

Ideal-based algorithms for VASSes and well-structured systems

LaBRI, LSV, and CRISAL

CONTENTS

SUMMARY	1
TABLE OF PARTICIPANTS	2
THE FULL PROPOSAL COMPARED TO THE PRE-PROPOSAL	3
1 PROPOSAL’S CONTEXT, POSITIONING AND OBJECTIVE(S)	3
1.1 OBJECTIVES AND SCIENTIFIC HYPOTHESES	3
1.2 ORIGINALITY AND RELEVANCE IN RELATION TO THE STATE OF THE ART	7
1.3 RISK MANAGEMENT AND METHODOLOGY	8
2 PROJECT ORGANISATION AND MEANS IMPLEMENTED	9
2.1 SCIENTIFIC COORDINATOR	9
2.2 CONSORTIUM AND RESOURCES	9
2.3 MEANS OF ACHIEVING THE OBJECTIVES	11
TASK A: ADHERENCE MEMBERSHIP: A NEW UNDERSTANDING OF THE KLM ALGORITHM	11
TASK B: HANDLING DOWNWARD-CLOSED SETS	13
TASK C: ACCELERATIONS AND COVERABILITY SETS FOR VASS EXTENSIONS	14
TASK D: COMPLEXITY ANALYSIS	16
TASK E: COORDINATION AND DISSEMINATION	17
3 IMPACT AND BENEFITS OF THE PROJECT	19
3.1 EXPECTED IMPACT OF BRAVAS	19
3.2 BRAVAS WITHIN THE ANR 2017 CALL	20
3.3 DISSEMINATION STRATEGY AND ITS EXPECTED IMPACT	20

SUMMARY

Vector additions systems, aka VASSes, have been investigated intensely since the 1960s, often under the form of Petri nets. Their reachability problems have been shown decidable in the early 1980s, opening the way to algorithmically decide many problems. However, the decidability proof and the associated KLM algorithm (for “Kosaraju–Lambert–Mayr”) is very complex, hard to adapt or extend. It is still being developed but does not yet provide an implementable paradigm. VASS reachability is EXPSPACE-hard, and this is all we know about it. Determining the complexity of that problem is the main unsolved question in the area.

Several VASS extensions exist and raise very challenging problems:

- For deciding data logics, Figueira *et al.* introduced an extension of VASSes where a clever mechanism enhances counter capabilities [DFP16]. *Unordered data nets* [Laz+08; HSS12] is another model that has been introduced for data-enriched logics. It is not known whether its reachability problems are decidable.
- More classically, questions on distributed algorithms or recursive programs have led to the introduction of *pushdown VASes*, a model whose reachability problem is still open after several decades.

- *Branching VASes*, aka BVASes, are VASSes extended with a weak form of alternation. BVASes have been introduced independently to attack problems in security protocols [VGL05], in the context of logics in data tree [JSD16], and in linear logic and lambda-calculi [GGS04] but it turned out that equivalent models had been developed in the '90s in computational linguistics [BJR91; Ram94; Sch10]. The decidability status of reachability in BVASes is still open.

Recently, Leroux and Schmitz showed in [LS15] that the KLM algorithm can be understood as computing a finite representation for the *downward closure* of the set of executions of a VASS from source to target, using a suitable ordering on the set of executions. This approach paves a way to adapt the KLM decomposition to various extended models or problems.

Nonelementary complexity. Bounding, from above or below, the running time of algorithms that—like KLM—rely on well-quasi-orderings is a difficult problem that is rarely tackled in the field of automata-theoretic verification. In recent years, Figueira, Schmitz, and Schnoebelen have developed new tools for bounding the length of bad sequences in well-quasi-orderings [Fig+11; SS11; Sch16b]. These techniques are currently demonstrating their relevance in a flurry of applications, e.g., [HSS12; KS13; HSS14; LPS14].

The work programme of the BRAVAS project comprises five tasks:

1. Adherence Membership: A New Understanding of the KLM Algorithm
2. Handling Downward-Closed Sets
3. Accelerations and Coverability Sets for VASS Extensions
4. Complexity Analysis
5. Coordination and Dissemination

The BRAVAS project is a 36-month public partnership collaborative project involving leading public research groups working on formal methods and verification of systems. The goal of the BRAVAS project is to develop *a new and powerful approach to decide the reachability problems for VASS-like models and to analyse their complexity*. The ambition here is to crack with a single hammer several long-lasting open problems that have all been identified as a barrier in different areas, but that are in fact closely related when seen as reachability. As in any fundamental research, the main risk is that of not being able to reach the goal, viz., establishing the desired decidability and complexity results. However, recent work by the proponents shows that we are on the right track, and intermediate milestones may prove to be interesting in their own right. Moreover, the level of risk is reduced since the project is considering many different counter models.

TABLE OF PARTICIPANTS

The participants to the BRAVAS project are listed in the table below. The third column states the involvement of the participant in man-months committed to the project. The last column indicates the coordinators as well as the tasks that the participant is responsible for.¹

¹It is understood that participants shall also contribute to other tasks.

Partner	Name & Position	Involvement	Role & Responsibilities
LaBRI	Jérôme Leroux <i>Research Director</i>	24 mm	Project Coordinator Tasks B and E
	Grégoire Sutre <i>Researcher</i>	12 mm	Task C
	Diego Figueira <i>Researcher</i>	8 mm	
LSV	Philippe Schnoebelen <i>Research Director</i>	16 mm	Local Coordinator Task D
	Alain Finkel <i>Professor</i>	16 mm	
	Jean Goubault-Larrecq <i>Professor</i>	8 mm	
	Sylvain Schmitz <i>Associate Professor</i>	8 mm	Task A
CRISAL	Sylvain Salvati <i>Professor</i>	8 mm	Local Coordinator

Our consortium gathers academic partners from three groups with complementary skills. It will benefit from already existing international collaborations, and in particular with the team of Ranko Lazić (University of Warwick in UK). During the last two years, mutual visits was funded by a 6 month sabbatical of S. Schmitz (LSV), and a bi-lateral project of the Royal Society (UK) between R. Lazić and J. Leroux. In the future we plan international collaborations thanks to additional bilateral projects.

THE FULL PROPOSAL COMPARED TO THE PRE-PROPOSAL

Jean Goubault-Larrecq joined the BRAVAS project between the pre-proposal and the full one.

1 PROPOSAL'S CONTEXT, POSITIONING AND OBJECTIVE(S)

1.1 OBJECTIVES AND SCIENTIFIC HYPOTHESES

1.1.1 CONTEXT

Vector Addition Systems. The main concept at the heart of the BRAVAS project is the abstract computational model known as *vector addition systems (with states)*, or VAS(S) for short, or equivalently Petri nets [Pet62]. A VASS is a finite-state machine equipped with a finite number of counters ranging over nonnegative integers. Operations on counters are limited to incrementation and decrementation. The latter is allowed only if the counter is positive (otherwise the corresponding transition is disabled). VASS find a wide range of applications in the modeling of concurrent, chemical, biological, or business processes. Their algorithmics, and in particular the decidability of their *reachability problem*, is the cornerstone of many decidability results in logic [Kan95; Boj+11; DFP16; CM14], automata [CM77; Boj+11], verification [GS92; GM12], etc.

Reachability Problem. One of the most celebrated results of theoretical computer science is the proof of decidability of the reachability problem in vector addition systems, first shown by Mayr in 1981. The result is famous for several reasons. First, the problem itself was open since the '60s [Pet62; KM69], and was soon identified as a key issue not solely for the formal verification of Petri nets and VASSes, but in many sometimes distantly related fields (see the survey in [Sch16a]). Second, the

problem resisted for a long time in spite of intense scrutiny, and only partial results and incomplete proofs could be obtained over the '70s [ST77; HP79]. Third, the solution proposed by Mayr [May81] is notoriously difficult to comprehend, as it relies on many seemingly ad-hoc concepts; at its core Mayr's algorithm proceeds by decomposing the problem into collections of progressively simpler instances, until either one of them satisfies a rather esoteric condition implying a positive answer, or we have exhausted all possibilities. Regarding this third aspect, two simplified versions of this decomposition algorithm were later offered by Kosaraju [Kos82] (with a full book by Reutenauer [Reu90] dedicated to its self-contained presentation) and Lambert [Lam92]. We call this approach nowadays the "KLM decomposition algorithm" after its main three inventors Kosaraju, Lambert, and Mayr. It is however fair to say that the intuition behind the structures and the termination criteria involved in the KLM algorithm remained elusive: while one can (with quite some work) convince oneself that the algorithm is correct, it was unclear how Kosaraju, Lambert, and Mayr had arrived to this particular approach, besides painstaking trial and error. Another sign of this lack of understanding is that no complexity upper bound were known for the algorithm before [LS15]. This was unsatisfying, not only from a mathematical elegance point of view, but more importantly because it hampered considerably any attempt to generalize the result to extensions of the basic VASS model.

VASS Extensions. Regarding VASS extensions, three reachability problems have now been in the spotlight for some time, namely those of *unordered data Petri nets* (UDPN), *pushdown VAS* (PVAS), and *branching VAS* (BVAS). Each of these problems has been eluding attempts for quite a while, and the general sentiment is that the problem should be decidable, but that we are lacking the proper tools to tackle it. For deciding data logics, Figueira et al. introduced an extension of VASSes where a clever mechanism enhances counter capabilities [DFP16]. *Unordered data nets* [Laz+08; HSS12] is another model that has been introduced for data-enriched logics. More classically, questions on distributed algorithms or recursive programs have led to the introduction of *pushdown VASes*. The third problem, namely BVAS reachability, is especially notorious [Boj14]. Several communities have indeed arrived to the same roadblock: it is equivalent to provability in multiplicative exponential linear logic [GGS04]—the sole fragment of propositional linear logic with unknown decidability status, open since 1990—, it was also defined independently in computational linguistics [Ram94], cryptographic protocol verification [VGL05], verification of APIs for distributed computing [BE13], and a slight extension of it is also equivalent to satisfiability of data logics for XML processing [JSD16; AFF17] and functional program equivalence in ML [CBMO17]; finally, it has the dubious honor of having an incorrect published decidability proof [Bim15].

Related Problems and Applications. Problems related to counter systems can be found in a wide variety of areas, including program verification, databases, logic, mathematical linguistics. Thus progress on the foundations of counter systems also impacts many seemingly unrelated problems. In other words, the intrinsic difficulty of decision problems related to counter systems is a sheer obstacle to progress in these fields. We present several examples here.

Verification of Concurrent Systems. The concurrency revolution is well underway, but writing reliable concurrent software remains a challenge. The analysis of concurrent systems with an unbounded number of processes classically uses a so-called *counter abstraction* [GS92]. The main idea is to forget about the identity of each process, so as to make processes indistinguishable. Assuming that each process is modeled by a finite-state automaton, it is then enough to count, for each state q , how many processes are in state q . The resulting model is a VASS, with no a priori bound on the values of the counters. The verification of a property written in linear temporal logic (including liveness conditions) on the original concurrent system translates into a *reachability* question on the VASS: Is it possible to reach a configuration that is component-wise larger than a given configuration? This reduction to VASS reachability has been extended to a class of recursive asynchronous programs [GM12].

