which additively combine tensor products of vectors representing roles (e.g., sequence positions) and vectors representing fillers (e.g., particular words). To test this hypothesis, we introduce Tensor Product Decomposition Networks (TPDNs), which use TPRs to approximate existing vector representations. We demonstrate using synthetic data that TPDNs can successfully approximate linear and tree-based RNN autoencoder representations, suggesting that these representations exhibit interpretable compositional structure; we explore the settings that lead RNNs to induce such structure-sensitive representations. By contrast, further TPDN experiments show that the representations of four models trained to encode naturally-occurring sentences can be largely approximated with a bag of words,

with only marginal improvements from more sophisticated structures. We conclude that TPDNs provide a powerful method for interpreting vector representations, and that standard RNNs can induce compositional sequence representations that are remarkably well approximated by TPRs; at the same time, existing training tasks for sentence representation learning may not be sufficient for inducing robust structural representations.

- LSTMs have proven very successful at language modeling. However, it remains unclear to what extent they are able to capture complex morphosyntactic structures. In [29], we examine whether LSTMs are sensitive to verb argument structures. We introduce a German grammaticality dataset in which ungrammatical sentences are constructed by manipulating case assignments (eg substituting nominative by accusative or dative). We find that LSTMs are better than chance in detecting incorrect argument structures and slightly worse than humans tested on the same dataset. Surprisingly, LSTMs are contaminated by heuristics not found in humans like a preference toward nominative noun phrases. In other respects they show human-similar results like biases for particular orders of case assignments.
- Pater (2019) proposes to use neural networks to model learning within existing grammatical frameworks. In [16] we argue that there is a fundamental gap to be bridged that does not receive enough attention : how can we use neural networks to examine whether it is possible to learn some linguistic representation (a tree, for example) when, after learning is finished, we cannot even tell if this is the type of representation that has been learned (all we see is a sequence of numbers)? Drawing a correspondence between an abstract linguistic representational system and an opaque parameter vector that can (or perhaps cannot) be seen as an instance of such a representation is an implementational mapping problem. Rather than relying on existing frameworks that propose partial solutions to this problem, such as harmonic grammar, we suggest that fusional research of the kind proposed needs to directly address how to ‘find’ linguistic representations in neural network representations.

6.4. Learnability relevant descriptions of linguistic corpora

Evidently, infants are acquiring their language based on whatever linguistic input is available around them. The extent of variation that can be found across languages, cultures and socio-economic background provides strong constraints (lower bounds on data, higher bounds on noise, and variation and ambiguity) for language learning algorithms.

- Previous computational modeling suggests it is much easier to segment words from child-directed (CDS) than adult-directed speech (ADS). However, this conclusion is based on data collected in the laboratory, with CDS from play sessions and ADS between a parent and an experimenter, which may not be representative of ecologically-collected CDS and ADS. In [15], fully naturalistic ADS and CDS collected with a non-intrusive recording device as the child went about her day were analyzed with a diverse set of algorithms. The difference between registers was small compared to differences between algorithms, it reduced when corpora were matched, and it even reversed under some conditions. These results highlight the interest of studying learnability using naturalistic corpora and diverse algorithmic definitions.
- A number of unsupervised learning algorithms have been proposed in the last 20 years for modeling early word learning, some of which have been implemented computationally, but whose results remain difficult to compare across papers. In [14], we created a tool that is open source, enables reproducible results, and encourages cumulative science in this domain. WordSeg has a modular architecture: It combines a set of corpora description routines, multiple algorithms varying in complexity and cognitive assumptions (including several that were not publicly available, or insufficiently documented), and a rich evaluation package. In the paper, we illustrate the use of this package by analyzing a corpus of child-directed speech in various ways, which further allows us to make recommendations for experimental design of follow-up work. Supplementary materials allow readers to reproduce every result in this paper, and detailed online instructions further enable them to go

beyond what we have done. Moreover, the system can be installed within container software that ensures a stable and reliable environment. Finally, by virtue of its modular architecture and transparency, WordSeg can work as an open-source platform, to which other researchers can add their own segmentation algorithms.

6.5. Test of the psychological validity of AI algorithms.

In this section, we focus on the utilisation of machine learning algorithms of speech and language processing to derive testable quantitative predictions in humans (adults or infants).

- In [24], we compare the performance of humans (English and French listeners) versus an unsupervised speech model in a perception experiment (ABX discrimination task). Although the ABX task has been used for acoustic model evaluation in previous research, the results have not, until now, been compared directly with human behaviour in an experiment. We show that a standard, well-performing model (DPGMM) has better accuracy at predicting human responses than the acoustic baseline. The model also shows a native language effect, better resembling native listeners of the language on which it was trained. However, the native language effect shown by the models is different than the one shown by the human listeners, and, notably, the models do not show the same overall patterns of vowel confusions.
- Word learning relies on the ability to master the sound contrasts that are phonemic (i.e., signal meaning difference) in a given language. Though the timeline of phoneme development has been studied extensively over the past few decades, the mechanism of this development is poorly understood. In [20], we take inspiration from computational modeling work in language grounding where phonetic and visual information is learned jointly. In this study, we varied the taxonomic distance of pairs of objects and tested how adult learners judged the phonemic status of the sound contrast associated with each of these pairs. We found that judgments were sensitive to gradients in the taxonomic structure, suggesting that learners use probabilistic information at the semantic level to optimize the accuracy of their judgements at the phonological level. The findings provide evidence for an interaction between phonological learning and meaning generalization in human learning.

6.6. Applications and tools for researchers

Some of CoMLs' activity is to produce speech and language technology tools that facilitate research into language development or clinical applications.

- Speech classifiers of paralinguistic traits traditionally learn from diverse hand-crafted low-level features, by selecting the relevant information for the task at hand. We explore an alternative to this selection, by learning jointly the classifier, and the feature extraction. Recent work on speech recognition has shown improved performance over speech features by learning from the waveform. In [24], we extend this approach to paralinguistic classification and propose a neural network that can learn a filterbank, a normalization factor and a compression power from the raw speech, jointly with the rest of the architecture. We apply this model to dysarthria detection from sentence-level audio recordings. Starting from a strong attention-based baseline on which mel-filterbanks out-perform standard low-level descriptors, we show that learning the filters or the normalization and compression improves over fixed features by 10% absolute accuracy. We also observe a gain over OpenSmile features by learning jointly the feature extraction, the normalization, and the compression factor with the architecture. This constitutes a first attempt at learning jointly all these operations from raw audio for a speech classification task.
- This paper [23] presents the problems and solutions addressed at the JSALT workshop when using a single microphone for speaker detection in adverse scenarios. The main focus was to tackle a wide range of conditions that go from meetings to wild speech. We describe the research threads we explored and a set of modules that was successful for these scenarios. The ultimate goal was to explore speaker detection; but our first finding was that an effective diarization improves detection,

and not having a diarization stage impoverishes the performance. All the different configurations of our research agree on this fact and follow a main backbone that includes diarization as a previous stage. With this backbone, we analyzed the following problems: voice activity detection, how to deal with noisy signals, domain mismatch, how to improve the clustering; and the overall impact of previous stages in the final speaker detection. In this paper, we show partial results for speaker diarization to have a better understanding of the problem and we present the final results for speaker detection.

RITS Project-Team

6. New Results

6.1. Multi-Task Cross-Modality Deep Learning for Pedestrian Risk Estimation

Participants: Danut Ovidiu Pop, Fawzi Nashashibi.

We want to solve the problem of multi-task pedestrian protection system (PPS) including not only pedestrian classification, detection and tracking, but also pedestrian action-unit classification and prediction, and finally pedestrian risk estimation. The goal of our research work is to develop an intelligent pedestrian protection component based only on single stereo vision system using an optimal cross-modality deep learning architecture in order to fulfill the prior requirements.

The system has to be able not only to detect all the pedestrians with high precision but also to track all the pedestrian paths, to classify the current pedestrian action and to predict their next actions and, finally, to estimate the pedestrian risk by the time to crossing for each pedestrian.

First, we investigate the classification component where we analyzed how learning representations from one modality would enable recognition for other modalitie(s) within various deep learning, which one terms as cross-modality learning. Second, we study how the cross modality learning improves an end-to-end the pedestrian action detection. Third, we analyze the pedestrian action prediction and the estimation of time to cross the street.

This work has been done in collaboration with Alexandrina Rogozan and Abdelaziz Bensrhair of INSA Rouen. More detail can be fund in [12], [13], [20], [11] and in the PhD manuscript of Danut Ovidiu Pop [6].

6.2. Study on the effect of rain on computer vision

Participants: Raoul de Charette, Fabio Pizzati.

Following the works initiated in past years, we have emphasized the need of developing for outdoor-applications to be robust to adverse weather conditions.

Three works were developed this year: two in the the context of the Samuel de Champlain Québec-France collaboration with Jean-François Lalonde from Univ. Laval (Canada) and another in the context of the new co-tutelle PhD thesis of Fabio Pizzati.

- We have first proposed a physically-based rain rendering pipeline for realistically inserting rain into clear weather images. Our research [16] was published at ICCV'19 and relies on a physical particle simulator, an estimation of the scene lighting and an accurate rain photometric modeling to augment images with arbitrary amount of realistic rain or fog. We validate our rendering with a user study, proving our rain is judged 40% more realistic that state-of-the-art. Using our generated weather augmented Kitti and Cityscapes dataset, we conduct a thorough evaluation of deep object detection and semantic segmentation algorithms and show that their performance decreases in degraded weather, on the order of 15% for object detection and 60% for semantic segmentation. Furthermore, we show refining existing networks with our augmented images improves the robustness of both object detection and semantic segmentation algorithms. We experiment on the popular nuScenes dataset and measure an improvement of 15% for object detection and 35% for semantic segmentation compared to original rainy performance.

Along with the research we have released the full augmented dataset on our project page ⁰ and the source code will be soon released.

- An alternative proposal is to use generative networks (GANs) to learn the translation of clear weather images to rainy images. This was achieved in the thesis of Fabio Pizzati and led to an accepted conference paper at WACV'20. To overcome the limitation of publicly available annotated datasets, we propose to learn the clear to rain mapping from datasets of different sources. Standard image-to-image translation architectures have limited effectiveness in such case due to the large source / target domain gap and usually fail to model typical traits of rain as water drops, which ultimately impacts the synthetic images realism. We proposed here a new type of domain bridge, that benefits from web-crawled data to reduce the domain gap.
- To circumvent the limitation of physics-based rendering and GANs rendering, we are currently working on extensions of [16] with Maxime Tremblay, PhD student at Univ. Laval. In this work, we are combining data-driven GAN approaches and physics-based driven learning.

6.3. Unsupervised Domain Adaptation

Participants: Raoul de Charette, Maximilian Jaritz, Fawzi Nashashibi, Fabio Pizzati.

