RESEARCH CENTRE

Inria Centre at Université Grenoble Alpes

IN PARTNERSHIP WITH:

Université de Grenoble Alpes

2024

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

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

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

B9.9. – Ethics

1 Team members, visitors, external collaborators

Research Scientists

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

Faculty Member

- Xavier Nicollin [GRENOBLE INP, Associate Professor]

Post-Doctoral Fellow

- Alexandre Honorat [FLORALIS, Post-Doctoral Fellow]

PhD Students

- Baptiste De Goer De Herve [INRIA]
- Giovanni Fabbretti [INRIA, until Sep 2024]
- Aurélie Kong Win Chang [INRIA]
- Pietro Lami [INRIA]
- Alexander Obeid Guzman [INRIA]

Interns and Apprentices

- Tristan Mathon [INRIA, Intern, from Jul 2024 until Sep 2024]

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 [32], makes them more pressing and more difficult to address: the targeted system properties – dynamic modularity, time-predictability, energy efficiency, and fault-tolerance – are largely antagonistic (*e.g.*, having a highly dynamic software structure is at variance with ensuring that resource and behavioral constraints are met). Tackling these questions together is crucial to address this antagonism, and constitutes a key point of the SPADES research program.

A few remarks are in order:

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

3 Research program

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 tasks 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 [56]. Witness component-based toolsets such as PTOLEMY [44], BIP [36], or the modular architecture frameworks used, for instance, in the automotive industry (AUTOSAR) [33]. 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 [39]), 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 [34, 43], 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. [34]. 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 [40, 54]. 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 [64], natural sciences, law [65], and statistics [67], 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 [60], to allow the diagnosis of faults in a complex concurrent system [55], or to enforce accountability [59], 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 [49]), or is broken (*e.g.*, by limiting fault propagation [70]).

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 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 2024. 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 5,6 t CO_2e. Commute travels sum up to 1,8 t CO_2e. We purchased new hardware (1 computer) for a total of 174 kg CO_2e. We roughly estimate our share of INRIA services and building usage to 6 t CO_2e. Based on the above estimations, our carbon footprint totals 13,6 t CO_2e for the team, or an average of 1,0 t CO_2e 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 4 Ateliers SEnS in 2024.

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 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.

Publication: hal-04222884

Contact: Aurélie Kong Win Chang

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

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). Four languages are addressed for the moment: The OASIS's TOSCA specification, the ETSI's NFV SOL001 specification, OpenStack's HOT (Heat Orchestration Template), and 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 Guillouard, Meryem Ouzzif, Frédéric Klamm, Jean-Luc Coulin

Partner: Orange Labs

6.1.3 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

6.1.4 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

7 New results

7.1 Dataflow models of computation

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 [47] and BPDF [66]), we have written a survey providing a comprehensive description of the existing parametric dataflow MoCs [41], we have studied *symbolic* analyses of dataflow graphs [42] and an original method to deal with lossy communication channels in dataflow graphs [48]. 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 focusing on optimizing the memory consumption of dataflow graphs. This objective is critical for embedded systems and neural networks which can be implemented in a dataflow fashion. We have first proposed new techniques to find optimal sequential schedules (optimal in the sense that their memory peak is minimal) of tasks graphs and SDF graphs. 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 [13, 28]. 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 provides 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.

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.

A natural extension of this work was to consider parallel schedules with shared memory. The optimal memory peak found for a sequential execution provides a lower bound for all its parallel versions. Using a list scheduling scheduling algorithm adapted to take into account memory requirements, we were able to derive dynamic parallel schedules from the optimal sequential schedules [21]. We produce parallel schedules that always meet the memory constraints (even the harshest ones) and enjoy relatively good speedups. As expected, the more relaxed the memory constraints, the better the speedups are. When memory constraints are close to their minimum, our approach always succeeds in finding a parallel schedule meeting the given constraints, whereas the previous state-of-the-art approaches mostly fail. Furthermore, our approach is faster and can deal with very large task graphs.

Finally, we have applied the Affine DataFlow Graph (ADFG) theory to the domain of reconfigurable processors (FPGA) [15]. 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 developped at INSA Rennes.

7.2 Real-Time Scheduling

Participants: Alain Girault.

