1. Excellence #@REL-EVA-RE@#

1.1 Quality and pertinence of the project's research and innovation objectives

The Internet of Things (IoT), and more largely reactive programs, are interconnected networks of programming devices that coordinate to achieve complex measurements and actuations. To give an idea of the size of such network, the number of "Things" (such as embedded devices, sensors, and actuators) connected on the "Internet" is expected to reach 41 billion in 2025[1], 5 times more than humans living on the planet. The recent emergence of low cost and powerful micro-controllers (e.g., ESP32 with RISC-V architecture) is another sign that the complexity of IoT networks is increasing. Those devices measure physical quantities (temperature, energy consumption), and also compute decisions that trigger physical actions and change the state of the world (smart cities, smart grid).

The new possibilities that offers the IoT network also comes with new challenges. The complexity to design IoT networks increases with the number of devices, sometimes even exponentially when each node is connected through the network to all other nodes. As a consequence, a local change in one of the nodes can have global implications on the emerging properties of the network. For instance, in the case of an automatic industrial process, machines operate in series on a product and the time and physical effect of each operation needs to be properly calibrated to get the expected result. There comes two main challenges: give a design framework with sufficient expressiveness to specify real time effects of programs (interaction with physical object); and provide formal tools to analyse the individual properties of each node, and the resulting property of the network.

Ideally, the design of an application should be close enough to the programmer's way of thinking, so that few errors are made during the specification. Then, a compiler automatically transforms a program to a binary that a machine can execute. The correctness of a compiler lays in the semantic preservation theorem: the generated binary has the same behaviour as the specification that the programmer wrote. The state of the art for system's language correct compiler mainly consists of two: CompCert[2] and CakeML[3].The semantic preservation theorem for both of those tools is solid, as it has been formally proved in a proof assistant, Coq[4] for the former and Isabelle[5] for the latter. However, none of the two compilers provide real time guarantees in their semantic preservation theorem, which is a critical concern in IoT networks.

The use of powerful categorical structures to model effects of programs on memory has been a very big step in compiler correctness for more advanced imperative languages. The focus, however, has mostly been done on the preservation of memory properties, but not on time properties. My **idea** is therefore to extend the theoretical and practical tools used in compiler correctness to provide timing guarantees in the semantic preservation theorem. This challenge is ambitious as it would provide a new spectrum of certifications for hard real time applications; and is seen as the next big step for current compiler design. I will use some advanced techniques from reactive programming and current knowledge in compiler correctness to extend the CompCert compiler with time guarantees.

I motivate the need for improving the state of the art on formal compiler with real time guarantees by linking to explicit challenges highlighted at European level. European union described in the Next Generation IoT (NGIoT) roadmap key research directions (R) for the coming years (Horizon 2027), from which I cite a few: deploy safe and resilient networks (R7.1), increase privacy by design of IoT networks (R5.2), and lower energetic impacts (R1.2/ R10.1)[6]. I list few reasons why a compiler that provides hard real time guarantees on program's execution will consolidate each of the research direction mentioned above. A network is **safe and resilient** if the timing for the program execution is precise and deterministic. A network has an increased **privacy by design** if each node is programmed in a programming language that supports formal reasoning to prove isolation properties (e.g., that a node cannot read information from another node). A network has a **lower energetic impact** if the time of execution of the running program can be provably optimized.

Aim and objectives of the project.

I detail the general aim and the three main objectives of the project in the next paragraphs, and discuss how such objectives relate and go beyond the state of the art.

The Aim: Advance the state of the art in compiler correctness to provide certifications on the real time guarantees of reactive applications.

To achieve **the aim**, I describe three objectives, one theoretical (O.1), one practical (O.2), and one applicative on a real use case (O.3).