There is an evident dead end to the paradigm of supervised learning, as it requires costly human labeling of millions of data frames to learn the appearance models of objects. As of today, the databases are recorded in very narrow conditions (e.g. only clear weather, only USA, only daytime). Adjusting to unseen conditions such as snow, hail, nighttime or unseen cities, require supervised algorithms to be retrained. Conversely, as humans we're capable of generalizing prior knowledge to new tasks. During this year, we initiated two works on transfer learning, typically Unsupervised Domain Adaptation (UDA) which is crucial to tackle the lack of annotations in a new domain. We have conducted two parallel projects on UDA: the first one in the scope of Maximilian Jaritz' thesis [27] (submitted), and the second one in the scope of Fabio Pizzati's work on rainy scenarios:

- **xMUDA:** In the first work, we explore how to learn from multi-modality and proposed cross-modal UDA (xMUDA) where we assumed the presence of 2D images and 3D point clouds for 3D semantic segmentation. This is challenging as the two input spaces are heterogeneous and can be impacted differently by domain shift. In xMUDA, modalities learn from each other through mutual mimicking, disentangled from the segmentation objective, to prevent the stronger modality from adopting false predictions from the weaker one. We evaluated on new UDA scenarios including day-to-night, country-to-country and dataset-to-dataset, leveraging recent autonomous driving datasets. xMUDA brings large improvements over uni-modal UDA on all tested scenarios, and is complementary to state-of-the-art UDA techniques.
- **Weighted Pseudo Labels:** The second work focus specifically on semantic segmentation in rainy scenarios. We benefited from our other work on GANs clear to rain translation to apply a self-supervised domain adaptation (aka UDA) that learns from the use of pseudo labels. Using pseudo labels enables the self-supervision of the learning reinforcing the network belief in its own predictions. To circumvent the use of hard-coded threshold, which is a common practice for pseudo labels, we proposed new Weighted Pseudo Labels that actively learn the ad-hoc threshold in a sort of region-growing techniques.

6.4. 3D completion and surface modeling

Participants: Raoul de Charette, Maximilian Jaritz, Manohar Kv.

Depth sensors (LiDARs, Time-of-flight cameras, stereo) gather geometrical knowledge about the scene which are rich and may be beneficial for many tasks. However, the depth information is usually sparse in nature and do not recover volumes and surfaces of objects.

⁰<https://team.inria.fr/rits/computer-vision/weather-augment/>

This year we have conducted three works on the topic: one work to densify the 3D point clouds generated from LiDAR sensors, another work to fuse 2D images and 3D point clouds, and finalized another work to reconstruct 3D deformable objects.

- The first work is in spirit a 3D point completion and was initiated with intern Manohar Kv. We developed a 3D pipeline to process point cloud and densify existing point clouds. It uses a modified version of the popular PointNet++ and it is thus able to reconstruct highly occluded 3D point clouds. The work is not yet published.
- In [17] we introduce a framework to fuse 2D multi-view images and 3D point clouds in an effective way by computing image features in 2D first, lifting them to 3D, and then fuse complementary geometry and image information in canonical 3D space. This work has been done while Maximilian Jaritz was visiting San Diego University.
- In [24], we propose a new algorithm to reconstruct 3D deformable objects heavily occluded. It uses an automatic registration of multiple depth sensors and Gaussian Mixture Modeling in the radial domain to detect and reconstruct object from their symmetrical properties. This research was applied in the context of pottery wheel for the preservation of the cultural heritage and conducted in collaboration with Mines ParisTech. It resulted that our method enabled reconstruction of challenging deformable objects with an average precision of 7.6mm.

6.5. 3D Surface Reconstruction from Voxel-based Lidar Data

Participants: Luis Roldao, Raoul de Charette, Anne Verroust-Blondet.

To achieve fully autonomous navigation, vehicles need to compute an accurate model of their direct surroundings. In fact, imprecise representations may lead to unexpected situations that could endanger the passengers. This year, we have proposed an algorithm capable to perform a fine and accurate 3D surface reconstruction of the environment from depth sensors. This representation keeps a high level of detail on the reconstruction, while maintaining a high density in the areas close to the vehicle.

Existing methods used for surface reconstruction from 3D data struggle to accommodate to the heterogeneous density of the input data while keeping the reconstruction accuracy. Conversely, our method is capable of handling this variable density by using an adaptive neighborhood kernel that perform local approximations of the data at different levels. This also permit to gain robustness against noise and output a smoother reconstruction. We also introduce a Gaussian confidence function capable to select the most adequate kernel for the local surface estimation. A Truncated Signed Distance Function (TSDF) is then globally estimated from the local surfaces to obtain the final mesh that represents the input scan.

The proposed method was evaluated in both simulated and real data. Reconstruction results show an improvement on the representation when compared with popular methods such as Implicit Moving Least Squares (IMLS), as the average error of our reconstruction is often 50% lower. Furthermore, almost 80% of vertices from our output mesh present an error below $0.2m$, while only 40% of vertices lie below the same threshold for IMLS. Our method is capable to output a higher level of detail on the reconstruction, while keeping a high density in vehicle surroundings, the mesh can be of special interest for both the robotics and the graphics community to perform different tasks, such as terrain traversability assessment or physical modeling.

More details can be found in [21]. This research is partially funded by AKKA Technology.

6.6. Attention mechanisms for vehicle trajectory prediction

Participants: Kaouther Messaoud, Fawzi Nashashibi, Anne Verroust-Blondet, Itheri Yahiaoui.

Scene understanding and future motion prediction of surrounding vehicles are crucial to achieve safe and reliable decision-making and motion planning for autonomous driving in a highway environment. This is a challenging task considering the correlation between the drivers behaviors. Two methods using attention mechanisms have been introduced in this context:

- In [18], we present a new approach based on an LSTM encoder-decoder that uses a social pooling mechanism to model the interactions between all the neighboring vehicles. This social pooling module combines both local and non-local operations: the non-local multi-head attention mechanism captures the relative importance of each vehicle despite the inter-vehicle distances to the target vehicle, while the local blocks represent nearby interactions between vehicles. Evaluations have been performed using two naturalistic driving datasets: Next Generation Simulation (NGSIM) and the highD Dataset⁰. The proposed method outperforms existing ones in terms of RMS values of prediction error, which shows the effectiveness of combining local and non-local operations in such a context.
- In [19] we propose an RRNNs based encoder-decoder architecture where the encoder analyzes the patterns underlying in the past trajectories and the decoder generates the future trajectory sequence. The originality of this network is that it combines the advantages of the LSTM blocks in representing the temporal evolution of trajectories and the attention mechanism to model the relative interactions between vehicles. The proposed method outperforms LSTM encoder decoder in terms of RMSE values of the predicted trajectories on the large scaled naturalistic driving highD dataset.

6.7. A unified framework for robust 2D/3D PML-SLAM

Participants: Kathia Melbouci, Fawzi Nashashibi.

Enhancing the outdoor mapping with SLAM based approaches is still an active research area. The main reason is that a consistent map of the vehicle's surrounding is one of the prerequisites for an effective vehicle interaction with this environment. In this context, and for the VALET project purpose, we have extended the PML-SLAM framework to handle 2D and 3D Lidars by replacing the localization module and designing a sparse pose graph optimizer. The sparse pose graph jointly optimizes the poses of the submaps generated by the local SLAM, which are already used for the mapping task, and the poses of the scans estimated following the scan matching process. This optimization is formulated as a non linear least square problem, and runs online in a background thread. The optimized poses are used to correct the vehicle's trajectory and to update the environment map. Furthermore, the graph-based PML-SLAM can deal with different sensors (IMU, GPS), that is, a sensor fusion "Kalman-filter" based is available to provide a good pose estimate for the local SLAM.

6.8. LIDAR-Based perception For Vehicle Localization in an HD Map

Participants: Farouk Ghallabi, Fawzi Nashashibi.

Self-vehicle localization is one of the fundamental tasks for autonomous driving. Most of current techniques for global positioning are based on the use of GNSS (Global Navigation Satellite Systems). However, these solutions do not provide a localization accuracy that is better than 2-3 m in open sky environments. Alternatively, the use of maps has been widely investigated for localization since maps can be pre-built very accurately. State of the art approaches often use dense maps or feature maps for localization. This year, we tackled to problems:

- In [14] we proposed a road sign perception system for vehicle localization within a third party map. This is challenging since third party maps are usually provided with sparse geometric features, which makes the localization task more difficult in comparison to dense maps. Experiments have been conducted on a Highway-like test track using GNSS/INS with RTK corrections as ground truth (GT).

⁰<https://www.highd-dataset.com/>

- In [15] High Reflective Landmarks (HRL) - such as lane markings, road signs and guard rail reflectors (GRR) - are detected from a 3D point cloud. A particle filtering algorithm estimates the position of the vehicle by matching observed HRLs with HD map attributes. Experiments have been conducted on a highway-like test track using GNSS/INS with RTK corrections as a ground truth (GT). Error evaluations are given as cross-track (CT) and along-track (AT) errors defined in the curvilinear coordinates related to the map. The obtained accuracies of the localization system is 18 cm for the cross-track error and 32 cm for the along-track error.

6.9. Motion planning in presence of highly dynamic obstacles with uncertain motion

Participants: Pierre de Beaucorps, Renaud Poncelet, Anne Verroust-Blondet, Fawzi Nashashibi.

Safe motion planning in a dynamic environment is of great importance in many robotics applications. This year, we have worked in two directions:

- The work on reachable interaction sets introduced in [37] has been extended to the case of dynamic obstacles with uncertain motions. We consider that the obstacles have stochastic motions and we use a probabilistic formulation to compute the RIS at each time step. Our approach improves existing methods in such a context (cf. Pierre de Beaucorps PhD thesis [7]).
- Focusing on autonomous vehicles, we begun to study scenarios with occluded dynamic obstacles.

6.10. A vehicle dynamic model corrector with side slip estimation for adding safety capabilities in autonomous vehicle

Participants: Imane Mahtout, Fawzi Nashashibi.

The ability to identify malfunctions on autonomous vehicles is critical for their deployment. As a matter of fact, systems able to identify when the positioning systems are not providing accurate data, or the perception algorithms are not properly detecting the environment, are extremely important to assure a certain safety level for automated vehicles. This is especially true since these systems are finally connected to the control module that provides the adequate commands to vehicle's actuators. For control algorithms to work properly, proper inputs are necessary to reduce noise, increase controllability and avoid system's malfunctions and instability. It is also critical for these algorithms to identify/consider vehicle physical limits for determining when is the automated system still capable of handling the vehicle. From the above, it is clear that automated vehicles are in need of proper inputs to control the vehicle, but also it is necessary to detect critical situations where the nominal control behavior is no longer assured, in order to take the vehicle to a safe state. Slide slip state is an example of a critical situation where the vehicle is no longer able to correct its trajectory. Thus, this part of my thesis work consists on developing a module for providing smooth signals to the controller and, at the same time, detect side slip situations. The lateral controller implemented in our automated driving (AD) system is based on the yaw error minimization between the desired yaw rate (obtained from the road layout information in function of the curvature) and the current vehicle yaw rate. From this, the first step is to provide a proper current yaw rate measurement. The proposed device compares the measured yaw rate value (coming from the vehicle sensor) with a model-based estimated yaw rate value. The idea is to identify vehicle model mismatches, correcting the model in real time. This permits to extend the nominal vehicle planar model to a road layout-independent model, where roll and pitch variations are considered. This first stage consists of a vehicle model compensator that includes unmodeled vehicle dynamics and parameter incertitude in real time when the vehicle is operating in autonomous mode. The vehicle lateral controller is then fed by the compensator output to allow robust performance in all road conditions. Once lateral control is fed with proper inputs, the second stage is the one detecting that vehicle handling physical limits are surpassed. The proposed system based on Youla-Kucera parametrization identifies the vehicle physical limits by estimating front and rear lateral forces, using as input the previous corrected model and on-board vehicle info. This permits to provide an accurate identification of slide slip vehicle states without adding any additional sensor to production vehicle's on-board sensors (see Fig. 1).

