RESEARCH CENTRE

Inria Center at Université Grenoble Alpes

IN PARTNERSHIP WITH:

Université de Grenoble Alpes

2022
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

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
- A2.5.2. – Component-based Design

Other research topics and application domains

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

1 Team members, visitors, external collaborators

Research Scientists

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

Faculty Member

- Xavier Nicollin [GRENOBLE INP, Associate Professor]

Post-Doctoral Fellow

- Alexandre Honorat [INRIA]

PhD Students

- Giovanni Fabbretti [INRIA]
- Aurélie Kong Win Chang [INRIA]
- Pietro Lami [INRIA]
- Thomas Mari [GRENOBLE INP]
- Aina Rasoldier [INRIA]

Technical Staff

- Roger Pissard-Gibollet [INRIA]

Interns and Apprentices

- Ludmila Courtillat-Piazza [ENS RENNES, from Aug 2022]

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

A few remarks are in order:

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

3 Research program

The SPADES research program is organized around three main themes, *Design and Programming Models*, *Certified real-time programming*, and *Fault management and causal analysis*, that seek to answer the three key questions identified in Section 2. We plan to do so by developing and/or building on programming languages and techniques based on formal methods and formal semantics (hence the use of *"sound programming"* in the project-team title). In particular, we seek to support design where correctness is obtained by construction, relying on proven tools and verified constructs, with programming languages and programming abstractions designed with verification in mind.

3.1 Design and Programming Models

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

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

On the design side, we focus on the definition of component-based models for software architectures combining distribution, dynamicity, real-time and fault-tolerant aspects. Component-based construction has long been advocated as a key approach to the "correct-by-construction" design of complex embedded systems [43]. Witness component-based toolsets such as PTOLEMY [33], BIP [27], or the modular architecture frameworks used, for instance, in the automotive industry (AUTOSAR) [25]. 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 [28]), 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 [26, 32], in particular in relation with mixed criticality systems. The main goal in this theme is to address several of these open questions.

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

Distributed reactive programming is rightly singled out as a major open issue in the recent, but heavily biased (it essentially ignores recent research in synchronous and dataflow programming), survey by Bainomugisha et al. [26]. For our part, we intend to focus on devising synchronous programming languages for distributed systems and precision-timed architectures.

3.3 Fault Management and Causal Analysis

Managing faults is a clear and present necessity in networked embedded systems. At the hardware level, modern multicore architectures are manufactured using inherently unreliable technologies [29, 41]. 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 [49], natural sciences, law [50], and statistics [52], 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 [45], to allow the diagnosis of faults in a complex concurrent system [42], or to enforce accountability [44], 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 [36]), or is broken (*e.g.*, by limiting fault propagation [55]).

4 Application domains

4.1 Industrial Applications

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

Industrial acceptance of formal techniques, as well as their deployment, goes necessarily through their usability by specialists of the application domain, rather than of the formal techniques themselves. Hence, we are looking to propose domain-specific (but generic) realistic models, validated through experience (*e.g.*, control tasks systems), based on formal techniques with a high degree of automation (*e.g.*, synchronous models), and tailored for concrete functionalities (*e.g.*, code generation). We also consider the development of formal tools that can certify the result of industrial applications (see *e.g.*, CertiCAN in Sec. 7.2.2).

4.2 Current Industrial Cooperations

Regarding applications and case studies with industrial end-users of our techniques, we cooperate with Orange Labs on software architecture for cloud services. We also collaborate with RTaW regarding the integration of our CAN-bus analysis certifier (CertiCAN) in the RTaW-Pegase program suite.

5 Social and environmental responsibility

5.1 Footprint of research activities

We have not yet computed the footprint of our research activities in 2022. A tool designed in collaboration with Labo 1.5 should be available next year. At the time being, it allows us to compute the carbon footprint due to individual travels (both professional and commute travels), due to new hardware, but not the footprint due to our share of the Inria services (data centers, networks, ...) nor the building usage. Finally, only the *carbon* footprint can be computed, but not the usage of other resources (water, raw materials) or the resulting pollution (impact on the bio-diversity).

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, this type of research may contribute to make more acceptable or even to promote many problematic systems such as IoT, drones, avionics, autonomous vehicles, ... with a potential strong 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. Around 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. SPADES organized its own Atelier SEnS in 2022. Sophie Quinton has facilitated 8 Ateliers SEnS in 2022.

