

RESEARCH CENTRE

**Inria Centre
at Université Grenoble Alpes**

IN PARTNERSHIP WITH:

Université de Grenoble Alpes

2023

ACTIVITY REPORT

Project-Team

SPADES

**Sound Programming of Adaptive
Dependable Embedded Systems**

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

DOMAIN

Algorithmics, Programming, Software and
Architecture

THEME

Embedded and Real-time Systems

Inria

Contents

Project-Team SPADES	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	2
3 Research program	3
3.1 Design and Programming Models	3
3.2 Certified Real-Time Programming	4
3.3 Fault Management and Causal Analysis	4
4 Application domains	5
4.1 Industrial Applications	5
4.2 Current Industrial Cooperations	5
5 Social and environmental responsibility	5
5.1 Footprint of research activities	5
5.2 Impact of research results	5
6 New software, platforms, open data	6
6.1 New software	6
6.1.1 CertiCAN	6
6.1.2 cloudnet	6
6.1.3 LDDL	7
6.1.4 MASTAG	7
7 New results	7
7.1 Design and Programming Models	7
7.1.1 Dynamicity in dataflow models	8
7.1.2 The ForeC time-predictable programming language	8
7.2 Certified Real-Time Programming	9
7.2.1 A Markov Decision Process approach for energy minimization policies	9
7.2.2 Formal proofs for schedulability analysis of real-time systems	9
7.3 Fault Management and Causal Analysis	10
7.3.1 Causal Explanations for Embedded Systems	10
7.3.2 Causal Explanations in Concurrent Programs	10
7.3.3 Reversibility for concurrent and distributed debugging	11
7.3.4 A certified fault-tolerance technique for SEMTs	11
7.4 Transversal activity: ICT and the Anthropocene	11
8 Bilateral contracts and grants with industry	12
8.1 Bilateral contracts with industry	12
9 Partnerships and cooperations	12
9.1 National initiatives	12
9.1.1 ANR	12
9.1.2 Défi Inria	13
9.2 Exploratory Actions	13
10 Dissemination	14
10.1 Promoting scientific activities	14
10.1.1 Scientific events: organisation	14
10.1.2 Scientific events: selection	14
10.1.3 Journal	14
10.1.4 Invited talks	15

10.1.5 Leadership within the scientific community	15
10.1.6 Research administration	15
10.2 Teaching - Supervision - Juries	15
10.2.1 Teaching	15
10.2.2 Supervision	16
10.2.3 Juries	16
10.3 Popularization	16
10.3.1 Education	16
10.3.2 Interventions	17
11 Scientific production	17
11.1 Major publications	17
11.2 Publications of the year	18
11.3 Cited publications	19

Project-Team SPADES

Creation of the Project-Team: 2015 July 01

Keywords

Computer sciences and digital sciences

- 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

- B3.1. – Sustainable development
- B4.5. – Energy consumption
- B6.3.3. – Network Management
- B6.4. – Internet of things
- B6.6. – Embedded systems
- B9. – Society and Knowledge
- B9.9. – Ethics

1 Team members, visitors, external collaborators

Research Scientists

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

Faculty Member

- Xavier Nicollin [GRENOBLE INP -UGA, Associate Professor]

Post-Doctoral Fellow

- Alexandre Honorat [INRIA, Post-Doctoral Fellow]

PhD Students

- Baptiste De Goer De Herve [INRIA, from Oct 2023]
- Giovanni Fabbretti [INRIA]
- Aurélie Kong Win Chang [INRIA]
- Pietro Lami [INRIA]
- Thomas Mari [INRIA (-01/2023), CNRS (02/2023-05/2023), until Nov 2023]
- Aina Rasoldier [INRIA]

Interns and Apprentices

- Wiame Karmouni Tlemcani [INRIA, Intern, from May 2023 until Aug 2023]
- Alexander Obeid Guzman [INRIA, Intern, from Nov 2023]

Administrative Assistant

- Julia Di Toro [INRIA]

2 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 [28], 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

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. 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.

3.1 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 and dynamic reconfigurations 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 [48]. Witness component-based toolsets such as PTOLEMY [37], BIP [31], or the modular architecture frameworks used, for instance, in the automotive industry (AUTOSAR) [29]. 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.

3.2 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 [32]), 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 [30, 36], 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. [30]. For our part, we intend to focus on devising synchronous programming languages for distributed systems and precision-timed architectures.

3.3 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 [33, 46]. 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 [54], natural sciences, law [55], and statistics [57], 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 [51], to allow the diagnosis of faults in a complex concurrent system [47], or to enforce accountability [50], 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 [40]), or is broken (*e.g.*, by limiting fault propagation [59]).

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). We also consider the development of formal tools that can certify the result of industrial applications (see *e.g.*, CertiCAN in Sec. 7.2.2).

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. We also collaborate with RTaW regarding the integration of our CAN-bus analysis certifier (CertiCAN) in the RTaW-Pegase program suite.

5 Social and environmental responsibility

5.1 Footprint of research activities

With the help of the **GES 1point5** tool we have estimated the direct carbon footprint of our research activities in 2023. Our estimation is based on data gathered in a non-automated manner, as no tool automating the data extraction is available yet.

Professional travels, including the coming of jury members, amount to a total of 4,0 t CO₂e. Commute travels sum up to 1,8 t CO₂e. We purchased new hardware (2 computers) for a total of 649 kg CO₂e. We roughly estimate our share of INRIA services and building usage to 6 t CO₂e. Based on the above estimations, our carbon footprint totals 12,5 t CO₂e for the team, or an average of 0,9 t CO₂e per team member.

5.2 Impact of research results

Our research on certification and fault-tolerance aims at making embedded systems safer. Certified systems tend also to be simpler, less depending on updates and therefore less prone to obsolescence. A potential major application of causality analysis is to help establish liability for accidents caused by software errors.

On the other hand, our research may contribute to make more acceptable or even to promote many problematic systems such as IoT, drones, avionics, autonomous vehicles, ... with a potential negative environmental impact.

Sophie Quinton and Éric Tannier (from the BEAGLE team in Lyon), with the help of many colleagues, including some in the SPADES team, have set up a series of one-day workshops called “**Ateliers SEnS**” (for Sciences-Environnements-Sociétés), which offer a venue for members of the research community (in particular, but not limited to, researchers) to reflect on the social and environmental implications of their research. More than 50 Ateliers SEnS have taken place so far, all across France and beyond INRIA and the computer science field. Participants to a workshop can replicate it, and quite a few have already done so. Sophie Quinton has facilitated 6 Ateliers SEnS in 2023.

Research into the connection between ICT (Information and Communication Technologies) and the environmental crisis has started in 2020 within the SPADES team, see Section 7.4.

6 New software, platforms, open data

6.1 New software

6.1.1 CertiCAN

Name: 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 light-weight 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).