Databases. Static analysis problems on semi-structured data over unbounded domains is often reduced to related problems of infinite state systems. In particular, Vector Addition Systems are closely related to

the satisfiability problem for query languages on data words and data trees. A *data word* is a finite string whose every position contains a letter from a finite alphabet and a data value from some unbounded domain. This model has been considered not only in the realm of semi-structured data, but also timed automata, extended temporal logics, program computations. Query languages on this model can address the finite part directly, while the infinite part can only be tested for equality. In many different scenarios, it has been shown that the static analysis such as the satisfiability problem on these languages are related to the emptiness of automata equipped with *registers* and equality tests, which can usually be reformulated as the reachability or coverability problems for counter systems [DL09; Boj+11; GT16]. Intuitively, the static analysis problems on logics for data words need to “count” and maintain the number of equivalence classes with a certain property along the word, thus the connection with counter systems. On data words, the satisfiability FO^2 , first-order logic with two variables, was shown to be equivalent (i.e., reducible from and to) to the reachability problem for VASS [Boj+11]. On the other hand, the satisfiability for a simple temporal logic on data words was shown to be equivalent to the coverability problem for VASS if future temporal modalities were allowed, or to the reachability problem if both future and past modalities are allowed [DFP16].

Further, connections between data logics and counter systems do not stop on data words, they extend also to trees. The extension of data words to trees plays a central role in semi-structured databases, since it is a common abstraction of XML documents—arguably the most widely-used format for structuring information nowadays. On *data trees*, the extension of data words to trees, FO^2 is related to reachability for *branching* VASS [Boj+09]. In fact it was shown to be equivalent to the reachability problem for a proper extension of branching VASS [JSD16]. On the other hand, temporal logics on data trees have also shown to be equivalent to the reachability and coverability problem for branching VASS in some cases or to an extension of branching VASS in other [AFF17]. For none of these branching VASS models the satisfiability status for the reachability problem is known—the only case known to be decidable heretofore is on dimension one [Göl+16]. Thus, any new insights on these branching models would have immediate implications on the satisfiability problem for query languages over XML documents.

Logic. Linear logic is a logic that has been introduced by Jean-Yves Girard [Gir87]. This logic comes from the observation that higher-order computations, i.e. computations of programs that can manipulate other programs, have invariants that can be decomposed in simpler objects which are fine-grained enough to represent the resource usage (time, memory, ...) of the programs. These objects can be assembled according to the rules of linear logic. Linear logic is then a logic which allows to express behaviors of programs with respect to resources. Therefore constructing a proof in linear logic not only allows to construct a program that satisfies a typing specification, but also that has a particular behavior with respect to its resources. This in particular illustrated by the work on implicit complexity that is based on linear logic and that captures complexity classes by means of typing [Gir98; Laf04].

A major problem that is left open in linear logic is the problem of deciding its Multiplicative and Exponential fragment MELL. Deciding this logic would allow to be able to generate programs that comply with a specification of its usage of resources. The formulae of MELL allow to represent processes that can be replicated at will and other that need to be used exactly once. The difficulty, is that processes that can be replicated can create resources and consume others, but also create new processes. Somehow reasoning within this logic requires to count processes that can be used exactly once and model the actions of the processes that can be used at will. This counting property makes reasoning in that logic close to reasoning with VASSes, but not exactly. An early version of BVASes has been defined in relation to this problem [GGS04]. It is shown that deciding MELL is equivalent to deciding the reachability problem for BVASes. The research direction of BRAVAS brings a fresh look at this problem and brings some hope about solving it.

Mathematical Linguistics. The work of Chomsky in his book about Syntactic Structures [Cho57] has started a long run research program concerned with the formal modeling of natural languages. Part of this program has been concerned with defining a class of languages that would be rich enough to model all natural languages but restricted enough so as to allow for modeling the human capacities related to

natural language. Part of these capacities are the ease to understand, and generate natural language utterances. A particular requirement of the formalisms related to these capacities is that they should have a decidable or even polynomial parsing. The generative approach the Chomsky has developed in his work is centered on models of syntactic structures. These structures are generally trees from which the surface structure can be read. However, as trees present some rigidity, soon enough other features have been added.

One of them, coming from Government and Binding theory [Cho93] and further developed in Chomsky's Minimalist program [Cho95], is the notion of trace. The idea is that a trace is supposed to relate constituents with positions in the structure where they play a syntactic or a semantic role. Stabler [Sta97] has proposed a formal definition of Minimalism which he called minimalist grammars. In this formalism, it has been shown that the parsing problem is equivalent to the reachability problem of BVASes [Sal11]. Somehow traces can be modeled by counters, each construct can have a certain number of traces that need to be used in other places of the derivation and thus as such traces accumulate and are required to be consumed we can model a BVAS.

Another feature comes from Natural Language Processing concerns. It comes from free word order languages. These languages come with a rich feature system which allow words to be placed in many possible orders without affecting the meaning of sentences. Among these languages we can find Latin, Russian, Polish, ... Rambow [Ram94] proposed a formalism extending grammars with tree derivations that deals with multisets of strings. The idea is that those strings represent constituents to be used in the construction of the sentence in arbitrary orders. Multisets are able to simulate counters and here again the branching nature of derivations make the formalism equivalent to BVASes [Sch10].

Though BVASes are complex machines they naturally appear within models of natural languages. It is then interesting to understand more the properties of their reachability sets and relate them to linguistic properties.

1.1.2 OBJECTIVES AND ROADMAP.

The goal of the BRAVAS project is to develop *a new and powerful approach to decide the reachability problems for VASS extensions and to analyse their complexity*. The ambition here is to crack with a single hammer several long-lasting open problems that have all been identified as a barrier in different areas, but that are in fact closely related when seen as reachability.

Our main insight for attacking these difficult problems comes from a new perception of the KLM algorithm that can be interpreted as a counter-example guided abstract refinement algorithm (CEGAR). The two authors of this milestone are partners of the BRAVAS project [LS15]. This presentation of the KLM algorithm agrees with the original construction, but it brings new key insights that are not apparent in the original one: the role of the *completed spaces* of well-quasi-ordered sets (wqo), and the use of an extension of the coverability algorithm to analyse the feasibility of a *downward-closed set of execution*, the so-called adherence property. Future investigations on the KLM algorithm will certainly take advantage of the powerful toolset recently developed for wqo-based techniques. This view also paves the way to generalizations to BVASes and other extensions.

The project is organized in an initial task followed by three main tasks.

A New Understanding of the KLM Algorithm. At its inception, the BRAVAS project is based on a new understanding on the KLM algorithm. An important goal is to identify a logical language that captures the kind of properties needed to express adherence properties, so that we have a clear and abstract problem for which decidability and complexity must be investigated for all considered VASS extensions. This effort will be led by the main experts with whom the idea originated (Leroux & Schmitz) and will profit from interaction with the other partners who can bring expertise on specific aspects like the algorithmic issues related to wqo, or the Fast-Growing complexity classes.

Handling Downward Closed Sets. This task is devoted to the development of data structures and algorithms for handling downward-closed subsets of a wqo via the *canonical ideal decomposition* proposed

in [Gou09; FG09]. This is a generic technique that supports polynomial-time set-theoretical operations. This representation is already well understood for words over a finite alphabet or vectors from \mathbb{N}^k but this task aims at developing it for more varied wqos. For example, the previous task needs to represent sets of executions. The ideal-based representation needs to be developed for trees for BVASes, but also for simpler wqos like sequences of sequences for executions of nets with linearly-ordered data tokens, or with unordered data tokens [HSS12; Laz+08].

Accelerations and Coverability Sets for VASS Extensions. In this task, we will develop a Karp and Miller approach [KM69] for VASS extensions. The original Karp and Miller algorithm is known to be extremely useful for deciding various problems on VASSes. Basically, it computes the downward closure of non-decreasing sequences of reachable configurations thanks to the precise computation of the transitive closure of sequences of actions, called *acceleration* [LS05]. The two main challenges to extend the Karp and Miller approach are (1) developing acceleration techniques, and (2) showing that they are complete for the effective computation of the downward-closure of reachability sets.

Complexity Analysis. This task will develop the tools originally proposed by Schmitz et al. [Fig+11; SS11; Sch16b] for the complexity analysis of wqo-based algorithms. These tools have proved well-suited for proving tight upper bounds via length-function theorems, and for proving matching lower-bounds via robust encodings in standard complexity classes. However the project BRAVAS will need specific developments that fit more closely the form of the KLM algorithm, and variants developed for VASS extensions.

1.2 ORIGINALITY AND RELEVANCE IN RELATION TO THE STATE OF THE ART

The starting point of this proposal is a series of recent milestone results from the BRAVAS participants that pave a promising way to solve long-standing open problems. Those milestone results are presented in the sequel in separate paragraphs:

- Decomposition theorem for solving reachability problems.
- Backward-chaining algorithm for optimal complexity bounds.
- Ideal decomposition for downward closed sets.
- Length-functions theorems for computational complexity.

Decomposition Theorem for Solving Reachability Problems. The Decomposition Theorem of Leroux and Schmitz [LS15] was one of the main result of the ANR funded project REACHARD (2011-2014). This theorem shows that the KLM algorithm can be understood as computing a finite representation for the *downward-closure* of the set of executions of a VASS from source to target, using a suitable ordering on the set of executions. This yields decidability, since the target is reachable from the source in a VASS if and only if the set of executions from source to target is non empty, if and only if the downward closure of this set of executions is non empty. The main ingredients of the proofs of Kosaraju, Lambert, and Mayr are recast in this light: the manipulated structures turn out to be representations for the particular *ideals* arising from VASS executions (see [Task B](#)), the decomposition approach is an instance of an abstraction refinement loop that gets closer and closer to the desired downward closure, and the termination condition is an instance of *adherence membership* for these ideals in the set of executions. This latter condition, of a topological flavor, asks whether there exists executions arbitrarily close to the limit denoted by the ideal. Goubault-Larrecq and Schmitz [GLS16] further show that, under mild conditions, computing downward closures and solving adherence membership are recursively equivalent, thus establishing adherence membership as the key problem in this context.

Backward-Chaining Algorithm for Optimal Complexity Bounds Well-structured, also called well-quasi-ordered, systems, aka “WSTS”, are a generic family of models with important explanatory and unifying power. Broadly speaking, these are models where monotonicity w.r.t. a well-quasi-ordering (a wqo) brings with it generic techniques for representing infinite sets and computing with them

[Fin87; FS01; SS13]. Alain Finkel and Philippe Schnoebelen published the theory in 2001 in the journal *Theoretical Computer Science*; that paper recently received an award as the *the top cited 2001 article in Theoretical Computer Science*. Many of the VASS extensions considered in the BRAVAS project are WSTS for some natural orders over the configurations (except for the pushdown VAS model). The WSTS framework comes with generic verification algorithms that are the main motivation behind the publication of new WSTS models. The most important and best known such algorithm is the backward-chaining algorithm for the decision of safety properties expressed in terms of “coverability”.