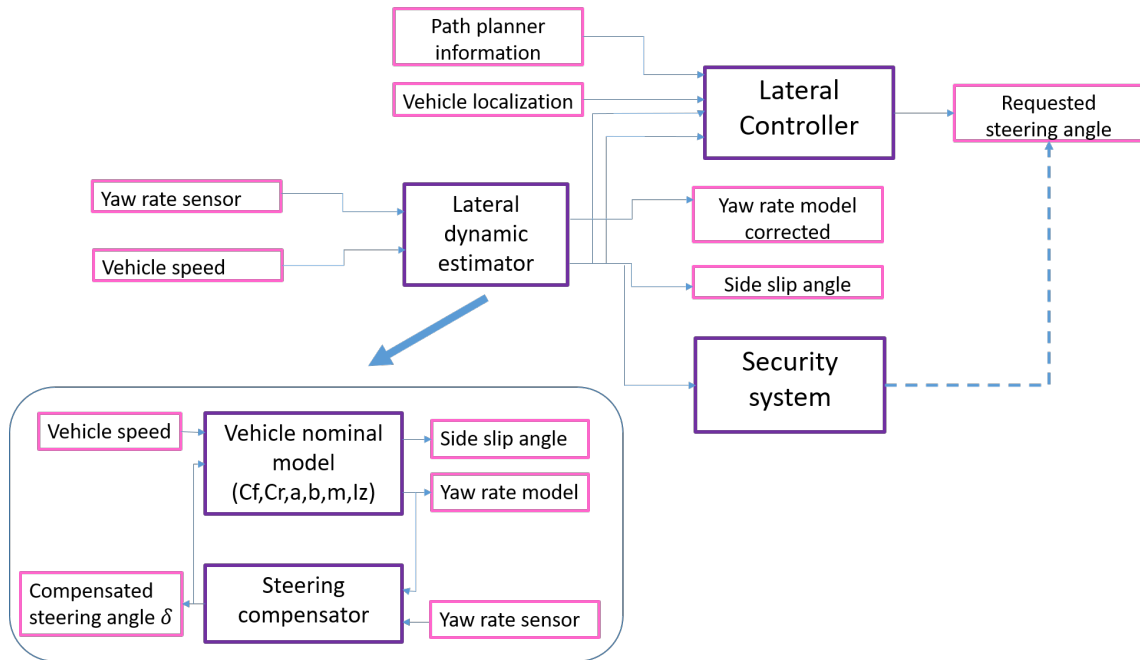


Figure 1. Lateral model compensator

6.11. Perception-adapted controller device for autonomous vehicles

Participants: Imane Mahtout, Fawzi Nashashibi.

Without loss of generality, let us consider a single camera for detecting a preceding vehicle in the road. It is clear that there are two parameters that impact the performance of the perception system:

- 1) the specific algorithm developed to detect and track the objects providing accurate measurement; and
- 2) the physical limitation of the sensor itself. For a camera, the number of pixels limits the resolution of the image so the farther away the vehicle is, the lower the accuracy in its detection. This implies a more inaccurate measurement that will degrade the ego-vehicle performance. From the vehicle response point of view, we cannot expect that a single control device can handle for example a camera-based car-following system for all detected vehicle distances.

For the sake of clarity, Figure 2 shows the speed of a preceding vehicle measured from a ground truth (solid blue line); and the measured speed from an on-board perception system. The speed of the preceding vehicle was computer controlled so we can assure a given response for it. In this example, it follows four consecutive reference speed changes from 0 to 5m/s, and finally to 8m/s. The ego-vehicle equipped with the on-board perception system was following that preceding vehicle at a speed dependent distance (i.e. as any on-the-market ACC system), meaning that the higher the speed, the higher the inter-vehicle distance. One can clearly see how the higher the speed, the higher the inaccuracy of the perception system, the more degraded the measurement.

We worked on a novel control device able to adapt its response to the perception system capabilities, modifying vehicle response accordingly to the level of accuracy of the perception system. This novel idea redefines and extends the capabilities of any ADAS or autonomous vehicle technology, not only because it improves vehicle's performance but also because we can, a-priori, understand the limitations of the full vehicle performance with a complete closed loop analysis from perception to vehicle control. As additional




Figure 2. Perception speed profile measurement

remark, there are ways of dealing with these problems by filtering the perception signal. This causes two problems: control response stability is no longer guaranteed, and it smooths the response but it's not possible to link the full system performance to that filtering. Figure 3 presents an overview of the method. The block Perception set-of-sensor represents the specific embedded perception system. It can consist either in a single sensor or a combination of them. The characterization of the specific perception setup can be done offline, calibrating the system performance accordingly to the on-board sensors. The module Offline calibration will contain a 3-D look-up table with object distance, preceding-vehicle speed, and as a third parameter the desired measurement inaccuracy (it can be either the speed as shown in Figure 2 or any other relevant parameter as the distance, yaw rate...). This offline calibration allows defining specific design parameters for the vehicle performance. Current control systems linked to perception don't consider this inaccuracy when designing the vehicle performance. We here include two different control design criteria blocks. Assuming that we keep the regular controller design that considers perfect measurement from the perception system; the First control design criteria block includes the current production system controller. On the contrary side, we have also included a Second control design criteria block that can be adapted in function of the specific interest of each application. Following with the example on Figure 2, let us assume that we are interested on developing an application between 0 and 10m/s and the inaccuracy of the perception system is the one presented in the plots. Having this in mind, we can design the second controller with the goal of minimizing the impact of that inaccuracy in the vehicle performance. The system also uses as input the real-time perception value coming from the Perception system measurement block (in the case of Figure 2 would provide the speed of the preceding vehicle in real-time). Then, this measurement feeds the Perception-adapted controller block and the Performance degradation block. This last, accordingly to the information from the Offline calibration block, determines the status of the on-board perception system. The output of the Performance degradation block with the output of the first and second design criteria blocks fed the core module of this work: the real-time vehicle performance adaptation module. It is composed by two main blocks: the response corrector block to adapt the vehicle performance and the Perception-adapted controller block that merges both designed controller in a single stable structure.

6.12. Cyberphysical constructs and mobile communications for fully automated networked vehicles

Participant: Gérard Le Lann.

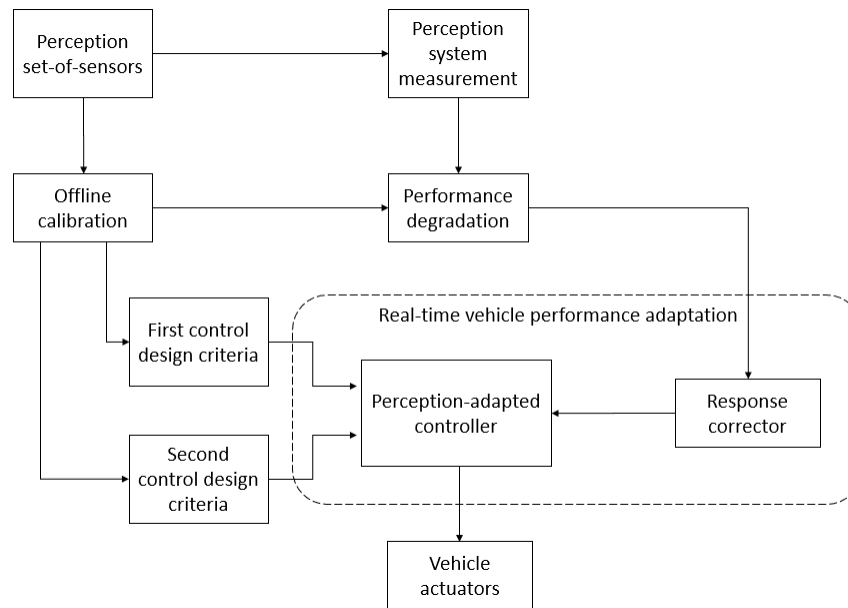


Figure 3. General overview of the block diagram for perception-adapted autonomous vehicle control device.

Safety, privacy, efficiency, and cybersecurity (SPEC) properties are key to the advent of self-forming and self-healing networks of fully automated (driverless) terrestrial vehicles. Such vehicles are referred to as Next-Gen Vehicles (NGVs) in order to avoid confusion with Connected Autonomous Vehicles (CAVs). NGVs prefigure SAE level 5 vehicles. CAVs and NGVs rest on robotics capabilities (sensors, motion control laws, actuators, onboard systems, etc.). CAVs are equipped with V2X (vehicle-to-everything) functionalities based on medium range WiFi radio communications. NGVs will be equipped with CMX (coordinated mobility for X) functionalities, X standing for S, P, E, and C, based on very short range communications (cellular radio and optics).

Work in 2019 has been devoted to defining the CMX framework and to comparing V2X and CMX functionalities. The outputs of this work have been published in [23].

Highest SE (safety and efficiency) is one of the most fundamental goals set to designers of onboard systems. It is surmised that onboard robotics must be supplemented with inter-vehicular communication (IVC) capabilities in order to achieve highest SE properties. Thus the question: which IVC capabilities? In the V2X framework, two distinct sets of IVC capabilities are considered, namely DSRC-V2X (WiFi radio) and C-V2X (4G LTE, 5G, cellular radio). IVC capabilities in the CMX framework encompass cellular radio, VLC and passive optics.

Since V2X functionalities rest on medium range radio communications, they are vulnerable to remote and local cyberattacks (message falsification, masquerading, Sybil attacks, injection of bogus messages, DDoS attacks, etc.). It has been amply demonstrated that such cyberattacks can compromise safety (collisions caused by remote and/or local attackers) as well as efficiency (congested roadways and cities). Furthermore, V2X functionalities break down when radio channels are noisy (messages get lost) or/and jammed (intentional remote and local cyberattacks). Finally, owing to decade-old design decisions, there are no privacy properties with V2X functionalities. For example, every CAV must periodically broadcast messages that carry vehicle-centric characteristics and unencrypted current GNSS space coordinates (referred to as beaconing, frequencies ranging between 1 Hz and 10 Hz. Despite certificate-based pseudonymisation, routes followed by vehicles can

be tracked and communications can be eavesdropped and recorded. Linkage with passengers) personal data is straightforward.

Therefore, in addition to degrading safety and efficiency properties achieved by onboard robotics, V2X functionalities do not meet elementary requirements regarding privacy and cybersecurity. Some proponents of the V2X approach assert that it is impossible to deliver road safety without breaching passengers' privacy. To be valid, that statement should be backed with an impossibility proof. Such a proof has not appeared yet and will never appear for the simple reason that safety and privacy properties can be achieved jointly, by design, proofs given, as demonstrated with the CMX approach.

