IN PARTNERSHIP WITH:
Institut polytechnique de Grenoble

Activity Report 2019

Project-Team SPADES

Sound Programming of Adaptive Dependable Embedded Systems

IN COLLABORATION WITH: Laboratoire d'Informatique de Grenoble (LIG)

RESEARCH CENTER
Grenoble - Rhône-Alpes

THEME
Embedded and Real-time Systems

Table of contents

Project-Team SPADES

Creation of the Team: 2013 January 01, updated into Project-Team: 2015 July 01

Keywords:

Computer Science and Digital Science:

A1.1.1. - Multicore, Manycore
A1.1.9. - Fault tolerant systems
A1.3. - Distributed Systems
A2.1.1. - Semantics of programming languages
A2.1.6. - Concurrent programming
A2.1.9. - Synchronous languages
A2.3. - Embedded and cyber-physical systems
A2.3.1. - Embedded systems
A2.3.2. - Cyber-physical systems
A2.3.3. - Real-time systems
A2.4.1. - Analysis
A2.4.3. - Proofs
A2.5.2. - Component-based Design

Other Research Topics and Application Domains:

B5.2.1. - Road vehicles
B6.3.3. - Network Management
B6.4. - Internet of things
B6.6. - Embedded systems

1. Team, Visitors, External Collaborators

Research Scientists

Gregor Goessler [Team leader, Inria, Researcher, HDR]
Alain Girault [Inria, Senior Researcher, HDR]
Pascal Fradet [Inria, Researcher, HDR]
Sophie Quinton [Inria, Researcher]
Jean-Bernard Stefani [Inria, Senior Researcher]

Faculty Member

Xavier Nicollin [Institut polytechnique de Grenoble, Associate Professor]

Post-Doctoral Fellow

Jia Jie Wang [Inria, Post-Doctoral Fellow]

PhD Students

Xiaojie Guo [Univ. Grenoble Alpes, PhD Student]
Maxime Lesourd [Univ. Grenoble Alpes, PhD Student]
Thomas Mari [Institut polytechnique de Grenoble, PhD Student, from Oct 2019]
Stephan Plassart [Univ. Grenoble Alpes, PhD Student]
Christophe Prévot [Inria, PhD Student, until Nov 2019]
Sihem Cherrared [Orange Labs, Univ. Rennes 1, PhD Student]
Arash Shafiei [Orange Labs, PhD Student]

Martin Vassor [Inria, PhD Student]

Technical staff

Souha Ben Rayana [Inria, Engineer, until Apr 2019]
Roger Pissard-Gibollet [Inria, Engineer, from Mar 2019]

Interns and Apprentices

Jonathan Julou [Univ. Grenoble Alpes, from Jun 2019 until Aug 2019]
Thomas Mari [Inria, from Feb 2019 until Jul 2019]
Martin Portalier [Univ. Grenoble Alpes, from Jun 2019 until Aug 2019]

Administrative Assistant

Helen Pouchot-Rouge-Blanc [Inria, Administrative Assistant]

2. Overall Objectives

2.1. Overall Objectives

The SPADES project-team aims at contributing to meet the challenge of designing and programming dependable embedded systems in an increasingly distributed and dynamic context. Specifically, by exploiting formal methods and techniques, SPADES aims to answer three key questions:

1. How to program open distributed embedded systems as dynamic adaptive modular structures?
2. How to program reactive systems with real-time and resource constraints?
3. How to program fault-tolerant and explainable embedded systems?

These questions above are not new, but answering them in the context of modern embedded systems, which are increasingly distributed, open and dynamic in nature [32], makes them more pressing and more difficult to address: the targeted system properties – dynamic modularity, time-predictability, energy efficiency, and fault-tolerance – are largely antagonistic (*e.g.*, having a highly dynamic software structure is at variance with ensuring that resource and behavioral constraints are met). Tackling these questions together is crucial to address this antagonism, and constitutes a key point of the SPADES research program.

A few remarks are in order:

- We consider these questions to be central in the construction of future embedded systems, dealing as they are with, roughly, software architecture and the provision of real-time and fault-tolerance guarantees. Building a safety-critical embedded system cannot avoid dealing with these three concerns.
- The three questions above are highly connected. For instance, composability along time, resource consumption and reliability dimensions are key to the success of a component-based approach to embedded systems construction.
- For us, “Programming” means any constructive process to build a running system. It can encompass traditional programming as well as high-level design or “model-based engineering” activities, provided that the latter are supported by effective compiling tools to produce a running system.
- We aim to provide semantically sound programming tools for embedded systems. This translates into an emphasis on formal methods and tools for the development of provably dependable systems.

3. Research Program

3.1. Introduction

The SPADES research program is organized around three main themes, *Design and Programming Models*, *Certified real-time programming*, and *Fault management and causal analysis*, that seek to answer the three key questions identified in Section 2.1. We plan to do so by developing and/or building on programming languages and techniques based on formal methods and formal semantics (hence the use of *"sound programming"* in the project-team title). In particular, we seek to support design where correctness is obtained by construction, relying on proven tools and verified constructs, with programming languages and programming abstractions designed with verification in mind.

3.2. Design and Programming Models

Work on this theme aims to develop models , languages and tools to support a "correct-by-construction" approach to the development of embedded systems.

On the programming side, we focus on the definition of domain specific programming models and languages supporting static analyses for the computation of precise resource bounds for program executions. We propose dataflow models supporting dynamicity while enjoying effective analyses. In particular, we study parametric extensions where properties such as liveness and boundedness remain statically analyzable.

On the design side, we focus on the definition of component-based models for software architectures combining distribution, dynamicity, real-time and fault-tolerant aspects. Component-based construction has long been advocated as a key approach to the "correct-by-construction" design of complex embedded systems [55]. Witness component-based toolsets such as PTOLEMY [47], BIP [38], or the modular architecture frameworks used, for instance, in the automotive industry (AUTOSAR) [30]. For building large, complex systems, a key feature of component-based construction is the ability to associate with components a set of *contracts*, which can be understood as rich behavioral types that can be composed and verified to guarantee a component assemblage will meet desired properties.

Formal models for component-based design are an active area of research. However, we are still missing a comprehensive formal model and its associated behavioral theory able to deal *at the same time* with different forms of composition, dynamic component structures, and quantitative constraints (such as timing, fault-tolerance, or energy consumption).

We plan to develop our component theory by progressing on two fronts: a semantical framework and domain-specific programming models. The work on the semantical framework should, in the longer term, provide abstract mathematical models for the more operational and linguistic analysis afforded by component calculi. Our work on component theory will find its application in the development of a COQ-based toolchain for the certified design and construction of dependable embedded systems, which constitutes our first main objective for this axis.

3.3. Certified Real-Time Programming

Programming real-time systems (*i.e.*, systems whose correct behavior depends on meeting timing constraints) requires appropriate languages (as exemplified by the family of synchronous languages [40]), but also the support of efficient scheduling policies, execution time and schedulability analyses to guarantee real-time constraints (*e.g.*, deadlines) while making the most effective use of available (processing, memory, or networking) resources. Schedulability analysis involves analyzing the worst-case behavior of real-time tasks under a given scheduling algorithm and is crucial to guarantee that time constraints are met in any possible execution of the system. Reactive programming and real-time scheduling and schedulability for multiprocessor systems are old subjects, but they are nowhere as mature as their uniprocessor counterparts, and still feature a number of open research questions [36], [45], in particular in relation with mixed criticality systems. The main goal in this theme is to address several of these open questions.

We intend to focus on two issues: multicriteria scheduling on multiprocessors, and schedulability analysis for real-time multiprocessor systems. Beyond real-time aspects, multiprocessor environments, and multicore ones in particular, are subject to several constraints *in conjunction*, typically involving real-time, reliability and energy-efficiency constraints, making the scheduling problem more complex for both the offline and the online cases. Schedulability analysis for multiprocessor systems, in particular for systems with mixed criticality tasks, is still very much an open research area.

Distributed reactive programming is rightly singled out as a major open issue in the recent, but heavily biased (it essentially ignores recent research in synchronous and dataflow programming), survey by Bainomugisha et al. [36]. For our part, we intend to focus on devising synchronous programming languages for distributed systems and precision-timed architectures.

3.4. Fault Management and Causal Analysis

Managing faults is a clear and present necessity in networked embedded systems. At the hardware level, modern multicore architectures are manufactured using inherently unreliable technologies [41], [51]. The evolution of embedded systems towards increasingly distributed architectures highlighted in the introductory section means that dealing with partial failures, as in Web-based distributed systems, becomes an important issue.

In this axis we intend to address the question of *how to cope with faults and failures in embedded systems?*. We will tackle this question by exploiting reversible programming models and by developing techniques for fault ascription and explanation in component-based systems.

