

2025 Activity Report


RESEARCH CENTRE: Inria Centre at Université Grenoble Alpes
IN PARTNERSHIP WITH: Université de Grenoble Alpes

Project-Team

SPADES

Sound Programming of Adaptive Dependable
Embedded Systems

In collaboration with Laboratoire d'Informatique de Grenoble (LIG)



Project-Team SPADES

Creation of the Project-Team: 2015 July 01

Each year, Inria research teams publish an Activity Report presenting their work and results over the reporting period. These reports follow a common structure, with some optional sections depending on the specific team. They typically begin by outlining the overall objectives and research programme, including the main research themes, goals, and methodological approaches. They also describe the application domains targeted by the team, highlighting the scientific or societal contexts in which their work is situated. The reports then present the highlights of the year, covering major scientific achievements, software developments, or teaching contributions. When relevant, they include sections on software, platforms, and open data, detailing the tools developed and how they are shared. A substantial part is dedicated to new results, where scientific contributions are described in detail, often with subsections specifying participants and associated keywords. Finally, the Activity Report addresses funding, contracts, partnerships, and collaborations at various levels, from industrial agreements to international cooperations. It also covers dissemination and teaching activities, such as participation in scientific events, outreach, and supervision. The document concludes with a presentation of scientific production, including major publications and those produced during the year.

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
- A2.3.5. – Cyber-physical systems
- A4.5.3. – Program proof

Other research topics and application domains

- B3.1. – Sustainable development
- B4.5. – Energy consumption
- B6.3.3. – Network Management
- B6.6. – Embedded systems
- B9. – Society and Knowledge

Contents

Project-Team SPADES	1
1 Team members, visitors, external collaborators	5
2 Overall objectives	5
3 Research program	6
3.1 Design and Programming Models	6
3.2 Certified Real-Time Programming	7
3.3 Causality and Reversibility	7
4 Application domains	8
4.1 Industrial Applications	8
4.2 Current Industrial Cooperations	8
5 Social and environmental responsibility	8
5.1 Footprint of research activities	8
5.2 Impact of research results	8
6 Latest software developments, platforms, open data	9
6.1 Latest software developments	9
6.1.1 CESAn	9
6.1.2 MASTAG	9
6.1.3 Tabvar	9
6.1.4 COSMetyc	10
7 New results	10
7.1 Dataflow models of computation	10
7.1.1 Memory minimization for dataflow graphs	10
7.2 Real-time scheduling	10
7.2.1 Learning the speed policy for minimizing the energy consumption of hard real-time jobs	11
7.3 Causality and reversibility	11
7.3.1 Causal explanations for embedded and concurrent systems	11
7.4 Transversal activity: ICT and the Anthropocene	12
7.5 Tooling up the teaching of proofs	13
8 Bilateral contracts and grants with industry	14
8.1 Bilateral contracts with industry	14
9 Partnerships and cooperations	14
9.1 National initiatives	14
9.1.1 Défi Inria	14
10 Dissemination	15
10.1 Promoting scientific activities	15
10.1.1 Scientific events: organisation	15
10.1.2 Scientific events: selection	15
10.1.3 Journal	16
10.1.4 Invited talks	16
10.1.5 Leadership within the scientific community	16
10.1.6 Scientific expertise	16
10.1.7 Research administration	16
10.2 Teaching - Supervision - Juries - Educational and pedagogical outreach	17

10.2.1 Teaching	17
10.2.2 Supervision	17
10.2.3 Juries	17
10.2.4 Educational and pedagogical outreach	17
10.3 Popularization	17
10.3.1 Specific official responsibilities in science outreach structures	17
10.3.2 Participation in Live events	17
11 Scientific production	18
11.1 Major publications	18
11.2 Publications of the year	19
11.3 Cited publications	20

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, Associate Professor]

Post-Doctoral Fellow

- Antoine Hardy [INRIA, Post-Doctoral Fellow, from Feb 2025]

PhD Students

- Baptiste De Goër De Herve [INRIA]
- Aurélie Kong Win Chang [INRIA, until May 2025]
- Alexander Obeid Guzman [INRIA]

Interns and Apprentices

- Timothe Baleras [INRIA, Intern, from Jun 2025 until Aug 2025]
- Henri Marion [UGA, Intern, from Jun 2025 until Aug 2025]
- Amaury Mazoyer [ENS DE LYON, Intern, from Jun 2025 until Jul 2025]
- Gregoire Stoupy [INRIA, Intern, from Jun 2025 until Aug 2025]

Administrative Assistants

- Julia Di Toro [INRIA]
- Nathalie Gillot [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 [26], 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 *Causality and reversibility*, 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. We also study the memory consumption of dataflow graphs. We study new techniques to find optimal sequential schedules (optimal in the sense that their memory peak is minimal) of task graphs and SDF graphs. We use these techniques to find parallel schedules that satisfy memory constraints, possibly the strongest ones.