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 and platforms

6.1 New software

6.1.1 CertiCAN

Name: Certifier of CAN bus analysis results

Keywords: Certification, CAN bus, Real time, Static analysis

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

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

Authors: Xiaojie Guo, Pascal Fradet, Sophie Quinton

Contact: Xiaojie Guo

6.1.2 cloudnet

Name: Cloudnet

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

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

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

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

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

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

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

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

Publication: hal-02940938v1

Contact: Philippe Merle

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

Partner: Orange Labs

7 New results

7.1 Design and Programming Models

Participants: Pascal Fradet, Alain Girault, Xavier Nicollin, Jean-Bernard Stefani.

7.1.1 Dynamicity in dataflow models

Dataflow Models of Computation (MoCs) are widely used in embeddedsystems, 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 lacks 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 static analyses for boundedness and liveness. We have proposed several parametric dataflow models of computation (MoCs) (SPDF [34] and BPDF [51]), we have written a survey providing a comprehensive description of the existing parametric dataflow MoCs [30], and we have studied *symbolic* analyses of dataflow graphs [31]. We have proposed an original method to deal with lossy communication channels in dataflow graphs [35].

More recently, we have studied models allowing *dynamic reconfigurations* of the *topology* of the dataflow graphs. This is required by many modern streaming applications that have a strong need for reconfigurability, for instance to accommodate changes in the input data, the control objectives, or the environment. We have proposed a new MoC called Reconfigurable Dataflow (RDF) [3]. RDF extends SDF with transformation rules that specify how the topology and actors of the graph may be reconfigured. Starting from an initial RDF graph and a set of *transformation rules*, an arbitrary number of new RDF graphs can be generated at runtime. The major feature and advantage of RDF is that it can be *statically analyzed* to guarantee that all possible graphs generated at runtime will be connected, consistent, and live. To the best of our knowledge, RDF is the only dataflow MoC allowing an arbitrary number of topological reconfigurations while remaining statically analyzable.

We have also worked on a practical way to integrate statically parameterized SDF graphs into the PREESM tool [18]. While this work does not provide a theoretical analysis of the evolution of the graph according to the parameters as in SPDF, it supports a wider range of parameterized expressions including complex mathematic operations. On the few tested applications, we have shown that a mere design-space exploration on a subset of all possible parameter configurations was providing reasonable approximations of the impact of each parameter on throughput, latency and energy.

Up to now, the main application domain of dataflow models has been streaming multimedia applications but they also seem particularly well suited to the efficient implementation of neural networks. We have started an exploratory action (see Section 9.2) to study the potential of dataflow MoCs for the implementation of neural networks. We expect advances in the form of a better time efficiency and a lower memory and energy consumption.

We are currently working on the reduction of the memory footprint of tasks graphs scheduled on unicore processors. This is motivated by the fact that some recent neural networks such as GPT-3, seen as tasks graphs, use too much memory and cannot fit on a single GPU. Using several exact and heuristic techniques, we are able to produce schedules that minimize the memory requirement of a sequential dataflow application, and therefore tasks graphs as well. Another technique used by memory greedy neural networks is activity and gradient checkpointing (a.k.a. rematerialization) which recompute intermediate values rather than keeping them in memory. We are now studying rematerialization in the more general dataflow context.

7.2 Certified Real-Time Programming

Participants: Pascal Fradet, Alain Girault, Xavier Nicollin, Sophie Quinton.

7.2.1 A Markov Decision Process approach for energy minimization policies

In 2022, we have completed our work on a very general model of real-time systems, made of a single-core processor equipped with DVFS and an infinite sequence of preemptive real-time jobs. Each job J_i is characterized by the triplet (τ_i, w_i, d_i), where τ_i is the *inter-arrival time* between J_i and J_{i-1}, w_i is the *actual size* of J_i, upper-bounded by the maximal size W, and d_i is the *relative deadline* of J_i, upper-bounded by Δ. The key point is that the system is *non-clairvoyant*, meaning that, at release time, w_i is not known until the job J_i actually terminates. What is available to the processor are the *statistical information* on the jobs' characteristics: release time, AET, and relative deadline. In this context, we have proposed a Markov Decision Process (MDP) solution to compute the optimal online speed policy guaranteeing that each job completes before its deadline and minimizing the energy consumption. To the best of our knowledge, our MDP solution is *the first to be optimal*. We have also provided counter examples to prove that the two previous state of the art algorithms, namely OA [56] and PACE [47], are both sub-optimal. Finally, we have proposed a new heuristic online speed policy called Expected Load (EL) that incorporates an aggregated term representing the future expected jobs into a speed equation similar to that of OA. A journal paper is currently under review.