From a more theoretical perspective, the V2X and the CMX frameworks can be contrasted as follows. Unquestionably, full asynchrony is the appropriate model for representing the vehicular network universe faithfully. Vehicles are started or stopped at arbitrary times, velocities change unpredictably, ditto for lane changes, on-ramp merging, concurrent traversals of intersections and roundabouts, and so on. Onboard processes that are life/safety critical are run in the presence of fortuitous failures, cyberattacks, and concurrency (due to resource sharing). It follows that even if one postulates the existence of finite bounds for process execution durations, it is impossible to assume any a priori knowledge of values taken by those bounds. That is precisely the definition of full asynchrony.

Numerous impossibility results relative to fully asynchronous systems have been published since the late-1970s. For example, problems akin to distributed consensus (terminating reliable broadcast, consistent multi-copied data structures, exact agreement, leader election, etc.) have no solutions in the presence of a single failure, even when communications are assumed to be perfect (no message losses). Since mobile wireless communications are unreliable, those results hold a fortiori in vehicular networks. Obviously, problems that involve termination in computable/predictable time bounds (a real-time property) have no solutions either.

The above-mentioned problems shall be solved in order to provide vehicles and vehicular networks with the SPEC properties. Knowing that solutions exist when considering synchrony models -such as e.g. partial synchrony, timed asynchrony, full synchrony- the challenge is to show how synchrony models could emerge from full asynchrony. This challenge is ignored in the V2X framework. Conclusion: since V2X designs are conducted considering full asynchrony, none of the SPEC properties may hold true.

The CMX framework results from addressing this challenge. NGVs are endowed with CMX functionalities which are based on specific cyberphysical constructs (cells, cohorts, flocks). These constructs serve to instantiate synchrony models within which it is possible to design protocols and algorithms (e.g., deterministic MAC protocols, time-bounded distributed algorithms for message dissemination, approximate agreement, and consensus) that are needed for establishing and proving the SPEC properties, while matching the real vehicular networks universe.

Concepts at the core of the CMX framework (cyberphysical levels, unfalsifiable vehicle profiles, proactive security modules, privacy-preserving naming, etc.) are detailed in [23]. Regarding SE properties, we show how to achieve theoretical absolute safety (no fatalities, no severe injuries) while keeping smallest safe gaps (highest efficiency) in cohort-structured vehicular networks, under assumptions of high coverage. As for PC properties, we show that passengers' privacy cannot be compromised via cyber eavesdropping and/or physical tracking of vehicles. This is due to the fact that messages do not carry vehicle-centric characteristics or GNSS space coordinates. CMX functionalities are shown to be immune to remote cyberattacks. Thanks to optical communications (in addition to very short range cellular radio), they can withstand radio channel jamming. Owing to controlled cohort admission, external local cyberattacks aimed at cohort members are inoperative. Local cyberattacks launched from the inside of a cohort, i.e. by cohort members themselves, can be thwarted. In the unlikely case of success, dishonest members would be involved in those collisions which they create. Conclusion: the only cyberattacks that may compromise safety in cohort-structured vehicular networks are due to irrational attackers.

6.13. Belief propagation inference for traffic prediction

Participant: Jean-Marc Lasgouttes.

This work [36], [35], in collaboration with Cyril Furtlehner (TAU, Inria), deals with real-time prediction of traffic conditions in a urban setting with incomplete data. The main focus is on finding a good way to encode available information (flow, speed, counts,...) in a Markov Random Field, and to decode it in the form of real-time traffic reconstruction and prediction. Our approach relies in particular on the Gaussian belief propagation algorithm.

Through our collaboration with PTV Sistema, we obtained extensive results on large-scale datasets containing 250 to 2000 detectors. The results show very good ability to predict flow variables and a reasonably good performance on speed or occupancy variables. Some element of understanding of the observed performance are given by a careful analysis of the model, allowing to some extent to disentangle modelling bias from intrinsic noise of the traffic phenomena and its measurement process.

This year we worked on code optimization and submitted our work to *Transportation Research: Part C*.

6.14. Stabilization of traffic through cooperative autonomous vehicles

Participants: Guy Fayolle, Carlos Flores, Jean-Marc Lasgouttes.

We investigate in [26] the transfer function emanating from the linearization of a car-following model, when taking into account a driver reaction time. This leads to stability conditions, which are explicitly given. We also show how this reaction time can introduce a *weak string instability*.

This paper is intended as a foundation of a larger work on traffic stabilization by means of a fleet of cooperative automated vehicles. Contrary to some earlier works, our approach is based on a car-following model with reaction-time delay, rather than on a first order fluid model. The continuation of these studies will concern shockwave analysis and adequate traffic-stabilizing control strategies.

6.15. Random walks in orthants and lattice path combinatorics

Participant: Guy Fayolle.

In the second edition of the book [2], original methods were proposed to determine the invariant measure of random walks in the quarter plane with small jumps (size 1), the general solution being obtained via reduction to boundary value problems. In this framework, number of difficult open problems related to lattice path combinatorics are currently being explored, in collaboration with A. Bostan and F. Chyzak (project-team SPECFUN, Inria-Saclay), both from theoretical and computer algebra points of view: concrete computation of the criteria, utilization of differential Galois theory, genus greater than 1 (i.e. when some jumps are of size ≥ 2), etc. A recent topic deals with the connections between simple product-form stochastic networks (so-called *Jackson networks*) and explicit solutions of functional equations for counting lattice walks, see [25].

6.16. Optimization of test case generation for ADAS via Gibbs sampling algorithms

Participant: Guy Fayolle.

Validating Advanced Driver Assistance Systems (ADAS) is a strategic issue, since such systems are becoming increasingly widespread in the automotive field.

But ADAS validation is a complex issue, particularly for camera based systems, because these functions may be facing a very high number of situations that can be considered as infinite. Building at a low cost level a sufficiently detailed campaign is thus very difficult. Indeed, test case generation faces the crucial question of *inherent combinatorial explosion*. An important constraint is to generate *almost all* situations in the most economical way. This task can be considered from two points of view: deterministic via binary search trees, or stochastic via Markov chain Monte Carlo (MCMC) sampling. We choose the latter probabilistic approach described below, which in our opinion seems to be the most efficient one. Typically, the problem is to produce samples of large random vectors, the components of which are possibly dependent and take a finite number of values with some given probabilities. The following flowchart is proposed.

1. In a first step, starting from the simulation graph generated by the toolboxes of MATLAB, we construct a so-called *Markov Random Field (MRF)*. When the parameters are locally dependent, this can be achieved from the user's specifications and by a systematic application of Bayes' formula.
2. Then, to cope with the combinatorial explosion, test cases are produced by implementing (and comparing) various *Gibbs samplers*, which are fruitfully employed for large systems encountered in physics. In particular, we strive to make a compromise between the convergence rate toward equilibrium, the percentage of generated duplicates and the path coverage, keeping in mind that the speed of convergence is exponential, a classical property deduced from the general theory of Markov chains.
3. The generation of rare events by mixing Gibbs samplers, large deviation techniques (LDT) and cross-entropy method is a work in progress.

The French car manufacturer *Groupe PSA* shows a great interest in these methods and has established a contractual collaboration involving ARMINES-Mines ParisTech (Guy Fayolle as associate researcher) and Can Tho University in Vietnam (Pr. Van Ly Tran).

VALDA Project-Team

7. New Results

7.1. Foundations of data management

We obtained a number of results on the foundations of data management, i.e., in database theory.

We worked on **knowledge bases**. In our work a knowledge base consists of an incomplete database together with a set of existential rules. We investigated the problem of query answering: computing the answers that are logically entailed from the knowledge base. This brings to light the fundamental chase tool, and its different variants that have been proposed in the literature. We studied the problem of chase termination, which has applications beyond query answering, and studied its complexity for restricted but useful classes of existential rules [27].

We worked on **data integration**. In our scenario a user can access data sitting in multiple sources by means of queries over a global schema, related to the sources via mappings. Data sources often contain sensitive information, and thus an analysis is needed to verify that a schema satisfies a privacy policy, given as a set of queries whose answers should not be accessible to users. We show that source constraints can have a dramatic impact on disclosure analysis [22]. Another work related to data integration is [16], where we connect the problem of answering queries under limited accesses (e.g., using Web forms) to two foundational issues: containment of Monadic datalog (MDL) programs, and containment problems involving regular tree languages. In particular, we establish a 2EXPTIME lower bound on the problem of containment of a MDL program into a conjunctive query, resolving an open problem from the early 1990s.

We also considered some other foundational topics, further from core database topics. In [18], we establish bounds on the height of maximal finite towers (a *tower* is a sequence of words alternating between two languages in such a way that every word is a subsequence of the following word) between two regular languages. In [17], we present an online $O(\sigma|y|)$ -time algorithm for finding approximate occurrences of a word x within a word y , where σ is the alphabet size.

Note that two other works in this theme will be described in the 2020 activity report, as they are published in 2020 conferences [25], [26].

7.2. Uncertainty and provenance of data

We have a strong focus on the uncertainty and provenance in databases. See [20] for a high-level introduction to the area.

In [15], we investigate the use of knowledge compilation, i.e., obtaining compact circuit-based representations of functions, for (Boolean) provenance. Some width parameters of the circuit, such as bounded treewidth or pathwidth, can be leveraged to convert the circuit to structured classes, e.g., deterministic structured NNFs (d-SDNNFs) or OBDDs. In [14], we investigate parameterizations of both database instances and queries that make query evaluation fixed-parameter tractable in combined complexity. We show that clique-frontier-guarded Datalog with stratified negation (CFG-Datalog) enjoys bilinear-time evaluation on structures of bounded treewidth for programs of bounded rule size. Such programs capture in particular conjunctive queries with simplicial decompositions of bounded width, guarded negation fragment queries of bounded CQ-rank, or two-way regular path queries. Our result is shown by translating to alternating two-way automata, whose semantics is defined via cyclic provenance circuits (cycluits) that can be tractably evaluated.

In previous work [39], [40], we have shown that the only restrictions to database instances that make probabilistic query evaluation tractable for a large class of queries is that of having a small treewidth. In [28], [32], we provide the first large-scale experimental study of treewidth and tree decompositions of real-world database instances (25 datasets from 8 different domains, with sizes ranging from a few thousand to a few million vertices). The goal is to determine which data, if any, has reasonably low treewidth. We also show that, even when treewidth is high, using partial tree decompositions can result in data structures that can assist algorithms.

To conclude on provenance management, in [23], [24], after investigating the complexity of satisfiability and query answering for attributed DL-LiteR ontologies, we propose a new semantics, based on provenance semirings, for integrating provenance information with query answering. Finally, we establish complexity results for satisfiability and query answering under this semantics.

We also consider **other notions of incompleteness**, such as in [13], where we study the complexity of query evaluation for databases whose relations are partially ordered; the problem commonly arises when combining or transforming ordered data from multiple sources. We focus on queries in a useful fragment of SQL, namely positive relational algebra with aggregates, whose bag semantics we extend to the partially ordered setting. Our semantics leads to the study of two main computational problems: the possibility and certainty of query answers. We show that these problems are respectively NP-complete and coNP-complete, but identify tractable cases depending on the query operators or input partial orders.