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 [40]. Witness component-based toolsets such as PROLEMY [33], BIP [29], or the modular architecture frameworks used, for instance, in the automotive industry (AUTOSAR) [27]. 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 RocQ-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 [30]), 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 [28, 32], in particular in relation with mixed criticality systems. The main goal in this theme is to address several of these open questions.

We 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. [28]. For our part, we focus on devising synchronous programming languages for distributed systems and precision-timed architectures.

3.3 Causality and Reversibility

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 [31, 38]. 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 [45], natural sciences, law [46], and statistics [47], 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 [43], to allow the diagnosis of faults in a complex concurrent system [39], or to enforce accountability [42], 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 [35]), or is broken (*e.g.*, by limiting fault propagation [52]).

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.

We are closely involved in Certisen, a startup project supported by ISS (Inria Startup Studio). The project focuses on Time-Sensitive Networking (TSN), which is the next generation of communication networks in critical sectors such as automotive, aerospace, and industrial systems. Certisen objective is to prove the correctness of TSN key properties (latency, priority, memory). A concrete tool, automatically extracted from Rocq/Coq proofs, will enable TSN configurations to be formally certified. This project is lead by Xiaojie Guo, a former PhD student of SPADES. Pascal Fradet participates to Certisen as a technical adviser and by working on the formal specifications of TSN and Rocq proofs.

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 2025. 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 725 kg CO₂e. Commute travels sum up to an estimated 1,5 t CO₂e. We have purchased one computer, amounting for 192 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 8,4 t CO₂e for the team, or an average of 765 kg 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 2 Ateliers SEnS in 2025.

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 Latest software developments, platforms, open data

6.1 Latest software developments

6.1.1 CESAn

Name: Core Erlang Semantics Analyzer

Keyword: Formal semantics

Functional Description: CESAn implements the semantics of a subset of Core Erlang. Given a Core Erlang program with finite semantics, it outputs its semantics in the form of a labeled transition system. The underlying small-step semantics faithfully represents message passing and signal handling in Erlang.

Publications: [hal-04222884](#), [hal-05353119](#)

Contact: Aurélie Kong Win Chang

Participants: Aurélie Kong Win Chang, Jerome Feret, Gregor Goessler

6.1.2 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 sequential 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.

From an optimal sequential schedule, MASTAG can derive a dynamic parallel schedule satisfying any practical memory constraints and achieving good speedups. MASTAG is efficient and always succeeds in finding a parallel schedule that respects the required memory constraints, unlike alternative tools.

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

Contact: Alexandre Honorat

6.1.3 Tabvar

Keywords: Coq, GUI (Graphical User Interface), Teaching

Functional Description: This is a prototype of a graphical interface for the Coq proof assistant that includes graphical elements (here, curve sketching tabs). These elements are meant to interact with the usual textual proof of Coq, and help students learning how what a mathematical proof is expected to be.

URL: <https://mbodin1.gitlabpages.inria.fr/td-math-coq/interface.html>

Publication: [hal-04096240](#)

Contact: Martin Bodin

Partner: LIG

6.1.4 COSMetyc

Keyword: OpenStreetMap

Functional Description: This library can import, store, manipulate, and export OpenStreetMap data from and to various format of the community (like GeoJSON and OSM XML).

URL: <https://gitlab.inria.fr/mbodin1/logic-osm>

Contact: Martin Bodin

7 New results

7.1 Dataflow models of computation

Participants: Pascal Fradet, Alain Girault.

7.1.1 Memory minimization for dataflow graphs

Dataflow Models of Computation (MoCs) are widely used in embedded systems, including multimedia processing, digital signal processing, telecommunications, and automatic control. The reason is twofold. On the one hand the dataflow MoC is very close to classical block diagrams familiar to automatic control engineers. On the other hand it allows to nicely express the parallelism of an application, thereby allowing its efficient implementation on modern parallel architectures, be they multi-core processors or HPC platforms.

In 2022, we started focusing on optimizing the memory consumption of dataflow graphs, more precisely on the memory peak, defined as the maximum of the memory space used during the whole execution of the dataflow graph. This objective is critical both for embedded systems (due to the very small memory space available) and neural networks implemented in a dataflow fashion (due to the huge memory requirement). We have first proposed new techniques to find optimal *sequential* schedules of dataflow graphs, optimal in the sense that their memory peak is minimal. Then, we have studied how to find parallel schedules that satisfy memory constraints possibly the strongest ones.

We have proposed graph transformations that compress a task graph while preserving its optimal memory peak [17, 34]. These transformations compress a large class of graphs into a single node representing their optimal schedule. In particular, we have formally proved that this is the case for all Series-Parallel Directed Acyclic Graphs (SP-DAGs). In addition, our graph transformations provide a simple characterization of optimal schedules for sets of independent tasks and allowed us to design an optimal compositional analysis dealing with single-source single-sink subgraphs (called S-T subgraphs). These S-T subgraphs can be analyzed separately and replaced by their best schedule. These results (correctness of the transformation rules, the compositional analysis, etc.) were formally proved. 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 task graphs, between 30 and 50 nodes.