One of the main features of the backward chaining algorithm makes it even more interesting for model-checking: it computes co-reachability sets of configurations, that is, it will actually compute the set of all safe configurations instead of just telling whether a given starting configuration is safe. In a recent paper, Sylvain Schmitz and Ranko Lazić revisited the backward-chaining algorithm for VASS, and showed directly the optimal time computational complexity of that algorithm. Moreover, they observed that the complexity of the coverability problem for VASS can be obtained by inspecting the computational complexity of the backward-chaining algorithm. Compared to the Rackoff approach based on combinatorial arguments, the new approach is based on a dual view of the algorithm by looking at downward closed sets (just obtained as the complement of the upward closed sets computed by the backward-chaining algorithm). This observation paves a new way to derive complexity results of the coverability problems for various WSTS, and in particular many VASS extensions considered in the BRAVAS project, just by inspecting the computational complexity of the backward-chaining algorithm; [LS16a] successfully applied this approach on data nets.

Ideal Decomposition for Downward Closed Sets. The decomposition theorem and the optimal complexity of the backward-chaining algorithm are two recent observations based on the ideal decomposition of downward closed sets. This is a classical notion—Fraïssé [Fra00, section 4.5] attributes finite ideal decompositions to Bonnet [Bon75]—which has been rediscovered in the study of well structured transition systems [FG09]. The basic theory of ideals can be found in [Bon75; Fra00; KP92; FG09]. Basically, the notion of ideals provide a canonical way for decomposing downward closed sets as finite union of ideals. Should we manage to find *effective* representations of wqo ideals, this will provide us with algorithmic means to manipulate downward-closed sets. This endeavour is the subject of [FG09; GL+17]. In that paper, an algebraic theory of ideals is presented for elementary operations on wqos like Cartesian products and finite sequences.

Length-Functions Theorems for Computational Complexity. Recently, and motivated in large part by VASS extensions, Schmitz and Schnoebelen have developed a complexity-theoretical framework for the analysis of wqo-based algorithms [Fig+11; SS11; SS13; Sch14; Sch16b]. This *Fast-Growing Hierarchy* framework introduces complexity classes that had not been previously identified by computer scientists. It aims at allowing precise complexity assessments for many algorithms and problems whose complexity ranges between “non-elementary” and “provably total in Peano arithmetic”, and has already proved quite successful in this respect, see [HSS12; LS15; LS16a; LS16b] and the catalogue of complete problems in [Sch16b].

1.3 RISK MANAGEMENT AND METHODOLOGY

Our new understanding of the KLM algorithm thanks to the decomposition theorem, the computational complexity of the backward-chaining algorithm to derive complexity analysis of coverability questions, and the length-function theorems for characterizing complexity of algorithms based on wqo are fundamental results for vector addition systems that are just waiting for extensions.

Whereas those problems interest various fields of computer science (see the impact in [Section 3](#)), all the recent results published over the past two years on this field are from authors and co-authors of the BRAVAS project. This contrast between the large number of researchers interested in those problems and the small group of researchers making progress in that field can be explained by major recent progress in this scientific field.

Nowadays, we strongly think that the time has come for solving long standing open problems about vector addition systems and extensions. As in any fundamental research, the main risk is that of not being able to reach the goal, viz., establishing the desired decidability and complexity results. However, recent work by the proponents shows that we are on the right track, and intermediate milestones may prove to be interesting in their own right.

The level of risk is reduced in two fold. First, the project is considering many different counter models. As a direct consequence, even if a deliverable of a task x is important for a task y , it is possible to start y before the publication of the deliverable by working on intermediate results of task x for some VASS extensions. Second, many important sub-problems can be considered as fall-back (like the small dimension problems [LPS14; LST15b; Blo+15] and the special case of flat control structures [LS05]).

In order to make quick progress on difficult problems, the BRAVAS project focuses his staff funding on two one-year post-doc researchers rather than some Ph.D. grants. Moreover, we plan to hire postdoc with a strong background on the verification of infinite state systems. Our request for a one-year postdoc position is slightly higher than usual to be able to hire junior researchers between 2 and 7 years after the Ph.D.

2 PROJECT ORGANISATION AND MEANS IMPLEMENTED

2.1 SCIENTIFIC COORDINATOR

The project coordinator, Jérôme Leroux, is a CNRS Directeur de recherche (DR2) at LaBRI and head of the team Méthodes Formelles since 2015. Over the past 15 years, he has been involved in several national and international research projects in the areas of his scientific interests, supported by research grants from governmental research agencies. In particular, his research has been supported by two Royal Society grants in UK for collaborations with the university of Warwick and the university of Liverpool, and ANR projects. His scientific interests are mainly focused on the algorithmic verification of counter systems and extensions. He has regularly published his results in international journals (I&C, LMCS, TCS, etc.) and international conferences (ICALP, POPL, LICS, CAV, TACAS, FOSSACS, CONCUR, etc), and he was invited to many international conferences (STACS, RP, ATVA, HIGHLIGHTS) for presenting his recent results on VASS. Moreover, he has been co-chair and organiser of the “12th International Workshop of Reachability Problems” (Bordeaux), as well as member of various PC of international conferences and workshops (STACS, HIGHLIGHTS, ATVA, RP, INFINITY etc). He has also participated to the following ANR projects: PERSÉE (2003-2006), SPACIFY (2006-2009), AVERISS (2006-2009), BINCOA (2009-2012), REACHARD (2011-2014). His main publications related to the projects are [Hof+16; LS16c; Esp+15; LST15b].

2.2 CONSORTIUM AND RESOURCES

The BRAVAS consortium gathers three partners: LaBRI (Laboratoire Bordelais de Recherche en Informatique — Univ. Bordeaux & CNRS), LSV (Laboratoire Spécification et Vérification — ENS Paris-Saclay & CNRS), and CRISAI (Centre de Recherche en Informatique, Signal et Automatique de Lille – Uni. Lille & CNRS & INRIA).

Complementary Expertise. This is an ideal combination of partners for this 36-month project, with expertise in vector addition systems (and extensions), well structured transition systems, and more generally in formal verification. The members of the project include full-time researchers, lecturers and professors. All of them have a very strong theoretical background with an expertise on the foundations of formal verification. Nevertheless, almost all of them have already co-published results in a recent past.

Current Collaboration. Collaborations between participants of the projects have existed for many years, even though not always on a regular pace; for instance between 2005 and 2009 participants from LaBRI and those from LSV has almost not performed any joint work. However, currently there is a

very strong desire to work together, and the team already co-published papers recently. Notice that S. Salvati used to be a member of LaBRI as junior INRIA researcher before he got a professor position in CRISTAL last year. As a former member of LaBRI, S. Salvati has started discussions on topics related to counter machines with G. Sutre and J. Leroux.

Resources. The consortium will contribute 100 man months (of permanent members) to the project, and the requested funding amounts to k€ 250. A one-year post-doctoral researcher will be appointed first by LSV and then by LaBRI for the next year in such a way the same junior researcher could be hired for two years in a row. Postdocs will mainly contribute to [Task C](#) for small dimensional problems (Deliverable [LC.2](#)).

The details of our request are provided below, for each partner. The *Mgmt/Struc.* line in the upcoming tables corresponds to “management fees and structural costs” (called « frais de gestion / frais de structure » in French). We base our travel expenses on the following common rules.

Project meetings: We estimate four project meetings per year, each taking place over two days, with 25% of them in Bordeaux and 75% of them in Paris. The cost of a mission for out-of-town participants is estimated at € 300.

Conferences: The scientific results of the project will be presented at international conferences (see [Task E](#)). Academic participants (except trainees) attend approximately 1.5 international conferences per year, with an average cost of € 2,500 per mission. We only ask for partial support: The mission costs of a participant are weighted by his implication level.

2.2.1 LABORATOIRE BORDELAIS DE RECHERCHE EN INFORMATIQUE

The LaBRI participants (D. Figueira, J. Leroux, G. Sutre) bring their expertise in infinite states model-checking, data-bases formal verification, and vector addition systems.

A one-year post-doctoral researcher will be appointed by LaBRI for this project. The post-doctoral researcher will be involved in the reachability problem for 1-dimensional pushdown VASS. In addition, we plan to recruit two trainees at Master or Engineer level, for research activities.

Budget. Staff funding consists of a one-year post-doc position (between 2 and 7 years after the Ph.D.), and two 5-month internships. Travel expenses include project meetings (€ 9,600) and participation to international conferences (€ 25,000). These amounts are obtained as follows. We plan to organize 12 project meetings involving all the participants. Eight of those meetings will be organized in Paris and the other ones will be organized in Bordeaux. We also plan three week visits of a LaBRI participant in one of the two other partner laboratories. The dissemination of our research results in our community is based on proceeding papers presented in top-ranked international conferences. We plan 10 presentations.

Equipment corresponds to the cost of four high-performance laptops with associated peripherals (€ 6,000).

Staff	Man mon.	k€	Grant k€	Expenses	k€	Grant k€
Permanent	44	303.8	0	Travel	34.6	34.6
Post-doc	12	65.0	65.0	Equipment	6	6
Trainees	10	5.6	5.6	Mgmt/Struc.	8.6	8.6

2.2.2 LABORATOIRE SPÉCIFICATION ET VÉRIFICATION

The LSV participants (A. Finkel, J. Goubault-Larrecq, S. Schmitz, P. Schnoebelen) bring their expertise in infinite states model-checking, well-quasi-order theory, complexity theory, and vector addition systems.

A one-year post-doc researcher will be appointed by LSV for this project. The post-doc researcher will be involved in extending the KLM decomposition theorem to 1-dimensional pushdown VASS. In addition, we plan to recruit one trainee at Master or Engineer level, for research activities.

Budget. Staff funding consists of a one-year post-doc position (between 2 and 7 years after the Ph.D.), and one 5-months internship. Unlike at LaBRI, whose budget includes two internships, some of the LSV internship students are funded by other scholarships, and we consider that one internship is sufficient for the BRAVAS project.

Travel expenses include project meetings (€ 6,000) and participation to international conferences (€ 32,500). These amounts are obtained as follows. We plan to organize 12 project meetings involving all the participants. Four of those meetings will be organized in Bordeaux. The other ones will be organized in Paris. We also plan three week visits of a LSV participant in one of the two other partner laboratories.

The dissemination of our research results in our community is based on proceeding papers presented in top-ranked international conferences. We plan 14 presentations.

Equipment corresponds to the cost of four high-performance laptops with associated peripherals (€ 6,000).

Staff	Man mon.	k€	Grant k€	Expenses	k€	Grant k€
Permanent	48	334.0	0	Travel	38	38
Post-doc	12	65.0	65.0	Equipment	6	6
Trainees	5	2.8	2.8	Mgmt/Struc.	8.9	8.9

2.2.3 CENTRE DE RECHERCHE EN INFORMATIQUE, SIGNAL ET AUTOMATIQUE DE LILLE

The CRISAL participant (S. Salvati) bring his expertise in logic, model-checking, and databases.

He plans to recruit one trainee at Master or Engineer level, for research activities.

Budget. Staff funding consists of one 5-month internships. Travel expenses include project meetings (€ 5,100) and participation to international conferences (€ 5,000). These amount, are obtained as follows. We plan to organize 12 project meetings involving all the participants. Those meetings will be organized in Bordeaux or in Paris. We also plan three week visits in one the two other partner laboratories.