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[1] 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. The only information available to the processor are the statistical information on the jobs' characteristics: release time, average execution 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 [14]. 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 [35] and PACE [61, 62], 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. 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 worst-case execution time (WCET). This was the topic of Stéphan Plassart's PhD [52, 51][50], who defended his PhD in June 2020 [68].

7.3 Causality and Reversibility

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

Following up on Thomas Mari's PhD thesis [63], we are currently working on a symbolic construction of robustness functions for real-time systems, allowing us to effectively generate explanations.

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 construction of causal explanations for Erlang programs. In [58] 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 Analizer (CESAn, 6.1.1). We are currently developing a tool that, based on this semantics, 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 infer causal dependencies from the single observation of an execution trace of a telecommunications network.

7.3.2 Reversibility for concurrent and distributed debugging

Concurrent and distributed debugging is a promising application of the notion of reversible computation [53]. As part of the ANR DCore project we contribute to the theory behind, and the developoment of the CauDEr reversible debugger for the Erlang programming language and system.

[1] DVFS=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.

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, even though the Erlang programming language was the main application target for the DCore project, Erlang is a real-world language with a fair number of features to handle. For this reason, we decided to develop a small process calculus conceived as an abstraction of the behaviour of distributed Erlang systems, featuring crash failures for nodes running Erlang systems, and for communication links between Erlang nodes, as well as recoveries for nodes and links. Node recovery in Erlang has a fairly weak semantics, namely, a node that recovers from a crash failure restarts with no memory from its past execution. Furthermore, as we discovered from experiments with Erlang, Erlang makes use of incarnation numbers to distinguish between the different instances of a node across failures (an incarnation number tells the number of failures a node has had in the past), and maintains for each node a (possibly erroneuous) view of the current status of nodes it is linked to. Unfortunately we did not find in the litterature any process calculus analysis of distributed systems with crash failures and recoveries with a weak recovery semantics, incarnation numbers and imperfect node views. Our first achievement in this area thus concerned the development of a new process calculus with the preceding features, called $D\pi FR$ (for distributed π-calculus with crash failures and recoveries), and the study of its behavioral theory, with a result of full abstraction, namely the characterization of contextual equivalence in $D\pi FR$ by means of a bisimilarity relation. This work has been accepted to LMCS with minor revisions [45]. Our second achievement is the development of two reversible semantics for $D\pi FR$, that constitutes two interesting design points in a range of possible reversible semantics for $D\pi FR$, depending on various assumptions pertaining e.g. to persistency of information across failures or the degree of cooperation between nodes. The first semantics, called *non-cooperative*, assumes that local causal dependency information is totally lost during a node crash and that nodes do not cooperate to exchange causal information during execution. The second semantics, called *cooperative*, still assumes that local causal dependecy information is lost during a node crash but that nodes cooperate during execution to exchange causal information, in a manner similar to causal logging checkpoint/rollback recovery schemes in the distributed algorithms literature. This work is reported in Giovanni Fabbretti's PhD thesis, which was successfully defended in October 2024.

Concerning reversibility for shared memory concurrency, we first considered imperative primitives from the Erlang environment providing access to a shared dictionary [16]. A shared dictionary is a form of shared memory and this work demonstrated the complexities involved in defining from scratch a causally consistent reversible semantics for a concurrent programming language. We thus decided to first concentrate on developing a modular operational framework for the definition of shared memory models. In our framework, a concurrent language with a shared memory model takes the form of a synchronous product of three components: a set of concurrent *threads* of execution that cooperate by means of a *memory* whose concurrent access is governed by a *scheduler*. We have shown that our framework is expressive and general enough to faithfully capture diverse memory models, including a sequentially consistent memory model, weak memory models with memory access buffering and reordering, as well as strong and weak transactional memory models. We have then developed a theory for the causally consistent reversing of synchronous products of labelled transition systems with independence. We show that the causal consistent reversibility of a synchronous product of LTSs meeting a set of conditions (square property, independence of backward transitions, well-foundedness, and unambiguity of transitions), can be obtained by the same synchronous product of the reversible variants of the composed LTSs. This provides us with a general recipe for reversing memory models described with our framework. Devising LTSs with independence meeting the required conditions can itself be automated using the approach we have developed [46] for generating a reversible semantics from an operational one defined as a reduction relation formalized in Maude (a logical framework for conditional rewriting logic). These results on reversing concurrent memory models are documented in PIetro Lami's PhD thesis, which was successfully defended in December 2024.