A natural extension of this work was to consider parallel schedules with shared memory, which we did in 2025. The optimal memory peak found for a sequential execution provides a lower bound for all its parallel versions. Thanks to a dynamic list scheduling algorithm under constraint, we are able to derive parallel schedules that are guaranteed to meet the desired memory peak constraint [21], even the harshest ones, whereas the previous state-of-the-art approaches mostly fail. To achieve this, we take advantage from the optimal sequential schedules from [17, 34]. Furthermore, our approach is faster and can deal with very large task graphs, up to 50,000 tasks. When the constraints are the most severe (equal to the minimal memory peak), we still obtain very good speedups, respectively $\times 2.68$ on 4 processors and $\times 5.76$ on 8 processors on average. As expected, the more relaxed the memory constraints, the better the speedups become.

7.2 Real-time scheduling

Participants: Alain Girault.

7.2.1 Learning the speed policy for minimizing the energy consumption of hard real-time jobs

Since 2017, we have been working with Bruno Gaujal and Stephan Plassart on a very general model of real-time systems, made of a single-core processor equipped with DVFS¹ and an infinite sequence of preemptive hard real-time jobs. Each real-time job is characterized by the triplet (inter-arrival time, size, relative deadline). The key point is that the system is *clairvoyant without statistical knowledge*, meaning that the job features are not known in advance, not even their probability distributions, but only revealed when the job is released. In this context, we wish to find the *optimal speed policy*, optimal meaning that all the jobs meet their deadline (hard real-time constraint) and that the average energy consumption is minimized. In 2025, we have proposed a *machine learning* method to learn the optimal speed policy, which consists of two successive phases. During the first phase, we observe the features of the incoming real-time jobs and we estimate their probability distribution. During this phase, the jobs are executed with the speed policy (OA) [53], which is known to be optimal when the future jobs are unknown. From this estimated probability distribution, we build and solve a constrained MDP (Markov Decision Process) and we compute the optimal speed policy (L), as in [37, 48, 36]. Then the second phase starts, during which the jobs are executed with the learned speed policy (L). We have established a stopping criterion that gives the duration of the learning phase. Our experiments show that, after a learning period of 105 steps, our learned speed policy (L) is extremely close to the offline optimal one, with an error margin less than 0.002%.

7.3 Causality and reversibility

Participants: Gregor Goessler, Jean-Bernard Stefani, Aurélie Kong Win Chang, Alexander Obeid Guzmán.

7.3.1 Causal explanations for embedded and concurrent systems

In order to pinpoint the root causes of system failures and explain *why* an observed outcome occurred, we have been developing techniques based on counterfactual analysis. Counterfactual analyses determine causal dependencies beyond those captured by classical causal semantics for concurrent systems. Given a system model, our goal is to extract the part of a failing execution trace that is causally relevant for the failure. We have been continuing our work on this topic in several directions.

We have developed a theory of explanations for discrete-event systems that allows us to analyze and compare different types of explanations by introducing a notion of semantics of an explanation. Based on this semantics we have formalised a set of expected properties for explanations, including *completeness* (all faulty traces can be explained) and *precision* (different faults have different explanations). We have then cast a frequently used type of explanation — called *causal compression* or *subsequence explanation* — in this framework and shown that it has no semantics that makes it both complete and precise. We have finally proposed a new type of explanations called *choice explanation* that focuses on branching choices in the execution of a system, and that satisfies our requirements.

Following up on Thomas Mari’s PhD thesis [44], we are currently working on a symbolic construction of robustness functions for real-time systems, and its implementation, allowing us to effectively generate such explanations.

As part of the DCORE project on causal debugging of concurrent programs, the goal of Aurélie Kong Win Chang’s PhD thesis was to investigate the construction of causal explanations for Erlang programs. In [41] we have formalized a small step semantics for a subset of Core Erlang that models, in particular, its monitoring and signal systems. We have implemented this semantics in the Core Erlang Semantics Analyzer

¹DVFS=Dynamic Voltage and Frequency Scaling. This is a feature commonly found on modern processors, which allows to decrease the (frequency,voltage) operating point in order to decrease the energy consumption.

(CESAn [19], Section 6.1.1). Based on the semantics, the tool explains the causes of a program execution violating an expected safety property.

So far we assumed a system model to be known in order to assess causality. However, for many applications a faithful system model is not available. The goal of Alexander Obeid Guzman’s PhD thesis is to leverage a measure of conditional algorithmic complexity in order to discover causal dependencies from the single observation of an execution trace coming from a distributed system.