[1] https://digital-strategy.ec.europa.eu/en/policies/next-generation-internet-things
[2] https://compcert.org/
[3] https://cakeml.org/
[4] https://coq.inria.fr/
[5] https://isabelle.in.tum.de/overview.html
[6] https://www.ngiot.eu/wp-content/uploads/sites/73/2020/09/D3.1.pdf

O.1: Develop a compositional and functional semantics that reflects timing effect of reactive programs.

What: The semantics of a programming language usually focuses on the logic of the program. A semantics is *functional* if the meaning of a program is a mathematical function. A functional semantics has many benefits, such as being compositional, i.e., the semantics of a program is the composition of the semantics of its parts; and enjoying powerful analysis methods. However, in case of direct interaction with the physical world (e.g., a time sensitive program), the same program can lead to different results if the physics change: the meaning of a program is, *a priori*, no longer functional. To overcome similar problems, the functional description of effects using categorical structures has been developed. The first objective (O.1) is to define a formal semantics that captures time of execution of a program as an effect. Recent research on semantics for reactive programs has led to notions of computations worth exploring for capturing effects and streaming application in a functional and verifiable way. Time as an effect is a new perspective, in line with recent advancements in functional languages. A formal semantics is therefore the first step to analyse runtime behaviour sensitive to time, and prove **safety and resilient** properties of time sensitive applications.
How: I will develop a categorical semantics that capture the hard-real time of reactive program (WP 2). The use of monads for modelling effects of programs at a semantic level enables powerful static analysis (using theorem provers) on runtime. The recent exploration of co-monads for streaming application is also a step in this direction. The categorical semantic will be defined within a proof assistant (such as Coq) in order to provably reason about timing guarantees of real world applications. I will study how program compositions (sequential, parallel) behave with respect to the composition of their timing effects (WP 2) and develop a type system to statically enforce timing properties.
Measure of success: The first objective will be successful once I formalize the semantics and succeed to prove that two communicating processes satisfy a time-sensitive protocol. Usually two processes communicate by sending signals on a wire with some precise hard real time frequencies. Given a specification of the sender, a specification of the receiver, and a specification of the time-sensitive protocol that rules their interaction, I will provide the proof in Coq that the two programs comply to the specification of the protocol. This result will already be a theoretical and practical contribution.

O.2: Provide certification that timing guarantees of reactive programs are preserved through compilation.

What: Current research, and more particularly at Inria, uses proof assistant (such as Coq) to verify the correctness of compilers for widely adopted programming languages (such as the C programming language). However, timing aspects of reactive programs are explicitly left out in the semantic preservation theorem. The reactive nature (dynamic input/output at runtime), the effect of a program on their physical environment (actuation and sensing), and the interconnection of reactive programs into a network are challenges that require new extensions of current tools to provide timing guarantee on compiled programs. Having a proof of the behaviour of a program at the design level and a compiler that preserves those invariants is key for having **privacy by design**.
How: I will extend the semantics of intermediate languages in CompCert to include hard real time information (WP 3), using the formal semantics defined in WP 1. Currently, the semantics used in CompCert is monadic and captures memory effects. I will extend the monadic semantics with structures that captures execution time compositionally. Next, I will define refinement techniques to provably certify the safety of runtime effects of reactive programs within their environment and architecture and extract a runnable binary with safety properties (WP 3). The certification ensures that the compiled binaries, running on the targeted architecture, satisfies the specification of an interaction protocol. The existence of tool-chains using CompCert demonstrates the feasibility of such approach.
Measure of success: The second objective will be successful once I can provably extend the semantic preservation theorem of CompCert for time sensitive applications.

O.3: Generate the first certified reference implementation for the Universal Asynchronous Receiver Transmitter protocol.