A common theme in this axis is the use and exploitation of causality information. Causality, *i.e.*, the logical dependence of an effect on a cause, has long been studied in disciplines such as philosophy [61], natural sciences, law [62], and statistics [63], but it has only recently emerged as an important focus of research in computer science. The analysis of logical causality has applications in many areas of computer science. For instance, tracking and analyzing logical causality between events in the execution of a concurrent system is required to ensure reversibility [58], to allow the diagnosis of faults in a complex concurrent system [54], or to enforce accountability [57], that is, designing systems in such a way that it can be determined without ambiguity whether a required safety or security property has been violated, and why. More generally, the goal of fault-tolerance can be understood as being to prevent certain causal chains from occurring by designing systems such that each causal chain either has its premises outside of the fault model (*e.g.*, by introducing redundancy [53]), or is broken (*e.g.*, by limiting fault propagation [65]).

4. Application Domains

4.1. Industrial Applications

Our applications are in the embedded system area, typically: transportation, energy production, robotics, telecommunications, the Internet of things (IoT), systems on chip (SoC). In some areas, safety is critical, and motivates the investment in formal methods and techniques for design. But even in less critical contexts, like telecommunications and multimedia, these techniques can be beneficial in improving the efficiency and the quality of designs, as well as the cost of the programming and the validation processes.

Industrial acceptance of formal techniques, as well as their deployment, goes necessarily through their usability by specialists of the application domain, rather than of the formal techniques themselves. Hence, we are looking to propose domain-specific (but generic) realistic models, validated through experience (*e.g.*, control tasks systems), based on formal techniques with a high degree of automation (*e.g.*, synchronous models), and tailored for concrete functionalities (*e.g.*, code generation).

4.2. Current Industrial Cooperations

Regarding applications and case studies with industrial end-users of our techniques, we cooperate with Orange Labs on software architecture for cloud services.

5. New Software and Platforms

5.1. pyCPA_TWCA

Analysis tool for weakly-hard real-time systems

KEYWORDS: Real time - Scheduling analyses

FUNCTIONAL DESCRIPTION: pyCPA_TWCA is a pyCPA plugin for Typical Worst-Case Analysis. pyCPA is an open-source Python implementation of Compositional Performance Analysis developed at TU Braunschweig, which allows in particular response-time analysis. pyCPA_TWCA is an extension of that tool that is co-developed by Sophie Quinton and Zain Hammadeh at TU Braunschweig. It allows in particular the computation of weakly-hard guarantees for real-time tasks, i.e. number of deadline misses out of a sequence of executions. So far, pyCPA_TWCA is restricted to uniprocessor systems of independent tasks. pyCPA_TWCA can handle the following scheduling policies: Fixed Priority Preemptive, Fixed Priority Non-Preemptive, Weighted Round-Robin, Earliest Deadline First.

- Contact: Sophie Quinton

5.2. CertiCAN

Certifier of CAN bus analysis results

KEYWORDS: Certification - CAN bus - Real time - Static analysis

FUNCTIONAL DESCRIPTION: CertiCAN is a tool, produced using the Coq proof assistant, allowing the formal certification of the correctness of CAN bus analysis results. Result certification is a process that is lightweight and flexible compared to tool certification, which makes it a practical choice for industrial purposes. The analysis underlying CertiCAN, which is based on a combined use of two well-known CAN analysis techniques, is computationally efficient. Experiments demonstrate that CertiCAN is able to certify the results of RTaW-Pegase, an industrial CAN analysis tool, even for large systems. Furthermore, CertiCAN can certify the results of any other RTA tool for the same analysis and system model (periodic tasks with offsets in transactions).

- Contact: Xiaojie Guo

6. New Results

6.1. Design and Programming Models

Participants: Pascal Fradet, Alain Girault, Gregor Goessler, Xavier Nicollin, Arash Shafiei, Jean-Bernard Stefani, Martin Vassor, Souha Ben Rayana.

6.1.1. Hypercells

The location graph framework we have introduced in [66] has evolved into the Hypercell framework presented in [18]. The Hypercell framework allows the definition of different component models for dynamic software architectures featuring both sharing and encapsulation. The basic behavioral theory of hypercells in the form of a contextual bisimulation has been developed and we are currently developing proofs of correctness for encapsulation policies based on this theory.

In collaboration with the Spirals team at Inria Lille – Nord Europe, and Orange, we have used hypercells as a pivot model for developing interpretations, formally defined with the Alloy specification language, of various languages and formalisms for the description of software configurations for cloud computing environments. Configuration languages considered include the TOSCA and OCCI standards, as well as the Open Stack Heat Orchestration Template (HOT), Docker Compose, and the Aeolus component model for cloud deployment. This work, developed as part of a bilateral contract with Orange, allowed the development of a verification tool for the correctness of HOT configurations, and helped uncover several flaws in the ETSI NFV standard.

6.1.2. Dynamicity in dataflow models

Recent dataflow programming environments support applications whose behavior is characterized by dynamic variations in resource requirements. The high expressive power of the underlying models (*e.g.*, Kahn Process Networks or the CAL actor language) makes it challenging to ensure predictable behavior. In particular, checking *liveness* (*i.e.*, no part of the system will deadlock) and *boundedness* (*i.e.*, the system can be executed in finite memory) is known to be hard or even undecidable for such models. This situation is troublesome for the design of high-quality embedded systems. In the past few years, we have proposed several parametric dataflow models of computation (MoCs) [49], [39], we have written a survey providing a comprehensive description of the existing parametric dataflow MoCs [42], and we have studied *symbolic* analyses of dataflow graphs [43]. More recently, we have proposed an original method to deal with lossy communication channels in dataflow graphs [48].

We are nowadays studying models allowing *dynamic reconfigurations* of the *topology* of the dataflow graphs. This is required by many modern streaming applications that have a strong need for reconfigurability, for instance to accommodate changes in the input data, the control objectives, or the environment.

We have proposed a new MoC called Reconfigurable Dataflow (RDF) [13]. RDF extends SDF with transformation rules that specify how the topology and actors of the graph may be reconfigured. Starting from an initial RDF graph and a set of *transformation rules*, an arbitrary number of new RDF graphs can be generated at runtime. Transformations can be seen as graph rewriting rules that match some sub-part of the dataflow graph and replace it by another one. Transformations can be applied an arbitrary number of times during execution and therefore can produce an arbitrary number of new graphs. The major feature and advantage of RDF is that it can be statically analyzed to guarantee that all possible graphs generated at runtime will be connected, consistent, and live. To the best of our knowledge, RDF is the only dataflow MoC allowing an arbitrary number of topological reconfigurations while remaining statically analyzable. It remains to complete the RDF implementation and to evaluate it on realistic case studies. Preliminary results indicate that dynamic reconfigurations can be implemented efficiently.

This is the research topic of Arash Shafiei's PhD, in collaboration with Orange Labs.

6.2. Certified Real-Time Programming

Participants: Pascal Fradet, Alain Girault, Gregor Goessler, Xavier Nicollin, Sophie Quinton, Xiaojie Guo, Maxime Lesourd.

6.2.1. Time predictable programming languages and architectures

Time predictability (PRET) is a topic that emerged in 2007 as a solution to the ever increasing unpredictability of today's embedded processors, which results from features such as multi-level caches or deep pipelines [46]. For many real-time systems, it is mandatory to compute a strict bound on the program's execution time. Yet, in general, computing a tight bound is extremely difficult [69]. The rationale of PRET is to simplify both the programming language and the execution platform to allow more precise execution times to be easily computed [35].

Within the CAPHCA project, we have proposed a new approach for predictable inter-core communication between tasks allocated on different cores. Our approach is based on the execution of synchronous programs written in the FOREC parallel programming language on PREcision Timed (hence deterministic) architectures [71], [72]. The originality resides in the time-triggered model of computation and communication that allows for a very precise control over the thread execution. Synchronization is done via configurable Time Division Multiple Access (TDMA) arbitrations (either physical or conceptual) where the optimal size and offset of the time slots are computed to reduce the inter-core synchronization costs. Results show that our model guarantees time-predictable inter-core communication, the absence of concurrent accesses (without relying on hardware mechanisms), and allows for optimized execution throughput [17]. This is a collaboration with Nicolas Hili and Eric Jenn, the postdoc of Nicolas Hili being funded by the CAPHCA project.

We have also proposed a *multi-rate* extension of FOREC [16]. Indeed, up to now FOREC programs were constrained to operate at a single rate, meaning that all the parallel threads had to share the same execution rate. While this simplified the semantics, it also represented a significant limitation.

Finally, we have extended the compiler of the PRET-C programming language [33], [34] in order to make it energy aware. PRET-C is a parallel programming language in the same sense as Esterel [44], meaning that the parallelism is "compiled away": the PRET-C compiler generates sequential code where the parallel threads from the source program are interleaved according to the synchronous semantics, and produces a classical Control Flow Graph (CFG). This CFG is then turned into a Timed Control Flow Graph (TCFG) by labeling each basic block with the number of clock cycles required to execute it on the chosen processor, based on its micro-architectural characteristics. From the TCFG, we use the method described in Section 6.2.5 to compute a Pareto front of non-dominated (worst-case execution time – WCET, worst-case energy consumption – WCEC) compromises.