7.4 Transversal activity: ICT and the Anthropocene

Participants: Martin Bodin, Baptiste De Goër De Herve, Pascal Fradet, Alain Girault, Gregor Goessler, Antoine Hardy, Xavier Nicollin, Sophie Quinton, 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”), so we started research in this axis by investigating the potential of ICT for supporting the ecological transition. More generally, we are interested in the role played by ICT in the Anthropocene as well as new approaches to their design. This raises many complex challenges: For example, 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)? These questions require a systemic approach, and at least some basic knowledge of the state of the art in many scientific disciplines, in particular in STS (Science and Technology Studies).

Following recent work on these topics in the SPADES team (e.g., the PhD of Aina Rasoldier defended in 2024 [50, 51], the internship of Ludmila Courtillat--Piazza in 2022 and the ongoing **AEx SIA project** in which we study how computer science should evolve to take into account the global changes brought about by the Anthropocene), two members of the SPADES team, namely Sophie Quinton and Jean-Bernard Stefani, have joined forces with other colleagues from Grenoble to launch a **new Inria team** named ADN for “Anthropocène et Décroissance du Numérique” (Anthropocene, Degrowth and ICT).

Work in the SIA project has progressed through the PhD thesis of Baptiste de Goër, which focuses on how ICT related sustainability issues are and could be taught in computer science classes. One major contribution was the study of how researchers problematize their work in the ICT4S conference community, and how nonepistemic values are involved [22]. Another contribution is still in progress and relates to how CS high school teachers approach the environmental impacts of ICT in class and how to help them teach this complex question (initial results will be presented at the **Didapro 2026** conference). In addition, in his postdoctoral work, Antoine Hardy has studied from a sociological perspective how the CS research community is coping with the ecological catastrophe. This work will be presented at the **Archipel** conference in 2026.

A second line of research on this topic relates to society’s dependence on ICT, and focuses on the relation between the resilience of ICT infrastructures and the resilience of our societies in a future more and more constrained by the destabilization of planetary boundaries and its consequences on our social organisations. There has been so far no research on the potential consequences on the network performance of possible long term and continuous disruptions such as semiconductor shortages, for example due to shortages of raw material or environmental policy constraints. In a **workshop paper**, we study the resilience of mobile networks in the face of hardware aging and approach mobile network resilience using the stochastic geometry framework, proposing a network model including multiple dependent techno-bands (TB) and deriving approximate closed-form expressions for new resilience metrics. Numerical experiments show the influence of parameters such as load, base station density and the number of TB on network resilience. Another **workshop paper** presents how we intend to articulate a sociological survey with this CS modeling work.

A part of the ADN research proposal is centered on the notion of convivial technology initially proposed by the philosopher Ivan Illich in his book *Tools for conviviality*. We have argued in a paper entitled *Conviviality for digital degrowth*, to appear in the journal *Philosophia Scientiae* (as a follow up of [49]), the relevance of conviviality for designing digital systems that are congruent with, and conducive to a degrowth

society, and highlighted different research directions towards *designing for conviviality* and rethinking digital infrastructures and workflow systems.

7.5 Tooling up the teaching of proofs

Participants: Martin Bodin, Alain Girault.

Proof is a key concept in early licence degree in mathematics. It is known to be difficult to teach, due to its level of abstraction, but also due to the various roles it is meant to fulfill: it serves first to certify the mathematical statements themselves, but it also has a communication role and it intervenes within the research process (exploring the consequences of hypotheses, coming up with conjectures, etc.). Each of these roles requires different skills, and the licence's year comes with an increase of the required formalism. This makes it difficult but key to teach at this level.

The SPADES team has long been working with the Rocq proof assistant, which provides a strong logical base for understanding proofs. It was thus natural for us to try to tool up licence's mathematical classes with Rocq. Mathematical classes have long been tooled up with computer software (like computer algebra systems or programming languages) to help providing intuitions or building conjectures. We want here to test whether a proof assistant can help *teaching* proof writing. This is the goal of the LIBERABACI project within Inria (which includes various other project-teams).

Within this project in SPADES, we approached the IREMI institute, which brings together researchers and teachers to do research on mathematical teaching. We are also working with Julien Puydt, teacher in the Pupilles de l'Air high school.

We focus in reducing the entry cost of Rocq. Indeed, this software has originally been designed by and for experts, and not for math students. We thus explore several aspects in which it can be improved to more closely match the practice in math courses:

- We experiment with its interface to incorporate graphical components.
- We aim to ease the user interaction when writing proofs.
- We explore new ways in which mathematical notations can be defined.
- We seek for mechanisms to support a library of courses that scales and meet the need of teachers.