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"), 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).

The PhD of Aina Rasoldier (defended in 2024 [27]) addressed the following question: can ICT significantly contribute to reducing the environmental footprint of other activity sectors? As a first contribution, the PhD detailed limitations of state of the art assessments of the claims stated in the previous paragraph. This is an important problem because it may lead to overestimating the current or potential benefits of digital solutions. The PhD then focused 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. Our evaluation of the ridesharing potential indicates that 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. It is based on a synthetic travel demand computed using the existing software from Hörl et al. [57] that we ran on the public data for the Grenoble metropolitan area. Based on this population synthesis, we have developed an original matching algorithm, tailored to our problem, in order to evaluate the maximum potential offered by ridesharing [69]. The analysis tool, called MEPCEL, is now publicly available.

The internship of Ludmila Courtillat--Piazza in 2022, and now the SIA project, explore the broader issue of how to choose a research topic in computer science given the ecological emergency, and taking into account the complex relations between computer science, the ICT it helps design, society (that both shapes and is shaped by its digital tools) and the environment. As a first step, the PhD thesis of Baptiste de Goër approaches this question by focusing on how ICT related sustainability issues are and could be taught in computer classes, which raises similar questions regarding the relation between ICT and sustainability [22, 31]. Specifically, one reason teaching ICT related sustainability issues is difficult is that the state of the art in academia is still incomplete, and there are many disagreements and even controversies in the research community regarding what scientific knowledge in computer science is relevant to tackle sustainability issues. New teaching approaches (typically inspired by the way earth science and biology classes teach socially acute questions such as GMOs), but also new research at the intersection of computer science and STS are needed. In complement, we study how to make explicit (and discuss) the social values that influence both computer science and IT design. We have in particular started investigating how Illich's notion of conviviality could offer a framework for rethinking ICT based on values consistent with a strong sustainability view [25].

A third line of research on this topic relates to society's dependance 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 paper currently under submission, we study the resilience of mobile networks in the face of hardware aging and approach mobile network resilience using the stochastic geometry framework, propose a network model including multiple dependent techno-bands (TB) and derive 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.

7.5 Tooling up the teaching of proofs

Participants: Martin Bodin.

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 search 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 COQ proof assistant, which provides a strong logical base for understanding proofs. It was thus natural to us to try to tool up licence's classes with COQ. 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 into 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 focused in reducing the entry cost of using COQ, especially through its interface. Indeed, COQ has originally been designed by and for experts and not for non-computer science students. We study how to incorporate graphical aspects into COQ proofs, as mathematical proofs in licence tend to include non-textual elements: in order to fit as much as possible the expected mindset for the student, we work in making the interface accepts such non-textual elements.

In order to understand what an interface for COQ should look like, we assisted mathematical courses in the university and studied the kind of reasoning that teachers and students made in practice: all such arguments should have a corresponding element within the interface. This helped us understand the constraints that apply to an interface for teaching proofs in this setting, and by studying the program taught in mathematics in licence we were able to identify a suitable target subject: curve sketching[2].

We asked about 50 students who got their licence to solve a proof sketching problem, and took their draft. These students were chosen as they already had the abilities that we aim to teach at the licence's level. We studied all the arguments they used within their draft and categorised them on whether they were textual or non-textual, as well as the role each argument had in the overall proof they were (implicitly) building. The results were unexpectedly diverse in both their arguments and method, showing that although apparently simple curve sketching can support a wide variety of reasoning. This constitutes a precious dataset that will later be used to validate our interface for COQ: students should be able to somehow express all these arguments within our interface. We published our approach on building and validating such an interface [38].

We developed a web interface (see Tabvar in Section 6.1.4) to manipulate curve sketching charts in association to a more usual textual-based interaction with a tool similar to COQ. The goal of this software is to help students getting connexions between textual and non-textual proofs, and help them understanding what is expected in their proof. We have shown this interface to both COQ experts and math teachers, but are still in the process of integrating their feedbacks.

In parallel, we designed non-textual activities to teach logic to high school students, in the form of a small board game [37] featuring pirate maps. We hope to eventually integrate such activities into our interface to provide alternative views of logical objects.