6.2.2. Synthesis of switching controllers using approximately bisimilar multiscale abstractions

The use of discrete abstractions for continuous dynamics has become standard in hybrid systems design (see *e.g.*, [67] and the references therein). The main advantage of this approach is that it offers the possibility to leverage controller synthesis techniques developed in the areas of supervisory control of discrete-event systems [64]. The first attempts to compute discrete abstractions for hybrid systems were based on traditional systems behavioral relationships such as simulation or bisimulation, initially proposed for discrete systems most notably in the area of formal methods. These notions require inclusion or equivalence of observed behaviors which is often too restrictive when dealing with systems observed over metric spaces. For such systems, a more natural abstraction requirement is to ask for closeness of observed behaviors. This leads to the notions of approximate simulation and bisimulation introduced in [50]. These approaches are based on sampling of time and space where the sampling parameters must satisfy some relation in order to obtain abstractions of a prescribed precision. In particular, the smaller the time sampling parameter, the finer the lattice used for approximating the state-space; this may result in abstractions with a very large number of states when the sampling period is small. However, there are a number of applications where sampling has to be fast; though this is generally necessary only on a small part of the state-space.

In previous work we have proposed an approach using mode sequences as symbolic states for our abstractions [59]. By using mode sequences of variable length we are able to adapt the granularity of our abstraction to the dynamics of the system, so as to automatically trade off precision against controllability of the abstract states [12]. We have shown the effectiveness of the approach on examples inspired by road traffic regulation.

6.2.3. A Markov Decision Process approach for energy minimization policies

In the context of independent real-time sporadic jobs running on a single-core processor equipped with Dynamic Voltage and Frequency Scaling (DVFS), we have proposed a Markov Decision Process approach (MDP) to minimize the energy consumption while guaranteeing that each job meets its deadline. The idea is to leverage on the *statistical information* on the jobs' characteristics available at design time: release time, worst-case execution time (WCET), and relative deadline. This is the topic of Stephan Plassart's PhD, funded by the CASERM Persyval project. We have considered several cases depending on the amount of information available at design time:

Offline case: In the offline case, all the information is known and we have proposed the first linear complexity offline scheduling algorithm that minimizes the total energy consumption [15]: our complexity is $\mathcal{O}(n)$ where n is the number of jobs to be scheduled, while the previously best known algorithms were in $\mathcal{O}(n^2)$ and $\mathcal{O}(n \log n)$ [60].

Clairvoyant case: In the clairvoyant case, the characteristics of the jobs are only known statistically, and each job's WCET and relative deadline are only known at release time. We want to compute the *optimal* online scheduling speed policy that minimizes the *expected* energy consumption while guaranteeing that each job meets its deadline. This general constrained optimization problem can be modeled as an unconstrained MDP by choosing a proper state space that also encodes the constraints of the problem. In the finite horizon case we use a dynamic programming algorithm, while in the infinite horizon case we use a value iteration algorithm [25].

Non-clairvoyant case: In the non-clairvoyant case, the actual execution time (AET) of a job is only known only when this job completes its execution. This AET is of course assumed to be less than the WCET, which is known at the job's release time. Again, by building an MDP for the system with a well chosen state, we compute the *optimal* online scheduling speed policy that minimizes the *expected* energy consumption [26].

Learning case: In the learning case, the only information known for the jobs are a bound on the jobs' WCETs and a bound on their deadlines. We have proposed two *reinforcement learning* algorithms, one that learns the optimal value of the expected energy (Q-learning), and another one that learns the probability transition matrix of the system, from which we derive the optimal online speed policy.

This work led us to compare several existing speed policies with respect to their feasibility. Indeed, the policies (OA) [70], (AVR) [70], and (BKP) [37] all assume that the maximal speed S_{max} available on the processor is infinite, which is an unrealistic assumption. For these three policies and for our (MDP) policy, we have established necessary and sufficient conditions on S_{max} guaranteeing that no job will ever miss its deadline [27].

6.2.4. Formal proofs for schedulability analysis of real-time systems

We contribute to Prosa [31], a Coq library of reusable concepts and proofs for real-time systems analysis. A key scientific challenge is to achieve a modular structure of proofs, *e.g.*, for response time analysis. Our goal is to use this library for:

1. a better understanding of the role played by some assumptions in existing proofs;
2. a formal verification and comparison of different analysis techniques; and
3. the certification of results of existing (*e.g.*, industrial) analysis tools.

We have further developed CertiCAN, a tool produced using Coq for the formal certification of CAN analysis results [14]. Result certification is a process that is light-weight and flexible compared to tool certification, which makes it a practical choice for industrial purposes. The analysis underlying CertiCAN is based on a combined use of two well-known CAN analysis techniques [68]. Additional optimizations have been implemented (and proved correct) to make CertiCAN computationally efficient. Experiments demonstrate that CertiCAN is able to certify the results of RTaW-Pegase, an industrial CAN analysis tool, even for large systems.

In addition, we have started investigating how to connect Prosa with implementations and less abstract models. Specifically, we have used Prosa to provide a schedulability analysis proof for RT-CertiKOS, a single-core sequential real-time OS kernel verified in Coq [20]. A connection with a timed-automata based formalization of the CAN specification is also in progress. Our objective with this line of research is to understand and bridge the gap between the abstract models used for real-time systems analysis and actual real-time systems implementation.

Finally, we contributed to a major refactoring of the Prosa library to make it more easily extendable and usable.

6.2.5. Scheduling under multiple constraints and Pareto optimization

We have completed a major work on embedded systems subject to multiple non-functional constraints, by proposing the first of its kind multi-criteria scheduling heuristics for a DAG of tasks onto an homogeneous multi-core chip [9], [23]. Given an application modeled as a Directed Acyclic Graph (DAG) of tasks and a multicore architecture, we produce a set of non-dominated (in the Pareto sense) static schedules of this DAG onto this multicore. The criteria we address are the execution time, reliability, power consumption, and peak temperature. These criteria exhibit complex antagonistic relations, which make the problem challenging. For instance, improving the reliability requires adding some redundancy in the schedule, which penalizes the execution time. To produce Pareto fronts in this 4-dimension space, we transform three of the four criteria into constraints (the reliability, the power consumption, and the peak temperature), and we minimize the fourth one (the execution time of the schedule) under these three constraints. By varying the thresholds used for the three constraints, we are able to produce a Pareto front of non-dominated solutions. Each Pareto optimum is a static schedule of the DAG onto the multicore. We propose two algorithms to compute static schedules. The first is a ready list scheduling heuristic called ERPOT (Execution time, Reliability, POwer consumption and Temperature). ERPOT actively replicates the tasks to increase the reliability, uses Dynamic Voltage and Frequency Scaling to decrease the power consumption, and inserts cooling times to control the peak temperature. The second algorithm uses an Integer Linear Programming (ILP) program to compute an optimal schedule. However, because our multi-criteria scheduling problem is NP-complete, the ILP algorithm is limited to very small problem instances, namely DAGs of at most 8 tasks. Comparisons showed that the schedules produced by ERPOT are on average only 9% worse than the optimal schedules computed by the ILP program, and that ERPOT outperforms the PowerPerf-PET heuristic from the literature on average by 33%. This is a joint work with Athena Abdi and Hamid Zarandi from Amirkabir University in Tehran, Iran.

In a related line of work, we have considered the bi-criteria minimization problem in the (worst-case execution time – WCET, worst-case energy consumption – WCEC) space for real-time programs. To the best of our knowledge, this is the first contribution of this kind in the literature.

A real-time program is abstracted as a Timed Control Flow Graph (TCFG), where each basic block is labeled with the number of clock cycles required to execute it on the chosen processor at the nominal frequency. This timing information can be obtained, for instance, with WCET analysis tools. The target processor is equipped with dynamic voltage and frequency scaling (DVFS) and offers several (frequency f, voltage V) operating points. The goal is to compute a set of non-dominated points in the (WCET, WCEC) plane, non-dominated in the Pareto sense. Each such point is an assignment from the set of basic blocks of the TCFG to the set of available (f,V) pairs.