URL: <https://team.inria.fr/spades/certican/>

Authors: Xiaojie Guo, Pascal Fradet, Sophie Quinton

Contact: Xiaojie Guo

6.1.2 cloudnet

Name: Cloudnet

Keywords: Cloud configuration, Tosca, Docker Compose, Heat Orchestration Template, Alloy

Scientific Description: The multiplication of models, languages, APIs and tools for cloud and network configuration management raises heterogeneity issues that can be tackled by introducing a reference model. A reference model provides a common basis for interpretation for various models and languages, and for bridging different APIs and tools. The Cloudnet Computational Model formally specifies, in the Alloy specification language, a reference model for cloud configuration management. The Cloudnet software formally interprets several configuration languages in it, including the TOSCA configuration language, the OpenStack Heat Orchestration Template and the Docker Compose configuration language.

The use of the software shoes, for examples, how the Alloy formalization allowed us to discover several classes of errors in the OpenStack HOT specification.

Functional Description: Application of the Cloudnet model developed by Inria to software network deployment and reconfiguration description languages.

The Cloudnet model allows syntax and type checking for cloud configuration templates as well as their visualization (network diagram, UML deployment diagram). Three languages are addressed for the moment with the modules:

* Cloudnet TOSCA toolbox for TOSCA including NFV description * cloudnet-hot for HOT (Heat Orchestration Template) from OpenStack * cloudnet-compose for Docker Compose

We can use directly the software from an Orange web portal: <https://toscatoolbox.orange.com>

URL: <https://github.com/Orange-OpenSource/Cloudnet-TOSCA-toolbox>

Publication: hal-02940938v1

Contact: Philippe Merle

Participants: Philippe Merle, Jean-bernard Stefani, Roger Pissard-Gibollet, Souha Ben Rayana, Karine Guilloard, Meryem Ouzzif, Frédéric Klamm, Jean-Luc Coulin

Partner: Orange Labs

6.1.3 LDDL

Name: Coq proofs of circuit transformations for fault-tolerance

Keywords: Fault-tolerance, Transformation, Coq, Semantics

Functional Description: We have developed a Coq-based framework to formally verify the functional and fault-tolerance properties of circuit transformations. Circuits are described at the gate level using LDDL, a Low-level Dependent Description Language inspired from muFP. Our combinator language, equipped with dependent types, ensures that circuits are well-formed by construction (gates correctly plugged, no dangling wires, no combinational loops, ...). Fault-tolerance techniques can be described as transformations of LDDL circuits.

The framework has been used to prove the correctness of three fault-tolerance techniques for SETs (Single Event Transients): TMR (the classic triple modular redundancy) and two new time redundancy techniques developed within the Spades team: TTR and DTR. More recently, LDDL has been used to prove the correctness of TMR+, a modified TMR able to tolerate SEMTs (Single Event Multiple Transients) a more involved type of faults.

The specifications of the framework (LDDL syntax and semantics, libraries, tactics) are made of 5000 lines of Coq (excluding comments and blank lines). The correctness proofs of fault-tolerance techniques are made of 700 lines of Coq for TMR, 700 for TMR+, 3500 for TTR and 7000 for DTR.

URL: <https://team.inria.fr/spades/fthwproofs/>

Authors: Pascal Fradet, Dmitry Burlyayev, Vincent Bonczak

Contact: Pascal Fradet

6.1.4 MASTAG

Name: Memory Analyzer and Scheduler for Task Graphs

Keyword: Task scheduling

Functional Description: The MASTAG software computes sequential schedules of a task graph or an SDF graph in order to minimize its memory peak.

MASTAG is made of several components: (1) a set of local transformations that compress a task graph while preserving its optimal memory peak, (2) an optimized branch and bound algorithm able to find optimal schedules for medium sized (30-50 nodes) task graphs, (3) support to accommodate SDF graphs in particular, their conversion into task graphs and a suboptimal technique to reduce their size.

MASTAG finds optimal schedules in polynomial time for a wide range of directed acyclic task graphs (DAG), including trees and series-parallel DAG. On classic benchmarks, MASTAG always outperforms the state-of-the-art.

URL: <https://gitlab.inria.fr/spades-pub/mastag>

Authors: Alexandre Honorat, Pascal Fradet, Alain Girault

Contact: Alexandre Honorat

7 New results

7.1 Design and Programming Models

Participants: Pascal Fradet, Alain Girault, Alexandre Honorat.

7.1.1 Dynamicity in dataflow models

Dataflow Models of Computation (MoCs) are widely used in embedded systems, including multimedia processing, digital signal processing, telecommunications, and automatic control. One of the first and most popular dataflow MoCs, Synchronous Dataflow (SDF), provides static analyses to guarantee boundedness and liveness, which are key properties for embedded systems. However, SDF and most of its variants lack the capability to express the dynamism needed by modern streaming applications.

For many years, the Spades team has been working on more expressive and dynamic models that nevertheless allow the static analyses of boundedness and liveness. We have proposed several parametric dataflow models of computation (MoCs) (SPDF [38] and BPDF [56]), we have written a survey providing a comprehensive description of the existing parametric dataflow MoCs [34], we have studied *symbolic* analyses of dataflow graphs [35] and an original method to deal with lossy communication channels in dataflow graphs [39]. We have also proposed the RDF (Reconfigurable Dataflow) MoC [3] which allows *dynamic reconfigurations* of the *topology* of the dataflow graphs. RDF extends SDF with transformation rules that specify how the topology and actors of the graph may be dynamically reconfigured. 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, which in turn guarantees that they can be executed in bounded time and bounded memory. To the best of our knowledge, RDF is the only dataflow MoC allowing an arbitrary number of topological reconfigurations while remaining statically analyzable.

In 2022, we started an exploratory action (see Section 9.2) to study the potential of dataflow MoCs for the implementation of neural networks. We started by working on the reduction of the memory footprint of tasks graphs scheduled on uncore processors. This is motivated by the fact that some recent neural networks such as GPT-3, seen as tasks graphs, use too much memory and cannot fit on a single GPU.

We have proposed graph transformations that compress the given task graph while preserving its optimal memory peak. We have proved that these transformations always compress Series-Parallel Directed Acyclic Graphs (SP-DAGs) to a single node representing their optimal schedule [18]. For graphs that cannot be compressed to a single node, we have designed an optimized branch and bound algorithm able to find optimal schedules for medium sized (30-50 nodes) task graphs. Our approach also applies to SDF graphs after converting them to task graphs. However, since that conversion may produce very large graphs, we also propose a new suboptimal method, similar to Partial Expansion Graphs, to reduce the problem size. We evaluated our approach on classic benchmarks, on which we always outperform the state-of-the-art.