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

This was the topic of Stephan Plassart's PhD [53][37, 39, 38], funded by the CASERM Persyval project, who defended his PhD in June 2020.

7.2.2 Formal proofs for schedulability analysis of real-time systems

We contribute to Prosa [23], a Coq library of reusable concepts and proofs for real-time systems analysis. A key scientific challenge is to achieve a modular structure of proofs, *e.g.*, for response time analysis. Our goal is to use this library for:

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

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

We have completed our work on a formal connection between Network Calculus (NC) and Response Time Analysis (RTA) [19]. This enables specialists of both formalisms to get increased confidence in their models (or to discover errors, as has happened). The presented mathematical results are all mechanically checked with the interactive theorem prover Coq, building on existing formalizations of RTA (namely Prosa) and NC (namely NCCoq). Establishing such a link between NC and RTA paves the way for improved real-time analyses obtained by combining both theories to enjoy their respective strengths (e.g., multicore analyses for RTA or clock drifts for NC).

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

7.3 Fault Management and Causal Analysis

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

7.3.1 Causal Explanations for Embedded Systems

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

In particular, as part of the SEC project, we are investigating how such techniques can be extended to classes of hybrid systems. As a first result we have studied the problem of explaining faults in real-time systems [48]. We have provided a formal definition of causal explanations on dense-time models, based on the well-studied formalisms of timed automata and zone-based abstractions. We have proposed a symbolic formalization to effectively construct such explanations, which we have implemented in a prototype tool. Basically, our explanations identify the parts of a run that move the system closer to the violation of an expected safety property, where safe alternative moves would have been possible.

We have recently generalized the work of [48] and defined *robustness functions* as a family of mappings from system states to a scalar that, intuitively, associate with each state its distance to the violation of a given safety requirement, e.g., in terms of the remaining number of bad system moves or of the time remaining to react. An explanation then summarizes the portions of the execution on which robustness decreases. However, as our instantiation of robustness in [48] is defined on a discrete abstraction, robustness may decrease in discrete steps once some timing threshold is crossed, thus exonerating the preceding absence of action. We are currently working on a truly hybrid definition of robustness functions that “anticipate” such thresholds, hence ensuring a smooth decrease indicating early when a dangerous event is approaching.

7.3.2 Causal Explanations in Concurrent Programs

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

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

7.3.3 Reversibility for concurrent and distributed debugging

Concurrent and distributed debugging is a promising application of the notion of reversible computation [40]. As part of the ANR DCore project, we have contributed to the theory behind, and the development of the CauDEr reversible debugger for the Erlang programming language. In [14], we have shown how to automate using the Maude rewriting logic environment the generation of a reversible semantics for a concurrent program such as Erlang, in effect implementing in Maude the theory developped by Lanese and Medic [46]. In the same paper we have also shown how to automatically generate the semantics of the imperative rollback primitive which is at the core of the CauDEr debugger. In [15], we have extended the reversible semantics of Erlanf and CauDEr to take into account of imperative constructs in Erlang allowing Erlang processes to access a shared map of process names and process identifiers. This is an first instance of dealing with reversibility in presence of a form of shared memory.

We have also started this year two novel threads of activity: studying reversibility for distributed programs and studying reversibilty for concurrent programs based on shared memory. For the time being we only have preliminary results on these two threads. On the first one, we have devised a small distributed process calculus featuring located processes with location and link failures with recovery and are currently working on its behavioral theory. This small calculus is intended as a reasonably faithful abstraction of the behavior of Erlang systems in presence of node and link failures. On the second thread, we have introduced a modular operational framework for the definition of shared memory concurrency models, and shown that we can capture in our framework several forms of weak memory models, including models of transactional memory.