The dissemination of our research results in our community is based on proceeding papers presented in top-ranked international conferences. We plan 2 presentations.

Equipment corresponds to the cost of one high-performance four laptops with associated peripherals (€ 1,500).

Staff	Man mon.	k€	Grant k€	Expenses	k€	Grant k€
Permanent	8	20.8	0	Travel	10.1	10.1
Trainees	5	2.8	2.8	Equipment	1.5	1.5
				Mgmt/Struc.	1.2	1.2

2.3 MEANS OF ACHIEVING THE OBJECTIVES

The expected schedule for the project's tasks is summarized in [Figure 1](#) page 18. The list of deliverables is given in [Table 1](#) page 18. The list of previous or ongoing projects and funding received in connection with this proposal is given in [Table 2](#) page 19.

TASK A: ADHERENCE MEMBERSHIP: A NEW UNDERSTANDING OF THE KLM ALGORITHM

Leader: Sylvain Schmitz (LSV)

$T_0 \rightarrow T_0 + 18$

Other participants:

Jean Goubault-Larrecq (LSV), Jérôme Leroux (LaBRI), and Sylvain Salvati (CRISAL).

Highlights:

- A new approach through abstraction refinement for the reachability problem, with a strong potential for VASS extensions.
- Challenge: extend the approach to other VASS and BVAS models.
- Challenge: endow execution structures in these extensions with suitable wqos.
- Challenge: identify good classes of execution ideals.

Context. The reachability problem for VASSes is known to be decidable thanks to the KLM algorithm. Intuition behind the structures and the termination criteria involved in the KLM algorithm remained elusive: while one can (with quite some work) convince oneself that the algorithm is correct, it was unclear how Kosaraju, Lambert, and Mayr had arrived to this particular approach, besides painstaking trial and error. The starting point of this proposal is the Decomposition Theorem of Leroux and Schmitz [LS15]. This theorem shows that the KLM algorithm can be understood as computing a finite representation for the *downward-closure* of the set of executions of a VASS from source to target, using a suitable ordering on the set of executions. This yields decidability, since the target is reachable from the source in a VASS if and only if the set of executions from source to target is non empty, if and only if the downward closure of this set of executions is non empty. In Leroux and Schmitz [LS15], the main ingredients of the proofs of Kosaraju, Lambert, and Mayr are recast in this light: the manipulated structures turn out to be representations for the particular *ideals* arising from VASS executions (see [Task B](#)), the decomposition approach is an instance of an abstraction refinement loop that gets closer and closer to the desired downward closure, and the termination condition is an instance of *adherence membership* for these ideals in the set of executions. This latter condition, of a topological flavor, asks whether there exists executions arbitrarily close to the limit denoted by the ideal. Goubault-Larrecq and Schmitz [GLS16] further show that, under mild conditions, computing downward closures and solving adherence membership are recursively equivalent, thus establishing adherence membership as the key problem in this context.

Problems. The Decomposition Theorem opens a new avenue for attempts at solving the reachability problem for VASS extensions, and we intend to explore it thoroughly. First, the set of executions in these extensions has to be endowed with a suitable wqo; here there are some natural candidates. Second, we need effective representations for the ideals arising from this wqo, which is the main topic of [Task B](#). Third, we need to be able to decide adherence membership for these ideals, which involves among other things being able to decide boundedness properties, and this is the main topic of [Task C](#).

In addition to the difficulties inherent to [Tasks B](#) and [C](#), the challenges preventing the generalization to VASS extensions should however not be taken lightly: already in the case of VASS, the naive ideals defined by the wqo solely have an undecidable adherence membership problem, and one has to restrict oneself to a restricted class of *execution ideals* with decidable adherence membership and preserved under decomposition steps. In turn, exhibiting such a class of ideals implies building a deep understanding of the interplay between wqos and actual execution structures in VASS extensions.

- Suitable wqos must be defined over the algebraic structures that underlie executions in VASS extensions. BVAS executions are trees, PVAS executions are nested sequences, and UDPN executions are sequences of multisets of vectors.
- Suitable classes of execution ideals must be invented. The proof of the Decomposition Theorem offers a hint on how to construct them, namely as the downward closures (over the set of all structures) of the set of executions greater than some execution. This is however a semantic characterization lacking the necessary effectiveness: how can we perform basic operations on those ideals?
- An abstraction refinement loop must be designed for each extension. The refinement step should in particular yield execution ideals.

Deliverables

LA.1: $T_0 + 6$: Wqo on executions of VASS extensions

LA.2: $T_0 + 18$: Algorithmic properties of execution ideals for VASS extensions

TASK B: HANDLING DOWNWARD-CLOSED SETS

Leader: Jérôme Leroux (LaBRI)

$T_0 \rightarrow T_0 + 36$

Other participants:

Alain Finkel (LSV), Jean Goubault-Larrecq (LSV), Sylvain Schmitz (LSV), and Philippe Schnoebelen (LSV).

Context. Well-structured, also called well-quasi-ordered, systems, aka “WSTS,” are a generic family of models with important explanatory and unifying power. Broadly speaking, these are models where monotonicity w.r.t. a well-quasi-ordering (a wqo) brings with it generic techniques for representing infinite sets and computing with them [Fin87; FS01; SS13]. The WSTS framework comes with generic verification algorithms that are the main motivation behind the publication of new WSTS models. The most important and best known such algorithm is the backward-chaining algorithm for the decision of safety properties expressed in terms of “coverability.” Basically, the coverability problem for a WSTS equipped with a wqo \leq over the configurations consists in deciding if a configuration c is coverable by a reachable configuration x , i.e. such that $c \leq x$. The backward-chaining algorithm basically computes by induction on k the set of configurations that can cover c in at most k steps. One of the main features of the backward chaining algorithm makes it even more interesting for model-checking: it computes co-reachability sets of configurations, that is, it will actually compute the set of all safe configurations instead of just telling whether a given starting configuration is safe.

Downward Closed Sets. Whereas upward closed sets form a natural class of sets for deciding problems on WSTS, two recent breakthroughs of members of the BRAVAS project showed that the class of downward closed sets plays a central role as well. Recently, Sylvain Schmitz and Ranko Lazić recently revisited the backward-chaining algorithm for VASS, and extracted directly the optimal time complexity of the algorithm. Moreover, they observed that the complexity of the coverability problem for VASS can be obtained by inspecting the computational complexity of the backward-chaining algorithm. Compared to the Rackoff approach based on combinatorial arguments, the new approach is based on a dual view of the algorithm by looking at downward closed sets (just obtained as the complement of the upward closed sets computed by the backward-chaining algorithm). This observation paves a new way to derive complexity results of the coverability problems for various WSTS, and in particular many VASS extensions considered in the BRAVAS project, just by inspecting the computational complexity of the backward-chaining algorithm; [LS16a] successfully applied this approach on data nets.

Ideal Decompositions. The main ingredient in the previous results as well as in the decomposition theorem described in [Task A](#) is the decomposition of downward closed sets into ideals. This is a classical notion—Fraïssé [Fra00, section 4.5] attributes finite ideal decompositions to Bonnet [Bon75]—which has been rediscovered in the study of well structured transition systems [FG09]. The basic theory of ideals can be found in [Bon75; Fra00; KP92; FG09]. Ideals also appear in some important results about WSTS:

- The structures used by the KLM algorithm for deciding the VASS reachability problem are ideals over runs [LS15].
- The nodes of the Karp and Miller tree for deciding various problems for VASS are labeled by ideals over the configurations [KM69].
- The set of languages introduced in [Abd+04] for analyzing communicating lossy channel systems, also called SRE in that context are ideals over the set of channel contents.

In those results the notion of ideals is not mentioned since the theory of wqo ideals attracted attention only recently as a central notion for many problems that involves wqo.

Problems. Wqo ideals are ideal for analyzing WSTS. In this task, our main objective is the development of that theory in the context of VASS extensions.

Extending the ideals decomposition of the KLM algorithm to VASS extensions (see [Task A](#) for more details) requires to develop the theory of ideals for complex wqo over the configurations or over the runs. In order to reach that ultimate goal, we identify some challenging intermediate problems:

- First of all, wqo ideals for a large class of wqo obtained algebraically from simpler ones must be identified. For instance, ideals of Cartesian products, or words are known. More general algebraic operations are required for VASS extensions. For instance data Petri nets require multisets of vectors, branching VASS and pushdown VASS require trees of vectors, and so on. Members of the BRAVAS project started working on that problem in [GL+17].
- The adherence property (see [Task A](#)) must be adapted to VASS extensions. Extracting properties satisfied by ideals that are the adherence of the set of runs of VASS extension is a challenging problem. For instance, that characterization for the VASS model provides exactly the theta condition introduced by Kosaraju for deciding the reachability problem. Looking at the form of the ideals for small dimension models (when the number of counters is small) should provide interesting insights to the general problems.
- For practical applications, we are interested in algorithmic properties satisfied by ideals. This problem is motivated by the good algorithmic properties satisfied by the SRE decomposition for lossy channel systems [Abd+04]. More generally, the inclusion of downward closed sets represented by ideals reduces in quadratic time to the inclusion problem for ideals. This last problem naturally depend on the structure of the ideals and must be studied to obtain efficient algorithms.

Deliverables

LB.1: $T_0 + 18$: An algebraic and algorithmic framework for wqo ideals.

LB.2: $T_0 + 36$: Adherence property for VASS extensions.

TASK C: ACCELERATIONS AND COVERABILITY SETS FOR VASS EXTENSIONS

Leader: Grégoire Sutre (LaBRI)

$T_0 + 6 \rightarrow T_0 + 36$

Other participants:

Diego Figueira (LaBRI), Alain Finkel (LSV), Jérôme Leroux (LaBRI), Sylvain Salvati (CRISTAL), and Sylvain Schmitz (LSV).

Context. The analysis of concurrent systems with unboundedly many processes classically uses the so-called *counter abstraction* [GS92]. The main idea is to forget about the identity of processes, so as to make them indistinguishable. The resulting model is a vector addition system (VASS), with no a priori bound on the values of the counters. The verification of a safety property on the original concurrent system (e.g., mutual exclusion) translates into a *coverability* question on the VASS: Is it possible to reach a configuration that is component-wise larger than a given configuration?

Karp-Miller Algorithm. The first decision procedure to solve coverability for VASS was given in 1967 by Karp and Miller [KM69]. In fact, the algorithm of Karp and Miller does much more than just solve coverability. It computes a finite representation of the downward closure of the reachability set. This provides a tight over-approximation of the reachability set that is extremely useful to decide various problems on VASS. In short, their algorithm unwinds the system to analyze into a reachability tree, and, at each step, pumps the positive cycles that have been traversed thus far. This pumping is represented by putting a special symbol ∞ , which intuitively denotes a very large number, in the components that are increased by the cycle. Termination of the algorithm is guaranteed by the fact that the component-wise extension of the usual order on natural numbers is a wqo. The Karp-Miller algorithm has been

adapted to BVASS [VGL05], to UDPN [Hof+16], and to several other VASS extensions [Val78; FMP04; Bon+12; Gee+15].