Finally, we also consider uncertainty through another angle, that of learning in a dynamic environment, using techniques from **reinforcement learning** and the **multi-armed bandit** field.

In [19], we tackle the problem of *influence maximization*: finding influential users, or nodes, in a graph so as to maximize the spread of information. We study a highly generic version of influence maximization, one of optimizing influence campaigns by sequentially selecting “spread seeds” from a set of influencers, a small subset of the node population, under the hypothesis that, in a given campaign, previously activated nodes remain persistently active. We introduce an estimator on the influencers’ remaining potential – the expected number of nodes that can still be reached from a given influencer – and justify its strength to rapidly estimate the desired value, relying on real data gathered from Twitter. We then describe a novel algorithm, GT-UCB, relying on probabilistic upper confidence bounds on the remaining potential.

In [21], we propose a Bayesian information-geometric approach to the exploration-exploitation trade-off in stochastic multi-armed bandits. The uncertainty on reward generation and belief is represented using the manifold of joint distributions of rewards and beliefs. Accumulated information is summarised by the barycentre of joint distributions, the pseudobelief-reward. While the pseudobelief-reward facilitates information accumulation through exploration, another mechanism is needed to increase exploitation by gradually focusing on higher rewards, the pseudobelief-focal-reward. Our resulting algorithm, BelMan, alternates between projection of the pseudobelief-focal-reward onto belief-reward distributions to choose the arm to play, and projection of the updated belief-reward distributions onto the pseudobelief-focal-reward.

In [29], we consider another form of bandits, *linear bandits*, in which the available actions correspond to arbitrary context vectors whose associated rewards follow a non-stationary linear regression model. In this setting, the unknown regression parameter is allowed to vary in time. To address this problem, we propose D-LinUCB, a novel optimistic algorithm based on discounted linear regression, where exponential weights are used to smoothly forget the past.

7.3. Web data management

We finally describe research more oriented towards applications.

The PhD of Karima Rafes [11] dealt with **semantic knowledge bases** and their applications to the management of scientific data, through the development of the LinkedWiki platform. Another practical work on semantic knowledge bases is [30], where we show how the edit history of a knowledge base can help correct constraint violations.

Finally, we investigate **transparency and bias** in data management and artificial intelligence. [12] presents to the data management community the challenges raised by new regulatory frameworks in this area. In [31], we discuss the possibility for artificial intelligence systems to be used in the practice of law.

WILLOW Team

7. New Results

7.1. 3D object and scene modeling, analysis, and retrieval

7.1.1. Learning joint reconstruction of hands and manipulated objects

Participants: Yana Hasson, Gül Varol, Dimitrios Tzionas, Igor Kalevatykh, Michael Black, Ivan Laptev, Cordelia Schmid.

Estimating hand-object manipulations is essential for interpreting and imitating human actions. Previous work has made significant progress towards reconstruction of hand poses and object shapes in isolation. Yet, reconstructing hands and objects during manipulation is a more challenging task due to significant occlusions of both the hand and object. While presenting challenges, manipulations may also simplify the problem since the physics of contact restricts the space of valid hand-object configurations. For example, during manipulation, the hand and object should be in contact but not interpenetrate. In [14] we regularize the joint reconstruction of hands and objects with manipulation constraints. We present an end-to-end learnable model that exploits a novel contact loss that favors physically plausible hand-object constellations. Our approach improves grasp quality metrics over baselines, using RGB images as input. To train and evaluate the model, we also propose a new large-scale synthetic dataset, ObMan, with hand-object manipulations. We demonstrate the transferability of ObMan-trained models to real data. Figure 1 presents some example results.

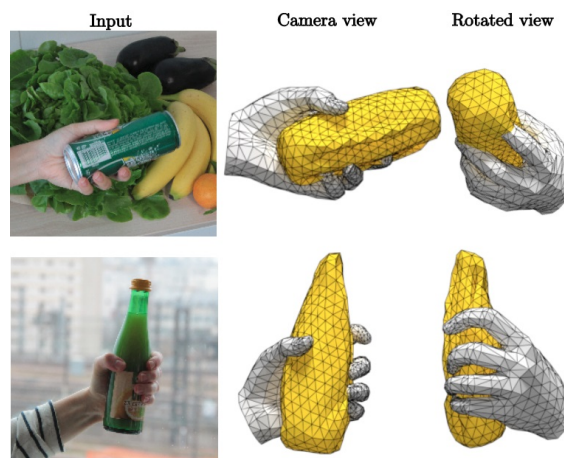


Figure 1. Our method jointly reconstructs hand and object meshes from a monocular RGB image. Note that the model generating the predictions for the above images, which we captured with an ordinary camera, was trained only on images from our synthetic dataset, ObMan.

7.1.2. D2-Net: A Trainable CNN for Joint Detection and Description of Local Features

Participants: Mihai Dusmanu, Ignacio Rocco, Tomas Pajdla, Marc Pollefeys, Josef Sivic, Akihiko Torii, Torsten Sattler.

In [13], we address the problem of finding reliable pixel-level correspondences under difficult imaging conditions. We propose an approach where a single convolutional neural network plays a dual role: It is simultaneously a dense feature descriptor and a feature detector, as illustrated in Figure 2. By postponing the detection to a later stage, the obtained keypoints are more stable than their traditional counterparts based on early detection of low-level structures. We show that this model can be trained using pixel correspondences extracted from readily available large-scale SfM reconstructions, without any further annotations. The proposed method obtains state-of-the-art performance on both the difficult Aachen Day-Night localization dataset and the InLoc indoor localization benchmark, as well as competitive performance on other benchmarks for image matching and 3D reconstruction.




Figure 2. A feature extraction CNN \mathcal{F} is used to extract feature maps that play a dual role: (i) local descriptors d_{ij} are simply obtained by traversing all the n feature maps D^k at a spatial position (i, j) ; (ii) detections are obtained by performing a non-local-maximum suppression on a feature map followed by a non-maximum suppression across each descriptor.

7.1.3. Is This The Right Place? Geometric-Semantic Pose Verification for Indoor Visual Localization

Participants: Hajime Taira, Ignacio Rocco, Jiri Sedlar, Masatoshi Okutomi, Josef Sivic, Tomas Pajdla, Torsten Sattler, Akihiko Torii.

Visual localization in large and complex indoor scenes, dominated by weakly textured rooms and repeating geometric patterns, is a challenging problem with high practical relevance for applications such as Augmented Reality and robotics. To handle the ambiguities arising in this scenario, a common strategy is, first, to generate multiple estimates for the camera pose from which a given query image was taken. The pose with the largest geometric consistency with the query image, e.g., in the form of an inlier count, is then selected in a second stage. While a significant amount of research has concentrated on the first stage, there is considerably less work on the second stage. In [21], we thus focus on pose verification. We show that combining different modalities, namely appearance, geometry, and semantics, considerably boosts pose verification and consequently pose accuracy, as illustrated in Figure 3. We develop multiple hand-crafted as well as a trainable approach to join into the geometric-semantic verification and show significant improvements over state-of-the-art on a very challenging indoor dataset.

7.1.4. An Efficient Solution to the Homography-Based Relative Pose Problem With a Common Reference Direction

Participants: Yaqing Ding, Jian Yang, Jean Ponce, Hui Kong.

In [12], we propose a novel approach to two-view minimal-case relative pose problems based on homography with a common reference direction. We explore the rank-1 constraint on the difference between the Euclidean homography matrix and the corresponding rotation, and propose an efficient two-step solution for solving both the calibrated and partially calibrated (unknown focal length) problems. We derive new 3.5-point, 3.5-point, 4-point solvers for two cameras such that the two focal lengths are unknown but equal, one of them is unknown,

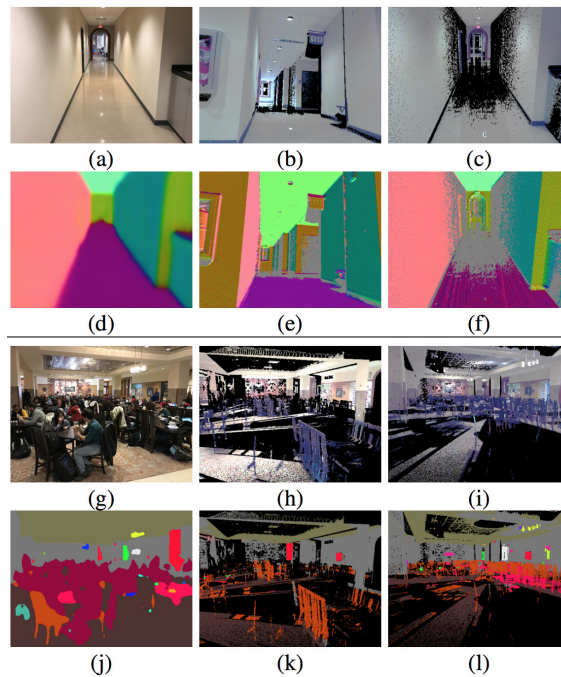


Figure 3. Given a set of camera pose estimates for a query image (a, g), we seek to identify the most accurate estimate. (b, h) Due to severe occlusion and weak textures, a state-of-the-art method fails to identify the correct camera pose. To overcome those difficulties, we use several modalities along with visual appearance: (top) surface normals and (bottom) semantics. (c, i) Our approach verifies the estimated pose by comparing the semantics and surface normals extracted from the query (d, j) and database (f, l).

and both are unknown and possibly different, respectively. We present detailed analyses and comparisons with existing 6-and 7-point solvers, including results with smart phone images.

7.1.5. *Coordinate-Free Carlsson-Weinshall Duality and Relative Multi-View Geometry*

Participants: Matthew Trager, Martial Hebert, Jean Ponce.

In [23], we present a coordinate-free description of Carlsson-Weinshall duality between scene points and camera pinholes and use it to derive a new characterization of primal/dual multi-view geometry. In the case of three views, a particular set of reduced trilinearities provide a novel parameterization of camera geometry that, unlike existing ones, is subject only to very simple internal constraints. These trilinearities lead to new “quasi-linear” algorithms for primal and dual structure from motion. We include some preliminary experiments with real and synthetic data.

7.1.6. *Build your own hybrid thermal/EO camera for autonomous vehicle*

Participants: Yigong Zhang, Yicheng Gao, Shuo Gu, Yubin Guo, Minghao Liu, Zezhou Sun, Zhixing Hou, Hang Yang, Ying Wang, Jian Yang, Jean Ponce, Hui Kong.

In [24], we propose a novel paradigm to design a hybrid thermal/EO (Electro-Optical or visible-light) camera, whose thermal and RGB frames are pixel-wisely aligned and temporally synchronized. Compared with the existing schemes, we innovate in three ways in order to make it more compact in dimension, and thus more practical and extendable for real-world applications. The first is a redesign of the structure layout of the thermal and EO cameras. The second is on obtaining a pixel-wise spatial registration of the thermal and RGB frames by a coarse mechanical adjustment and a fine alignment through a constant homography warping. The third innovation is on extending one single hybrid camera to a hybrid camera array, through which we can obtain wide-view spatially aligned thermal, RGB and disparity images simultaneously. The experimental results show that the average error of spatial-alignment of two image modalities can be less than one pixel. Some results of our method are illustrated in Figure 4.