What: The last objective (O.3) consists in applying the methodology to verify an implementation of the Universal Asynchronous Receiver Transmitter (UART), a widely used protocol in hardware communication. The protocol is used to communicate between two components that do not share the same clock (hence, asynchronous). The UART protocol ensures that the receiver and transmitters agree on a shared frequency and can proceed to exchange data. The difficulty in verifying that two programs (a receiver and a transmitter) implement the UART protocol comes from the time-sensitive nature of the protocols. Moreover, the application on a real case is also a good opportunity to

compare the efficiency of two programs. A program that does fewer steps can be preferred to a program that does more steps as it may **reduce the energetic impact**.
How: The formalization of a programming language with timing effects in Coq provides the basis for specifying the UART receiver and transmitter formally (WP 4), and proving that their composition implements the specification of the protocol (WP 4).
Measure of success: The third objective will lead to a reference implementation that runs on a dedicated architecture (xtensa for ESP32 chips). The outcome is a proof of concept that formally captures time as an effect, and enables proof that the time sensitive communication between a receiver and of a transmitter satisfies the UART protocol.

The TEA (Time, Event, Architecture) team of the Inria Centre at Rennes University has expertise in leveraging formal methods to program engineering, with recent research publications in top conferences on end-to-end verified programming of operating system services using the theorem prover Coq, in the context of Inria Challenge RIOT-fp[7]. The advanced algebraic concepts (co-induction, formal cyber-physical coordination) developed during my PhD will be valuable for the TEA projects that are oriented towards proving safety of interactive reactive systems.

Beyond state of the art.
I split the state of the art under three categories, each dedicated to improve safety, privacy by design, and efficiency of time sensitive application.

Formal compiler and correctness. The field of formal verification mainly considers the logical effects of programs. Compilers such as CompCert or CakeML are a formalization of a C compiler and ML compiler respectively. The semantic preservation theorem of CompCert and CakeML shows that a functional description can be compiled to an imperative program while preserving the input/output behaviour. The two compilers do not yet deliver certificates about time guarantees. As CompCert explicitly says in its documentation, the semantic preservation theorem ensures that observable behaviours of the source and target programs are the same, and define observable behaviour as "*everything the user of the program, or the physical world in which it executes, can "see" about the actions of the program, with the notable exception of execution time and memory consumption.*". The project I propose goes beyond the state of the art by extending the CompCert semantic preservation theorem by making time an observable behaviour. Several works have studied extensions of the CompCert semantic preservation theorem for constant time complexity preservation[8] and worst case analysis[9]. At the operating system level, a major work of certification has been done in the seL4 project. However, "*what seL4 cannot (yet) do, and no other OS can either, is to provide temporal isolation guarantees."*[10]. They continue with "*we have not yet developed the formal framework for reasoning about timing guarantees on top of the Mixed Critical System model*" to finish with "*the world's emerging cyberphysical systems need more.*". The project I propose is ambitious and would be a first step towards certification of time guarantees on critical systems. The formalisation of time as an effect, and the use of categorical framework can be the "formal framework" that will be used as a standard for other formalisations for compositional verification. Frameworks used by industrials such as Frama-C or KataOS (Google) would also benefit from the project.

Space and timing guarantees. Most proof systems focus on space guarantees, i.e., safe allocation of memory at runtime. Languages such as Rust are popular since they enforce at the design level some properties about memories. The type theory underlying the type system of Rust is a recent work that focuses on memory ownerships and management. Logics such as separation logic[11] is one instance of a logic that provides primitives to prove memory isolation and heap invariants. The project I propose deals with time properties instead. To reason about time at the type system level is a new challenge that would also benefit the ecosystem of programming languages. Until now, time is essentially abstracted as a logical step in synchronous languages such as LUSTRE[12]. However, the presence of physical components interacting with a program renders logical time inadequate for analysis. Instead, safety for cyber-physical systems requires precise physical time[13]. The formal semantics developed in this project will provide an intermediate representation for more abstract descriptions such as in Platzer with its formal framework for analysis of cyber-physical systems[14], Naijun Zhang with the formalization of hybrid systems[15], and Edward Lee with an actor-