The use of graphical (and more generally non-textual) support within proofs in math courses is frequent, yet totally absent from Rocq. We thus sketched a new interface for the software that enables to incorporate such non-textual elements, and to be able to interact with them within a proof. The goal here is to be as close as possible to the current practice of teachers and students to reduce the entry cost of using the software. Our goal is not for the student to be able to use a proof assistant, but to help them produce better pen-and-paper proofs: the use of the proof assistant is here to guide them and help them understand what it expected and where. We are working on this subject with Emmanuel Beffara (LIG, team MeTAH) and Rémi Molinier (Institut Fourier). We are hoping to get an in-class experiment in 2026.

Proof assistants are heavily based on tactics, which are programs that can write parts of a proof: most of the user interaction with these tools are based on these tactics. We asked about 50 students at an early master's degree to solve a math problem and took their draft with their raw arguments. These students were chosen as they already had the abilities that we want to teach at the licence's level. They were also chosen by their relative diversity. We studied the arguments they used within their draft and categorised them on whether they were textual or not, on what the role each argument had in the overall proof they were (implicitly) building. We are designing tactics that follow as close as possible the students's arguments, and that could potentially guide them through the proof.

Math heavily relies on notations, that are frequently overloaded. For instance the \times symbol is both used as number multiplication, matrix multiplication, and cartesian product, although these operations are quite different. There already exists numerous mechanisms in Rocq to define notations and infer which instance of an overloaded symbol was meant to be used at a given point. We find these mechanisms incredible and they can be very useful in development. However, we fear that their error messages are difficult to understand by a

non-expert: these systems tend to internally produce complex candidates that weren't always expected by the user, and their error messages tend to mention unfulfilled constraints for these complex candidates. We are thus working on much simpler mechanisms for notations that we expect to be closer to teachers's intuition. Such systems will be much less powerful than the state of the art, but we expect them to be more aligned with teaching. We are working in collaboration with Julien Puydt (Pupilles de l'Air) and Arthur Charguéraud (Inria, Camus team).

Still with the help of Julien Puydt, we are looking for mechanisms to build library content that could scale. Building libraries is known to be an intensive task, and current mathematical libraries in Rocq are from time to time subject to large-scale refactoring. This is due to their structure from the ground-up: each mathematical notion is built on smaller ones, and each change in an inner structure can impact all the ones on top of it. We fear that this library structure isn't the perfect fit for teachers. First, teachers usually lack the time to maintain content up to date with external changes. Second, this library structure doesn't follow the natural order of a course: monoids are taught after groups, but monolithic libraries tend to define groups with the help of monoids, in the reverse order. We are thus currently assessing what features teachers would need to design a mathematical library for teaching. Once identified, we plan to implement these features, or to relate them with existing features.

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 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 Rocq proof assistant in a Mathematical course in the University to help teaching proofs. At SPADES, Martin Bodin is working with the IREMI at Grenoble to involve math teachers and didactic researchers to the project.

SmartNet

Participants: Gregor Goessler, Alexander Obeid Guzman.

The SmartNet project, kicked off in 2024, aims to develop network management techniques to handle the increasing complexity of networks. In particular, SmartNet seeks to provide comprehensive insights of the network and its subsystems by identifying cause-effect relationships, enabling strategic interventions for malfunction prevention. We collaborate in this project with the APTIKAL team from LIG, and Nokia Research. This project funds Alexander Guzmán's PhD thesis as well as the postdoc position of Nishchal Prasad, who is hosted by the APTIKAL team and co-advised by Gregor Goessler.

SIA

Participants: Baptiste De Goër, 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, Pascal Fradet, Alain Girault, Gregor Goessler, Xavier Nicollin, Sophie Quinton, Jean-Bernard Stefani.

10.1 Promoting scientific activities

10.1.1 Scientific events: organisation

General chair, scientific chair

- Alain Girault was co-chair of the International Workshop on Synchronous Programming (SYN-CHRON'25).

Member of the organizing committees

- Alain Girault is a member of the Steering Committee of the Embedded Systems Week (ESWEEK).
- Martin Bodin is a member of the Steering Committee for Didapro'11.
- In connection to the **CIS-PEN** working group and the **Persyval-Lab ICT-S** research axis, Sophie Quinton co-organized two national scientific events in Grenoble, "Numérique et soutenabilité, quels usages de la modélisation ?" and "Regards croisés sur les infrastructures numériques et numérisées", as well as a panel "Regards croisés sur les enjeux et limites de la quantification environnementale du numérique" at the Journées Scientifiques du CIS.

10.1.2 Scientific events: selection

Chair of conference program committees Sophie Quinton is a co-chair for the **Undone Computer Science 2026** conference.

Member of the conference program committees

- Alain Girault was TPC member of the Forum on Description Languages international conference (FDL'25) and the Time-Centric Reactive Software international workshop (TCRS'25).
- Sophie Quinton was a TPC member of the SCW 2026 workshop.

10.1.3 Journal

Member of the editorial boards

- Alain Girault is associate editor for the EURASIP Journal on Embedded Systems and for the Real-Time Systems Journal.
- Gregor Goessler has been a guest editor of the IST special issue on *Causal modeling and inference in software engineering* [24].