Acceleration. One of the main ingredients of the Karp-Miller algorithm is the above-mentioned pumping of cycles. The idea is to replace by ∞ the components that the cycle makes larger and larger. This approach is fine for the computation of the downward closure of the reachability set, but it is not precise enough to decide the reachability problem. A refinement of Karp and Miller’s pumping of cycles, called *acceleration*, consists in computing a formula describing the effect of taking the cycle arbitrarily many times. This approach has been developed for various classes of infinite-state systems [BW94; BH99; CJ98; FPS03; Bar+08; BIK10], and has been implemented in successful model-checkers such as LASH, TREX, FAST, and FLATA. Using acceleration instead of Karp and Miller’s pumping of cycles leads to an exact representation of the reachability set, but the resulting procedure is only a semi-algorithm, i.e., it is not guaranteed to terminate in general. It turns out that this semi-algorithm is in fact an algorithm for many subclasses of VASS [LS05].

Problems. The algorithm of Karp and Miller [KM69] is extremely useful to decide various problems on VASS, and it has inspired a lot of research in the larger area of infinite-state systems. The main goal of this task is to develop a Karp and Miller approach for VASS extensions.

Karp and Miller for VASS Extensions. The “KLM decomposition algorithm” for VASS reachability relies on the Karp-Miller algorithm to decide boundedness properties on VASS. Similarly, as mentioned in [Task A](#), turning the Decomposition Theorem of Leroux and Schmitz [LS15] into a decision procedure for VASS extensions requires solving boundedness properties for these extensions. So the first objective of this task is to generalize the Karp-Miller algorithm to VASS extensions. This generalization needs to be powerful enough to decide the required adherence membership questions identified in [Task A](#). It follows that this part of the task will receive more man power after Deliverable [LA.2](#). The main focus of the project is on branching VASS (BVASS), pushdown VASS (PVASS) and unordered data Petri nets (UDPN), for which the reachability problem is still open. The existing adaptations of Karp and Miller’s algorithm to BVASS [VGL05] and UDPN [Hof+16] are a first step towards solving adherence membership questions.

Both BVASS and UDPN are subclasses of the large class of *well-structured transition systems* (WSTS) introduced in [Fin87] and further investigated in [Abd+96; FS01]. The Karp-Miller algorithm has been generalized to the class of WSTS in [Fin87] and more recently in [BFGL17]. This recent generalization yields a very generic semi-algorithm that relies on the completion via ideals [FG12b] of the system under analysis. This gives us a solid foundation — on which we intend to build — to devise new algorithms to solve adherence membership questions. However, several algorithmic challenges need to be addressed in order to turn the generalization of [BFGL17] into an effective algorithm for BVASS and UDPN, including the handling of ideals and downward-closed sets (see [Task B](#)) and the computation of accelerations à la Karp and Miller.

The case of PVASS is different since they do not fall in the class of WSTS. In fact, the coverability and reachability problems for PVASS are easily seen to be irreducible in polynomial time. This is in sharp contrast with the case of VASS where coverability is known to be complete for exponential space [Lip76; Rac78] while the complexity of reachability is still open. An adaptation of Karp and Miller’s algorithm can be used to solve basic boundedness questions on PVASS [LPS14]. However, the complexity of these basic boundedness questions is largely open.

Small Dimensions. In parallel to the development of a Karp-Miller algorithm for VASS extensions of arbitrary dimension, we will study the reachability problem for the case of small dimensions. One of the first subclasses of VASS for which reachability was shown decidable is the class of 2-dimensional VASS. For this class, Hopcroft and Pansiot devised an algorithm that computes a finite description (more precisely, a *semilinear* presentation) of the reachability set [HP79]. As an immediate consequence, they obtained that reachability is decidable for this class. In fact, the algorithm of Hopcroft and Pansiot can

be viewed as a refinement of the Karp-Miller algorithm where the abstract pumping of cycles (putting ∞ in some components) is replaced by an exact acceleration of cycles (adding new vectors to the current set of periods). We will try to push this idea further for VASS extensions in small dimensions. The two one-year postdoc positions will be affected to the small dimension problems. The postdoc will join LSV for one year and then LaBRI for another year in such a way the same junior researcher could be hired for two years in a row.

The idea of cycle acceleration was also used recently to solve coverability in 1-dimensional PVASS [LST15b] and reachability in 1-dimensional BVASS [Göl+16]. But the decidability status of these two problems is still open in dimension 2. We will investigate these two questions. Our conjecture is that these two problems are decidable and that they can be solved using cycle acceleration.

Deliverables.

LC.1: $T_0 + 24$: Karp and Miller algorithm for VASS extensions.

LC.2: $T_0 + 36$: Reachability for VASS extensions (BVASS, PVASS, UDPN) in small dimensions.

TASK D: COMPLEXITY ANALYSIS

Leader: Philippe Schnoebelen (LSV)

$T_0 \rightarrow T_0 + 36$

Other participants:

Diego Figueira (LaBRI), Sylvain Schmitz (LSV), and Grégoire Sutre (LaBRI).

Highlights:

- The complexity analysis of the KLM and related wqo-based algorithms is at the source of recent complexity-theoretical developments.
- These developments must be consolidated and turned into a comprehensive and easy-to-use toolbox.
- Challenges: develop Length-Function Theorems for elementary wqos and beyond, refine them for powersets and multisets, apply them to adherence computations and other algorithms with ideal-based termination arguments.

Context. The verification of models based on VASSes and their extensions raise specific challenges for computational complexity. Indeed, several key algorithms rely on well-quasi-ordering (wqo) theory for their termination arguments, but these do not translate easily into running-time bounds.

Recently, and motivated in large part by VASS extensions, Schmitz and Schnoebelen have developed a complexity-theoretical framework for the analysis of wqo-based algorithms [Fig+11; SS11; SS13; Sch14; Sch16b]. This *Fast-Growing Hierarchy* framework introduces complexity classes that had not been previously identified by computer scientists. It aims at allowing precise complexity assessments for many algorithms and problems whose complexity ranges between “non-elementary” and “provably total in Peano arithmetic”, and has already proved quite successful in this respect, see [HSS12; LS15; LS16a; LS16b] and the catalog of complete problems in [Sch16b].

The general objective of this task is to develop the FGH framework along directions required by the VASS extensions targeted in the project.

Subtask: Length-Function Theorems for elementary wqos and beyond For complexity upper-bound, the FGH framework allows establishing so-called *Length-function Theorems* that provide upper bounds for the length of bad sequences inside a wqo. In order to use these bounds in complexity-theoretical statements, we need an asymptotic analysis of lengths as a function of some input, further parameterized by the complexity of the sequence-generating process, and sometimes parameterized for a family of wqos rather than a fixed one. In practice, this is a kind of resource-conscious and parameterized computation of maximal order types for wqos. This analysis has been carried in [SS11] for all exponential wqos, that is, wqos that can be embedded in arbitrary combinations of subword orderings.

Regarding Length-Function theorems, our main goals are:

Elementary wqos: The main objective is to extend the existing analysis from exponential to elementary wqos, i.e., allowing nesting of subword orderings. This is needed for reasoning on sets of runs (as in [Task A](#) and [Task B](#)) and more generally will account for models like data nets, bounded-depth trees and higher-order channel systems [HSS12; HSS14].

Powersets and multisets: in principle, length-function theorems for powerset and multiset constructions can be accounted for by embedding in subword extensions. However this comes at a clear loss of precision and it would be beneficial to derive sharper, special-cased, length function theorems for these simpler constructs. This objective is probably within reach of our current methods for multisets but the situation is not so clear for powersets. We note in particular that there is no known method for the maximal order type of the (finitary) powerset of a wqo X , and we have already proved that it is not a function of the order type of X .

Beyond elementary wqos: length-function theorems for elementary wqos are too restricted for wqos with maximal order type above ε_0 , including most classes of ordered trees and graphs that are ubiquitous in computer science. In particular, runs of BVASes are trees. This is a quite challenging problem (one just has to consider how complex is the maximal order type of trees), that can probably be best tackled by identifying easy subclasses but we have no candidate at the moment.

Subtask: Fast-Growing Complexity for Ideal-Based Algorithms Algorithms for manipulation of ideals have not yet been scrutinized under the light of computational complexity. In particular, several constructions are just known to be effective, with no particular complexity bound. For example, complementing ideals can be done by enumerating candidate members until an oracle declares that we have gathered a finite basis [GL+17]. At a higher level, the PTL separability algorithm of Goubault-Larrecq and Schmitz [GLS16] combines two semi-decision procedures and relies on arbitrary enumerations.

This lack of complexity bound is in part due to the fact that the aforementioned algorithms have been developed in a generic setting: e.g., one considers a generic wqo X with some effectiveness assumptions. But when in [Task A](#) one looks at adherence questions on sets of runs of VASSes or extensions, the generic algorithms have to be instantiated with sequences of vectors (or richer structures for VASS extensions) and the FGH framework, in particular the Length-Function Theorems, should be able to provide upper bounds.

Our objective here is to provide fully explicit descriptions for ideal-based algorithms. For algorithms on simple wqos (e.g., on VASSes and UDPNs) it is expected that the above work will lead to precise complexity upper-bounds.

Deliverables

LD.1: $T_0 + 24$: “Length-function theorems for elementary wqos”

LD.2: $T_0 + 36$: “A toolbox for fast-growing complexity”

TASK E: COORDINATION AND DISSEMINATION

Leader: Jérôme Leroux (LaBRI)

$T_0 \rightarrow T_0 + 36$

Other participants: All members.

This task focuses on the overall project management and the dissemination of the project’s results.

Special care will be given to the coordination of the project. This includes project meetings and work sessions between partners (see [Section 2.2](#)), human resources (recruitment of post-doc researchers and trainees), as well as project progression and deliverables (notably the final project report). In addition to the organisation presented in [Section 2.2](#), this task will rely on a collaborative work space to collect and exchange results.