From the TCFG we extract the longest execution path, therefore deriving the WCET and the WCEC for a chosen fixed (f,V) pair. By construction, all the other execution paths are shorter, so this WCET and this WCEC hold for the whole program. This ensures that each single-frequency assignment is a non-dominated point. Then, we study two frequencies assignments, still for the longest execution path. When the frequency switching costs in time and in energy are assumed to be negligible, we demonstrate that each two frequencies (say with f_i and f_j) assignment is a point in the segment between the single frequency assignment at f_i and the single frequency assignment at f_j. We also propose a linear time heuristic to assign a (f,V) pair to all the other blocks (*i.e.*, those not belonging to the longest path) such that all the other execution paths have a shorter WCET and a lesser WCEC. A key result is that we demonstrate that any two frequencies assignment where the two frequencies are not contiguous is dominated either by a single frequency assignment or by a two frequencies assignment with contiguous frequencies. A corollary is that the Pareto front is a continuous piece-wise affine function. Finally, we generalize these results to the case where the frequency switching costs are not negligible. This is the topic of Jia Jie Wang's postdoc.

We evaluate our method and heuristic on a set of hard real time benchmark programs and we show that they perform extremely well. Our DVFS assignment algorithm can also be used as a back-end for the compiler of the PRET-C programming language [33], [34] in order to make it energy aware, thanks to the ability of this compiler to generate TCFGs (see Section 6.2.1).

6.3. Fault Management and Causal Analysis

Participants: Gregor Goessler, Jean-Bernard Stefani, Sihem Cherrared, Thomas Mari, Martin Vassor.

6.3.1. Fault Ascription in Concurrent Systems

Fault ascription is a precise form of fault diagnosis that relies on counterfactual analysis for pinpointing the causes of system failures. Research on counterfactual causality has been marked, until today, by a succession of definitions of causation that are informally validated against human intuition on mostly simple examples. This approach suffers from its dependence on the tiny number and incompleteness of examples in the literature, and from the lack of objective correctness criteria [52].

We have defined in [28] a set of expected properties for counterfactual analysis, and presented a refined analysis that conforms to our requirements. As an early study of the behavior of our analysis under abstraction we have established its monotony under refinement.

6.3.2. Causal Explanations in Discrete Event Systems

Model-Based Diagnosis of discrete event systems (DES) usually aims at detecting failures and isolating faulty event occurrences based on a behavioural model of the system and an observable execution log. The strength of a diagnostic process is to determine *what* happened that is consistent with the observations. In order to go a step further and explain *why* the observed outcome occurred, we borrow techniques from causal analysis.

In [21] we have presented two constructions of explanations that are able to extract the relevant part of a property violation that can be understood by a human operator. Both support partial observability of events. The first construction is based on minimal sub-sequences of the traces of the log that entail a violation of the property. The second approach is based on a construction of layers similar to [56], in which the explanation is constructed from the choices that definitely move the system closer to the violation of the property. Both approaches are complementary: while subsequence-based explanations are well suited to “condense” the execution trace in sequential portions of the model but are prone to keep non-pertinent parts such as initialisation sequences in the explanation, effective choice explanations highlight the “fateful” choices in an execution, as well as alternative events that would have helped avoid the outcome. Effective choice explanations are therefore able to explain failures stemming from non-deterministic choices, such as concurrency bugs.

6.3.3. Fault Management in Virtualized Networks

From a more applied point of view we have been investigating, in the context of Sihem Cherrared’s PhD thesis, approaches for fault explanation and localization in virtualized networks. In essence, Network Function Virtualization (NFV), widely adopted by the industry and the standardization bodies, is about running network functions as software workloads on commodity hardware to optimize deployment costs and simplify the life-cycle management of network functions. However, it introduces new fault management challenges including dynamic topology and multi-tenant fault isolation.

In [29] we have proposed a model-based root cause analysis framework called SAKURA. In order to overcome the lack of accurate previous knowledge, SAKURA features a self-modeling algorithm that models the dependencies within and between layers of virtual networks, including auto-recovery and elasticity aspects. Model-based diagnosis is performed using constraint solving on the previous and acquired knowledge. As an illustration we have applied SAKURA to the virtual IpMultimedia Subsystem (vIMS).

Finally, in our survey on fault management in network virtualization environments [11] we have addressed the impact of virtualization on fault management, proposed a new classification of the recent fault management research achievements in network virtualization environments, and compared their major contributions and shortcomings.

7. Bilateral Contracts and Grants with Industry

7.1. Bilateral Contracts with Industry

- Inria and Orange Labs have established in 2015 a joint virtual research laboratory, called I/O LAB. We have been heavily involved in the creation of the laboratory and are actively involved in its operation (Jean-Bernard Stefani is one of the two co-directors of the lab). I/O LAB focuses on the network virtualization and cloudification. As part of the work of I/O LAB, we have cooperated with Orange Lab, as part of a cooperative research contract funded by Orange, on defining architectural principles and frameworks for network cloud infrastructures encompassing control and management of computing, storage and network resources.

7.2. Bilateral Grants with Industry

With Orange:

- Fault Management in Multi-Tenant Programmable Networks. This CIFRE grant funds the PhD of Sihem Cherrared.
- Dynamic dataflow models of computation. This CIFRE grant funds the PhD of Arash Shafiei.

8. Partnerships and Cooperations

8.1. Regional Initiatives

8.1.1. CASERM (Persyval-Lab project)

Participants: Pascal Fradet, Alain Girault, Gregor Goessler, Xiaojie Guo, Maxime Lesourd, Xavier Nicollin, Stephan Plassart, Sophie Quinton, Jean-Bernard Stefani, Martin Vassor.

The CASERM project represents a significant effort towards a COQ-based design method for reconfigurable multi-view embedded systems, in order to formalize the structure and behavior of systems and to prove their main properties. The use of a proof assistant to support such a framework is motivated by the fact that the targeted systems are both extremely complex and critical. The challenges addressed are threefold:

1. to model software architectures for embedded systems taking into account their dynamicity and multiple constraints (functional as well as non functional);
2. to propose novel scheduling techniques for dynamically reconfiguring embedded systems; and
3. to advance the state of the art in automated proving for such systems.

The objectives of CASERM that address these challenges are organized in three tasks. They consist respectively in designing an architecture description framework based on a process calculus, in proposing online optimization methods for dynamic reconfiguration systems (this is the topic of Stephan Plassart's PhD), and in developing a formal framework for real-time analysis in the COQ proof assistant (this is the topic of Xiaojie Guo's and Maxime Lesourd's PhD).

The CASERM consortium gathers researchers from the LIG and VERIMAG laboratories who are reknowned specialists in these fields. The project started in November 2016 and was completed in November 2019.

8.1.2. SEC: Construction of Safe Explainable Cyber-physical systems

Participants: Gregor Goessler, Thomas Mari.

In cyber-physical systems (CPS), software interacts with physical processes so as achieve desired functionalities. CPS are usually subject to safety and reliability requirements. Depending on the application, their failure may have unacceptable consequences, it is therefore crucial to ensure their correctness at design time. In addition, explainability of increasingly autonomous CPS is becoming crucial in order for the CPS to be socially acceptable.

The goal of this project is twofold. First, we will investigate a contract-based design approach for safe CPS in which different aspects – such as functional requirements, real-time constraints, and continuous behaviors – are modeled and verified separately. Second, we will leverage the contracts in order to ensure explainability of the system behavior by construction. By explainability we understand, informally, that for any behavior of the system we can automatically construct, from a log generated by the execution, an excerpt that retains only the events that causally contributed to the outcome, and that is easy to understand by a human expert.

The SEC project is supported by the "Initiatives de Recherche Stratégiques (IRS)" program of the IDEX UGA. It funds the PhD thesis of Thomas Mari, who will be co-advised by Gregor Gössler and Thao Dang (VERIMAG).

8.2. National Initiatives

8.2.1. ANR

8.2.1.1. RT-proofs

Participants: Pascal Fradet, Xiaojie Guo, Maxime Lesourd, Sophie Quinton.

RT-proofs is an ANR/DFG project between Inria, MPI-SWS, Onera, TU Braunschweig and Verimag, running from 2018 until 2022.

The overall objective of the RT-proofs project is to lay the foundations for computer-assisted formal verification of timing analysis results. More precisely, the goal is to provide:

1. a strong formal basis for schedulability, blocking, and response-time analysis supported by the Coq proof assistant, that is as generic, robust, and modular as possible;
2. correctness proofs for new and well-established generalized response-time analysis results, and a better, precise understanding of the role played by key assumptions and formal connections between competing analysis techniques;
3. an approach for the generation of proof certificates so that analysis results – in contrast to analysis tools – can be certified.

The results obtained in 2019 in connection with the RT-proofs project are described in Section 6.2.4.

8.2.1.2. DCore

Participants: Gregor Goessler, Jean-Bernard Stefani.

DCORE is an ANR project between Inria project teams ANTIQUE, FOCUS and SPADES, and the IRIF lab, running from 2019 to 2023.