Reviewer - reviewing activities

- Alain Girault reviewed a paper for ACM Transactions on Parallel Computing.
- Martin Bodin reviewed a paper for ICFP.

10.1.4 Invited talks

Sophie Quinton was invited to panels at GRETSI 2025, Réparer le Futur (rencontres inter-AtEcoPol 2025) and the Journée d'Étude "Enjeux environnementaux du numérique au prisme des SHS". She gave an invited talk at a workshop on "IA et recherche en santé".

10.1.5 Leadership within the scientific community

Sophie Quinton co-chairs a **working group of the GDR CIS** associated with the Center for Internet and Society focused on environmental issues. She also co-chairs the **"ICT and sustainability" Persyval-lab axis**.

10.1.6 Scientific expertise

Sophie Quinton is a member of the scientific board of Grenoble Métropole and the Agence d'urbanisme de la région grenobloise.

10.1.7 Research administration

- Since October 2013, Pascal Fradet has been head of the committee for doctoral studies ("Responsable du comité des études doctorales") of the Inria Grenoble research center. He is also the local correspondent for the young researchers Inria mission ("Mission jeunes chercheurs") and serves as the substitute of the director of the Inria Grenoble research center at the doctoral school council (MSTII).
- Up to June 2025, Alain Girault was Deputy Scientific Director at INRIA, in charge of the domain "Algorithmics, Programming, Software and Architecture" (51 research teams over 9 INRIA research centers).
- Gregor Goessler is a member of the *Scientific policy committee (COS)* of the Inria Grenoble center, of the council of the MSTIC department of Grenoble University, and of the scientific jobs committee at Inria Grenoble.
- Sophie Quinton facilitated until the end of 2025 the SEnS-GRA group which hosts discussions and proposes actions regarding the environmental and societal impact of our research at Inria Grenoble.

10.2 Teaching - Supervision - Juries - Educational and pedagogical outreach

10.2.1 Teaching

- Licence: Pascal Fradet, Modèles de Calcul: λ -calcul, CM & TD, 30 H eqTD, level L3, Univ. Grenoble Alpes, France.
- Licence: Alain Girault, Modèles de Calcul: λ -calcul, TD, 18 H eqTD, level L3, Univ. Grenoble Alpes, France.
- Licence: Martin Bodin, Modèles de Calcul: λ -calcul, TD, 18 h eqTD, level L3, Univ. Grenoble Alpes, France.
- Licence: Martin Bodin, Langage et traducteurs, Programmation Fonctionnelle, TP, 30 h eqTD, level M1, Polytech Grenoble, France.
- Master : Sophie Quinton, Numérique responsable, 15 HeqTD, level M1, Grenoble INP (Ensimag), France.
- Sophie Quinton gave a one-day course at the EJCP Summerschool, contributed to the scientific program and gave two lectures at the AI² autumn school, gave one lecture and attended one panel at the Ensimag engineering school, and gave a lecture at the ENS Lyon CS departement seminar.

10.2.2 Supervision

- Gregor Goessler, PhD completed: Aurélie Kong Win Chang, “Causal explanations for concurrent programs in Erlang”; co-advised by Gregor Goessler and Jérôme Feret.
- Gregor Goessler, PhD in progress: Alexander Obeid Guzmán, “Inference of causal models for networks from single observations”; since January 2024.
- Sophie Quinton, PhD in progress: Baptiste de Goër, “Teaching ICT-related sustainability issues in computer science courses”.
- Sophie Quinton, PhD in progress: Ludmila Courtillat–Piazza, “Dépendance au numérique et vulnérabilités dans un contexte d’urgence écologique, abordées sous l’angle de la résilience des réseaux mobiles”.

10.2.3 Juries

Sophie Quinton, member of the PhD committee of Simon Delarue, Telecom Paris.

10.2.4 Educational and pedagogical outreach

Sophie Quinton contributed to the serious game [PhoneImpact](#).

10.3 Popularization

10.3.1 Specific official responsibilities in science outreach structures

- Martin Bodin serves as a local correspondent for the Chiche! program.

10.3.2 Participation in Live events

- Martin Bodin animated an activity on logic during the science festival « Fête de la science ».
- Sophie Quinton gave a lecture on the environmental impacts of ICT at a high school event.