We shall disseminate the scientific results of the project through the usual means in academic research: presentations of main results to international conferences, workshops, summer schools and

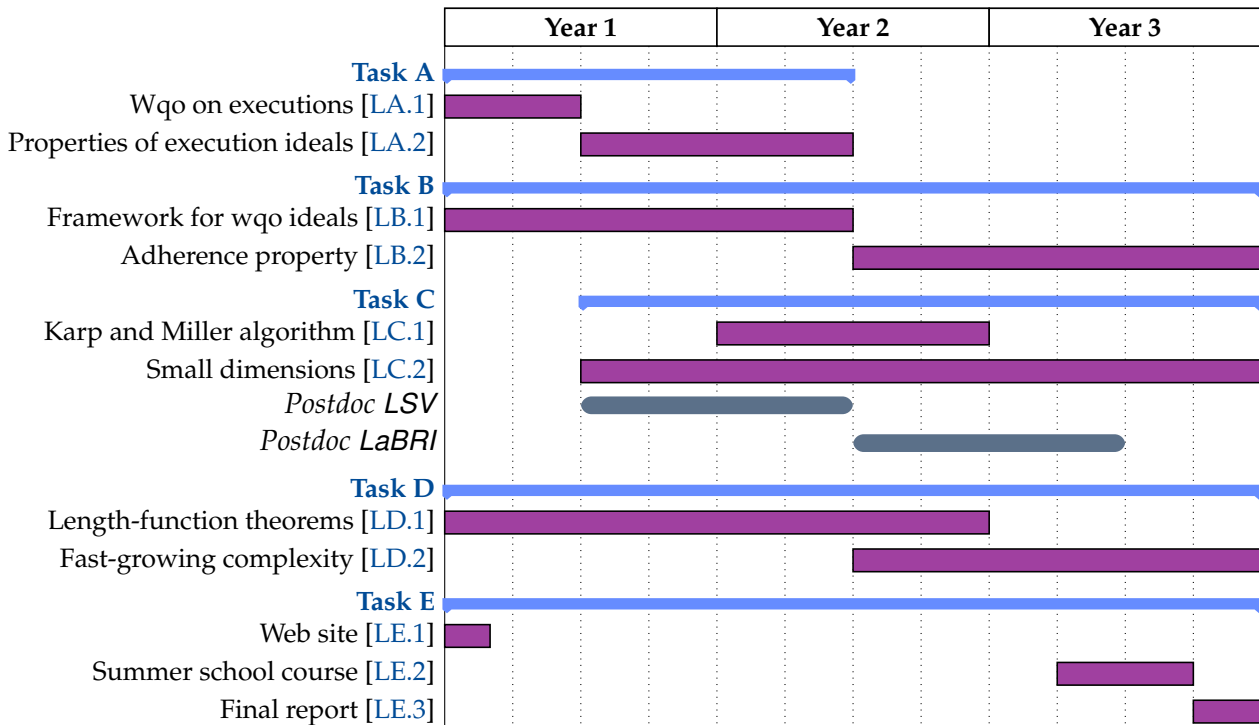


Figure 1: Schedule of the BRAVAS project.

Id	Title	Type	Date
LE.1	Project web site	Web	$T_0 + 2$
LA.1	Wqo on executions of VASS extensions	Report	$T_0 + 6$
LA.2	Algebraic properties of execution ideals of VASS extensions	Report	$T_0 + 18$
LB.1	An algebraic and algorithmic framework for wqo ideals	Report	$T_0 + 18$
LC.1	Karp and Miller algorithm for VASS extensions.	Report	$T_0 + 24$
LD.1	Length-function theorems for elementary wqos	Report	$T_0 + 24$
LE.2	Summer school course on reachability for VASS extensions	Lecture	$T_0 + 33$
LB.2	Adherence property for VASS extensions	Report	$T_0 + 36$
LC.2	Reachability for VASS extensions in small dimensions	Report	$T_0 + 36$
LD.2	A toolbox for fast-growing complexity	Report	$T_0 + 36$
LE.3	Final project report	Report	$T_0 + 36$

Table 1: Deliverables of the BRAVAS project.

publication in international journals. Moreover, we will design and maintain a web site for the project. This web site will be dedicated to the dissemination of our results, and will provide easy access to the deliverables (both reports and software). We also plan to propose a course at a summer school so as to disseminate the tools developed during the project in the direction of younger researcher. The summer school we have in mind is ESSLLI² which is a yearly summer school that gathers young researcher interested in logic, language and computer science. We hope that the broad audience of ESSLLI will allow us to have access to a large audience.

Deliverables

LE.1: Project web site (Web, $T_0 + 2$)

LE.2: Summer school course on reachability for VASS extensions (Lecture, $T_0 + 33$)

LE.3: Final project report (Report, $T_0 + 36$)

²S. Schmitz and P. Schnoebelen provided a lecture on the algorithmic aspects of wqo theory in ESSLLI 2016.

Call	Project Title	Name of coordinator	start date	end date	grant amount	Part.	Name	mm
						of the person involved in this proposal		
ANR PRC	REACHARD	A. Finkel	2011	2014	k€ 262	LSV	A. Finkel	18
						LaBRI	J. Leroux	21
						LSV	S. Schmitz	9
						LSV	P. Schnoebelen	15
						LaBRI	G. Sutre	15
Royal Society		R. Lazić	2015	2017	k£ 12	LaBRI	J. Leroux	2
						LaBRI	G. Sutre	2
						LaBRI	D. Figueira	2
ANR JCJC	PRODAQ	S. Schmitz	2015	2019	k€ 145	LSV	S. Schmitz	38

Table 2: Previous or ongoing projects and funding received in connection with this proposal.

3 IMPACT AND BENEFITS OF THE PROJECT

3.1 EXPECTED IMPACT OF BRAVAS

Decidability for VASS Extensions. The idea at the origin of the BRAVAS proposal is that new insights in the KLM algorithm for VASS will allow solving major open problems for VASS extensions.

If the project is a complete success, we will show decidability of reachability for UDPN, PVAS, and BVAS and this will have major and long-lasting impact on many areas where these and related models are central. As explained in [Section 1.1.1](#), this encompasses several areas in logic (linear logics, denotational semantics, ..), in computer-aided verification (in particular with models for concurrent programs), in database theory, in computational linguistics, etc.

If we do not manage to solve the aforementioned open problems in the course of the project, we expect that our work will identify hard bottlenecks, putting us closer to the heart of the problem. These bottlenecks will inspire new, more focused, questions and suggest new avenues of investigation, for us and for the whole community. They may also show the way to some undecidability proofs that we cannot envision now.

It is possible that some of the above problems, e.g., reachability for BVAS, are proven undecidable as an outcome of the project. In that case it seems that the impact will be less important than with positive answers. However, since the general opinion is that these problems should be decidable, a negative result may spur a paradigm shift in some areas. In any case this will help focus on decidable, as yet unidentified, subclasses.

Verification for VASSes. One clear outcome of the BRAVAS project, in particular its [Task A](#), is an improved understanding of the KLM algorithm for VASS reachability. This is likely to open new avenues for implementation.

Actually, the KLM algorithm has never been implemented simply because of its sheer conceptual complexity. By breaking it down into well-identified components, by using new concepts and data structures for downward-closed sets of runs, by introducing a logical language for handling boundedness and adherence questions, [Task A](#) will make it possible to envision an implementation, identify the most costly components, design heuristics for them or identify special cases where the combinatorial explosion can be tamed.

Let's make it clear right now that such an implementation is not part of our plans at the moment: first, we are interested in other algorithms, not based on the KLM scheme (see [FL15] and [Blo+16]), and second, many other groups have more practical experience with the problem than us, and more pressing real-life applications for such a tool. Nevertheless, such an implementation will have been made possible by the BRAVAS project.

Computing Downward Closures. [Task B](#) and [Task C](#) will provide new algorithms and data structures for ideals and downward-closed sets. This is likely to have impact beyond our concerns with the KLM algorithm and its generalizations to VASS extensions.

Observe that the computability of downward-closures has recently been identified as an important feature, e.g., for logical questions like separability [GLS16] or for decidability in distributed computing problems [LMW15]. It has also long been used for more practical questions like approximations in model-checking [HMW10].

The scope of potential applications is quite large. G. Zetsche recently reduced the computability of downward-closures of traces to the SUP problem [Zet15]—a specialized version of adherence membership—, opening the way to many new positive results (e.g. [HKO16; Cle+16]), but these are still theoretical results in need of efficient data structures and algorithms.

Fast-growing complexity. The bulk of [Task D](#) tackles problems of general interest, not limited to project BRAVAS.

The FGH framework is already put to good use in several groups unrelated to the originating one. There, it allows separating, or sometimes relating, models by comparing their computing power. The potential for applications is very large: any problem that is today just described as “non-elementary”, can be better described within the FGH framework.

These applications will profit from the development of the framework and the dissemination of its underlying concepts via our planned deliverables [LD.1](#) and [LD.2](#).

3.2 BRAVAS WITHIN THE ANR 2017 CALL

As we have explained in the previous section, the project BRAVAS is a foundational project. The challenges BRAVAS addresses may find application in the security of software and systems, in data management, in linguistics, ... This large spectrum of potential applications is related to the orientations 27 and 28 of the SNR, which mentions connected objects and big data. The project, though being clearly in the field of theoretical computer science, also has a multi-disciplinary dimension via its connection to linguistics and to logic. It therefore enters clearly into the Défi 7.

The project BRAVAS is targeting a conceptual breakthrough concerning counter systems. This breakthrough is two folds. First, it concerns understanding the KLM algorithm. The systematic presentation of reachability algorithm by means of downward closed sets in suitable wqos will provide a general architecture to support, design and understand these algorithms. Second, the definition and exploration of new complexity classes that capture these problems can open not only the way to solve major problems in the field but also enhance the understanding as to why these problems are difficult. In particular, this may lead to the development of tailor-made heuristic adapted to practical cases. The ambitious character of BRAVAS and the groundbreaking nature of the research we plan to follow clearly puts the project in Axe 1 of Défi 7.

3.3 DISSEMINATION STRATEGY AND ITS EXPECTED IMPACT

The results of BRAVAS will be mainly disseminated by publications in major conferences and in major journals. Moreover, the website of the project will be another mean of dissemination. It will collect the publications made during the project and make it easy for the scientific community to access the bleeding edge results on counter systems. We also plan to propose a course at a summer school so as to disseminate the tools developed during the project in the direction of younger researcher. The summer school we have in mind is ESSLLI which is a yearly summer school that gathers young researcher interested in logic, language and computer science. The broad audience of ESSLLI will allow us to have access to a large audience.

A CV

A.1 CV OF THE PROPOSAL COORDINATOR

The project coordinator, Jérôme Leroux, is a CNRS Directeur de recherche (DR2) at LaBRI and head of the team Méthodes Formelles since 2015. Over the past 15 years, he has been involved in several national and international research projects in the areas of his scientific interests, supported by research grants from governmental research agencies. In particular, his research has been supported by two Royal Society grants in UK for collaborations with the university of Warwick and the university of Liverpool, and ANR projects. His scientific interests are mainly focused on the algorithmic verification of counter systems and extensions. He has regularly published his results in international journals (I&C, LMCS, TCS, etc.) and international conferences (ICALP, POPL, LICS, CAV, TACAS, FOSSACS, CONCUR, etc), and he was invited to many international conferences (STACS, RP, ATVA, HIGHLIGHTS) for presenting his recent results on VASS. Moreover, he has been co-chair and organiser of the “12th International Workshop of Reachability Problems” (Bordeaux), as well as member of various PC of international conferences and workshops (STACS, HIGHLIGHTS, ATVA, RP, INFINITY etc). He has also participated to the following ANR projects: PERSÉE (2003-2006), SPACIFY (2006-2009), AVERISS (2006-2009), BINCOA (2009-2012), REACHARD (2011-2014). His main publications related to the projects are [Hof+16; LS16c; Esp+15; LST15b].

A.2 CV OF PARTICIPANTS

Figueira, D.: (8 man-months) is a CNRS Chargé de Recherche (CR2) at LaBRI since 2014. His work is centered around automata theory, and logics for words and trees in relation to databases, with a particular interest in the application of methods for infinite state systems in this context.

His main publications related to the project are [DFP16; AFF17; Fig+11].