7.4 Transversal activity: ICT and the Anthropocene

Participants: Martin Bodin, Ludmila Courtillat-Piazza, 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"). In [9], we have detailed limitations of state of the art assessments of such claims. First, most papers do not provide enough details on the scenarios underlying their evaluations: which hypotheses they are based on and why, and why specific scenarios are chosen as the baseline. This is a key point because it may lead to overestimating the current or potential benefits of digital solutions. Second, results are rarely discussed in the context of global strategies for greenhouse gas (GHG) emissions reduction. These leaves open how the proposed technologies would fit into a realistic plan for meeting current GHG reduction goals. To overcome the underlined limitations, we propose a set of guidelines that all studies on digital solutions for mitigating GHG emissions should satisfy, point out overlooked research directions, and provide concrete examples and initial results for the specific case of ridesharing. We are now working on estimating the potential of ridesharing as a solution for reducing the GHG emissions of commuting.

During her short internship in the group, Ludmila Courtillat-Piazza addressed the issue of how to choose a research topic in computer science taking into account the systemic nature of the environmental impact of ICT. In addition to a literature review, she experimented with a case study on web design.

The SPADES team has started working together on a project proposal to investigate the current role played by ICT in the Anthropocene as well as new approaches to their design. We have identified the following main challenges: How do local measures meant to reduce the environmental impact of ICT relate (or not) to global effects? What can we learn from, and what are the limits of, current quantitative approaches for environmental impact assessment and their use for public debate and policy making? Which criteria could/should we take into account to design more responsible computer systems (other than efficiency, which is already well covered and subject to huge rebound effects in the case of digital technologies)? To come up with a solid research agenda, we are thus studying the state of the art of many new topics, including STS (Science and Technology Studies), low tech software and hardware, lifecyle assessment, (digital) commons... A new network of collaborations is also in the making, in particular with colleagues from social sciences.

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

RT-proofs

Participants: Pascal Fradet, Sophie Quinton.

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

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

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

The results obtained in 2022 in connection with the RT-proofs project are described in Section 7.2.2.

DCORE

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

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

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

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

9.1.2 Défi Inria

LiberAbaci

Participants: Yves Bertot, Jean-Marie Madiot, Arthur Charguéraud, Nicolas Tabareau, Hugo Herbelin, Martin Bodin, Sylvie Boldo, Micaela Mayero.

LIBERABACI is a project between Inria project teams CAMBIUM, CAMUS, GALLINETTE, πr^2, SPADES, STAMP, TOCCATA, and the Laboratoire d'Informatique de Paris-Nord.

The overall objective is to study how one could use the COQ proof assistant in a Mathematical course in the University to help teaching proofs. At Spades, Martin Bodin is working with IREM de Grenoble to involve math teachers and didactic researchers to the project.

9.2 Exploratory Actions

DF4DL

Participants: Pascal Fradet, Alain Girault, Alexandre Honorat.

The DF4DL action is funded by Inria's DGDS. It aims at exploring the use of the dataflow model of computation to better program deep neural networks, in particular a variant called dynamic neural networks. As a first step, we have studied the problem of minimizing the peak memory requirement for the execution of a dataflow graph. This is of paramount importance for deep neural networks since the largest ones cannot fit on a single core due to their very high memory requirement. It happens that this was an open problem since 1995 [54], and the new algorithms we proposed in 2022 have allowed us to significantly lower the memory peak as well as the compute time required to find it. For instance, we managed to lower the optimal memory peak of the satellite benchmark in [54] from 1.920 units (obtained after 4 days of compute time) to 1.680 units (obtained after 10 milliseconds).

10 Dissemination

10.1 Promoting scientific activities

10.1.1 Scientific events: organisation

- Sophie Quinton co-organized a workshop (journée d'étude) on the connection between digital commons and the ecological transition.
- Sophie Quinton co-organized a panel at the Archipel conference on the relevance (or lack thereof) of using science and technologies for sustainability.

General chair, scientific chair

- Alain Girault belongs to the Steering Committee of ESWEEK (since 2022).

10.1.2 Scientific events: selection

Member of the conference program committees

- Alain Girault was a PC member of EMSOFT 2022 and FDL 2022.
- Sophie Quinton was a PC member of ECRTS 2022.

Reviewer

- Alain Girault reviewed articles for ECRTS 2022.

10.1.3 Journal

Member of the editorial boards

- Alain Girault is associate editor for EURASIP Journal on Embedded Systems (since 2005) and Real-Time Systems Journal (since 2020).
- Alain Girault was guest editor of a special issue of ACM TECS titled “Specification and Design Languages”.