7.2. Category-level object and scene recognition

7.2.1. *Detecting unseen visual relations using analogies*

Participants: Julia Peyre, Ivan Laptev, Cordelia Schmid, Josef Sivic.

In [19], we seek to detect visual relations in images of the form of triplets $t = (\text{subject}, \text{predicate}, \text{object})$, such as “person riding dog”, where training examples of the individual entities are available but their combinations are unseen at training. This is an important set-up due to the combinatorial nature of visual relations: collecting sufficient training data for all possible triplets would be very hard. The contributions of this work are three-fold. First, we learn a representation of visual relations that combines (i) individual embeddings for subject, object and predicate together with (ii) a visual phrase embedding that represents the relation triplet. Second, we learn how to transfer visual phrase embeddings from existing training triplets to unseen test triplets using analogies between relations that involve similar objects. Third, we demonstrate the benefits of our approach on three challenging datasets : on HICO-DET, our model achieves significant improvement over a strong baseline for both frequent and unseen triplets, and we observe similar improvement for the retrieval of unseen triplets with out-of- vocabulary predicates on the COCO-a dataset as well as the challenging unusual triplets in the UnRel dataset. Figure 5 presents an illustration of the approach.

7.2.2. *SFNet: Learning Object-aware Semantic Correspondence*

Participants: Junghyup Lee, Dohyung Kim, Jean Ponce, Bumsub Ham.




Figure 4. Results of alignment between the thermal and RGB frames of three sets of hybrid cameras before and after homography warping, respectively. (a), (b) and (c) are the alignment results before the homography warping, respectively. In each sub-figure, the layout of images is arranged as follows. Top-left: the aligned RGB image. Top-middle and bottom-left: the same aligned thermal image. Bottom-middle: the fusion image. (d), (e) and (f) are the alignment results after the homography warping, respectively. Likewise, the layout of images in each sub-figure is the same as those of (a), (b) and (c). To show the effect of homography rectification, we have overlaid red dotted lines horizontally and vertically onto the each sub-figure. In addition, the right column of each sub-figure zooms in four selected image regions to help us to view the warping result.

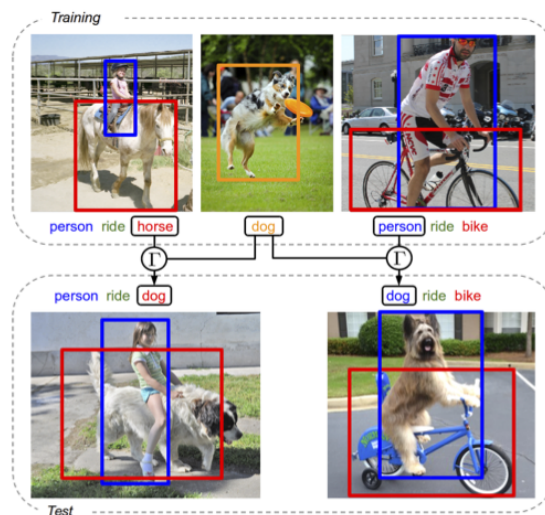


Figure 5. Illustration of transfer by analogy. We transfer visual representations of relations seen in the training set such as “person ride horse” to represent new unseen relations in the test set such as “person ride dog”.

In [15], we address the problem of semantic correspondence, that is, establishing a dense flow field between images depicting different instances of the same object or scene category. We propose to use images annotated with binary foreground masks and subjected to synthetic geometric deformations to train a convolutional neural network (CNN) for this task. Using these masks as part of the supervisory signal offers a good compromise between semantic flow methods, where the amount of training data is limited by the cost of manually selecting point correspondences, and semantic alignment ones, where the regression of a single global geometric transformation between images may be sensitive to image-specific details such as background clutter. We propose a new CNN architecture, dubbed SFNet, which implements this idea. It leverages a new and differentiable version of the argmax function for end-to-end training, with a loss that combines mask and flow consistency with smoothness terms. Experimental results demonstrate the effectiveness of our approach, which significantly outperforms the state of the art on standard benchmarks. Figure 6 presents an illustration of the approach.




Figure 6. We use pairs of warped foreground masks obtained from a single image (left) as a supervisory signal to train our model. This allows us to establish object-aware semantic correspondences across images depicting different instances of the same object or scene category (right). No masks are required at test time.

7.2.3. Hyperpixel Flow: Semantic Correspondence with Multi-layer Neural Features

Participants: Juhong Min, Jongmin Kim, Jean Ponce, Minsu Cho.

In [17], we establish visual correspondences under large intra-class variations requires analyzing images at different levels, from features linked to semantics and context to local patterns, while being invariant to instance-specific details. To tackle these challenges, we represent images by "hyper-pixels" that leverage a small number of relevant features selected among early to late layers of a convolutional neural network. Taking advantage of the condensed features of hyperpixels, we develop an effective real-time matching algorithm based on Hough geometric voting. The proposed method, hyperpixel flow, sets a new state of the art on three standard benchmarks as well as a new dataset, SPair-71k, which contains a significantly larger number of image pairs than existing datasets, with more accurate and richer annotations for in-depth analysis. Figure 7 presents an illustration of the approach.

7.2.4. Exploring Weight Symmetry in Deep Neural Networks

Participants: Xu Shell Hu, Sergey Zagoruyko, Nikos Komodakis.

In [7], we propose to impose symmetry in neural network parameters to improve parameter usage and make use of dedicated convolution and matrix multiplication routines. Due to significant reduction in the number of parameters as a result of the symmetry constraints, one would expect a dramatic drop in accuracy. Surprisingly,

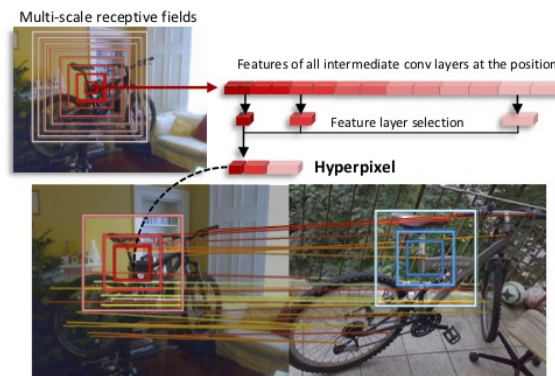


Figure 7. Hyperpixel flow. Top: The hyperpixel is a multi-layer pixel representation created with selected levels of features optimized for semantic correspondence. It provides multi-scale features, resolving local ambiguities. Bottom: The proposed method, hyperpixel flow, establishes dense correspondences in real time using hyperpixels.

we show that this is not the case, and, depending on network size, symmetry can have little or no negative effect on network accuracy, especially in deep overparameterized networks. We propose several ways to impose local symmetry in recurrent and convolutional neural networks, and show that our symmetry parameterizations satisfy universal approximation property for single hidden layer networks. We extensively evaluate these parameterizations on CIFAR, ImageNet and language modeling datasets, showing significant benefits from the use of symmetry. For instance, our ResNet-101 with channel-wise symmetry has almost 25% less parameters and only 0.2% accuracy loss on ImageNet.

7.2.5. Bilinear image translation for temporal analysis of photo collections

Participants: Théophile Dalens, Mathieu Aubry, Josef Sivic.

In [5], we propose an approach for analyzing unpaired visual data annotated with time stamps by generating how images would have looked like if they were from different times. To isolate and transfer time dependent appearance variations, we introduce a new trainable bilinear factor separation module. We analyze its relation to classical factored representations and concatenation-based auto-encoders. We demonstrate this new module has clear advantages compared to standard concatenation when used in a bottleneck encoder-decoder convolutional neural network architecture. We also show that it can be inserted in a recent adversarial image translation architecture, enabling the image transformation to multiple different target time periods using a single network. We apply our model to a challenging collection of more than 13,000 cars manufactured between 1920 and 2000 and a dataset of high school yearbook portraits from 1930 to 2009, as illustrated in Figure 8. This allows us, for a given new input image, to generate a "history-lapse video" revealing changes over time by simply varying the target year. We show that by analyzing the generated history-lapse videos we can identify object deformations across time, extracting interesting changes in visual style over decades.

7.3. Image restoration, manipulation and enhancement

7.3.1. Deformable Kernel Networks for Joint Image Filtering

Participants: Beomjun Kim, Jean Ponce, Bumsu Ham.



Figure 8. Our method takes as input an image of an object (in green), such as a car, and generates what it would have looked like in another time-period (in blue). Each row shows temporal translation for a different input car image (in green). The translation model is trained on an unpaired dataset of cars with time stamps. We show that analyzing changes between the generated images reveal structural deformations in car shape and appearance over time.

Joint image filters are used to transfer structural details from a guidance picture used as a prior to a target image, in tasks such as enhancing spatial resolution and suppressing noise. Previous methods based on convolutional neural networks (CNNs) combine nonlinear activations of spatially-invariant kernels to estimate structural details and regress the filtering result. In this paper, we instead learn explicitly sparse and spatially-variant kernels. In [28], we propose a CNN architecture and its efficient implementation, called the deformable kernel network (DKN), that outputs sets of neighbors and the corresponding weights adaptively for each pixel. The filtering result is then computed as a weighted average. We also propose a fast version of DKN that runs about four times faster for an image of size 640×480 . We demonstrate the effectiveness and flexibility of our models on the tasks of depth map upsampling, saliency map upsampling, cross-modality image restoration, texture removal, and semantic segmentation. In particular, we show that the weighted averaging process with sparsely sampled 3×3 kernels outperforms the state of the art by a significant margin.

7.3.2. Revisiting Non Local Sparse Models for Image Restoration

Participants: Bruno Lecouat, Jean Ponce, Julien Mairal.

In [29], we propose a differentiable algorithm for image restoration inspired by the success of sparse models and self-similarity priors for natural images. Our approach builds upon the concept of joint sparsity between groups of similar image patches, and we show how this simple idea can be implemented in a differentiable architecture, allowing end-to-end training. The algorithm has the advantage of being interpretable, performing sparse decompositions of image patches, while being more parameter efficient than recent deep learning methods. We evaluate our algorithm on grayscale and color denoising, where we achieve competitive results, and on demosaicking, where we outperform the most recent state-of-the-art deep learning model with 47 times less parameters and a much shallower architecture. Figure 9 shows results of the proposed approach.

7.4. Human activity capture and classification

7.4.1. Video Face Clustering with Unknown Number of Clusters

Participants: Makarand Tapaswi, Marc T. Law, Sanja Fidler.



Figure 9. Demosaicking result obtained by our method. Top right: Ground truth. Middle: Image demosaicked with our sparse coding baseline without non-local prior. Bottom: demosaicking with sparse coding and non-local prior. The reconstruction does not exhibit any artefact on this image which is notoriously difficult for demosaicking.

Understanding videos such as TV series and movies requires analyzing who the characters are and what they are doing. We address the challenging problem of clustering face tracks based on their identity. Different from previous work in this area, we choose to operate in a realistic and difficult setting where: (i) the number of characters is not known a priori; and (ii) face tracks belonging to minor or background characters are not discarded.