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] S. D. Djoko, R. Douence and P. Fradet. ‘Aspects preserving properties’. In: *Science of Computer Programming* 77.3 (2012), pp. 393–422.
- [4] P. Fradet, A. Girault and A. Honorat. ‘Graph Transformations for Memory Peak Minimization by Scheduling’. In: *ACM Transactions on Embedded Computing Systems (TECS)* (2024), pp. 1–36. URL: <https://inria.hal.science/hal-04816271>. In press.
- [5] 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>.
- [6] P. Fradet, X. Guo, J.-F. Monin and 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, Apr. 2019, pp. 1–10. DOI: [10.1109/RTAS.2019.00023](https://doi.org/10.1109/RTAS.2019.00023). URL: <https://hal.archives-ouvertes.fr/hal-02119024>.
- [7] 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>.
- [8] 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>.
- [9] 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>.
- [10] G. Gössler and D. Le Métayer. ‘A general framework for blaming in component-based systems’. In: *Science of Computer Programming* 113, Part 3 (2015). DOI: [10.1016/j.scico.2015.06.010](https://doi.org/10.1016/j.scico.2015.06.010). URL: <https://hal.inria.fr/hal-01211484>.
- [11] 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>.
- [12] 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>.
- [13] S. Andalam, P. S. Roop, A. Girault and C. Traulsen. ‘A Predictable Framework for Safety-Critical Embedded Systems’. In: *TC* 63.7 (July 2014), pp. 1600–1612.
- [14] 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>.
- [15] 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

- [16] G. Fabbretti, I. Lanese and J.-B. Stefani. ‘A Behavioral Theory for Distributed Systems with Weak Recovery’. In: *Logical Methods in Computer Science* 21.3 (1st July 2025). DOI: [10.46298/lmcs-21\(3:1\)2025](https://doi.org/10.46298/lmcs-21(3:1)2025). URL: <https://inria.hal.science/hal-05272017>.
- [17] P. Fradet, A. Girault and A. Honorat. ‘Graph Transformations for Memory Peak Minimization by Scheduling’. In: *ACM Transactions on Embedded Computing Systems (TECS)* (23rd Jan. 2025), pp. 1–36. DOI: [10.1145/3707206](https://doi.org/10.1145/3707206). URL: <https://inria.hal.science/hal-04816271> (cit. on p. 10).
- [18] 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)* 24.3 (2025), p. 45. DOI: [10.1145/3615586](https://doi.org/10.1145/3615586). URL: <https://hal.science/hal-04200195>.

Invited conferences

- [19] A. Kong Win Chang, J. Feret and G. Gössler. ‘CESAn: A Core Erlang Semantics Analyser’. In: *Components Operationally: Reversibility and System Engineering. Essays Dedicated to Jean-Bernard Stefani on the Occasion of His 65th Birthday*. DisCoTec 2025 - 20th International Federated Conference on Distributed Computing Techniques / CORSE - Components Operationally: Reversibility and System Engineering. Vol. LNCS-16065. Lecture Notes in Computer Science. Lille, France: Springer Nature Switzerland, 18th Oct. 2025, pp. 103–115. DOI: [10.1007/978-3-031-99717-4_6](https://doi.org/10.1007/978-3-031-99717-4_6). URL: <https://inria.hal.science/hal-05353119> (cit. on p. 12).

International peer-reviewed conferences

- [20] B. Diarra, K. Guillooard, M. Ouzzif, P. Merle and J.-B. Stefani. ‘Demo: A customizable toolchain to validate Kubernetes manifests for robust deployment’. In: *IC2E 2025 - 13th IEEE International Conference on Cloud Engineering*. Rennes, France: IEEE, 2025, pp. 1–2. URL: <https://hal.science/hal-05204020>.
- [21] P. Fradet, A. Girault and A. Honorat. ‘Parallel scheduling of task graphs with minimal memory requirements’. In: *IPDPS 2025 - 39th IEEE International Parallel and Distributed Processing Symposium*. Milano, Italy, 2025, pp. 1–12. URL: <https://inria.hal.science/hal-04879748> (cit. on p. 10).
- [22] B. de Goër and S. Quinton. ‘The Political Dimension of Problematizing ICT-related Sustainability Issues’. In: *ICT4S 2025 - International Conference on Information and Communications Technology for Sustainability*. Dublin, Ireland, 2025, pp. 1–12. URL: <https://inria.hal.science/hal-05064545> (cit. on p. 12).

National peer-reviewed Conferences

- [23] T. Baleras, M. Bodin and U. Comignani. ‘COSMetyc : OpenStreetMap en OCaml’. In: *JFLA 2026 – 37es Journées Francophones des Langages Applicatifs*. Vol. JFLA 2026 – 37es Journées Francophones des Langages Applicatifs. Oberbronn, Alsace, France, 27th Jan. 2026. URL: <https://hal.science/hal-05427963>.

Edition (books, proceedings, special issue of a journal)

- [24] *Special issue on causal modeling and inference in SE* 183 (Apr. 2025). DOI: [10.1016/j.infsof.2025.107754](https://doi.org/10.1016/j.infsof.2025.107754). URL: <https://inria.hal.science/hal-05353350> (cit. on p. 16).

Doctoral dissertations and habilitation theses

- [25] A. Kong Win Chang. ‘Causal explanations for concurrent programs in Erlang’. Université Grenoble Alpes [2020-....], 16th May 2025. URL: <https://theses.hal.science/tel-05262517>.