[2] In French, 'Tableaux de variations', which is a particular kind of curve sketching thoroughly taught in France. It doesn't seem to be taught as thoroughly in English-speaking countries, hence the lack of a good translation. We shall use curve sketching in the following for this particular notion.

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, Aurélie Kong Win Chang, Pietro Lami.

DCORE has been 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 was 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 encompasses 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 the IREMI de 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.

9.2 Exploratory Actions

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 served as general chair of the International conference ESWEEK'2024.

Member of the organizing committees

- Sophie Quinton was a member of the program committee of the JSI 2024, for which she organized the session on the (non)neutrality of science.
- Martin Bodin was a member of the program committee of VSTTE'24.

10.1.2 Scientific events: selection

Member of the conference program committees

- Alain Girault was TPC member of the international conferences FDL'2024 and CPSAT'2024.
- Sophie Quinton was TPC member of the Undone Computer Science conference 2024.

10.1.3 Journal

Member of the editorial boards

- Alain Girault has served as guest editor of ACM TECS for a special issue on the best papers from FDL'2020 and FDL'2021.
- Alain Girault is associated editor of Real-Time Systems Journal.
- Gregor Goessler has served as a guest editor of the Elsevier IST special issue on "Application of causal modeling and inference in software engineering".

Reviewer - reviewing activities

- Alain Girault has reviewed articles for ACM TECS.
- Gregor Goessler has reviewed articles for the European Journal of Control and IEEE Trans. on Industrial Informatics.

10.1.4 Invited talks

- Alain Girault gave a talk titled "Synchronous programming of reactive systems (cyber-physical, real-time,embedded)" during the Orange Dev Days, Grenoble, France, December 2024.
- Sophie Quinton gave talks related to her SIA research project in Rennes (Séminaire d'éthique des mathématiques), Paris (Étudiant.es en lutte de Sorbonne Université) and Lyon (AG du CITI), and presented the Ateliers SEnS to Labos1point5's GT réflexion.

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 doctorale") 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).
- In 2024, Pascal Fradet was member of the Inria Grenoble hiring committee for junior researchers (CRCN).
- Since January 2019, Alain Girault has been Deputy Scientific Director at INRIA, in charge of the "Algorithmics, Programming, Software and Architecture" research domain, which encompasses 46 research teams. In 2024, he launched two Inria research challenges (LLM4Code and CocoRISCo), organized a scientific workshop on RISC-V, and supported/audited the creation of four new Inria research teams. He was also one of the three Inria sherpas for the "Hardware Components of AI" French program, which will kick-off in 2025.
- Alain Girault was member of the INRIA Junior Researcher Admission jury in 2024.
- Gregor Goessler is member of the *scientific jobs committee* at Inria Grenoble.

- Sophie Quinton facilitates 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 a hiring committee for 2 McF positions 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, Algorithimque 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 : Alain Girault, Modèles de Calcul : λ-calcul, 12 HeqTD, niveau L3, Univ. Grenoble Alpes, France
- Master : Sophie Quinton, Numérique responsable, 15 HeqTD, niveau M1, Grenoble INP (Ensimag), France
- Master : Sophie Quinton, Kaléidoscope, 3 HeqTD, niveau M1, Grenoble INP, France
- École doctorale: Sophie Quinton gave a 3h course "Sciences, environnements, sociétés" at the College des Écoles Doctorales.
- Master : Martin Bodin, Kaléidoscope, 1 HeqTD, niveau M1, Grenoble INP, France
- Licence : Martin Bodin, Modèles de Calcul : λ-calcul, 12 HeqTD, niveau L3, Univ. Grenoble Alpes, France
- Master : Martin Bodin, LTPF, 27 HeqTD, niveau M1, Polytech (UGA), France

10.2.2 Supervision

- Gregor Goessler: PhD in progress: Aurélie Kong Win Chang, "Causal explanations for concurrent programs in Erlang"; since January 2021; co-advised by Gregor Goessler and Jérôme Feret.
- Gregor Goessler: PhD in progress: Alexander Obeid Guzman, "Inference of causal models for networks from single observations"; since January 2024.
- Sophie Quinton et Alain Girault: PhD completed: Aina Rasoldier, "Comment évaluer le potentiel d'une solution numérique face à l'urgence écologique ? Application aux plateformes de covoiturage régulier à l'échelle locale".