Reviewer - reviewing activities

- Pascal Fradet has reviewed an article for ACM Trans. on Embedded Computing Systems.
- Alain Girault reviewed articles for IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems and for IEEE Transactions on Parallel and Distributed Systems.
- Gregor Gössler reviewed articles for ACM Transactions on Computer Systems, Elsevier Artificial Intelligence, and Mathematical Structures in Computer Science.

10.1.4 Invited talks

- Martin Bodin gave an invited talk about Coq at the Séminaire de l’IREM de Grenoble.
- Sophie Quinton gave a keynote on ICT and sustainability at the 2nd Inria-DFKI European Summer School on AI (IDESSAI 2022).

10.1.5 Leadership within the scientific community

- Sophie Quinton was a member of the ECRTS Executive Committee until September 2022. She is now a member of the ECRTS Advisory Board.
- Sophie Quinton co-chairs a working group of the GDR CIS associated with the Center for Internet and Society focused on environmental issues.

10.1.6 Research administration

- Pascal Fradet is head of the committee for doctoral studies (“Responsable du comité des études doctorales”) of the Inria Grenoble research center. He is the local correspondent for the young researchers Inria mission (“Mission jeunes chercheurs”) and the substitute of the director of the Inria Grenoble research center at the doctoral school council (MSTII).
- Alain Girault is Deputy Scientific Director for the domain “Algorithmics, Programming, Software and Architecture” (since 2019).
- Alain Girault is Scientific Manager of the Cyber-Physical Systems axis of the labex Persyval-Lab.
- Gregor Gössler is member of the *commission of scientific jobs* of the Inria Grenoble research center.
- Sophie Quinton leads the SEnS-GRA group which hosts discussions and proposes actions regarding the environmental and societal impact of our research at Inria Grenoble Rhône-Alpes.
- Sophie Quinton was a member of the CRCN hiring committee in Grenoble.

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, 12 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 : Martin Bodin, Modèles de Calcul : λ-calcul, 12 HeqTD, niveau L3, Univ. Grenoble Alpes, France
- Licence : Martin Bodin, Projet Génie logiciel, 55 HeqTD, niveau M1, Grenoble INP (Ensimag), France
- Master: Jean-Bernard Stefani, Formal Aspects of Component Software, 9h, MOSIG, Univ. Grenoble Alpes, France.
- École doctorale: Sophie Quinton gave a 3h course "Sciences, environnements, sociétés" at the MSTII doctoral school.
- Sophie Quinton gave an introduction to the environmental impact of ICT and participated in the associated panel as part of a lecture in the "Entangled futures" course at the University of the Arts London.

10.2.2 Supervision

- Gregor Gössler: PhD in progress: Thomas Mari, “Construction of Safe Explainable Cyber-physical systems”; Grenoble INP; since October 2019; co-advised by Gregor Gössler and Thao Dang.
- Gregor Gössler: PhD in progress: Aurélie Kong Win Chang, "Abstractions for causal analysis and explanations in concurrent programs"; since January 2021; co-advised by Gregor Gössler and Jérôme Feret.
- Sophie Quinton and Alain Girault: PhD in progress: Aina Rasoldier, ICT in the Anthropocene: Technical and social challenges at the local scale.
- Sophie Quinton: PhD completed: Leonie Köhler "A Compositional Performance Analysis for Embedded Computing Systems with Weakly-Hard Real-Time Constraints", TU Braunschweig; defended in January 2022; co-advized with Rolf Ernst.
- Martin Bodin: intership student: Jean Abou-Samra, “Hiérarchies de monades”.

- Jean-Bernard Stefani: PhD in progress: Giovanni Fabbretti, "Reversibility in distributed systems with crash faults and recovery", UGA.
- Jean-Bernard Stefani: PhD in progress: Pietro Lami, "Reversibility in concurrent systems with shared memory", co-advised with Ivan Lanese (U. Bologne), UGA.

10.2.3 Juries

- Sophie Quinton was a member of the PhD committee of Lucien Rakotomalala, "Preuve Formelle en calcul réseau", defended in February 2022 (Université de Toulouse).
- Martin Bodin was a member of the PhD committee of Julien Braine, “The Data-abstraction Framework : abstracting unbounded data-structures in Horn clauses, the case of arrays”, defended in May 2022 (ENS Lyon).
- Jean-Bernard Stefani was a rapporteur on the PhD of Louis Noizet, "Necro, la sémantique sans y laisser les os", defended in September 2022 (U. Rennes).