The overall objective of the project is to develop a semantically well-founded, novel form of concurrent debugging, which we call *causal debugging*, that aims to alleviate the deficiencies of current debugging techniques for large concurrent software systems. The causal debugging technology developed by DCORE will comprise and integrate two main novel engines:

1. *a reversible execution engine* that allows programmers to backtrack and replay a concurrent or distributed program execution, in a way that is both precise and efficient (only the exact threads involved by a return to a target anterior or posterior program state are impacted);
2. a *causal analysis engine* that allows programmers to analyze concurrent executions, by asking questions of the form “what caused the violation of this program property?”, and that allows for the precise and efficient investigation of past and potential program executions.

8.2.2. Institute of Technology (IRT)

8.2.2.1. CAPHCA

Participants: Alain Girault, Nicolas Hili.

CAPHCA is a project within the Antoine de Saint Exupéry IRT in Toulouse. The general objective of the project is to provide methods and tools to achieve both performance and determinism on modern, high-performance, multi-core and FPGA-enabled SOCs. Our specific contribution lies withing work packages dedicated to the design of novel PRET architectures and programming languages (see Section 6.2.1). This contract has yielded two publications so far [17], [16].

8.3. European Initiatives

8.3.1. Collaborations in European Programs, Except FP7 & H2020

Program: Celtic-Plus

Project acronym: SENDATE

Project title: Secure Networking for a Data center cloud in Europe

Duration: April 2016 - March 2019

Coordinator: Nokia France

Other partners: Nokia, Orange, IMT, Inria

Abstract: The SENDATE project aims to develop a clean-slate architecture for converged telecommunications networks and distributed data centers supporting 5G cellular networks and the needs from the Industrial Internet and the Internet of Things. It aims to provide scientific and technical solutions for intra and inter data centrers security, control, management and orchestration, placement and management of virtual network functions, as well as high-speed transport networks for data centers access and interconnection.

8.3.2. Collaborations with Major European Organizations

We have a strong collaboration with the Technische Universität Braunschweig in Germany and the MPI-SWS in Kaiserslautern (Germany) on formal proofs for the analysis real-time systems. This collaboration is formalized by the ANR-PRCI project called RT-proofs started in 2018, which involves MPI-SWS, TU Braunschweig, Inria, and Onera.

8.4. International Initiatives

8.4.1. Inria Associate Teams Not Involved in an Inria International Labs

8.4.1.1. Quasar

Title: Quantitative systems formal verification

International Partner (Institution - Laboratory - Researcher):

CAS (China) - Department of Informatics - Lijun Zhang

Start year: 2019

The general scientific objectives are to extend formal analysis and verification methods such as model checking, process algebra and interactive theorem proving (Coq) to quantitative systems, more specifically probabilistic and quantum computing systems. Application fields include compositional modeling for dynamic real-time probabilistic software architectures and risk analysis. The collaboration will involve active scientists on all these fields not only from Inria and Inst Soft. CAS, but also from CWI, Verimag Grenoble, ECNU Shanghai, and partners of CWI (VU Amsterdam and Twente).

9. Dissemination

9.1. Environmental and societal responsibility

NB: We support the idea of an additional section in future Activity Reports that would be dedicated to actions in connection with our environmental and societal responsibility. We write below content that would fit into it for 2019.

- A discussion on computer ethics was scheduled at our last team seminar. There have been several meetings between permanent researchers to discuss how we can better align our research agenda with the current needs of our society. Our discussions cover aspects such as the environmental and societal impact of ICT (Information and Communication Technologies), and more broadly the role that ICT play in our lives, our role as researchers in computer science, etc. These discussions have led us to consider these issues as a research topic that we should investigate in the near future.
- Sophie Quinton and Jean-Bernard Stefani contributed to the so-called "MakeSEnS" working group appointed by the CEO of Inria to propose a list of concrete actions that the institute could take to tackle the current environmental crisis [24].
- Sophie Quinton has been mandated by Patrick Gros (Director of the Inria Grenoble – Rhône-Alpes research center) to organize discussions and actions regarding the environmental and societal impact of our research at Inria Grenoble Rhône-Alpes.
- Sophie Quinton is a member of the GDS EcoInfo (https://ecoinfo.cnrs.fr/). Her actions as member of EcoInfo include: leading a discussion group on the environmental impact of collaborative platforms as part of a two-day seminar organized at Mines ParisTech [1]; and contributing to a roadmap regarding the environmental and societal impact of ICT prepared by the Conseil National du Numérique (https://cnnumerique.fr/) at the government's request.
- Sophie Quinton co-chairs a working group of the GDR CIS associated with the Center for Internet and Society (http://cis.cnrs.fr/) focused on environmental issues.
- Sophie Quinton is part of the scientific committee of the upcoming "COP2 étudiante" (https://cop2etudiante.org/).

9.2. Promoting Scientific Activities

9.2.1. Scientific Events: Organisation

9.2.1.1. General Chair, Scientific Chair

- Alain Girault is member of the steering committee of the International Federated Conference on Distributed Computing Techniques (DISCOTEC) and of the ACM International Conference on Embedded Software (EMSOFT).
- Gregor Gössler is member of the steering committee of the International Workshop on Causal Reasoning for Embedded and Safety-critical Systems Technologies (CREST).
- Jean-Bernard Stefani is the current chair of the steering committee of the IFIP FORTE international conference series, a member of the steering committee of the IFIP DISCOTEC conference series, and the current chair of the IFIP Working Group 6.1.

9.2.1.2. Member of the Organizing Committees

- Alain Girault was co-organizer of the Workshop on Synchronous Programming (SYNCHRON'19) http://synchron19.org.
- Gregor Gössler was co-organizer of a Shonan seminar on *Causal reasoning in systems* https://project.inria.fr/shonan139.

9.2.2. Scientific Events: Selection

9.2.2.1. Chair of Conference Program Committees

- Sophie Quinton was the program chair of 31st Euromicro Conference on Real-Time Systems (ECRTS'19) https://www.ecrts.org/archives/fileadmin/WebsitesArchiv/ecrts2019/index.html.

9.2.2.2. Member of the Conference Program Committees

[1] http://www.mines-paristech.fr/Actualites/Ingenieurs-et-transitions-environnementales/4116

- Alain Girault served in the program committees of the Forum on specification and Design Languages (FDL'19) and the Conference on Applications of Concurrency to System Design (ACSD'19).
- Gregor Gössler served in the program committees of the ACM/IEEE International Conference on Embedded Software (EMSOFT'19), the 26th SPIN Symposium on Model Checking of Software (SPIN'19), the 17th ACM-IEEE International Conference on Formal Methods and Models for System Design (MEMOCODE'19), and the 4th international Workshop on Formal Reasoning about Causation, Responsibility, and Explanations in Science and Technology (CREST'19).

9.2.2.3. Reviewer

- Alain Girault reviewed an article for the International Conference on Fundamental Approaches to Software Engineering (FASE'19).
- Sophie Quinton reviewed an article for the ACM International Conference on Embedded Software (EMSOFT'19).
- Pascal Fradet reviewed an article for the ACM International Conference on Embedded Software (EMSOFT'19).

9.2.3. Journal

9.2.3.1. Member of the Editorial Boards

- Alain Girault is a member of the editorial board of the Journal on Embedded Systems.

9.2.3.2. Reviewer - Reviewing Activities

- Alain Girault reviewed articles for IEEE Transactions on Service Computing (TSC) and for the International Journal on Software Tools for Technology Transfer (STTT).
- Gregor Gössler reviewed an article for IEEE Transactions on Automatic Control (TAC).

9.2.4. Leadership within the Scientific Community

- Sophie Quinton is a member of the ACM SIGBED Executive Committee and Associate Editor of the SIGBED Review.

9.2.5. Research Administration

- Pascal Fradet is head of the committee for doctoral studies ("Responsable du comité des études doctorales") of the Inria Grenoble – Rhône-Alpes research center and local correspondent for the young researchers Inria mission ("Mission jeunes chercheurs").
- Alain Girault is Deputy Scientific Director at Inria in charge of the domain "Algorithmics, Programming, Software and Architecture". He is also Scientific Manager of the Cyber-Physical Systems axis of the Persyval-Lab labex https://persyval-lab.org.
- Xavier Nicollin is member of the committee for computing resources users ("Comité des Utilisateurs des Moyens Informatiques") of the Inria Grenoble – Rhône-Alpes research center.
- Jean-Bernard Stefani is Head of Science (délégué scientifique) of the Inria Grenoble – Rhône-Alpes research center and a member of the Inria Evaluation Committee.