Finkel, A.: (16mm) is a full professor (CExp) at LSV, where he heads the INFINI group. He invented WSTS (his 1987 ICALP paper) and a variant of lossy channel systems (1990). He initiated the French modern school on counter systems. He is a prominent researcher in WSTS, counter and channel systems, totalling close to 5,000 citations. In recent years, he developed, first with Jean Goubault-Larrecq and later with Michael Blondin and Pierre McKenzie, new forward-chaining approaches for WSTS verification. His main publications related to the projects are [FG09; FG12b; FG12a; BFM14; BFM16; BFGL17].

Goubault-Larrecq, J.: (8mm) is a full professor (CExp) at ENS Paris-Saclay, head of the SECSI (security of information systems) group at LSV, and was awarded the CNRS Silver Medal. He has worked in several domains of computer science: automated deduction and logic, computer security (verification of cryptographic protocols, intrusion detection), semantics of computer languages, static analysis, the λ -calculus, domain theory and topology, concurrency theory and directed algebraic topology. With Alain Finkel, he developed new forward-chaining approaches to the verification of WSTS, based on novel ideas on the representation of downwards-closed subsets of states, as finite sets of points taken in an adequate completion. He also showed that those ideas arose naturally from topological ideas, in turn leading to generalisations of WSTS to the richer class of topological WSTS, which are based on Noetherian spaces instead of well quasi-orders. Before that, he had introduced the branching VASS model with his student K. Neeraj Verma (independently found by de Groote, Salvati, and Guillaume), and had designed and proved a coverability tree algorithm for that model. His main publications related to the project are [VGL05; FG09; FG12b; FG12a; GLS16].

Salvati, S.: (coordinator for CRISAL, 8mm) is a full professor at the Université of Lille 1 and a member of the INRIA team Links. He has a 10-year experience in λ -calculus, logic and formal language theory. In particular he has connect automata theory with the denotational semantics of λ -calculus and proof theory, initiating the first works involving intersection types to extend the notion of recognizability to higher-order functions. His main publications related to the projects are [GGS04; Sal10; Sal11; SW16; Cle+16].

Schmitz, S.: (8mm) is an associate professor (MCF) at LSV with a 8-year expertise in logic and verification and an already impressive publication record (including a LICS paper each year since 2011). He is the key researcher behind the length function theorems that provide upper-bounds on WSTS verification [Fig+11; SS11], and he recently designed a complexity framework for the classification of problems of nonelementary complexity [Sch16b]. His main publications related to the projects are [Sch10; LS15; LS16a; GLS16].

Schoebelen, Ph.: (coordinator for LSV, 16mm) is a CNRS senior researcher (DR1) at LSV with 26-year expertise in temporal logic model-checking and complexity of verification, totalling approx. 3,000 citations. In recent years he developed, in collaboration with S. Schmitz, a series of new techniques for assessing the complexity of wqo-based algorithms, and in particular algorithms for the verification of VASSes and their extensions. His main publications related to the projects are [Fig+11; SS11; HSS12; LS14; GL+17].

Sutre, G.: (12mm) is a CNRS research fellow (CR1) at LaBRI with a 19-year expertise in formal modelling and model-checking of infinite-state systems, totalling approx. 2,300 citations. His current research activities focus on algorithmic verification of concurrent systems, including FIFO and counter systems. His main publications related to the projects are [LS05; LPS14; LST15b; LST15a].

REFERENCES

- [Abd+96] P. A. Abdulla, K. Čerāns, B. Jonsson, and Y.-K. Tsay. “General Decidability Theorems for Infinite-State Systems”. In: *LICS '96*. IEEE, 1996, pp. 313–321.
- [Abd+04] P. A. Abdulla, A. Collomb-Annichini, A. Bouajjani, and B. Jonsson. “Using Forward Reachability Analysis for Verification of Lossy Channel Systems”. In: *Formal Methods in System Design 25.1* (2004), pp. 39–65.
- [AFF17] S. Abriola, D. Figueira, and S. Figueira. “Logics of repeating values on data trees and branching counter systems”. In: *FOSSACS 2017*. To appear. 2017.
- [Bar+08] S. Bardin, A. Finkel, J. Leroux, and L. Petrucci. “FAST: Acceleration from theory to practice”. In: *International Journal on Software Tools for Technology Transfer 10.5* (2008), pp. 401–424.
- [BJR91] T. Becker, A. K. Joshi, and O. Rambow. “Long-distance scrambling and tree adjoining grammars”. In: *EACL'91*. ACL Press, 1991, pp. 21–26.
- [Bim15] K. Bimbó. “The decidability of the intensional fragment of classical linear logic”. In: *Theoretical Computer Science 597* (2015), pp. 1–17.
- [BFGL17] M. Blondin, A. Finkel, and J. Goubault-Larrecq. *Forward Analysis for WSTS, Part III: Model Checking*. Submitted. 2017.
- [BFM14] M. Blondin, A. Finkel, and P. McKenzie. “Handling Infinitely Branching WSTS”. In: *ICALP 2014*. Springer, 2014, pp. 13–25.
- [Blo+16] M. Blondin, A. Finkel, Ch. Haase, and S. Haddad. “Approaching the Coverability Problem Continuously”. In: *TACAS 2016*. Springer, 2016, pp. 480–496.
- [BFM16] M. Blondin, A. Finkel, and P. McKenzie. *Well Behaved Transition Systems*. Research Report 1608.02636 [cs.LO]. 18 pages. Computing Research Repository, Aug. 2016.
- [Blo+15] M. Blondin, A. Finkel, S. Göller, Ch. Haase, and P. McKenzie. “Reachability in Two-Dimensional Vector Addition Systems with States is PSPACE-Complete”. In: *LICS 2015*. IEEE, 2015, pp. 32–43.
- [BW94] B. Boigelot and P. Wolper. “Symbolic Verification with Periodic Sets”. In: *CAV 2004*. Springer, 1994, pp. 55–67.
- [Boj+11] M. Bojańczyk, C. David, A. Muscholl, Th. Schwentick, and L. Segoufin. “Two-variable logic on data words”. In: *ACM Transactions on Computational Logic 12.4* (2011), p. 27.
- [Boj14] M. Bojańczyk. “Some open problems in automata and logic”. In: *ACM SIGLOG News 1.2* (2014), pp. 3–12.
- [Boj+09] M. Bojańczyk, A. Muscholl, Th. Schwentick, and L. Segoufin. “Two-variable logic on data trees and XML reasoning”. In: *Journal of the ACM 56.3* (2009), pp. 1–48.
- [Bon+12] R. Bonnet, A. Finkel, J. Leroux, and M. Zeitoun. “Model Checking Vector Addition Systems with one zero-test”. In: *Logical Methods in Computer Science 8.2:11* (2012).

- [Bon75] R. Bonnet. "On the cardinality of the set of initial intervals of a partially ordered set". In: *Infinite and finite sets: to Paul Erdős on his 60th birthday*, Vol. 1. Coll. Math. Soc. János Bolyai. North-Holland, 1975, pp. 189–198.
- [BH99] A. Bouajjani and P. Habermehl. "Symbolic reachability analysis of FIFO-channel systems with nonregular sets of configurations". In: *Theoretical Computer Science* 221.1–2 (1999), pp. 211–250.
- [BE13] A. Bouajjani and M. Emmi. "Analysis of Recursively Parallel Programs". In: *ACM Transactions on Programming Languages and Systems* 35.3 (2013), 10:1–10:49.
- [BIK10] M. Bozga, R. Iosif, and F. Konečný. "Fast Acceleration of Ultimately Periodic Relations". In: *CAV 2010*. Springer, 2010, pp. 227–242.
- [Cho57] N. Chomsky. *Syntactic Structures*. Mouton, de Gruyter, 1957.
- [Cho93] N. Chomsky. *Lectures on government and binding: The Pisa lectures*. 9. Walter de Gruyter, 1993.
- [Cho95] N. Chomsky. *The minimalist program*. Vol. 28. MIT Press, 1995.
- [Cle+16] L. Clemente, P. Parys, S. Salvati, and I. Walukiewicz. "The Diagonal Problem for Higher-Order Recursion Schemes is Decidable". In: *LICS 2016*. ACM Press, 2016, pp. 96–105.
- [CM14] Th. Colcombet and A. Manuel. "Generalized Data Automata and Fixpoint Logic". In: *Proc. FST&TCS 2014*. LZI, 2014, pp. 267–278.
- [CJ98] H. Comon and Y. Jurski. "Multiple Counters Automata, Safety Analysis, and Presburger Arithmetic". In: *CAV '98*. Springer, 1998, pp. 268–279.
- [CBMO17] C. Cotton-Barratt, A. Murawski, and C.-H. L. Ong. "ML and Extended BVASS". In: *ESOP 2017*. To appear. 2017.
- [CM77] S. Crespi-Reghizzi and D. Mandrioli. "Petri nets and Szilard languages". In: *Information and Control* 33.2 (1977), pp. 177–192.
- [DFP16] S. Demri, D. Figueira, and M. Praveen. "Reasoning about Data Repetitions with Counter Systems". In: *Logical Methods in Computer Science* 12.3 (2016).
- [DL09] S. Demri and R. Lazic. "LTL with the freeze quantifier and register automata". In: *ACM Transactions on Computational Logic* 10.3 (2009).
- [Esp+15] J. Esparza, P. Ganty, J. Leroux, and R. Majumdar. "Verification of Population Protocols". In: *CONCUR 2015*. LZI, 2015, pp. 470–482.
- [Fig+11] D. Figueira, S. Figueira, S. Schmitz, and Ph. Schnoebelen. "Ackermannian and Primitive-Recursive Bounds with Dickson's Lemma". In: *LICS 2011*. IEEE, 2011, pp. 269–278.
- [Fin87] A. Finkel. "A Generalization of the Procedure of Karp and Miller to Well Structured Transition Systems". In: *ICALP '87*. Springer, 1987, pp. 499–508.
- [FG12a] A. Finkel and J. Goubault-Larrecq. "The Theory of WSTS: The Case of Complete WSTS". In: *PETRI NETS 2012*. Springer, 2012, pp. 3–31.
- [FL15] A. Finkel and J. Leroux. "Recent and simple algorithms for Petri nets". In: *Software and System Modeling* 14.2 (2015), pp. 719–725.
- [FPS03] A. Finkel, S. Purushothaman Iyer, and G. Sutre. "Well-abstracted transition systems: application to FIFO automata". In: *Information and Computation* 181.1 (2003), pp. 1–31.
- [FS01] A. Finkel and Ph. Schnoebelen. "Well-Structured Transition Systems Everywhere!" In: *Theoretical Computer Science* 256.1–2 (2001), pp. 63–92.
- [FG09] A. Finkel and J. Goubault-Larrecq. "Forward Analysis for WSTS, Part I: Completions". In: *STACS 2009*. LZI, 2009, pp. 433–444.
- [FG12b] A. Finkel and J. Goubault-Larrecq. "Forward Analysis for WSTS, Part II: Complete WSTS". In: *Logical Methods in Computer Science* 8.3:28 (2012).
- [FMP04] A. Finkel, P. McKenzie, and C. Picaronny. "A Well-Structured Framework for Analysing Petri Nets Extensions". In: *Information and Computation* 195.1–2 (2004), pp. 1–29.
- [Fra00] R. Fraïssé. *Theory of Relations*. Vol. 145. Studies in Logic and the Foundations of Mathematics. Elsevier, 2000.
- [GM12] P. Ganty and R. Majumdar. "Algorithmic verification of asynchronous programs". In: *ACM Transactions on Programming Languages and Systems* 34.1 (2012), 6:1–6:48.
- [Gee+15] G. Geeraerts, A. Heußner, M. Praveen, and J.-F. Raskin. " ω -Petri Nets: Algorithms and Complexity". In: *Fundamenta Informaticae* 137.1 (2015), pp. 29–60.
- [GS92] S. M. German and A. P. Sistla. "Reasoning about Systems with Many Processes". In: *Journal of the ACM* 39.3 (1992), pp. 675–735.
- [Gir87] J.-Y. Girard. "Linear Logic". In: *Theoretical Computer Science* 50.1 (1987), pp. 1–102.
- [Gir98] J.-Y. Girard. "Light Linear Logic". In: *Information and Computation* 143.2 (1998), pp. 174–204.