Another technique used by memory greedy neural networks is activity and gradient checkpointing (a.k.a. rematerialization), which recomputes intermediate values rather than keeping them in memory. We are currently studying rematerialization in the more general dataflow framework.

We have published a comprehensive paper about the Affine DataFlow Graph (ADFG) theory and software [13]. ADFG synthesizes task periods of real-time embedded applications modeled by SDF graphs. This paper concludes 10 years of work on the **ADFG** open-source software.

We have applied the ADFG theory to the domain of reconfigurable processors (FPGA) [12]. With the help of a few new equations, the theory of ADFG is adapted to minimize the buffer sizes of dataflow applications modeled by SDF graphs and executed on FPGA. This is particularly important for FPGAs which have a limited embedded memory. The corresponding open-source software **PREESM** is developed at INSA Rennes.

7.1.2 The ForeC time-predictable programming language

Embedded real-time systems are tightly integrated with their physical environment. Their correctness depends both on the outputs and timeliness of their computations. The increasing use of multi-core processors in such systems is pushing embedded programmers to be parallel programming experts. However, parallel programming is challenging because of the skills, experiences, and knowledge needed to avoid common parallel programming traps and pitfalls. We have proposed the ForeC synchronous multi-threaded programming language for the deterministic, parallel, and reactive programming of embedded multi-cores. The synchronous semantics of ForeC is designed to greatly simplify the understanding and debugging of parallel programs. ForeC ensures that ForeC programs can be compiled efficiently for parallel execution and be amenable to static timing analysis. ForeC's main innovation is its shared variable

semantics that provides thread isolation and deterministic thread communication. All ForeC programs are correct by construction and deadlock free because no non-deterministic constructs are needed. We have benchmarked our ForeC compiler with several medium-sized programs (e.g., a 2.274-line ForeC program with up to 26 threads and distributed on up to 10 cores, which was based on a 2.155-line non-multi-threaded C program). These benchmark programs show that ForeC can achieve better parallel performance than Esterel, a widely used imperative synchronous language for concurrent safety-critical systems, and is competitive in performance to OpenMP, a popular desktop solution for parallel programming (which implements classical multi-threading, hence is intrinsically non-deterministic). We also demonstrate that the worst-case execution time of ForeC programs can be estimated to a high degree of precision [15].

This topic has been a long-run effort, since we started working on ForeC in 2013 in the context of the PhD of Eugene Yip [61]. It took time to finalize this work, with the ultimate contribution in 2019 on multi-clock ForeC programs [45], paving the way for the long version article published in 2023 [15].

7.2 Certified Real-Time Programming

Participants: Pascal Fradet, Alain Girault, Sophie Quinton.

7.2.1 A Markov Decision Process approach for energy minimization policies

Since 2017 we have been working on a very general model of real-time systems, made of a single-core processor equipped with DVFS and an infinite sequence of preemptive real-time jobs. Each job J_i is characterized by the triplet (τ_i, w_i, d_i) , where τ_i is the *inter-arrival time* between J_i and J_{i-1} , w_i is the *actual size* of J_i , upper-bounded by the maximal size W , and d_i is the *relative deadline* of J_i , upper-bounded by Δ . The key point is that the system is *non-clairvoyant*, meaning that, at release time, w_i is not known until the job J_i actually terminates. What is available to the processor are the *statistical information* on the jobs' characteristics: release time, AET, and relative deadline. In this context, we have proposed a Markov Decision Process (MDP) solution to compute the optimal online speed policy guaranteeing that each job completes before its deadline and minimizing the energy consumption. To the best of our knowledge, our MDP solution is *the first to be optimal*. We have also provided counter examples to prove that the two previous state of the art algorithms, namely OA [60] and PACE [52], are both sub-optimal. Finally, we have proposed a new heuristic online speed policy called Expected Load (EL) that incorporates an aggregated term representing the future expected jobs into a speed equation similar to that of OA. A journal paper is currently under review.

Simulations show that our MDP solution outperforms the existing online solutions (OA, PACE, and EL), and can be very attractive in particular when the mean value of the execution time distribution is far from the WCET.

This was the topic of Stephan Plassart's PhD [58][41, 43, 42], funded by the CASERM Persyval project, who defended his PhD in June 2020.

7.2.2 Formal proofs for schedulability analysis of real-time systems

We contribute to Prosa [27], 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. the formal specification of real-time concepts;
2. a better understanding of the role played by some assumptions in existing proofs;
3. a formal verification and comparison of different analysis techniques; and
4. the certification of (results of) existing analysis techniques or tools.

We have developed CertiCAN, a tool produced using the Coq proof assistant, allowing the formal certification of CAN bus analysis results. CertiCAN is able to certify the results of industrial CAN analysis tools, even for large systems. We have described this work in a long journal article [11].

The work on the formalization in Prosa of Compositional Performance Analysis is still ongoing.

7.3 Fault Management and Causal Analysis

Participants: Gregor Goessler, Jean-Bernard Stefani, Aurélie Kong Win Chang, Thomas Mari, Giovanni Fabbretti, Pietro Lami, Pascal Fradet.

7.3.1 Causal Explanations for Embedded 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. We are currently exploring techniques that are able to extract, from an execution trace, the causally relevant part for a property violation.

In particular, as part of the SEC project, we are investigating how such techniques can be extended to classes of hybrid systems. As a first result we have studied the problem of explaining faults in real-time systems [53]. We have provided a formal definition of causal explanations on dense-time models, based on the well-studied formalisms of timed automata and zone-based abstractions. We have proposed a symbolic formalization to effectively construct such explanations, which we have implemented in a prototype tool. Basically, our explanations identify the parts of a run that move the system closer to the violation of an expected safety property, where safe alternative moves would have been possible.

We have recently generalized the work of [53] and defined *robustness functions* as a family of mappings from system states to a scalar that, intuitively, associate with each state its distance to the violation of a given safety requirement, e.g., in terms of the remaining number of bad system moves or of the time remaining to react. An explanation then summarizes the portions of the execution on which robustness decreases. However, as our instantiation of robustness in [53] is defined on a discrete abstraction, robustness may decrease in discrete steps once some timing threshold is crossed, thus exonerating the preceding absence of action. We are currently working on a truly hybrid definition of robustness functions that “anticipate” such thresholds, hence ensuring a smooth decrease indicating early when a dangerous event is approaching.

7.3.2 Causal Explanations in Concurrent Programs

As part of the DCore project on causal debugging of concurrent programs, the goal of Aurélie Kong Win Chang’s PhD thesis is to investigate the use of abstractions to construct causal explanations for Erlang programs. We are interested in developing abstractions that “compose well” with causal analyses, and understanding precisely how explanations found on the abstraction relate to explanations on the concrete system. It is worth noting that the presence of abstraction, which inherently comes with some induction and extrapolation processes, completely recasts the issue of reasoning about causality. Causal traces do no longer describe only potential scenarios in the concrete semantics, but also mix some approximation steps coming from the computation of the abstraction itself. Therefore, not all explanations are replayable counter-examples: they may contain some steps witnessing some lack of accuracy in the analysis. Vice versa, a research question to be addressed is how to define causal analyses that have a well understood behavior under abstraction.