10.3 Popularization

10.3.1 Articles and contents

- Sophie Quinton co-authored an article entitled "The crisis of the scientific mind : an investigation, a tragedy and a collective redistribution of roles" [13] in connection with the Ateliers SEnS.

10.3.2 Interventions

- Sophie Quinton presented the Ateliers SEnS at a panel at the conference "En route pour la sobriété numérique ?" in Geneva.
- Sophie Quinton gave a talk about the challenges of developing a "responsible ICT" approach at a Planet Tech’Care workshop and participated in a panel on the environmental impact of ICT at EPFL.
- Sophie Quinton attended a "Forum des métiers" organized by the Lycée du Grésivaudan.
- Martin Bodin and Alain Girault participated in the MathC2+ event.
- Martin Bodin gave a talk at the Séminaire sport-étude de l’ENS Lyon.

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] 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.

[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.

[4] 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. URL: https://hal.archives-ouvertes.fr/hal-02119024.

[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] 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.

11.2 Publications of the year

International journals

[10] 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.

[11] P. Fradet, X. Guo and S. Quinton. 'CertiCAN : Certifying CAN Analyses and Their Results'. In: *Real-Time Systems* (31st Dec. 2022). URL: https://hal.inria.fr/hal-03941096.

[12] M. Perrin, A. Mostefaoui, G. Bonin and L. Courtillat-Piazza. 'Extending the Wait-free Hierarchy to Multi-Threaded Systems'. In: *Distributed Computing* 35 (Aug. 2022), pp. 375–398. DOI: 10.1007/s00446-022-00425-x. URL: https://hal.science/hal-03819422.

[13] E. Tannier, V. Daubin and S. Quinton. 'The crisis of the scientific mind : an investigation, a tragedy and a collective redistribution of roles'. In: *Les Cahiers de Framespa : e-Storia* 40 (30th June 2022). DOI: 10.4000/framespa.13150. URL: https://hal.archives-ouvertes.fr/hal-03714886.

International peer-reviewed conferences

[14] G. Fabbretti, I. Lanese and J.-B. Stefani. 'Generation of a Reversible Semantics for Erlang in Maude'. In: *Formal Methods and Software Engineering : 23rd International Conference on Formal Engineering Methods, ICFEM 2022, Madrid, Spain, October 24–27, 2022, Proceedings*. ICFEM 2022 - 23rd International Conference on Formal Engineering Methods. Vol. LNCS13478. Lecture Notes in Computer Science. Madrid, Spain: Springer International Publishing, 10th Oct. 2022, pp. 106–122. DOI: 10.1007/978-3-031-17244-1_7. URL: https://hal.inria.fr/hal-03916227.

[15] P. Lami, I. Lanese, J.-B. Stefani, C. Sacerdoti Coen and G. Fabbretti. 'Reversibility in Erlang: Imperative Constructs'. In: *Reversible Computation : 14th International Conference, RC 2022, Urbino, Italy, July 5–6, 2022, Proceedings*. RC 2022 - 14th International Conference on Reversible Computation. Vol. LNCS-13354. Lecture Notes in Computer Science. Urbino, Italy: Springer International Publishing, 28th June 2022, pp. 187–203. DOI: 10.1007/978-3-031-09005-9_13. URL: https://hal.inria.fr/hal-03915947.

[16] 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.

National peer-reviewed Conferences

[17] J. Abou-Samra, Y. Zakowski and M. Bodin. ‘Effectful Programming across Heterogeneous Computations -Work in Progress’. In: *Journées Francophones des Langages Applicatifs*. JFLA 2023 - 34èmes Journées Francophones des Langages Applicatifs. Praz-sur-Arly, France, 31st Jan. 2023, pp. 7–23. URL: https://hal.science/hal-03886975.

Conferences without proceedings

[18] A. Honorat, T. Bourgoin, H. Miomandre, K. Desnos, D. Menard and J.-F. Nezan. ‘Influence of Dataflow Graph Moldable Parameters on Optimization Criteria’. In: DASIP 2022 - Workshop on Design and Architectures for Signal and Image Processing. Vol. 13425. Lecture Notes in Computer Science. Budapest, Hungary: Springer International Publishing, 20th June 2022, pp. 83–95. DOI: 10.1007/978-3-031-12748-9_7. URL: https://hal.science/hal-03752645.