To this end, we propose Ball Cluster Learning (BCL), a supervised approach to carve the embedding space into balls of equal size, one for each cluster (see Figure 10). The learned ball radius is easily translated to a stopping criterion for iterative merging algorithms. This gives BCL the ability to estimate the number of clusters as well as their assignment, achieving promising results on commonly used datasets. We also present a thorough discussion of how existing metric learning literature can be adapted for this task. This work has been published in [22].

7.4.2. Cross-task weakly supervised learning from instructional videos

Participants: Dimitri Zhukov, Jean-Baptiste Alayrac, Ramazan Gokberk Cinbis, David Fouhey, Ivan Laptev, Josef Sivic.

In [25], we investigate learning visual models for the steps of ordinary tasks using weak supervision via instructional narrations and an ordered list of steps instead of strong supervision via temporal annotations. At the heart of our approach is the observation that weakly supervised learning may be easier if a model shares components while learning different steps: “pour egg” should be trained jointly with other tasks involving “pour” and “egg”. We formalize this in a component model for recognizing steps and a weakly supervised learning framework that can learn this model under temporal constraints from narration and the list of steps. Past data does not permit systematic studying of sharing and so we also gather a new dataset, CrossTask, aimed at assessing cross-task sharing. Our experiments demonstrate that sharing across tasks improves performance, especially when done at the component level and that our component model can parse previously unseen tasks by virtue of its compositionality. Figure 11 illustrates the idea of sharing step components between different tasks.




Figure 10. Ball Cluster Learning carves the feature space into balls of equal radius. The number of samples in the cluster does not affect the ball radius or minimum separation to other balls.



Figure 11. Our method begins with a collection of tasks, each consisting of an ordered list of steps and a set of instructional videos from YouTube. It automatically discovers both where the steps occur and what they look like. To do this, it uses the order, narration and commonalities in appearance across tasks (e.g., the appearance of pour in both making pancakes and making meringue).

7.4.3. Leveraging the Present to Anticipate the Future in Videos

Participants: Antoine Miech, Ivan Laptev, Josef Sivic, Heng Wang, Lorenzo Torresani, Du Tran.

Anticipating actions before they are executed is crucial for a wide range of practical applications including autonomous driving and the moderation of live video streaming. While most prior work in this area requires partial observation of executed actions, in the paper we focus on anticipating actions seconds before they start (see Figure 12). Our proposed approach is the fusion of a purely anticipatory model with a complementary model constrained to reason about the present. In particular, the latter predicts present action and scene attributes, and reasons about how they evolve over time. By doing so, we aim at modeling action anticipation at a more conceptual level than directly predicting future actions. Our model outperforms previously reported methods on the EPIC-KITCHENS and Breakfast datasets. This paper was presented at the CVPR 2019 precognition workshop [34] and ranked second at the EPIC-KITCHENS action anticipation challenge.

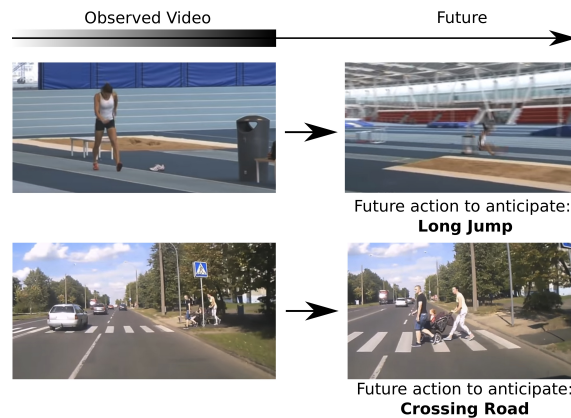


Figure 12. Examples of action anticipation in which the goal is to anticipate future actions in videos seconds before they are performed.

7.4.4. HowTo100M: Learning a Text-Video Embedding by Watching Hundred Million Narrated Video Clips

Participants: Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, Josef Sivic.

Learning text-video embeddings usually requires a dataset of video clips with manually provided captions. However, such datasets are expensive and time consuming to create and therefore difficult to obtain on a large scale. In this work, we propose instead to learn such embeddings from video data with readily available natural language annotations in the form of automatically transcribed narrations (see Figure 13). The contributions of this work are three-fold. First, we introduce HowTo100M: a large-scale dataset of 136 million video clips sourced from 1.22M narrated instructional web videos depicting humans performing and describing over 23k different visual tasks. Our data collection procedure is fast, scalable and does not require any additional manual annotation. Second, we demonstrate that a text-video embedding trained on this data leads to state-of-the-art results for text-to-video retrieval and action localization on instructional video datasets such as YouCook2 or CrossTask. Finally, we show that this embedding transfers well to other domains: fine-tuning on generic Youtube videos (MSR-VTT dataset) and movies (LSMDC dataset) outperforms models trained on these datasets alone. This work was presented at ICCV 2019 [16].




Figure 13. We learn a joint text-video embedding by watching millions of narrated video clips of people performing diverse visual tasks. The learned embedding transfers well to other instructional and non-instructional text-video datasets.

7.4.5. Are Large-Scale 3D Models Really Necessary for Accurate Visual Localization?

Participants: Akihiko Torii, Hajime Taira, Josef Sivic, Marc Pollefeys, Masatoshi Okutomi, Tomas Pajdla, Torsten Sattler.

Accurate visual localization is a key technology for autonomous navigation. 3D structure-based methods, as illustrated in Figure 14, employ 3D models of the scene to estimate the full 6 degree-of-freedom (DOF) pose of a camera very accurately. However, constructing (and extending) large-scale 3D models is still a significant challenge. In contrast, 2D image retrieval-based methods only require a database of geo-tagged images, which is trivial to construct and to maintain. They are often considered inaccurate since they only approximate the positions of the cameras. Yet, the exact camera pose can theoretically be recovered when enough relevant database images are retrieved. In [8], we demonstrate experimentally that large-scale 3D models are not strictly necessary for accurate visual localization. We create reference poses for a large and challenging urban dataset. Using these poses, we show that combining image-based methods with local reconstructions results in a higher pose accuracy compared to state-of-the-art structure-based methods, albeit at higher run-time costs. We show that some of these run-time costs can be alleviated by exploiting known database image poses. Our results suggest that we might want to reconsider the need for large-scale 3D models in favor of more local models, but also that further research is necessary to accelerate the local reconstruction process.

7.4.6. End-to-End Learning of Visual Representations from Uncurated Instructional Videos

Participants: Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, Andrew Zisserman.

Annotating videos is cumbersome, expensive and not scalable. Yet, many strong video models still rely on manually annotated data. With the recent introduction of the HowTo100M dataset, narrated videos now offer the possibility of learning video representations without manual supervision. In this work we propose a new learning approach, MIL-NCE, capable of addressing misalignments inherent to narrated videos (see Figure 15). With this approach we are able to learn strong video representations from scratch, without the need for any manual annotation. We evaluate our representations on a wide range of four downstream tasks over eight datasets: action recognition (HMDB-51, UCF-101, Kinetics-700), text-to-video retrieval (YouCook2, MSR-VTT), action localization (YouTube-8M Segments, CrossTask) and action segmentation (COIN). Our method outperforms all published self-supervised approaches for these tasks as well as several fully supervised baselines. This preprint [32] is currently under review.




Figure 14. The state-of-the-art for large-scale visual localization. 2D image-based methods (bottom) use image retrieval and return the pose of the most relevant database image. 3D structure-based methods (top) use 2D-3D matches against a 3D model for camera pose estimation. Both approaches have been developed largely independently of each other and never compared properly before.




Figure 15. We describe an efficient approach to learn visual representations from highly misaligned and noisy narrations automatically extracted from instructional videos. Our video representations are learnt from scratch without relying on any manually annotated visual dataset yet outperform all self-supervised and many fully-supervised methods on several video recognition benchmarks.

7.4.7. Synthetic Humans for Action Recognition from Unseen Viewpoints

Participants: Gul Varol, Ivan Laptev, Cordelia Schmid, Andrew Zisserman.

In [35], the goal is to improve the performance of human action recognition for viewpoints unseen during training by using synthetic training data. Although synthetic data has been shown to be beneficial for tasks such as human pose estimation, its use for RGB human action recognition is relatively unexplored. We make use of the recent advances in monocular 3D human body reconstruction from real action sequences to automatically render synthetic training videos for the action labels. We make the following contributions: (i) we investigate the extent of variations and augmentations that are beneficial to improving performance at new viewpoints. We consider changes in body shape and clothing for individuals, as well as more action relevant augmentations such as non-uniform frame sampling, and interpolating between the motion of individuals performing the same action; (ii) We introduce a new dataset, SURREACT, that allows supervised training of spatio-temporal CNNs for action classification; (iii) We substantially improve the state-of-the-art action recognition performance on the NTU RGB+D and UESTC standard human action multi-view benchmarks; Finally, (iv) we extend the augmentation approach to in-the-wild videos from a subset of the Kinetics dataset to investigate the case when only one-shot training data is available, and demonstrate improvements in this case as well. Figure 16 presents an illustration of the approach.

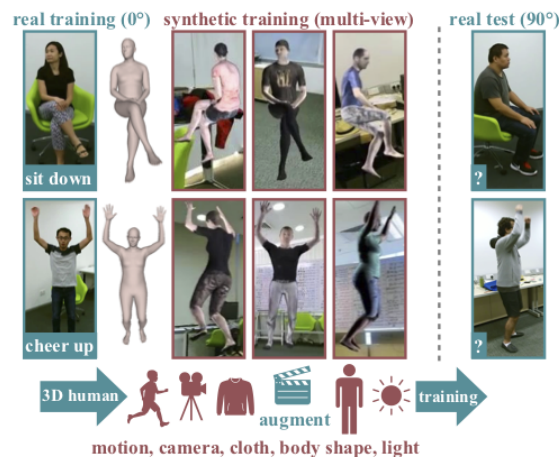


Figure 16. We estimate 3D shape from real videos and automatically render synthetic videos with action labels. We explore various augmentations for motions, viewpoints, and appearance. Training temporal CNNs with this data significantly improves the action recognition from unseen viewpoints.

7.5. Learning embodied representations and robotics

7.5.1. Roboticians and Reporters

Participants: Celine Pieters, Emmanuelle Danblon, Jean-Paul Laumond.

This paper reports on an experiment organized at the Cité des Sciences et de l'Industrie (CSI) of Paris in order to assess the importance of language in the representation and the integration of robots into the human culture. The experiment gathered specialized reporters and experts in robotics around a practical exercise of rhetoric. The objective of this work is to show that rhetoric is not a matter of communication, but a technique that allows to better understand the way roboticians understand their own discipline.

7.5.2. Robots

Participants: Jean-Paul Laumond, Denis Vidal.

What is a robot? How does it work? How is research progressing, what are the challenges and the economic and social questions posed by robotics in the twenty-first century? Today, as robots are becoming increasingly present in our professional, public and private lives, it is vital to understand their technological capabilities. We must more fully comprehend how they can help us and master their uses. Robots continue to fascinate us but our idea of them, stemming from literature and cinema, is often a purely imaginary one. This illustrated book accompanies the Robots exhibition at the Cité des sciences et de l'industrie. Figure 17 presents the front page of the exhibition.



Figure 17. Front page of the permanent exhibition at Cité des Sciences et de l'Industrie about Robotics.