In [19] we have formalized a small step semantics for a subset of Core Erlang that models, in particular, its monitoring and signal systems. Having a precise representation of these aspects is crucial to explain unexpected behaviors such as *concurrency bugs* stemming from non-determinism in the handling of messages.

We are currently working on a formalization of an abstract Erlang semantics that allows for a finite abstraction while still accounting for the exchanges of messages and signals between processes.

7.3.3 Reversibility for concurrent and distributed debugging

Concurrent and distributed debugging is a promising application of the notion of reversible computation [44]. As part of the ANR DCore project we contribute to the theory behind, and the development of the CauDEr reversible debugger for the Erlang programming language and system.

We have continued this year our work on two main themes: studying reversibility for distributed programs in presence of node and link failures with recovery, and studying reversibility for concurrent programs using a shared memory concurrency model.

Concerning reversibility for distributed programs, we have developed a novel process calculus, called $D\pi$ FR [26]. $D\pi$ FR provides a good basis for formally modeling Erlang distribution, including the behaviour of Erlang systems in presence of crash failure and recovery for nodes and links. We have developed a full behavioral theory for $D\pi$ FR, in the form of a weak observational equivalence, which we have proved fully abstract with respect to the contextual equivalence for the calculus. This work is under submission for publication. We have also started studying reversibility in $D\pi$ FR, considering in particular the difficult case where node failures imply the loss of causality information in the reversible operational semantics.

Concerning reversibility for shared memory concurrency, we have developed a modular operational semantics framework for defining different shared memory concurrency models, including various lock-based weak memory models and transactional memory models. We have proved strong equivalence results between the original formal operational semantics of these different memory models and the operational semantics obtained using our framework. We have also started working on a general theory for reversing synchronization products of transition systems with independence with the hope to directly apply it to our shared memory framework.

7.3.4 A certified fault-tolerance technique for SEMTs

Digital circuits are subject to transient faults caused by high-energy particles. As technology scales down, a single particle becomes likely to induce transients faults in several adjacent components. These single-event multiple transients (SEMTs) are becoming a major issue for modern digital circuits.

We have studied how to formalize SEMTs and how the standard triple modular redundancy (TMR) technique can be modified so that, along with some placement constraints, it completely masks SEMTs [25]. We specified this technique, denoted by TMR+, as a circuit transformation on the LDDL syntax (see 6.1.3) and the fault models for SEMTs as particular semantics of LDDL. We show that, for any circuit, its transformation by TMR+ masks all faults of the considered SEMT fault model. All this development was formalized in the Coq proof assistant where fault-tolerance properties are expressed and formally proved.

7.4 Transversal activity: ICT and the Anthropocene

Participants: Martin Bodin, Baptiste De Goer De Herve, Pascal Fradet, Alain Girault, Gregor Goessler, Xavier Nicollin, Roger Pissard, Sophie Quinton, Aina Rasoldier, Jean-Bernard Stefani.

Digital technologies are often presented as a powerful ally in the fight against climate change (see *e.g.*, the discourse around the “convergence of the digital and the ecological transitions”). The SPADES team has started working together on a project proposal to investigate the current role played by ICT in the Anthropocene as well as new approaches to their design. We have identified the following main challenges: How do local measures meant to reduce the environmental impact of ICT relate (or not) to global effects? What can we learn from, and what are the limits of, current quantitative approaches for environmental impact assessment and their use for public debate and policy making? Which criteria could/should we take into account to design more responsible computer systems (other than efficiency, which is already well covered and subject to huge rebound effects in the case of digital technologies)? To come up with a solid research agenda, we are thus studying the state of the art of many new topics [14],

including STS (Science and Technology Studies), low tech software and hardware, lifecycle assessment, (digital) commons... A new network of collaborations is also in the making, in particular with colleagues from social sciences. See [23] for a possible topic of interdisciplinary research. Besides, Baptiste de Goër has just started a PhD focusing on how to integrate ICT-related sustainability issues in computer science courses [22].

In the context of Aina Rasoldier's PhD, we have been working on estimating the potential of ridesharing as a solution for reducing the GHG emissions of commuting. Ridesharing is one of the solutions put forward by local authorities to reduce the carbon footprint of individual travel. But it is far from granted that this solution can achieve the long term objectives stated by the French government in its "Stratégie Nationale Bas Carbone", and declined locally in the "Plan de Déplacements Urbains" of the Grenoble metropolitan area. We have focused on the *daily peer-to-peer ridesharing* (also called car-pooling), in which people travel using the personal vehicle of one of them. Moreover, ridesharing is *prearranged* (also called static, or organized) ridesharing, which supposes that people know in advance their travel needs for the entire day and use digital platforms finding a match (i.e., finding passengers when one is driving her/his own car, or finding a car when one is a passenger). We have considered two matching schemes between drivers and passengers: on the one hand *identical ridesharing*, where drivers and passengers can only carpool if their origins (and destinations) are *close*, and on the other hand *inclusive ridesharing*, where passengers can be picked up and dropped off along the driver's route if the passenger's origin and destination are *close* to the driver's route. In both cases, *close* refers to a maximal walking distance for the passenger to reach the driver, and to a maximal time between her or his desired starting time and the driver's actual starting time. Our evaluation of the ridesharing potential is based on a synthetic travel demand computed using the existing software from Hörl et al. [49] that we ran on the public data for the Grenoble metropolitan area. Based on this population synthesis, we have developed an ad-hoc matching algorithm to evaluate the maximum potential offered by ridesharing. Extensive simulations performed with our algorithm show that to reach the goals stated in the Grenoble PDU would require at least 55% of the local population to adopt ridesharing on a daily basis, a ratio that seems completely out of reach in the near future (this ratio was obtained with the following parameters: maximal walking distance for the passengers equal to 1 km and maximal delay equal to 15 min). This preliminary study shows that betting solely on digital solutions (here, digital ridesharing platforms) to reduce our carbon footprint will not be sufficient [20].

8 Bilateral contracts and grants with industry

Participants: Jean-Bernard Stefani.

8.1 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 was one of the two co-directors of the lab, till Feb. 2020). 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 the verification of system configurations in cloud computing environments and software-defined networks.

9 Partnerships and cooperations

9.1 National initiatives

9.1.1 ANR

DCORE

Participants: Gregor Goessler, Jean-Bernard Stefani, Giovanni Fabbretti, Pietro Lami, Aurélie Kong Win Chang.

DCORE is an ANR project between Inria project teams ANTIQUE, FOCUS and SPADES, and the IRIF lab, running from 2019 to 2024.