9.3. Teaching - Supervision - Juries

9.3.1. Teaching

Licence : Pascal Fradet, Théorie des Langages 1 & 2, 36 HeqTD, niveau L3, Grenoble INP (Ensimag), France

Licence : Pascal Fradet, Modèles de Calcul : λ-calcul, 12 HeqTD, niveau L3, Univ. Grenoble Alpes, France

Master : Pascal Fradet, Langages et Traducteurs, 16 HeqTD, niveau M1, Polytech Grenoble, Univ. Grenoble Alpes, France

Master : Xavier Nicollin, Analyse de Code pour la Sûreté et la Sécurité, 45 HeqTD, niveau M1, Grenoble INP (Ensimag), France

Licence : Xavier Nicollin, Théorie des Langages 1, 48 HeqTD, niveau L3. Grenoble INP (Ensimag), France

Licence : Xavier Nicollin, Théorie des Langages 2, 37,5 HeqTD, niveau L3, Grenoble INP (Ensimag), France

Licence : Xavier Nicollin, Bases de la Programmation Impérative, 30 HeqTD, niveau L3, Grenoble INP (Ensimag), France

Licence : Sophie Quinton, Théorie des Langages 2, 20 HeqTD, niveau L3, Grenoble INP (Ensimag), France

Master : Sophie Quinton, Performance and Quantitative Properties, 8 HeqTD, MOSIG, Univ. Grenoble Alpes, France

Master: Jean-Bernard Stefani, Formal Aspects of Component Software, 9h, MOSIG, Univ. Grenoble Alpes, France.

9.3.2. Supervision

- PhD in progress: Sihem Cherrared, "Fault Management in Multi-Tenant Programmable Networks"; Univ. Rennes 1; since October 2016; co-advised by Eric Fabre and Gregor Gössler.
- PhD in progress: Thomas Mari, "Construction of Safe Explainable Cyber-physical systems"; Grenoble INP; since October 2019; co-advised by Gregor Gössler and Thao Dang.
- PhD: Christophe Prévot, "Early Performance assessment for evolving and variable Cyber-Physical Systems", Univ. Grenoble Alpes; defended on November 15th, 2019; co-advised by Alain Girault and Sophie Quinton.
- PhD: Zain A. H. Hammadeh, "Deadline Miss Models for Temporarily Overloaded Systems", TU Braunschweig; defended on May 28th, 2019; co-advised by Rolf Ernst and Sophie Quinton.
- PhD in progress: Stephan Plassart, "On-line optimization in dynamic real-time systems"; Univ. Grenoble Alpes; since September 2016; co-advised by Bruno Gaujal and Alain Girault.
- PhD in progress: Xiaojie Guo, "Formal Proofs for the Analysis of Real-Time Systems in COQ"; Univ. Grenoble Alpes; since December 2016; co-advised by Pascal Fradet, Jean-François Monin, and Sophie Quinton.
- PhD in progress: Maxime Lesourd, "Generic Proofs for the Analysis of Real-Time Systems in COQ"; Univ. Grenoble Alpes; since September 2017; co-advised by Pascal Fradet, Jean-François Monin, and Sophie Quinton.
- PhD in progress: Arash Shafiei, "RDF: A reconfigurable dataflow MoC supporting dynamic topological transformations and static analyzability"; Univ. Grenoble Alpes; since September 2017; co-advised by Pascal Fradet, Alain Girault, and Xavier Nicollin.
- PhD in progress: Martin Vassor, "Analysis and types for safe dynamic software reconfigurations"; Univ. Grenoble Alpes; since November 2017; co-advised by Pascal Fradet and Jean-Bernard Stefani.
- Internships: Jonathan Julou and Martin Portalier, "Formal proofs for real-time systems analysis"; ENSIMAG; co-advised by Pascal Fradet, Xiaojie Guo, Maxime Lesourd and Sophie Quinton.

9.3.3. Juries

- Alain Girault was referee for the Habilitation thesis of Julien De Antoni (Université Côte d'Azur) and was member of the PhD thesis of Pierre Donat-Bouillud (Sorbonne Université). He was also vice-president of the Inria Starting Research Position (SRP) Advanced Research Position (ARP) jurys.
- Jean-Bernard Stefani was examiner for the Habilitation (HDR) jury of Thomas Ledoux (Université de Nantes).

- Sophie Quinton was a member of the PhD thesis committee of Hugo Daigmorte (Université de Toulouse).

9.4. Popularization

9.4.1. Articles and contents

Software is a fundamental pillar of modern scientific research, across all fields and disciplines. However, there is a lack of adequate means to cite and reference software due to the complexity of the problem in terms of authorship, roles and credits. This complexity is further increased when it is considered over the lifetime of a software that can span up to several decades. Building upon the internal experience of Inria, we have provided a contribution to the ongoing efforts in order to develop proper guidelines and recommendations for software citation and reference. Namely, we recommend: (1) a richer taxonomy for software contributions with a qualitative scale; (2) to put humans at the heart of the evaluation; and (3) to distinguish citation from reference [10]. This has been a joint work between Alain Girault and colleagues from Software Heritage (Roberto Di Cosmo), from Inria's Evaluation Committee (Pierre Alliez, Benjamin Guedj, Mohand-Saïd Hacid, Nicolas Rougier), and a specialist in reproducible research (Arnaud Legrand).

9.4.2. Interventions

Sophie Quinton played the role of one of the experts in the fictional trial of a robot during the Transfo Festival in Grenoble. She also gave a lecture on the design and verification of real-time systems as part of the ISN curriculum for high school teachers.

Alain Girault played the role of one of the experts in the fictional trial of a robot during the Fête de la Science in Grenoble.

9.4.3. Creation of media or tools for science outreach

All SPADES members have contributed to the scenario of an animated movie, commissioned by the communication service of Inria, in order to popularise the research activities of SPADES.

10. Bibliography

Major publications by the team in recent years

[1] S. ANDALAM, P. S. ROOP, A. GIRAULT, C. TRAULSEN. *A Predictable Framework for Safety-Critical Embedded Systems*, in "IEEE Trans. on Computers", July 2014, vol. 63, n^{o} 7, pp. 1600–1612

[2] A. BOUAKAZ, P. FRADET, A. GIRAULT. *A Survey of Parametric Dataflow Models of Computation*, in "ACM Trans. Design Autom. Electr. Syst.", 2017, vol. 22, n^{o} 2, pp. 38:1–38:25, https://doi.org/10.1145/2999539

[3] S. DJOKO DJOKO, R. DOUENCE, P. FRADET. *Aspects preserving properties*, in "Science of Computer Programming", 2012, vol. 77, n^{o} 3, pp. 393-422

[4] G. FREHSE, A. HAMANN, S. QUINTON, M. WÖHRLE. *Formal Analysis of Timing Effects on Closed-loop Properties of Control Software*, in "35th IEEE Real-Time Systems Symposium 2014 (RTSS)", Rome, Italy, December 2014, https://hal.inria.fr/hal-01097622

[5] A. GIRARD, G. GÖSSLER, S. MOUELHI. *Safety Controller Synthesis for Incrementally Stable Switched Systems Using Multiscale Symbolic Models*, in "IEEE Transactions on Automatic Control", 2016, vol. 61, n^{o} 6, pp. 1537-1549 [*DOI :* 10.1109/TAC.2015.2478131], https://hal.archives-ouvertes.fr/hal-01197426

[6] G. GÖSSLER, D. LE MÉTAYER. *A general framework for blaming in component-based systems*, in "Science of Computer Programming", 2015, vol. 113, Part 3 [*DOI :* 10.1016/J.SCICO.2015.06.010], https://hal.inria.fr/hal-01211484

[7] I. LANESE, C. A. MEZZINA, J.-B. STEFANI. *Reversibility in the higher-order π-calculus*, in "Theoretical Computer Science", 2016, vol. 625, pp. 25-84 [*DOI :* 10.1016/J.TCS.2016.02.019], https://hal.inria.fr/hal-01303090

[8] S. QUINTON, M. HANKE, R. ERNST. *Formal analysis of sporadic overload in real-time systems*, in "2012 Design, Automation & Test in Europe Conference & Exhibition, DATE 2012, Dresden, Germany, March, 2012", 2012, pp. 515–520, http://dx.doi.org/10.1109/DATE.2012.6176523

Publications of the year

Articles in International Peer-Reviewed Journals

[9] A. ABDI, A. GIRAULT, H. ZARANDI. *ERPOT: A Quad-Criteria Scheduling Heuristic to Optimize Execution Time, Reliability, Power Consumption and Temperature in Multicores*, in "IEEE Transactions on Parallel and Distributed Systems", October 2019, vol. 30, n^{o} 10, pp. 2193-2210 [*DOI :* 10.1109/TPDS.2019.2906172], https://hal.inria.fr/hal-02400019

[10] P. ALLIEZ, R. DI COSMO, B. GUEDJ, A. GIRAULT, M.-S. HACID, A. LEGRAND, N. P. ROUGIER. *Attributing and Referencing (Research) Software: Best Practices and Outlook from Inria*, in "Computing in Science & Engineering", 2019, pp. 1-14, https://arxiv.org/abs/1905.11123 [*DOI :* 10.1109/MCSE.2019.2949413], https://hal.archives-ouvertes.fr/hal-02135891