- [Göl+16] S. Göller, C. Haase, R. Lazić, and P. Totzke. “A Polynomial-Time Algorithm for Reachability in Branching VASS in Dimension One”. In: *ICALP 2016*. LZI, 2016, 105:1–105:13.
- [GL+17] J. Goubault-Larrecq, S. Halfon, P. Karandikar, K. Narayan Kumar, and Ph. Schnoebelen. *The Ideal Approach to Computing Closed Subsets in Well-Quasi-Orderings*. In preparation. Mar. 2017.
- [Gou09] J. Goubault-Larrecq. *On a Generalization of a Result by Valk and Jantzen*. Research Report LSV-09-09. Laboratoire Spécification et Vérification, ENS Cachan, France, 2009.
- [GLS16] J. Goubault-Larrecq and S. Schmitz. “Deciding Piecewise Testable Separability for Regular Tree Languages”. In: *ICALP 2016*. LZI, 2016, 97:1–97:15.
- [GT16] R. Grigore and N. Tzevelekos. “History-Register Automata”. In: *Logical Methods in Computer Science* 12.1 (2016).
- [GGS04] P. de Groote, B. Guillaume, and S. Salvati. “Vector Addition Tree Automata”. In: *LICS 2004*. IEEE, 2004, pp. 64–73.
- [HSS14] Ch. Haase, S. Schmitz, and Ph. Schnoebelen. “The Power of Priority Channel Systems”. In: *Logical Methods in Computer Science* 10.4:4 (2014), pp. 1–39.
- [HMW10] P. Habermehl, R. Meyer, and H. Wimmel. “The Downward-Closure of Petri Net Languages”. In: *ICALP 2010*. Springer, 2010, pp. 466–477.
- [HSS12] S. Haddad, S. Schmitz, and Ph. Schnoebelen. “The Ordinal-Recursive Complexity of Timed-Arc Petri Nets, Data Nets, and Other Enriched Nets”. In: *LICS 2012*. IEEE, 2012, pp. 355–364.
- [HKO16] M. Hague, J. Kochems, and C.-H. L. Ong. “Unboundedness and downward closures of higher-order pushdown automata”. In: *POPL 2016*. ACM Press, 2016, pp. 151–163.
- [Hof+16] P. Hofman, S. Lasota, R. Lazić, J. Leroux, S. Schmitz, and P. Totzke. “Coverability Trees for Petri Nets with Unordered Data”. In: *FOSSACS 2016*. Springer, 2016, pp. 445–461.
- [HP79] J. Hopcroft and J.-J. Pansiot. “On the reachability problem for 5-dimensional vector addition systems”. In: *Theoretical Computer Science* 8.2 (1979), pp. 135–159.
- [JSD16] F. Jacquemard, L. Segoufin, and J. Dimino. “ $\text{FO}^2(<, +1, \sim)$ on data trees, data tree automata and branching vector addition systems”. In: *Logical Methods in Computer Science* 12.2 (2016).
- [KP92] M. Kabil and M. Pouzet. “Une extension d’un théorème de P. Jullien sur les âges de mots”. In: *RAIRO Theoretical Informatics and Applications* 26.5 (1992), pp. 449–482.
- [Kan95] M. I. Kanovich. “Petri Nets, Horn Programs, Linear Logic and Vector Games”. In: *Annals of Pure and Applied Logic* 75.1-2 (1995), pp. 107–135.
- [KS13] P. Karandikar and S. Schmitz. “The Parametric Ordinal-Recursive Complexity of Post Embedding Problems”. In: *FOSSACS 2013*. Springer, 2013, pp. 273–288.
- [KM69] R. M. Karp and R. E. Miller. “Parallel Program Schemata”. In: *Journal of Computer and System Sciences* 3.2 (1969), pp. 147–195.
- [Kos82] S. R. Kosaraju. “Decidability of reachability in vector addition systems”. In: *STOC ’82*. ACM Press, 1982, pp. 267–281.
- [LMW15] S. La Torre, A. Muscholl, and I. Walukiewicz. “Safety of Parametrized Asynchronous Shared-Memory Systems is Almost Always Decidable”. In: *CONCUR 2015*. LZI, 2015, pp. 72–84.
- [Laf04] Y. Lafont. “Soft linear logic and polynomial time”. In: *Theoretical Computer Science* 318.1-2 (2004), pp. 163–180.
- [Lam92] J. L. Lambert. “A structure to decide reachability in Petri nets”. In: *Theoretical Computer Science* 99.1 (1992), pp. 79–104.
- [Laz+08] R. Lazić, T. Newcomb, J. Ouaknine, A. W. Roscoe, and J. Worrell. “Nets with Tokens which Carry Data”. In: *Fundamenta Informaticae* 88.3 (2008), pp. 251–274.
- [LS16a] R. Lazić and S. Schmitz. “The Complexity of Coverability in ν -Petri Nets”. In: *LICS 2016*. ACM Press, 2016, pp. 467–476.
- [LS15] J. Leroux and S. Schmitz. “Demystifying Reachability in Vector Addition Systems”. In: *LICS 2015*. IEEE, 2015, pp. 56–67.
- [LS16b] J. Leroux and S. Schmitz. “Ideal Decompositions for Vector Addition Systems (Invited Talk)”. In: *STACS 2016*. LZI, 2016.
- [LS14] J. Leroux and Ph. Schnoebelen. “On Functions Weakly Computable by Petri Nets and Vector Addition Systems”. In: *RP 2014*. Springer, 2014, pp. 190–202.
- [LS05] J. Leroux and G. Sutre. “Flat counter systems almost everywhere!” In: *ATVA 2005*. Springer, 2005, pp. 489–503.
- [LST15a] J. Leroux, G. Sutre, and P. Totzke. “On Boundedness Problems for Pushdown Vector Addition Systems”. In: *RP 2015*. Springer, 2015, pp. 101–113.

- [LST15b] J. Leroux, G. Sutre, and P. Totzke. “On the Coverability Problem for Pushdown Vector Addition Systems in One Dimension”. In: *ICALP 2015*. Springer, 2015, pp. 324–336.
- [LPS14] J. Leroux, M. Praveen, and G. Sutre. “Hyper-Ackermannian bounds for pushdown vector addition systems”. In: *CSL-LICS 2014*. ACM Press, 2014, 63:1–63:10.
- [LS16c] J. Leroux and S. Schmitz. “Ideal Decompositions for Vector Addition Systems (Invited Talk)”. In: *STACS 2016*. LZI, 2016, 1:1–1:13.
- [Lip76] R. J. Lipton. *The Reachability Problem Requires Exponential Space*. Tech. rep. 63. Department of Computer Science, Yale University, Jan. 1976.
- [May81] E. W. Mayr. “An algorithm for the general Petri net reachability problem”. In: *STOC’81*. ACM Press, 1981, pp. 238–246.
- [Pet62] C. A. Petri. “Kommunikation mit Automaten”. PhD thesis. Universität Bonn, 1962.
- [Rac78] Ch. Rackoff. “The covering and boundedness problems for vector addition systems”. In: *Theoretical Computer Science* 6.2 (1978), pp. 223–231.
- [Ram94] O. Rambow. “Multiset-valued linear index grammars: imposing dominance constraints on derivations”. In: *ACL ’94*. ACL Press, 1994, pp. 263–270.
- [Reu90] C. Reutenauer. *The mathematics of Petri nets*. Masson and Prentice-Hall, 1990.
- [ST77] G. S. Sacerdote and R. L. Tenney. “The decidability of the reachability problem for vector addition systems”. In: *STOC’77*. ACM Press, 1977, pp. 61–76.
- [Sal10] S. Salvati. “A Note on the Complexity of Abstract Categorical Grammars”. In: *The Mathematics of Language*. Springer, 2010, pp. 266–271.
- [Sal11] S. Salvati. “Minimalist Grammars in the Light of Logic”. In: *Logic and Grammar*. Vol. 6700. Lecture Notes in Computer Science. Springer, 2011, pp. 81–117.
- [SW16] S. Salvati and I. Walukiewicz. “Simply typed fixpoint calculus and collapsible pushdown automata”. In: *Mathematical Structures in Computer Science* 26.7 (2016), pp. 1304–1350.
- [SS11] S. Schmitz and Ph. Schnoebelen. “Multiply-Recursive Upper Bounds with Higman’s Lemma”. In: *ICALP 2011*. Springer, 2011, pp. 441–452.
- [SS13] S. Schmitz and Ph. Schnoebelen. “The Power of Well-Structured Systems (Invited Talk)”. In: *CONCUR 2013*. Springer, 2013, pp. 5–24.
- [Sch10] S. Schmitz. “On the Computational Complexity of Dominance Links in Grammatical Formalisms”. In: *ACL 2010*. ACL Press, 2010, pp. 514–524.
- [Sch14] S. Schmitz. “Complexity Bounds for Ordinal-Based Termination (Invited Talk)”. In: *RP 2014*. Springer, 2014, pp. 1–19.
- [Sch16a] S. Schmitz. “Automata column: The complexity of reachability in vector addition systems”. In: *ACM SIGLOG News* 3.1 (2016), pp. 3–21.
- [Sch16b] S. Schmitz. “Complexity Hierarchies Beyond Elementary”. In: *ACM Transactions on Computation Theory* 8.1 (2016), 3:1–3:36.
- [Sta97] E. P. Stabler. “Derivational Minimalism”. In: *LACL ’96*. Springer, 1997, pp. 68–95.
- [Val78] R. Valk. “Self-modifying Nets, a natural extension of Petri nets”. In: *ICALP ’78*. Springer, 1978, pp. 464–476.
- [VGL05] K. N. Verma and J. Goubault-Larrecq. “Karp-Miller Trees for a Branching Extension of VASS”. In: *Discrete Mathematics and Theoretical Computer Science* 7 (2005), pp. 217–230.
- [Zet15] G. Zetsche. “An Approach to Computing Downward Closures”. In: *Proc. ICALP 2015*. Springer, 2015, pp. 440–451.