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.

9.1.2 Défi Inria

LiberAbaci

Participants: Martin Bodin.

LIBERABACI is a project between Inria project teams CAMBIUM, CAMUS, GALLINETTE, πr^2 , SPADES, STAMP, TOCCATA, and the Laboratoire d’Informatique de Paris-Nord. The overall objective is to study how one could use the COQ proof assistant in a Mathematical course in the University to help teaching proofs. At Spades, Martin Bodin is working with IREM de Grenoble to involve math teachers and didactic researchers to the project.

9.2 Exploratory Actions

DF4DL

Participants: Pascal Fradet, Alain Girault, Alexandre Honorat.

The DF4DL action is funded by Inria’s DGDS. It aims at exploring the use of the dataflow model of computation to better program deep neural networks.

As a first step, we have studied the problem of minimizing the peak memory requirement for the execution of a dataflow graph. This is of paramount importance for deep neural networks since the largest ones cannot fit on a single core due to their very high memory requirement. We have proposed different techniques in order to find a sequential schedule minimizing the memory peak (see 7.1.1).

Another technique used by memory greedy neural networks is rematerialization which recomputes intermediate values rather than keeping them in memory. We are currently studying rematerialization in the dataflow framework.

SIA

Participants: Baptiste De Goer De Herve, Sophie Quinton.

The SIA Exploratory Research project, supported by INRIA's DGDS, funds the PhD work of Baptiste de Goër and provides funding for an upcoming postdoctoral fellow in Sciences and Technology Studies.

The goal of the project is to provide interdisciplinary foundations for studying the complex relationship between computer science, information and communication technologies (ICT), society and the environment. We approach the problem from three complementary perspectives: 1) by contributing to an interdisciplinary overview of the state of knowledge on the environmental impacts of ICT; 2) by studying the complex connection between computer science and the Anthropocene through the way it is and could be taught in secondary schools; 3) by exploring, at a local scale, the possibility to deploy frugal or low tech alternatives to existing digital systems, following a participatory approach.

10 Dissemination

Participants: Martin Bodin, Fradet Pascal, Girault Alain, Gregor Goessler, Jean-Bernard Stefani, Nicollin Xavier.

10.1 Promoting scientific activities

10.1.1 Scientific events: organisation

General chair, scientific chair

- Alain Girault, vice-general chair of the [ESWEEK 2023](#) international conference.

Member of the organizing committees

- Gregor Gössler co-organized the colloquium [FunCausal](#) on Fundamental Challenges in Causality.
- Sophie Quinton co-organized a workshop (journée d'étude) on teaching the environmental consequences of ICT.

10.1.2 Scientific events: selection

Member of the conference program committees

- Gregor Gössler served in the PC of the ETAPS workshop CREST'23.
- Sophie Quinton served in the PC of the Undone Computer Science 2024 conference.

Reviewer

- Sophie Quinton was an external reviewer for DATE 2024.

10.1.3 Journal

Member of the editorial boards

- Alain Girault, EURASIP Journal on Embedded Systems (since 2005); Real-Time Systems Journal (since 2020).

Reviewer - reviewing activities

- Alain Girault, ACM Trans. on Embedded Computing Systems.
- Gregor Gössler, Elsevier Artificial Intelligence.

10.1.4 Invited talks

- Sophie Quinton gave invited talks at the DATE 2023 conference, at the SICT summer school as well as at the EEATS doctoral school PhD day and at the Institut Néel in Grenoble.

10.1.5 Leadership within the scientific community

- Sophie Quinton is a member of the ECRTS Advisory Board.
- Sophie Quinton co-chairs a working group of the GDR CIS associated with the Center for Internet and Society focused on environmental issues.

10.1.6 Research administration

- Pascal Fradet is head of the committee for doctoral studies (“Responsable du comité des études doctorales”) of the Inria Grenoble research center. He is the local correspondent for the young researchers Inria mission (“Mission jeunes chercheurs”) and serve as the substitute of the director of the Inria Grenoble research center at the doctoral school council (MSTII).
- Alain Girault is Deputy Scientific Director at Inria for the domain “Algorithmics, Programming, Software and Architecture” (since 2019).
- Alain Girault was president of the Inria Senior Researchers Admission 2023 jury (DR2).
- Gregor Gössler is member of the “commission of scientific jobs” of the Inria Grenoble research center.
- Jean-Bernard Stefani was member of the Inria Grenoble Junior Researches Admissibility 2023 jury (CRCN).
- Sophie Quinton leads the SENs-GRA group which hosts discussions and proposes actions regarding the environmental and societal impact of our research at Inria Grenoble.
- Sophie Quinton was a member of the CRCN 2023 hiring committee in Rennes.

10.2 Teaching - Supervision - Juries

10.2.1 Teaching

- Licence : Pascal Fradet, Théorie des Langages 1, 18 HeqTD, niveau L3, Grenoble INP (Ensimag), France
- Licence : Pascal Fradet, Modèles de Calcul : λ -calcul, CM & TD, 30 HeqTD, niveau L3, Univ. Grenoble Alpes, France
- Licence : Xavier Nicollin, Théorie des Langages 1, 40,5 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, Modèles de Calcul : Machines de Turing, 30 HeqTD, niveau L3, Univ. Grenoble Alpes, France
- Master : Xavier Nicollin, Analyse de Code pour la Sûreté et la Sécurité, 45 HeqTD, niveau M1, Grenoble INP (Ensimag), France
- Master : Xavier Nicollin, Algorithmique et Optimisation Discrète, 18 HeqTD, niveau M1, Grenoble INP (Ensimag), France
- Master : Xavier Nicollin, Fondements Logiques pour l’Informatique, 19,5 HeqTD, niveau M1, Grenoble INP (Ensimag), France

- Licence : Martin Bodin, Modèles de Calcul : λ -calcul, 12 HeqTD, niveau L3, Univ. Grenoble Alpes, France
- Licence : Martin Bodin, Théorie des Langages 2, 18 HeqTD, niveau L3, Grenoble INP (Ensimag), France
- Licence : Alain Girault, Modèles de Calcul : λ -calcul, 12 HeqTD, niveau L3, Univ. Grenoble Alpes, France
- Licence : Sophie Quinton contributed to two 3h workshops for first-year students at the Ensimag Engineering School to introduced them to the environmental impacts of ICT.
- École doctorale: Sophie Quinton gave a 3h course "Sciences, environnements, sociétés" at the College des Écoles Doctorales.
- Sophie Quinton co-supervised a one-week sociological study by students from the École des Mines de Paris on water management in the Grenoble area, with a focus on ICT related aspects.