[7] https://future-proof-iot.github.io/RIOT-fp/publications
[8] Gilles Barthe et al, *Formal Verification of a Constant-Time Preserving C Compiler*
[9] André Oliveira Maroneze, *Certified Compilation and Worst-Case Execution Time Estimation.*
[10] https://sel4.systems/About/more-research.pml
[11] Peter W. O'Hearn, *Resources, concurrency, and local reasoning*
[12] https://www-verimag.imag.fr/The-Lustre-Programming-Language-and
[13] Lion, Arbab, Talcott, *A semantic model for interacting cyber-physical systems*
[14] André Platzer: *Logical Foundations of Cyber-Physical Systems.*
[15] Xu, Talpin, Wang, Zhan, Zhan. *Semantics Foundation for Cyber-Physical Systems Using Higher-Order UTP.*

based implementation of hybrid systems[16]. The quantum physicist Nicolas Gisin also witnesses that using computational model for explaining physics is a region rich of theoretical and practical questions[17], fertile for discoveries. The project goes beyond the state of the art as it provides a formal framework to analyse cyber-physical systems with real time guarantees in Coq, and gives certification that the properties are preserved by compilation.

Embedded systems. Instead of a general theory, analysis on the worst case execution time have been performed on annotated programs, and CompCert has been used to preserve the real time guarantees of such annotated programs. Research in the field has lead to the AbsInt company[18], that develops techniques to analyse the worst case execution time (WCET) of dedicated software. However, as the director and lead researcher said in a recent article[19], the analysis of the worst-case execution time is not compositional: given a composition A;B, the execution time of instruction B may depend on the state of the architecture after execution of instruction A. I propose to develop a compositional model of execution time at design level using similar advanced mathematics as those developed for compositional reasoning on memory manipulation. The results will be formalized and goes beyond the state of the art as it offers new avenues for compositional verification. The application on the certification that a receiver and a transmitter follow the UART protocol is also a contribution that offers new avenue to verify other time sensitive protocols, such as Precise Time Protocol.

1.2 Soundness of the proposed methodology.

Overall methodology:

The proposed research naturally decomposes into both theoretical and practical components. The research methodology is to interleave both developments, by using practical tools to implement theoretical results, prove their soundness, and demonstrate their applicability. As a result, the theory will be restricted to one that can be implemented and constructively verified, and the use of theorem prover and process extraction will confirm and demonstrate the validity and applicability of the theory.

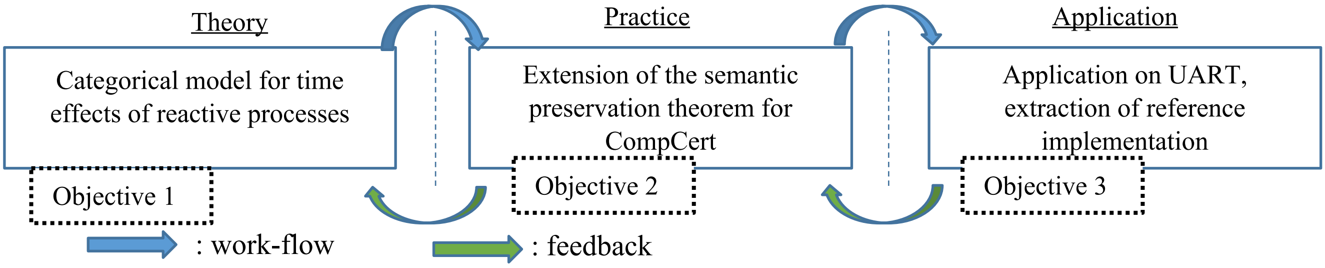
Theory
Categorical model for time effects of reactive processes
Objective 1
Practice
Extension of the semantic preservation theorem for CompCert
Objective 2
Application
Application on UART, extraction of reference implementation
Objective 3
: work-flow
: feedback

Figure 1: The methodology as an alternation of theoretical developments, practical formalization, and application on a critical case.