[11] S. CHERRARED, S. IMADALI, E. FABRE, G. GÖSSLER, I. G. B. YAHIA. *A Survey of Fault Management in Network Virtualization Environments: Challenges and Solutions*, in "IEEE Transactions on Network and Service Management", October 2019, pp. 1-15 [*DOI :* 10.1109/TNSM.2019.2948420], https://hal.inria.fr/hal-02370378

[12] A. GIRARD, G. GÖSSLER. *Safety Synthesis for Incrementally Stable Switched Systems using Discretization-Free Multi-Resolution Abstractions*, in "Acta Informatica", 2019 [*DOI :* 10.1007/S00236-019-00341-X], https://hal.archives-ouvertes.fr/hal-02286661

International Conferences with Proceedings

[13] P. FRADET, A. GIRAULT, R. KRISHNASWAMY, X. NICOLLIN, A. SHAFIEI. *RDF: Reconfigurable Dataflow*, in "DATE 2019 - Design, Automation & Test in Europe Conference & Exhibition", Florence, Italy, March 2019, pp. 1709-1714 [*DOI :* 10.23919/DATE.2019.8714987], https://hal.inria.fr/hal-01960788

[14] P. FRADET, X. GUO, J.-F. MONIN, S. QUINTON. *CertiCAN: A Tool for the Coq Certification of CAN Analysis Results*, in "RTAS 2019 - 25th IEEE Real-Time and Embedded Technology and Applications Symposium", Montreal, Canada, IEEE, April 2019, pp. 1-10 [*DOI :* 10.1109/RTAS.2019.00023], https://hal.archives-ouvertes.fr/hal-02119024

[15] B. GAUJAL, A. GIRAULT, S. PLASSART. *A Linear Time Algorithm for Computing Off-line Speed Schedules Minimizing Energy Consumption*, in "MSR 2019 - 12ème Colloque sur la Modélisation des Systèmes Réactifs", Angers, France, November 2019, pp. 1-14, https://hal.archives-ouvertes.fr/hal-02372136

[16] A. Girault, N. Hili, É. Jenn, E. Yip. *A Multi-Rate Precision Timed Programming Language for Multi-Cores*, in "FDL 2019 - Forum for Specification and Design Languages", Southampton, United Kingdom, IEEE, September 2019, pp. 1-8 [*DOI :* 10.1109/FDL.2019.8876950], https://hal.inria.fr/hal-02399998

[17] N. Hili, A. Girault, É. Jenn. *Worst-Case Reaction Time Optimization on Deterministic Multi-Core Architectures with Synchronous Languages*, in "RTCSA2019 2019 - 25th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA)", Hangzhou, China, IEEE, August 2019, pp. 1-11 [*DOI :* 10.1109/RTCSA.2019.8864570], https://hal.inria.fr/hal-02400009

[18] J.-B. Stefani, M. Vassor. *Encapsulation and Sharing in Dynamic Software Architectures: The Hypercell Framework*, in "FORTE 2019 - 39th International Conference on Formal Techniques for Distributed Objects, Components, and Systems (FORTE)", Copenhagen, Denmark, J. A. Pérez, N. Yoshida (editors), Formal Techniques for Distributed Objects, Components, and Systems, Springer International Publishing, 2019, vol. LNCS-11535, pp. 242-260, Part 1: Full Papers [*DOI :* 10.1007/978-3-030-21759-4_14], https://hal.inria.fr/hal-02313751

Conferences without Proceedings

[19] B. Gaujal, A. Girault, S. Plassart. *A Linear Time Algorithm for Computing Off-line Speed Schedules Minimizing Energy Consumption*, in "MSR 2019 - 12ème Colloque sur la Modélisation des Systèmes Réactifs", Angers, France, November 2019, pp. 1-14, https://hal.archives-ouvertes.fr/hal-02432311

[20] X. Guo, M. Lesourd, M. Liu, L. Rieg, Z. Shao. *Integrating Formal Schedulability Analysis into a Verified OS Kernel*, in "Computer Aided Verification", New York, United States, July 2019, pp. 496-514 [*DOI :* 10.1007/978-3-030-25543-5_28], https://hal.archives-ouvertes.fr/hal-02289494

[21] G. Gössler, T. Mari, Y. Pencolé, L. Travé-Massuyès. *Towards Causal Explanations of Property Violations in Discrete Event Systems*, in "DX'19 - 30th International Workshop on Principles of Diagnosis", Klagenfurt, Austria, November 2019, pp. 1-8, https://hal.inria.fr/hal-02369014

[22] A. N. Sylla, K. Guillouard, F. Klamm, M. Ouzzif, P. Merle, S. Ben Rayana, J.-B. Stefani. *Formal Verification of Orchestration Templates for Reliable Deployment with OpenStack Heat*, in "CNSM 2019 - 15th International Conference on Network and Service Management", Halifax, Canada, October 2019, pp. 1-5, https://hal.inria.fr/hal-02375386

Research Reports

[23] A. Abdi, A. Girault, H. Zarandi. *ERPOT: A quad-criteria scheduling heuristic to optimize the execution time, failure rate, power consumption and temperature in multicores*, Inria ; 37, March 2019, n^{o} RR-9196, pp. 1-37, https://hal.inria.fr/hal-01848087

[24] F. Berthoud, P. Guitton, L. Lefèvre, S. Quinton, A. Rousseau, J. Sainte-Marie, C. Serrano, J.-B. Stefani, P. Sturm, E. Tannier. *Sciences, Environnements et Sociétés : Rapport long du groupe de travail MakeSEnS d'Inria*, Inria, October 2019, https://hal.inria.fr/hal-02340948

[25] B. Gaujal, A. Girault, S. Plassart. *A Discrete Time Markov Decision Process for Energy Minimization Under Deadline Constraints*, Grenoble Alpes ; Inria Grenoble Rhône-Alpes, Université de Grenoble, December 2019, n^{o} RR-9309, 46 p. , https://hal.inria.fr/hal-02391948

[26] B. GAUJAL, A. GIRAULT, S. PLASSART. *Exploiting Job Variability to Minimize Energy Consumption under Real-Time Constraints*, Inria Grenoble Rhône-Alpes, Université de Grenoble ; Université Grenoble - Alpes, November 2019, n^{o} RR-9300, 23 p. , https://hal.inria.fr/hal-02371742

[27] B. GAUJAL, A. GIRAULT, S. PLASSART. *Feasibility of on-line speed policies in real-time systems*, Inria Grenoble Rhône-Alpes, Université de Grenoble ; Univ. Grenoble Alpes, November 2019, n^{o} RR-9301, 38 p. , https://hal.inria.fr/hal-02371996

[28] G. GÖSSLER, J.-B. STEFANI. *Causality Analysis and Fault Ascription in Component-Based Systems*, Inria - Research Centre Grenoble – Rhône-Alpes, June 2019, n^{o} RR-9279, pp. 1-27, https://hal.inria.fr/hal-02161534

Other Publications

[29] S. CHERRARED, S. IMADALI, E. FABRE, G. GÖSSLER. *SAKURA a Model Based Root Cause Analysis Framework for vIMS*, ACM Press, June 2019, pp. 594-595, MobiSys 2019 - 17th ACM International Conference on Mobile Systems, Applications, and Services, Poster, https://hal.inria.fr/hal-02291163

References in notes

[30] *Automotive Open System Architecture*, 2003, http://www.autosar.org

[31] *A Library for formally proven schedulability analysis*, http://prosa.mpi-sws.org/

[32] ARTEMIS JOINT UNDERTAKING. *ARTEMIS Strategic Research Agenda*, 2011

[33] S. ANDALAM, P. S. ROOP, A. GIRAULT. *Predictable Multithreading of Embedded Applications Using PRET-C*, in "International Conference on Formal Methods and Models for Codesign, MEMOCODE'10", Grenoble, France, IEEE, July 2010, pp. 159–168

[34] S. ANDALAM, P. S. ROOP, A. GIRAULT, C. TRAULSEN. *A Predictable Framework for Safety-Critical Embedded Systems*, in "IEEE Transactions on Computers", July 2014, 13 p. , https://hal.inria.fr/hal-01095468

[35] P. AXER, R. ERNST, H. FALK, A. GIRAULT, D. GRUND, N. GUAN, B. JONSSON, P. MARWEDEL, J. REINEKE, C. ROCHANGE, M. SEBATIAN, R. VON HANXLEDEN, R. WILHELM, W. YI. *Building Timing Predictable Embedded Systems*, in "ACM Trans. Embedd. Comput. Syst.", 2014, To appear

[36] E. BAINOMUGISHA, A. CARRETON, T. VAN CUTSEM, S. MOSTINCKX, W. DE MEUTER. *A Survey on Reactive Programming*, in "ACM Computing Surveys", 2013, vol. 45, n^{o} 4