10.2.2 Supervision

- Alain Girault and Sophie Quinton: PhD in progress: Aina Rasoldier, ICT in the Anthropocene: Technical and social challenges at the local scale.
- Gregor Gössler: PhD completed: Thomas Mari, "Construction of Safe Explainable Cyber-physical systems"; Grenoble INP; defended in November 2023; co-advised by Gregor Gössler and Thao Dang.
- Gregor Gössler: PhD in progress: Aurélie Kong Win Chang, "Abstractions for causal analysis and explanations in concurrent programs"; since January 2021; co-advised by Gregor Gössler and Jérôme Feret.
- Gregor Gössler: pre-thesis contract: Alexander Obeid Guzman, "Inference of causal models for networks from single observations"; since November 2023.
- Jean-Bernard Stefani: PhD in progress: Giovanni Fabbretti on reversibility for distributed programs (UGA), Pietro Lami on reversibility for shared memory concurrent programs (UGA and U. Bologna), Boubacar Diarra on verification of Kubernetes configurations (U. Lille).
- Sophie Quinton: PhD in progress: Baptiste de Goër, "Teaching ICT-related sustainability issues in computer science courses".

10.2.3 Juries

- Alain Girault, President of the PhD jury of Kevin Zagalo, Sorbonne Université; 2023.
- Alain Girault, Invited to the PhD jury of Baptiste Pauget, ENS-PSL; 2023.
- Alain Girault, Examiner in the PhD jury of Lou Grimal, UTT; 2023.

10.3 Popularization

10.3.1 Education

- Martin Bodin and Alain Girault, "Le théorème des quatre couleurs", Lecture Maths C2+ for high-school pupils, Grenoble, June 2023.

10.3.2 Interventions

- Martin Bodin was interviewed on the Twitch channel Chercheur-es de montagne.
- Martin Bodin with Emmanuel Beffara: creation and experimentation of an activity about Logic for the Fête de la science.
- Sophie Quinton was a member of the scientific committee of the GAES (groupe artistique d'exploration scientifique) 2024.
- Sophie Quinton and Baptiste de Goër collaborate with two teachers of the Lycée Stendhal on their project teaching about the environmental impacts of ICT.

11 Scientific production

11.1 Major publications

- [1] A. Abdi, A. Girault and 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* 30.10 (1st Oct. 2019), pp. 2193–2210. DOI: [10.1109/TPDS.2019.2906172](https://doi.org/10.1109/TPDS.2019.2906172). URL: <https://hal.inria.fr/hal-02400019>.
- [2] A. Bouakaz, P. Fradet and A. Girault. 'A Survey of Parametric Dataflow Models of Computation'. In: *ACM Trans. Design Autom. Electr. Syst.* 22.2 (2017), 38:1–38:25. DOI: [10.1145/2999539](https://doi.org/10.1145/2999539).
- [3] P. Fradet, A. Girault, R. Krishnaswamy, X. Nicollin and A. Shafiei. 'RDF: A Reconfigurable Dataflow Model of Computation'. In: *ACM Transactions on Embedded Computing Systems (TECS)* (19th Dec. 2022). DOI: [10.1145/3544972](https://doi.org/10.1145/3544972). URL: <https://hal.inria.fr/hal-03940615>.
- [4] P. Fradet, X. Guo and S. Quinton. 'CertiCAN : Certifying CAN Analyses and Their Results'. In: *Real-Time Systems* 59.2 (14th Mar. 2023), pp. 160–198. DOI: [10.1007/s11241-023-09393-2](https://doi.org/10.1007/s11241-023-09393-2). URL: <https://inria.hal.science/hal-03941096>.
- [5] G. Frehse, A. Hamann, S. Quinton and 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, Dec. 2014. URL: <https://hal.inria.fr/hal-01097622>.
- [6] A. Girard, G. Gössler and S. Mouelhi. 'Safety Controller Synthesis for Incrementally Stable Switched Systems Using Multiscale Symbolic Models'. In: *IEEE Transactions on Automatic Control* 61.6 (2016), pp. 1537–1549. DOI: [10.1109/TAC.2015.2478131](https://doi.org/10.1109/TAC.2015.2478131). URL: <https://hal.archives-ouvertes.fr/hal-01197426>.
- [7] G. Gössler and J.-B. Stefani. 'Causality analysis and fault ascription in component-based systems'. In: *Theoretical Computer Science* 837 (2020), pp. 158–180. DOI: [10.1016/j.tcs.2020.06.010](https://doi.org/10.1016/j.tcs.2020.06.010). URL: <https://hal.inria.fr/hal-02927216>.
- [8] I. Lanese, C. A. Mezzina and J.-B. Stefani. 'Reversibility in the higher-order π -calculus'. In: *Theoretical Computer Science* 625 (2016), pp. 25–84. DOI: [10.1016/j.tcs.2016.02.019](https://doi.org/10.1016/j.tcs.2016.02.019). URL: <https://hal.inria.fr/hal-01303090>.
- [9] A. Rasoldier, J. Combaz, A. Girault, K. Marquet and S. Quinton. 'How realistic are claims about the benefits of using digital technologies for GHG emissions mitigation?' In: *LIMITS 2022 - Eighth Workshop on Computing within Limits*. Virtual, France, 21st June 2022. URL: <https://hal.inria.fr/hal-03949261>.
- [10] P. Roux, S. Quinton and M. Boyer. 'A Formal Link Between Response Time Analysis and Network Calculus'. In: *ECRTS 2022 - 34th Euromicro Conference on Real-Time Systems*. Modene, Italy, 5th July 2022. DOI: [10.4230/DARTS.8.1.3](https://doi.org/10.4230/DARTS.8.1.3). URL: <https://hal.science/hal-03770727>.

11.2 Publications of the year

International journals

- [11] P. Fradet, X. Guo and S. Quinton. ‘CertiCAN : Certifying CAN Analyses and Their Results’. In: *Real-Time Systems* 59.2 (14th Mar. 2023), pp. 160–198. DOI: [10.1007/s11241-023-09393-2](https://doi.org/10.1007/s11241-023-09393-2). URL: <https://inria.hal.science/hal-03941096>.
- [12] A. Honorat, M. Dardaillon, H. Miomandre and J.-F. Nezan. ‘Automated Buffer Sizing of Dataflow Applications in a High-Level Synthesis Workflow’. In: *ACM Transactions on Reconfigurable Technology and Systems (TRETs)* (29th Sept. 2023), pp. 1–26. DOI: [10.1145/3626103](https://doi.org/10.1145/3626103). URL: <https://hal.science/hal-04237266>.
- [13] A. Honorat, H. N. Tran, T. Gautier, L. Besnard, S. S. Bhattacharyya and J.-P. Talpin. ‘Real-Time Fixed Priority Scheduling Synthesis using Affine DataFlow Graphs: from Theory to Practice’. In: *ACM Transactions on Embedded Computing Systems (TECS)* (18th Aug. 2023), pp. 1–30. DOI: [10.1145/3615586](https://doi.org/10.1145/3615586). URL: <https://hal.science/hal-04200195>.
- [14] G. Roussilhe, A.-L. Ligozat and S. Quinton. ‘A long road ahead: a review of the state of knowledge of the environmental effects of digitization’. In: *Current Opinion in Environmental Sustainability* 62 (June 2023), p. 101296. DOI: [10.1016/j.cosust.2023.101296](https://doi.org/10.1016/j.cosust.2023.101296). URL: <https://hal.science/hal-04448683>.
- [15] E. Yip, A. Girault, P. S. Roop and M. Biglari-Abhari. ‘Synchronous Deterministic Parallel Programming for Multi-Cores with ForeC’. In: *ACM Transactions on Programming Languages and Systems (TOPLAS)* 45.2 (26th June 2023), pp. 1–74. DOI: [10.1145/3591594](https://doi.org/10.1145/3591594). URL: <https://hal.science/hal-04338823>.