- 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”.
- Jean-Bernard Stefani: PhD completed: Giovanni Fabbretti on reversibility for distributed programs (UGA), Pietro Lami on reversibility for shared memory concurrent programs (UGA and U. Bologna).
- Jean-Bernard Stefani: PhD in progress: Boubacar Diarra on verification of Kubernetes configurations (U. Lille).

10.2.3 Juries

- Gregor Goessler, president of the PhD jury of Lei Zan, UGA.
- Sophie Quinton, member of the PhD committee of Vincent Giraud, ENS.

10.3 Popularization

10.3.1 Specific official responsibilities in science outreach structures

- Martin Bodin is local referent for Chiche!, a nationwide program to raise awareness of digital science and technology among second-grade students.

10.3.2 Productions (articles, videos, podcasts, serious games, ...)

- Sophie Quinton contributed to the serious game PhoneImpact.
- Sophie Quinton co-authored an article about water related issues of the cloud [12].

10.3.3 Participation in Live events

- Alain Girault gave an outreach talk titled “Régulateur Intelligent/Adaptatif de Vitesse” at the Conférence Sociéte Numérique en Question, Grenoble, March 2024.
- Sophie Quinton gave a conference on the social and environmental implications of ICT at the MathC2+ event and co-organized a workshop on ICT and water at the Fête de la science.
- Martin Bodin gave a talk at the MathC2+ event with Alain Girault.
- Martin Bodin animated an activity for high school students during the MathC2+ 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. URL: https://hal.inria.fr/hal-02400019.

[2] 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.

[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. URL: https://hal.inria.fr/hal-03940615 (cit. on p. 8).

[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. 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. 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. 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. URL: https://hal.inria.fr/hal-01303090.

[9] 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.

[10] 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.

[11] 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. URL: https://hal.science/hal-03770727.

11.2 Publications of the year

International journals

[12] S. Bouveret, A. Bugeau, A.-C. Orgerie and S. Quinton. 'De l'eau dans les nuages'. In: *Annales des Mines - Enjeux Numériques* 27 (Sept. 2024), pp. 40–47. URL: https://hal.science/hal-04698568 (cit. on p. 17).

[13] 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 (cit. on p. 8).

[14] 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* (5th Dec. 2024), p. 47. DOI: 10.1007/s11241-024-09433-5. URL: https://inria.hal.science/hal-02371742 (cit. on p. 9).

[15] 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)* 17.1 (31st Mar. 2024), pp. 1–26. DOI: 10.1145/3626103. URL: https://hal.science/hal-04237266 (cit. on p. 8).

[16] P. Lami, I. Lanese, J.-B. Stefani, C. Sacerdoti Coen and G. Fabbretti. 'Reversible debugging of concurrent Erlang programs: Supporting imperative primitives'. In: *Journal of Logical and Algebraic Methods in Programming* 138 (Apr. 2024), p. 100944. DOI: 10.1016/j.jlamp.2024.100944. URL: https://inria.hal.science/hal-04575594 (cit. on p. 10).

National journals

[17] E. Beffara and M. Bodin. 'Un jeu de plateau pour comprendre la dualité en logique'. In: *Adjectif : analyses et recherches sur les TICE* 2024.1 (9th Apr. 2024), pp. 1–5. URL: https://hal.science/hal-04748800.

Invited conferences

[18] I. Lanese and G. Gössler. 'Causal Debugging for Concurrent Systems'. In: RC 2024 - 16th International Conference on Reversible Computation. Vol. LNCS-14680. Reversible Computation : 16th International Conference, RC 2024, Toruń, Poland, July 4–5, 2024, Proceedings. Torun, Poland: Springer Nature Switzerland, 29th May 2024, pp. 3–9. DOI: 10.1007/978-3-031-62076-8_1. URL: https://inria.hal.science/hal-04610282.

International peer-reviewed conferences

[19] B. Diarra, K. Guillouard, 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.

[20] G. Fabbretti, I. Lanese and J.-B. Stefani. 'Reversibility with Holes'. In: *Lecture notes in computer science*. RC 2024 - 16th International Conference on Reversible Computation. Vol. LNCS-14680. Reversible Computation : 16th International Conference, RC 2024, Toruń, Poland, July 4–5, 2024, Proceedings. Torun, Poland: Springer, 2024, pp. 69–74. DOI: 10.1007/978-3-031-62076-8_5. URL: https://inria.hal.science/hal-04610283.