[19] 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.

Reports & preprints

[20] P. Ciblat, J. Combaz, M. Coupechoux, K. Marquet and A.-C. Orgerie. *Impacts environnementaux de la 5G: Partie 1 : La technologie 5G*. EcoInfo, 11th Oct. 2022, pp. 1–12. URL: https://hal.science/hal-03810501.

[21] G. Fabbretti, I. Lanese and J.-B. Stefani. *Generation of a reversible semantics for Erlang in Maude*. RR-9468. Inria - Research Centre Grenoble – Rhône-Alpes, 5th Apr. 2022, pp. 1–22. URL: https://hal.inria.fr/hal-03630407.

[22] P. Lami, I. Lanese, J.-B. Stefani, C. Sacerdoti Coen and G. Fabbretti. *Reversibility in Erlang: Imperative Constructs -Technical Report*. Inria - Research Centre Grenoble – Rhône-Alpes, 5th July 2022, pp. 1–28. URL: https://hal.archives-ouvertes.fr/hal-03655372.

11.3 Cited publications

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

[24] ARTEMIS Joint Undertaking. *ARTEMIS Strategic Research Agenda*. 2011.

[25] *Automotive Open System Architecture*. 2003. URL: http://www.autosar.org.

[26] 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).

[27] 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).

[28] 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).

[29] S. Borkar. ‘Designing Reliable Systems from Unreliable Components: The Challenges of Transistor Variability and Degradation’. In: *IEEE Micro* 25.6 (2005).

[30] 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.

[31] 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.

[32] R. Davis and A. Burns. ‘A Survey of Hard Real-Time Scheduling for Multiprocessor Systems’. In: *ACM Computing Surveys* 43.4 (2011).

[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).

[34] 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.

[35] 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.

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

[37] 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.

[38] 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.

[39] 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.

[40] 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.

[41] 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.

[42] 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.

[43] T. Henzinger and J. Sifakis. 'The Embedded Systems Design Challenge'. In: *Formal Methods 2006*. Vol. 4085. Lecture Notes in Computer Science. Springer, 2006.

[44] 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.

[45] 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.

[46] I. Lanese and D. Medic. 'A General Approach to Derive Uncontrolled Reversible Semantics'. In: *31st International Conference on Concurrency Theory, CONCUR 2020*. Vol. 171. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020, 33:1–33:24.

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

[48] T. Mari, T. Dang and G. Gössler. 'Explaining Safety Violations in Real-Time Systems'. In: *FORMATS 2021 - Formal Modeling and Analysis of Timed Systems*. Paris, France, Aug. 2021, pp. 100–116. DOI: 10.1007/978-3-030-85037-1_7. URL: https://hal.inria.fr/hal-03348010.

[49] 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.

[50] M. Moore. *Causation and Responsibility*. Oxford, 1999.

[51] 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.

[52] J. Pearl. ‘Causal inference in statistics: An overview’. In: *Statistics Surveys* 3 (2009), pp. 96–146.

[53] S. Plassart. ‘Online optimization in dynamic real-time systems’. Theses. Université Grenoble Alpes [2020-....], June 2020. URL: https://tel.archives-ouvertes.fr/tel-02990646.

[54] S. Ritz, M. Willems and H. Meyr. ‘Scheduling for optimum data memory compaction in block diagram oriented software synthesis’. In: *1995 International Conference on Acoustics, Speech, and Signal Processing, ICASSP ’95, Detroit, Michigan, USA, May 08-12, 1995*. IEEE Computer Society, 1995, pp. 2651–2654. DOI: 10.1109/ICASSP.1995.480106. URL: https://doi.org/10.1109/ICASSP.1995.480106.

[55] J. Rushby. *Partitioning for Safety and Security: Requirements, Mechanisms, and Assurance*. Tech. rep. CR-1999-209347. NASA Langley Research Center, 1999.

[56] F. Yao, A. Demers and S. Shenker. ‘A scheduling model for reduced CPU energy’. In: *Proceedings of lEEE Annual Foundations of Computer Science*. 1995, pp. 374–382.