International peer-reviewed conferences

- [16] E. Beffara, M. Bodin, N. Mandran and R. Molinier. ‘Instrumentation de l’association de registres sémiotiques dans un assistant de preuve’. In: EIAH2023 - 11ème Conférence sur les Environnements Informatiques pour l’Apprentissage Humain. Brest, France, 12th June 2023, pp. 1–5. URL: <https://hal.science/hal-04096240>.
- [17] B. Diarra, K. Guillooard, M. Ouzzif, P. Merle and J.-B. Stefani. ‘In-depth analysis of Kubernetes manifest verification tools for robust CNF deployment’. In: ICIN 2024 - Conference on Innovation in Clouds, Internet and Networks. Paris, France, 2024, pp. 1–8. URL: <https://inria.hal.science/hal-04421758>.
- [18] P. Fradet, A. Girault and A. Honorat. ‘Sequential Scheduling of Dataflow Graphs for Memory Peak Minimization’. In: *Proceedings of the 24th ACM SIGPLAN/SIGBED International Conference on Languages, Compilers, and Tools for Embedded Systems (LCTES '23)*. LCTES 2023 - 24th ACM SIGPLAN/SIGBED International Conference on Languages, Compilers, and Tools for Embedded Systems. Orlando (FL), United States: ACM, 13th June 2023, pp. 76–86. DOI: [10.1145/3589610.3596280](https://doi.org/10.1145/3589610.3596280). URL: <https://hal.science/hal-04163123>.
- [19] A. Kong Win Chang, J. Feret and G. Gössler. ‘A Semantics of Core Erlang with Handling of Signals’. In: *ACM Digital Library. Erlang 2023 - 22nd ACM SIGPLAN International Workshop on Erlang*. Seattle WA, United States: ACM, 2023, pp. 31–38. DOI: [10.1145/3609022.3609417](https://doi.org/10.1145/3609022.3609417). URL: <https://hal.science/hal-04222884>.
- [20] A. Rasoldier, A. Girault, S. Quinton, J. Combaz and K. Marquet. ‘Assessing the Potential of Carpooling for Reducing Vehicle Kilometers Traveled’. In: ICT4S 2023 - 9th International Conference on Information and Communications Technology for Sustainability. Rennes, France: IEEE, 2023, pp. 120–131. DOI: [10.1109/ICT4S58814.2023.00021](https://doi.org/10.1109/ICT4S58814.2023.00021). URL: <https://inria.hal.science/hal-04401006>.

National peer-reviewed Conferences

- [21] J. Abou-Samra, Y. Zakowski and M. Bodin. ‘Effectful Programming across Heterogeneous Computations -Work in Progress’. In: *Journées Francophones des Langages Applicatifs*. JFLA 2023 - 34èmes Journées Francophones des Langages Applicatifs. Praz-sur-Arly, France, 2023, pp. 7–23. URL: <https://hal.science/hal-03886975>.

Conferences without proceedings

- [22] B. de Goër, M. Hersch and S. Quinton. ‘Informatique et durabilité, une difficile transposition didactique’. In: *Didapro 10 - Didactique de l’informatique et des STIC*. Louvain-la-Neuve, Belgium, 30th Jan. 2024. URL: <https://hal.science/hal-04448717>.
- [23] S. Quinton and J.-B. Stefani. ‘Taking conviviality seriously (extended abstract)’. In: *Undone Computer Science 2024*. Nantes, France, 5th Feb. 2024. URL: <https://hal.science/hal-04448759>.

Reports & preprints

- [24] C. K. Assaad, E. Devijver, E. Gaussier, G. Gössler and A. Meynaoui. *Identifiability of total effects from abstractions of time series causal graphs*. 24th Oct. 2023. URL: <https://hal.science/hal-04250602>.
- [25] V. Bonczak and P. Fradet. *A formally verified circuit transformation to tolerate SEMTs*. RR-9523. Inria Grenoble - Rhône-Alpes, 9th Oct. 2023, pp. 1–25. URL: <https://inria.hal.science/hal-04236869>.
- [26] G. Fabbretti, I. Lanese and J.-B. Stefani. *A Behavioral Theory For Crash Failures and Erlang-style Recoveries In Distributed Systems*. RR-9511. Inria, 9th June 2023. URL: <https://hal.science/hal-04123758>.

11.3 Cited publications

- [27] *A Library for formally proven schedulability analysis*. URL: <http://prosa.mpi-sws.org/>.
- [28] ARTEMIS Joint Undertaking. *ARTEMIS Strategic Research Agenda*. 2011.
- [29] *Automotive Open System Architecture*. 2003. URL: <http://www.autosar.org>.
- [30] E. Bainomugisha, A. Carreton, T. Van Cutsem, S. Mostinckx and W. De Meuter. ‘A Survey on Reactive Programming’. In: *ACM Computing Surveys* 45.4 (2013).
- [31] A. Basu, S. Bensalem, M. Bozga, J. Combaz, M. Jaber, T.-H. Nguyen and J. Sifakis. ‘Rigorous Component-Based System Design Using the BIP Framework’. In: *IEEE Software* 28.3 (2011).
- [32] A. Benveniste, P. Caspi, S. A. Edwards, N. Halbwachs, P. Le Guernic and R. de Simone. ‘The synchronous languages 12 years later’. In: *Proceedings of the IEEE* 91.1 (2003).
- [33] S. Borkar. ‘Designing Reliable Systems from Unreliable Components: The Challenges of Transistor Variability and Degradation’. In: *IEEE Micro* 25.6 (2005).
- [34] A. Bouakaz, P. Fradet and A. Girault. ‘A Survey of Parametric Dataflow Models of Computation’. In: *ACM Transactions on Design Automation of Electronic Systems (TODAES)* (Jan. 2017). URL: <https://hal.inria.fr/hal-01417126>.
- [35] A. Bouakaz, P. Fradet and A. Girault. ‘Symbolic Analyses of Dataflow Graphs’. In: *ACM Transactions on Design Automation of Electronic Systems (TODAES)* (Jan. 2017). URL: <https://hal.inria.fr/hal-01417146>.
- [36] R. Davis and A. Burns. ‘A Survey of Hard Real-Time Scheduling for Multiprocessor Systems’. In: *ACM Computing Surveys* 43.4 (2011).
- [37] J. Eker, J. W. Janneck, E. A. Lee, J. Liu, X. Liu, J. Ludvig, S. Neuendorffer, S. Sachs and Y. Xiong. ‘Taming heterogeneity - the Ptolemy approach’. In: *Proceedings of the IEEE* 91.1 (2003).
- [38] P. Fradet, A. Girault and P. Polpavko. ‘SPDF: A schedulable parametric data-flow MoC’. In: *Design, Automation and Test in Europe, DATE’12*. IEEE, 2012.