11.3 Cited publications

- [26] ARTEMIS Joint Undertaking. *ARTEMIS Strategic Research Agenda*. 2011 (cit. on p. 6).
- [27] *Automotive Open System Architecture*. 2003. URL: <http://www.autosar.org> (cit. on p. 6).
- [28] 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) (cit. on p. 7).
- [29] 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) (cit. on p. 6).
- [30] 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) (cit. on p. 7).
- [31] S. Borkar. ‘Designing Reliable Systems from Unreliable Components: The Challenges of Transistor Variability and Degradation’. In: *IEEE Micro* 25.6 (2005) (cit. on p. 7).
- [32] R. Davis and A. Burns. ‘A Survey of Hard Real-Time Scheduling for Multiprocessor Systems’. In: *ACM Computing Surveys* 43.4 (2011) (cit. on p. 7).
- [33] 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) (cit. on p. 6).
- [34] 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)*. Orlando (FL), United States: ACM, June 2023, pp. 76–86. DOI: [10.1145/3589610.3596280](https://doi.org/10.1145/3589610.3596280). URL: <https://hal.science/hal-04163123> (cit. on p. 10).
- [35] F. C. Gärtner. ‘Fundamentals of Fault-Tolerant Distributed Computing in Asynchronous Environments’. In: *ACM Computing Surveys* 31.1 (1999) (cit. on p. 7).
- [36] B. Gaujal, A. Girault and S. Plassart. ‘An MDP-Based Solution for the Energy Minimization of Non-Clairvoyant Hard Real-Time Systems’. In: *Real-Time Systems* (Dec. 2024), p. 47. DOI: [10.1007/s11241-024-09433-5](https://doi.org/10.1007/s11241-024-09433-5). URL: <https://inria.hal.science/hal-02371742> (cit. on p. 11).
- [37] 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> (cit. on p. 11).
- [38] 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 (cit. on p. 7).
- [39] 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 (cit. on p. 7).
- [40] T. Henzinger and J. Sifakis. ‘The Embedded Systems Design Challenge’. In: *Formal Methods 2006*. Vol. 4085. Lecture Notes in Computer Science. Springer, 2006 (cit. on p. 6).
- [41] A. Kong Win Chang, J. Feret and G. Gössler. ‘A Semantics of Core Erlang with Handling of Signals’. In: *Erlang 2023 - 22nd ACM SIGPLAN International Workshop on Erlang*. Seattle WA, United States: ACM, Sept. 2023, pp. 31–38. DOI: [10.1145/3609022.3609417](https://doi.org/10.1145/3609022.3609417). URL: <https://hal.science/hal-04222884> (cit. on p. 11).
- [42] 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 (cit. on p. 7).
- [43] 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 (cit. on p. 7).
- [44] T. Mari. ‘Causal explanations for reactive real-time systems’. Theses. Université Grenoble - Alpes, Nov. 2023. URL: <https://inria.hal.science/tel-04848168> (cit. on p. 11).

- [45] 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> (cit. on p. 7).
- [46] M. Moore. *Causation and Responsibility*. Oxford, 1999 (cit. on p. 7).
- [47] J. Pearl. ‘Causal inference in statistics: An overview’. In: *Statistics Surveys* 3 (2009), pp. 96–146 (cit. on p. 7).
- [48] S. Plassart. ‘Online optimization in dynamic real-time systems’. Theses. Université Grenoble Alpes [2020-....], June 2020. URL: <https://tel.archives-ouvertes.fr/tel-02990646> (cit. on p. 11).
- [49] S. Quinton and J.-B. Stefani. ‘Taking conviviality seriously (extended abstract)’. In: *2024 - 1st conference on Undone Science in Computer Science*. Nantes, France, Feb. 2024, pp. 1–4. URL: <https://hal.science/hal-04448759> (cit. on p. 12).
- [50] A. Rasoldier. ‘Comment évaluer le potentiel d’une solution numérique face à l’urgence écologique ? Application aux plateformes de covoiturage régulier à l’échelle locale’. Theses. Université Grenoble Alpes [2020-....], Feb. 2024. URL: <https://theses.hal.science/tel-04702139> (cit. on p. 12).
- [51] A. Rasoldier, A. Girault, S. Quinton, J. Combaz and K. Marquet. ‘Assessing the Potential of Carpooling for Reducing Vehicle Kilometers Traveled’. In: *International Conference on Information and Communications Technology for Sustainability, ICT4S’23*. Rennes, France: IEEE, June 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> (cit. on p. 12).
- [52] J. Rushby. *Partitioning for Safety and Security: Requirements, Mechanisms, and Assurance*. Tech. rep. CR-1999-209347. NASA Langley Research Center, 1999 (cit. on p. 7).
- [53] 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 (cit. on p. 11).