7.5.3. Learning to Augment Synthetic Images for Sim2Real Policy Transfer

Participants: Alexander Pashevich, Robin Strudel, Igor Kalevtykh, Ivan Laptev, Cordelia Schmid.

Vision and learning have made significant progress that could improve robotics policies for complex tasks and environments. Learning deep neural networks for image understanding, however, requires large amounts of domain-specific visual data. While collecting such data from real robots is possible, such an approach limits the scalability as learning policies typically requires thousands of trials. In this paper [18], we attempt to learn manipulation policies in simulated environments. Simulators enable scalability and provide access to the underlying world state during training. Policies learned in simulators, however, do not transfer well to real scenes given the domain gap between real and synthetic data. We follow recent work on domain randomization and augment synthetic images with sequences of random transformations. Our main contribution is to optimize the augmentation strategy for sim2real transfer and to enable domain-independent policy learning. We design an efficient search for depth image augmentations using object localization as a proxy task. Given the resulting sequence of random transformations, we use it to augment synthetic depth images during policy learning. Our augmentation strategy is policy-independent and enables policy learning with no real images. We demonstrate our approach to significantly improve accuracy on three manipulation tasks evaluated on a real robot. Figure 18 presents an illustration of the approach.

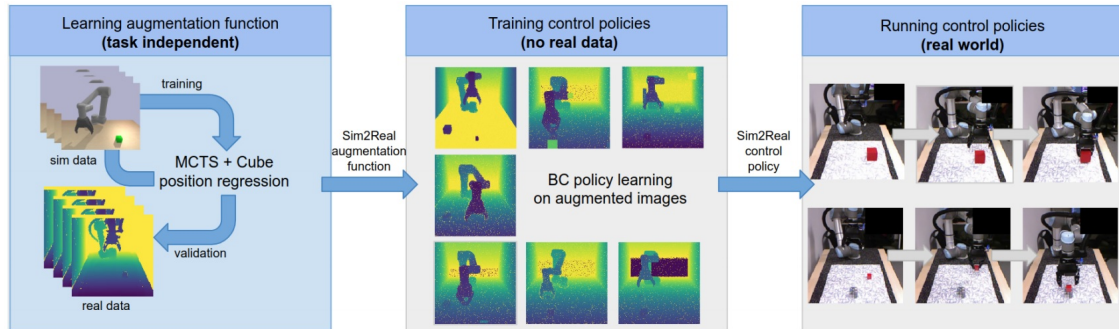


Figure 18. Overview of the method. Our contribution is the policy-independent learning of depth image augmentations (left). The resulting sequence of augmentations is applied to synthetic depth images while learning manipulation policies in a simulator (middle). The learned policies are directly applied to real robot scenes without finetuning on real images (right).

7.5.4. Learning to combine primitive skills: A step towards versatile robotic manipulation

Participants: Robin Strudel, Alexander Pashevich, Igor Kalevatykh, Ivan Laptev, Josef Sivic, Cordelia Schmid.

Manipulation tasks such as preparing a meal or assembling furniture remain highly challenging for robotics and vision. Traditional task and motion planning (TAMP) methods can solve complex tasks but require full state observability and are not adapted to dynamic scene changes. Recent learning methods can operate directly on visual inputs but typically require many demonstrations and/or task-specific reward engineering. In this paper [20], we aim to overcome previous limitations and propose a reinforcement learning (RL) approach to task planning that learns to combine primitive skills. First, compared to previous learning methods, our approach requires neither intermediate rewards nor complete task demonstrations during training. Second, we demonstrate the versatility of our vision-based task planning in challenging settings with temporary occlusions and dynamic scene changes. Third, we propose an efficient training of basic skills from few synthetic demonstrations by exploring recent CNN architectures and data augmentation. Notably, while all of our policies are learned on visual inputs in simulated environments, we demonstrate the successful transfer and high success rates when applying such policies to manipulation tasks on a real UR5 robotic arm. Figure 19 presents an illustration of the approach.

7.5.5. Monte-Carlo Tree Search for Efficient Visually Guided Rearrangement Planning

Participants: Sergey Zagoruyko, Yann Labbé, Igor Kalevatykh, Ivan Laptev, Justin Carpentier, Mathieu Aubry, Josef Sivic.

We address the problem of visually guided rearrangement planning with many movable objects, i.e., finding a sequence of actions to move a set of objects from an initial arrangement to a desired one, while relying on visual inputs coming from RGB camera. To do so, we introduce a complete pipeline relying on two key contributions. First, we introduce an efficient and scalable rearrangement planning method, based on a Monte-Carlo Tree Search exploration strategy. We demonstrate that because of its good trade-off between exploration and exploitation our method (i) scales well with the number of objects while (ii) finding solutions which require a smaller number of moves compared to the other state-of-the-art approaches. Note that on the contrary to many approaches, we do not require any buffer space to be available. Second, to precisely localize movable objects in the scene, we develop an integrated approach for robust multi-object workspace state estimation from a single uncalibrated RGB camera using a deep neural network trained only with synthetic data. We validate our multi-object visually guided manipulation pipeline with several experiments on a real




Figure 19. Illustration of our approach. (Left): Temporal hierarchy of master and skill policies. The master policy π_m is executed at a coarse interval of n time-steps to select among K skill policies $\pi_s^1 \dots \pi_s^K$. Each skill policy generates control for a primitive action such as grasping or pouring. (Right): CNN architecture used for the skill and master policies.

UR-5 robotic arm by solving various rearrangement planning instances, requiring only 60 ms to compute the complete plan to rearrange 25 objects. In addition, we show that our system is insensitive to camera movements and can successfully recover from external perturbation. Figure 20 shows an example of the problems we consider. This work is under-review and an early pre-print is available [37].




Figure 20. **Visually guided rearrangement planning.** Given a source (a) and target (b) RGB images depicting a robot and multiple movable objects, our approach estimates the positions of objects in the scene without the need for explicit camera calibration and efficiently finds a sequence of robot actions (c) to re-arrange the scene into the target scene. Final object configuration after re-arrangement by the robot is shown in (d).

7.5.6. Estimating the Center of Mass and the Angular Momentum Derivative for Legged Locomotion — A recursive approach

Participants: François Bailly, Justin Carpentier, Mehdi Benallegue, Bruno Watier, Philippe Soueres.

Estimating the center of mass position and the angular momentum derivative of legged systems is essential for both controlling legged robots and analyzing human motion. In this paper[4], a novel recursive approach to concurrently and accurately estimate these two quantities together is introduced. The proposed method employs kinetic and kinematic measurements from classic sensors available in robotics and biomechanics, to effectively exploits the accuracy of each measurement in the spectral domain. The soundness of the proposed approach is first validated on a simulated humanoid robot, where ground truth data is available, against an Extend Kalman Filter. The results demonstrate that the proposed method reduces the estimation error on the center of mass position with regard to kinematic estimation alone, whereas at the same time, it provides an accurate estimation of the derivative of angular momentum. Finally, the effectiveness of the proposed method is illustrated on real measurements, obtained from walking experiments with the HRP-2 humanoid robot. Figure 21 presents an illustration of the approach.




Figure 21. Illustration of the measurement apparatus. The several physical quantities involved in the estimation framework are displayed, as well as a simplified sketch of the estimation algorithm.

7.5.7. Dynamics Consensus between Centroidal and Whole-Body Models for Locomotion of Legged Robots

Participants: Rohan Budhiraja, Justin Carpentier, Nicolas Mansard.

It is nowadays well-established that locomotion can be written as a large and complex optimal control problem. Yet, current knowledge in numerical solver fails to directly solve it. A common approach is to cut the dimensionality by relying on reduced models (inverted pendulum, capture points, centroidal). However it is difficult both to account for whole-body constraints at the reduced level and also to define what is an acceptable trade-off at the whole-body level between tracking the reduced solution or searching for a new one. The main contribution of this paper [9] is to introduce a rigorous mathematical framework based on the Alternating Direction Method of Multipliers, to enforce the consensus between the centroidal state dynamics at reduced and whole-body level. We propose an exact splitting of the whole-body optimal control problem between the centroidal dynamics (under-actuation) and the manipulator dynamics (full actuation), corresponding to a rearrangement of the equations already stated in previous works. We then describe with details how alternating descent is a good solution to implement an effective locomotion solver. We validate this approach in simulation with walking experiments on the HRP-2 robot. Figure 22 presents a resulting motion of the proposed approach.

7.5.8. The Pinocchio C++ library – A fast and flexible implementation of rigid body dynamics algorithms and their analytical derivatives

Participants: Justin Carpentier, Guilhem Saurel, Gabriele Buondonno, Joseph Mirabel, Florent Lamiroux, Olivier Stasse, Nicolas Mansard.




Figure 22. Walking sequence generated for HRP-2 robot using the proposed ADMM solver.

In this paper [10], we introduce Pinocchio, an open-source software framework that implements rigid body dynamics algorithms and their analytical derivatives. Pinocchio does not only include standard algorithms employed in robotics (e.g., forward and inverse dynamics) but provides additional features essential for the control, the planning and the simulation of robots. In this paper, we describe these features and detail the programming patterns and design which make Pinocchio efficient. We evaluate the performances against RBDL, another framework with broad dissemination inside the robotics community. We also demonstrate how the source code generation embedded in Pinocchio outperforms other approaches of state of the art. Figure 23 presents the logo of Pinocchio.



Figure 23. Logo of Pinocchio.

7.5.9. Crocodyl: An Efficient and Versatile Framework for Multi-Contact Optimal Control

Participants: Carlos Mastalli, Rohan Budhiraja, Wolfgang Merkt, Guilhem Saurel, Bilal Hammoud, Maximilien Naveau, Justin Carpentier, Sethu Vijayakumar, Nicolas Mansard.

In this paper [31], we introduce Crocodyl (Contact RObot COntrol by Differential DYnamic Library), an open-source framework tailored for efficient multi-contact optimal control. Crocodyl efficiently computes the state trajectory and the control policy for a given predefined sequence of contacts. Its efficiency is due to the use of sparse analytical derivatives, exploitation of the problem structure, and data sharing. It employs differential geometry to properly describe the state of any geometrical system, e.g. floating-base systems. We have unified dynamics, costs, and constraints into a single concept-action-for greater efficiency and easy prototyping. Additionally, we propose a novel multiple-shooting method called Feasibility-prone Differential Dynamic Programming (FDDP). Our novel method shows a greater globalization strategy compared to classical Differential Dynamic Programming (DDP) algorithms, and it has similar numerical behavior to state-of-the-art multiple-shooting methods. However, our method does not increase the computational complexity typically encountered by adding extra variables to describe the gaps in the dynamics. Concretely, we propose two modifications to the classical DDP algorithm. First, the backward pass accepts infeasible state-control trajectories. Second, the rollout keeps the gaps open during the early "exploratory" iterations (as expected in multiple-shooting methods). We showcase the performance of our framework using different tasks. With our

method, we can compute highly-dynamic maneuvers for legged robots (e.g. jumping, front-flip) in the order of milliseconds. Figure 24 presents a resulting motion of the proposed approach.




Figure 24. *Crocodyl: an efficient and versatile framework for multi-contact optimal control. Highly-dynamic maneuvers needed to traverse an obstacle with the ANYmal robot.*