[37] N. BANSAL, T. KIMBREL, K. PRUHS. *Speed Scaling to Manage Energy and Temperature*, in "Journal of the ACM", 2007, vol. 54, n^{o} 1

[38] A. BASU, S. BENSALEM, M. BOZGA, J. COMBAZ, M. JABER, T.-H. NGUYEN, J. SIFAKIS. *Rigorous Component-Based System Design Using the BIP Framework*, in "IEEE Software", 2011, vol. 28, n^{o} 3

[39] V. BEBELIS, P. FRADET, A. GIRAULT, B. LAVIGUEUR. *BPDF: A Statically Analyzable Dataflow Model with Integer and Boolean Parameters*, in "International Conference on Embedded Software, EMSOFT'13", Montreal, Canada, ACM, September 2013

[40] A. BENVENISTE, P. CASPI, S. A. EDWARDS, N. HALBWACHS, P. LE GUERNIC, R. DE SIMONE. *The synchronous languages 12 years later*, in "Proceedings of the IEEE", 2003, vol. 91, n^{o} 1

[41] S. BORKAR. *Designing Reliable Systems from Unreliable Components: The Challenges of Transistor Variability and Degradation*, in "IEEE Micro", 2005, vol. 25, n^{o} 6

[42] A. BOUAKAZ, P. FRADET, A. GIRAULT. *A Survey of Parametric Dataflow Models of Computation*, in "ACM Transactions on Design Automation of Electronic Systems (TODAES)", January 2017, https://hal.inria.fr/hal-01417126

[43] A. BOUAKAZ, P. FRADET, A. GIRAULT. *Symbolic Analyses of Dataflow Graphs*, in "ACM Transactions on Design Automation of Electronic Systems (TODAES)", January 2017, https://hal.inria.fr/hal-01417146

[44] F. BOUSSINOT, R. DE SIMONE. *The Esterel Language*, in "Proceedings of the IEEE", September 1991, vol. 79, n^{o} 9, pp. 1293–1304

[45] R. DAVIS, A. BURNS. *A Survey of Hard Real-Time Scheduling for Multiprocessor Systems*, in "ACM Computing Surveys", 2011, vol. 43, n^{o} 4

[46] S. A. EDWARDS, E. A. LEE. *The Case for the Precision Timed (PRET) Machine*, in "44th Design Automation Conference (DAC)", IEEE, 2007

[47] J. EKER, J. W. JANNECK, E. A. LEE, J. LIU, X. LIU, J. LUDVIG, S. NEUENDORFFER, S. SACHS, Y. XIONG. *Taming heterogeneity - the Ptolemy approach*, in "Proceedings of the IEEE", 2003, vol. 91, n^{o} 1

[48] P. FRADET, A. GIRAULT, L. JAMSHIDIAN, X. NICOLLIN, A. SHAFIEI. *Lossy channels in a dataflow model of computation*, in "Principles of Modeling, Festschrift in Honor of Edward A. Lee", Berkeley, United States, Lecture Notes in Computer Science, Springer, October 2017, https://hal.inria.fr/hal-01666568

[49] P. FRADET, A. GIRAULT, P. POLPAVKO. *SPDF: A schedulable parametric data-flow MoC*, in "Design, Automation and Test in Europe, DATE'12", IEEE, 2012

[50] A. GIRARD, G. PAPPAS. *Approximation metrics for discrete and continuous systems*, in "IEEE Trans. on Automatic Control", 2007, vol. 52, n^{o} 5, pp. 782–798

[51] D. GIZOPOULOS, M. PSARAKIS, S. V. ADVE, P. RAMACHANDRAN, S. K. S. HARI, D. SORIN, A. MEIXNER, A. BISWAS, X. VERA. *Architectures for Online Error Detection and Recovery in Multicore Processors*, in "Design Automation and Test in Europe (DATE)", 2011

[52] C. GLYMOUR, D. DANKS, B. GLYMOUR, F. EBERHARDT, J. RAMSEY, R. SCHEINES, P. SPIRTES, C. M. TENG, J. ZHANG. *Actual causation: a stone soup essay*, in "Synthese", 2010, vol. 175, n^{o} 2, pp. 169–192

[53] F. C. GÄRTNER. *Fundamentals of Fault-Tolerant Distributed Computing in Asynchronous Environments*, in "ACM Computing Surveys", 1999, vol. 31, n^{o} 1

[54] S. HAAR, E. FABRE. *Diagnosis with Petri Net Unfoldings*, in "Control of Discrete-Event Systems", Lecture Notes in Control and Information Sciences, Springer, 2013, vol. 433, chap. 15

[55] T. HENZINGER, J. SIFAKIS. *The Embedded Systems Design Challenge*, in "Formal Methods 2006", Lecture Notes in Computer Science, Springer, 2006, vol. 4085

[56] H. JIN, K. RAVI, F. SOMENZI. *Fate and free will in error traces*, in "STTT", 2004, vol. 6, n^{o} 2, pp. 102–116, http://dx.doi.org/10.1007/s10009-004-0146-9

[57] R. KÜSTERS, T. TRUDERUNG, A. VOGT. *Accountability: definition and relationship to verifiability*, in "ACM Conference on Computer and Communications Security", 2010, pp. 526-535

[58] I. LANESE, C. A. MEZZINA, J.-B. STEFANI. *Reversing Higher-Order Pi*, in "21th International Conference on Concurrency Theory (CONCUR)", Lecture Notes in Computer Science, Springer, 2010, vol. 6269

[59] E. LE CORRONC, A. GIRARD, G. GÖSSLER. *Mode Sequences as Symbolic States in Abstractions of Incrementally Stable Switched Systems*, in "CDC - 52nd Conference on Decision and Control - 2013", IEEE, December 2013, pp. 3225-3230, http://hal.inria.fr/hal-00924815

[60] M. LI, F. YAO, H. YUAN. *An $O(n^2)$ Algorithm for Computing Optimal Continuous Voltage Schedules*, in "Annual Conference on Theory and Applications of Models of Computation, TAMC'17", Bern, Switzerland, LNCS, April 2017, vol. 10185, pp. 389–400

[61] P. MENZIES. *Counterfactual Theories of Causation*, in "Stanford Encyclopedia of Philosophy", E. ZALTA (editor), Stanford University, 2009, http://plato.stanford.edu/entries/causation-counterfactual

[62] M. MOORE. *Causation and Responsibility*, Oxford, 1999

[63] J. PEARL. *Causal inference in statistics: An overview*, in "Statistics Surveys", 2009, vol. 3, pp. 96-146

[64] P. RAMADGE, W. WONHAM. *Supervisory Control of a Class of Discrete Event Processes*, in "SIAM Journal on control and optimization", January 1987, vol. 25, n^{o} 1, pp. 206–230

[65] J. RUSHBY. *Partitioning for Safety and Security: Requirements, Mechanisms, and Assurance*, NASA Langley Research Center, 1999, n^{o} CR-1999-209347

[66] J.-B. STEFANI. *Components as Location Graphs*, in "11th International Symposium on Formal Aspects of Component Software", Bertinoro, Italy, Lecture Notes in Computer Science, September 2014, vol. 8997, https://hal.inria.fr/hal-01094208

[67] P. TABUADA. *Verification and Control of Hybrid Systems - A Symbolic Approach*, Springer, 2009

[68] K. TINDELL. *Using offset information to analyse static priority pre-emptively scheduled task sets*, Technical report YCS 182, University of York, Department of Computer Science, 1992, https://books.google.fr/books?id=qARQHAAACAAJ

[69] R. WILHELM, J. ENGBLOM, A. ERMEDAHL, N. HOLSTI, S. THESING, D. B. WHALLEY, G. BERNAT, C. FERDINAND, R. HECKMANN, T. MITRA, F. MUELLER, I. PUAUT, P. P. PUSCHNER, J. STASCHULAT, P. STENSTRÖM. *The Determination of Worst-Case Execution Times — Overview of the Methods and Survey of Tools*, in "ACM Trans. Embedd. Comput. Syst.", April 2008, vol. 7, n^{o} 3

[70] F. YAO, A. DEMERS, S. SHENKER. *A scheduling model for reduced CPU energy*, in "Proceedings of lEEE Annual Foundations of Computer Science", 1995, pp. 374–382

[71] E. YIP, P. S. ROOP, M. BIGLARI-ABHARI, A. GIRAULT. *Programming and Timing Analysis of Parallel Programs on Multicores*, in "International Conference on Application of Concurrency to System Design, ACSD'13", Barcelona, Spain, IEEE, July 2013, pp. 167–176, https://hal.inria.fr/hal-00842402

[72] E. YIP, P. S. ROOP, A. GIRAULT, M. BIGLARI-ABHARI. *Synchronous Deterministic Parallel Programming for Multicores with ForeC*, Inria - Research Centre Grenoble – Rhône-Alpes, August 2016, n^{o} RR-8943, https://hal.inria.fr/hal-01351552