[21] P. Fradet, A. Girault and A. Honorat. 'Parallel scheduling of task graphs with minimal memory requirements'. In: 39th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2025). Milano, Italy, 3rd June 2025. URL: https://inria.hal.science/hal-04879748 (cit. on p. 8).

[22] B. de Goër, H. Micha and S. Quinton. 'Informatique et durabilité, une difficile transposition didactique'. In: *Actes du colloque Didapro 10 sur la Didactique de l'informatique et des STIC. Volume 1*. DidaSTIC 2024 - Colloque Didapro 10 sur la Didactique de l'informatique et des STIC. Louvain-La-Neuve, Belgium, 2024, pp. 23–32. URL: https://hal.science/hal-04482133 (cit. on p. 11).

[23] P. Lami, I. Lanese and J.-B. Stefani. 'A Small-Step Semantics for Janus'. In: *Lecture notes in computer science*. RC 2024 - 16th International Conference on Reversible Computation. Vol. LNCS-14680. Reversible Computation : 16th International Conference, RC 2024, Toruń, Poland, July 4–5, 2024, Proceedings. Torun, Poland: Springer, 2024, pp. 105–123. DOI: 10.1007/978-3-031-62076-8_8. URL: https://inria.hal.science/hal-04610285.

Conferences without proceedings

[24] C. Assaad, E. Devijver, É. Gaussier, G. Gössler and A. Meynaoui. 'Identifiability of total effects from abstractions of time series causal graphs'. In: 40th Conference on Uncertainty in Artificial Intelligence. Barcelone, Spain, 2024. URL: https://hal.science/hal-04250602.

[25] S. Quinton and J.-B. Stefani. 'Taking conviviality seriously (extended abstract)'. In: 2024 - 1st conference on Undone Science in Computer Science. Nantes, France, 2024, pp. 1–4. URL: https://hal.science/hal-04448759 (cit. on p. 11).

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

[26] *Introduction to the Special Issue on Specification and Design Languages (FDL 2021)* 23.5 (14th Aug. 2024). DOI: 10.1145/3677316. URL: https://inria.hal.science/hal-04682174.

Doctoral dissertations and habilitation theses

[27] A. Rasoldier. 'Assessing the potential of a digital Solution in the context of ecological urgency : application to regular carpooling platforms at local scale'. Université Grenoble Alpes [2020-....], 12th Feb. 2024. URL: https://theses.hal.science/tel-04702139 (cit. on p. 11).

Reports & preprints

[28] P. Fradet, A. Girault and A. Honorat. *Graph Transformations for Memory Peak Minimization by Scheduling*. RR-9554. Inria - Research Centre Grenoble – Rhône-Alpes, 5th Sept. 2024, pp. 1–42. URL: https://inria.hal.science/hal-04688921 (cit. on p. 8).

Scientific popularization

[29] P. Ciblat, J. Combaz, M. Coupechoux, K. Marquet and A.-C. Orgerie. 'Effets environnementaux de la 5G (Partie 1) : La technologie 5G'. In: *1024 : Bulletin de la Société Informatique de France* 23 (Apr. 2024), pp. 53–72. DOI: 10.48556/SIF.1024.23.53. URL: https://hal.science/hal-03810501.

[30] P. Ciblat, J. Combaz, M. Coupechoux, K. Marquet and A.-C. Orgerie. 'Effets environnementaux de la 5G (partie 2) : Applications envisagées et acteurs impliqués'. In: *1024 : Bulletin de la Société Informatique de France* 24 (Nov. 2024), pp. 99–127. DOI: 10.48556/SIF.1024.24.99. URL: https://hal.science/hal-04828952.

[31] B. de Goër. *Parler de durabilité en cours d'informatique au lycée*. 10th Apr. 2024. URL: https://inria.hal.science/hal-04644556 (cit. on p. 11).

11.3 Cited publications

[32] ARTEMIS Joint Undertaking. *ARTEMIS Strategic Research Agenda*. 2011 (cit. on p. 3).

[33] *Automotive Open System Architecture*. 2003. URL: http://www.autosar.org (cit. on p. 3).

[34] 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. 4).