- [39] P. Fradet, A. Girault, L. Jamshidian, X. Nicollin and 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, Oct. 2017. URL: <https://hal.inria.fr/hal-01666568>.
- [40] F. C. Gärtner. ‘Fundamentals of Fault-Tolerant Distributed Computing in Asynchronous Environments’. In: *ACM Computing Surveys* 31.1 (1999).
- [41] B. Gaujal, A. Girault and S. Plassart. ‘A Pseudo-Linear Time Algorithm for the Optimal Discrete Speed Minimizing Energy Consumption’. In: *Discrete Event Dynamic Systems* 31 (2021), pp. 163–184. DOI: [10.1007/s10626-020-00327-9](https://doi.org/10.1007/s10626-020-00327-9). URL: <https://hal.science/hal-03030416>.
- [42] B. Gaujal, A. Girault and S. Plassart. ‘Dynamic Speed Scaling Minimizing Expected Energy Consumption for Real-Time Tasks’. In: *Journal of Scheduling* (July 2020), pp. 1–25. DOI: [10.1007/s10951-020-00660-9](https://doi.org/10.1007/s10951-020-00660-9). URL: <https://hal.inria.fr/hal-02888573>.
- [43] B. Gaujal, A. Girault and S. Plassart. ‘Feasibility of on-line speed policies in real-time systems’. In: *Real-Time Systems* (Apr. 2020). DOI: [10.1007/s11241-020-09347-y](https://doi.org/10.1007/s11241-020-09347-y). URL: <https://hal.inria.fr/hal-02557148>.
- [44] E. Giachino, I. Lanese and C. A. Mezzina. ‘Causal-Consistent Reversible Debugging’. In: *17th International Conference Fundamental Approaches to Software Engineering (FASE)*. Vol. 8411. Lecture Notes in Computer Science. 2014, pp. 370–384.
- [45] A. Girault, N. Hili, É. Jenn and 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, Sept. 2019, pp. 1–8. DOI: [10.1109/FDL.2019.8876950](https://doi.org/10.1109/FDL.2019.8876950). URL: <https://hal.inria.fr/hal-02399998>.
- [46] D. Gizopoulos, M. Psarakis, S. V. Adve, P. Ramachandran, S. K. S. Hari, D. Sorin, A. Meixner, A. Biswas and X. Vera. ‘Architectures for Online Error Detection and Recovery in Multicore Processors’. In: *Design Automation and Test in Europe (DATE)*. 2011.
- [47] S. Haar and E. Fabre. ‘Diagnosis with Petri Net Unfoldings’. In: *Control of Discrete-Event Systems*. Vol. 433. Lecture Notes in Control and Information Sciences. Springer, 2013. Chap. 15.
- [48] T. Henzinger and J. Sifakis. ‘The Embedded Systems Design Challenge’. In: *Formal Methods 2006*. Vol. 4085. Lecture Notes in Computer Science. Springer, 2006.
- [49] S. Hörl and M. Balac. ‘Synthetic population and travel demand for Paris and Île-de-France based on open and publicly available data’. In: *Transportation research. Part C, Emerging technologies* 130 (Sept. 2021), p. 103291. DOI: [10.1016/j.trc.2021.103291](https://doi.org/10.1016/j.trc.2021.103291). URL: <https://hal.science/hal-03286361>.
- [50] R. Küsters, T. Truderung and A. Vogt. ‘Accountability: definition and relationship to verifiability’. In: *ACM Conference on Computer and Communications Security*. 2010, pp. 526–535.
- [51] I. Lanese, C. A. Mezzina and J.-B. Stefani. ‘Reversing Higher-Order Pi’. In: *21th International Conference on Concurrency Theory (CONCUR)*. Vol. 6269. Lecture Notes in Computer Science. Springer, 2010.
- [52] J. Lorch and A. Smith. ‘PACE: A New Approach to Dynamic Voltage Scaling’. In: *IEEE Trans. on Computers* 53.7 (2004), pp. 856–869.
- [53] T. Mari, T. Dang and G. Gössler. ‘Explaining Safety Violations in Real-Time Systems’. In: *FORMATS 2021 - Formal Modeling and Analysis of Timed Systems*. Paris, France, Aug. 2021, pp. 100–116. DOI: [10.1007/978-3-030-85037-1_7](https://doi.org/10.1007/978-3-030-85037-1_7). URL: <https://hal.inria.fr/hal-03348010>.
- [54] P. Menzies. ‘Counterfactual Theories of Causation’. In: *Stanford Encyclopedia of Philosophy*. Ed. by E. Zalta. Stanford University, 2009. URL: <http://plato.stanford.edu/entries/causation-counterfactual>.
- [55] M. Moore. *Causation and Responsibility*. Oxford, 1999.
- [56] V. Bebelis, P. Fradet, A. Girault and B. Lavigueur. ‘BPDF: A Statically Analyzable Dataflow Model with Integer and Boolean Parameters’. In: *International Conference on Embedded Software, EMSOFT’13*. Montreal, Canada: ACM, Sept. 2013.

-
- [57] J. Pearl. 'Causal inference in statistics: An overview'. In: *Statistics Surveys* 3 (2009), pp. 96–146.
- [58] S. Plassart. 'Online optimization in dynamic real-time systems'. Theses. Université Grenoble Alpes [2020-....], June 2020. URL: <https://tel.archives-ouvertes.fr/tel-02990646>.
- [59] J. Rushby. *Partitioning for Safety and Security: Requirements, Mechanisms, and Assurance*. Tech. rep. CR-1999-209347. NASA Langley Research Center, 1999.
- [60] F. Yao, A. Demers and S. Shenker. 'A scheduling model for reduced CPU energy'. In: *Proceedings of IEEE Annual Foundations of Computer Science*. 1995, pp. 374–382.
- [61] E. Yip, P. S. Roop, M. Biglari-Abhari and 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. URL: <https://hal.inria.fr/hal-00842402>.