[35] N. Bansal, T. Kimbrel and K. Pruhs. 'Speed Scaling to Manage Energy and Temperature'. In: *Journal of the ACM* 54.1 (2007) (cit. on p. 9).

[36] 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. 3).

[37] E. Beffara and M. Bodin. 'Un jeu de plateau pour comprendre la dualité en logique'. In: *Adjectif : analyses et recherches sur les TICE (Adjectif)*. 2024 (cit. on p. 12).

[38] E. Beffara, M. Bodin, N. Mandran and R. Molinier. 'Instrumentation de l'association de registres sémiotiques dans un assistant de preuve'. In: *Conférence sur les Environnements Informatiques pour l'Apprentissage Humain (EIAH)*. 2023 (cit. on p. 12).

[39] 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. 4).

[40] S. Borkar. 'Designing Reliable Systems from Unreliable Components: The Challenges of Transistor Variability and Degradation'. In: *IEEE Micro* 25.6 (2005) (cit. on p. 4).

[41] 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 (cit. on p. 8).

[42] 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 (cit. on p. 8).

[43] 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. 4).

[44] 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. 3).

[45] G. Fabbretti, I. Lanese and J.-B. Stefani. *A Behavioral Theory for Distributed Systems with Weak Recovery*. 2024. arXiv: 2406.12574 [cs.LO]. URL: https://arxiv.org/abs/2406.12574 (cit. on p. 10).

[46] G. Fabbretti, I. Lanese and J.-B. Stefani. 'Generation of a Reversible Semantics for Erlang in Maude'. In: *ICFEM 2022 - 23rd International Conference on Formal Engineering Methods*. Vol. LNCS13478. Lecture Notes in Computer Science. Madrid, Spain: Springer International Publishing, Oct. 2022, pp. 106–122. DOI: 10.1007/978-3-031-17244-1_7. URL: https://inria.hal.science/hal-03916227 (cit. on p. 10).

[47] 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 (cit. on p. 8).

[48] 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 (cit. on p. 8).

[49] F. C. Gärtner. 'Fundamentals of Fault-Tolerant Distributed Computing in Asynchronous Environments'. In: *ACM Computing Surveys* 31.1 (1999) (cit. on p. 5).

[50] 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. URL: https://hal.science/hal-03030416 (cit. on p. 9).

[51] 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. URL: https://hal.inria.fr/hal-02888573 (cit. on p. 9).

[52] 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. URL: https://hal.inria.fr/hal-02557148 (cit. on p. 9).

[53] 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 (cit. on p. 9).

[54] 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. 4).

[55] 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. 4).

[56] 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. 3).

[57] S. Hörl and M. Balac. 'Open Synthetic Travel Demand for Paris and Île-de-France: Inputs and Output Data'. In: *Data in Brief* 39 (Dec. 2021), p. 107622. DOI: 10.1016/j.dib.2021.107622. URL: https://hal.science/hal-03448125 (cit. on p. 11).

[58] 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. URL: https://hal.science/hal-04222884 (cit. on p. 9).

[59] 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. 4).

[60] 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. 4).

[61] J. Lorch and A. Smith. 'Improving Dynamic Voltage Scaling Algorithms with PACE'. In: *Joint International Conference on Measurements and Modeling of Computer Systems, SIGMETRICS'01*. Cambridge (MA), USA: ACM, June 2001, pp. 50–61 (cit. on p. 9).

[62] J. Lorch and A. Smith. 'PACE: A New Approach to Dynamic Voltage Scaling'. In: *IEEE Trans. on Computers* 53.7 (2004), pp. 856–869 (cit. on p. 9).

[63] 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. 9).

[64] 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. 4).

[65] M. Moore. *Causation and Responsibility*. Oxford, 1999 (cit. on p. 4).

[66] 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 (cit. on p. 8).

[67] J. Pearl. 'Causal inference in statistics: An overview'. In: *Statistics Surveys* 3 (2009), pp. 96–146 (cit. on p. 4).

[68] 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. 9).

[69] 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. URL: https://inria.hal.science/hal-04401006 (cit. on p. 11).

[70] J. Rushby. *Partitioning for Safety and Security: Requirements, Mechanisms, and Assurance*. Tech. rep. CR-1999-209347. NASA Langley Research Center, 1999 (cit. on p. 5).