As shown in Figure 1 with the blue arrows, the theoretical framework of objective 1 will be formalized in the Coq proof assistant and used as input for objective 2, and then applied on extracting a reference implementation for the UART protocol. Additionally, the insights that the use of the formalization gives on the theory and the need that the application provides on the formalization will be feedback to update the theoretical framework and the formalization (green arrows). The overall methodology is reflected in the Gantt diagram, as an alternation of two blocks, each decomposed into a theoretical, practical, and applicative component. The alternation ensures that feedback loops (green arrows) appear in due time. Each objective has its own methodology, detailed hereafter.

Methodology towards Objective 1. I will start from a series of theoretical works that model effects in functional programming languages[20]. The use of monads to capture effects of programs is widely accepted in the programming community, and the use of co-monad to capture context-dependency is theoretically elegant[21], but not yet widely adopted in practice. The proof of compositionality for a semantics combining both effects (as a monad) and context-dependency (as a co-monad) will be the starting point of the theoretical framework. The research will extend the compositionality result to enlarge the semantic model in order to model time sensitive programming constructs. The methodology for objective 1 will be to chose a monad and a co-monad that capture the time and I/O effects of a reactive program, and prove that such semantics is compositional (WP2). The theory will be formalized in the Coq proof assistant. I will collaborate with David Nowak from X2S team at CNRS Lille on the formalization of the theory, as David has strong mathematical background in category theory and in using the Coq proof assistant.

[16] Lohstroh, Menard, Bateni, A. Lee: *Toward a Lingua Franca for Deterministic Concurrent Systems.*
[17] Gisin, Nicolas. *Indeterminism in Physics, Classical Chaos and Bohmian Mechanics: Are Real Numbers Really Real?*
[18] https://www.absint.com/ait/
[19] https://cacm.acm.org/magazines/2020/10/247596-real-time-spent-on-real-time/fulltext#R6
[20] Eugenio Moggi, *Notion of computation and monads*
[21] Tarmo Uutsalu, Varmo Vene, *Comonadic notions of computation*

Methodology towards Objective 2. I will start from the Coq formalization of the time sensitive semantics developed in WP1 to give a time semantics of an intermediate language used in the CompCert compiler. I will work in close collaboration with Frédéric Besson, in the Epicure team in Rennes. Frédéric has extensive knowledge on the CompCert compiler, as he already provided several extensions. Weekly meeting with Frédéric will be the opportunity to discuss details on the extension of the formal semantics of CompCert with timing information. I will then prove an extension of the semantic preservation theorem for the RISC-V processor architecture, as it is broadly supported by recent processors.

Methodology towards Objective 3. I will start from a well known protocol that is used in most embedded system to initially flash a program. An embedded system usually contains a micro-controller that boots, when turned on, with a small program called bootloader. A micro-controller, if connected with a serial link to a machine, may have its code updated. The update is usually performed through a protocol, called UART. In the UART protocol, a program runs on both ends of the serial link, i.e., one receiver program on the micro-controller, and one transmitter program on the machine that writes to the controller. The protocol ensures that the two machines agree with a clock signal, and the transmitter then starts to send signals to the receiver. The receiver decodes the signals and writes the corresponding bits to its memory. The protocol goes on until the transmitter sends the terminating sequence. For objective 3, I will consider an implementation in C of the receiver and transmitter, and formally verify that their composition satisfies the UART protocol, namely that the receiver decodes precisely the messages that the transmitter sent. To achieve this goal, I will use WP2 as a formal framework for the receiver and transmitter programs, prove the properties specified by the UART protocol, and generate a binary with the certified compiler of WP3. To the best of my knowledge, this application will be the first verified implementation of the UART protocol.

Integration of method and discipline:

This project combines three disciplines: programming engineering, applied mathematics for computer science, and physics. For all work packages in each of those disciplines, experts will act as advisor to support and collaborate towards the success of the project.

Objective 1 is theoretical and the work will be supported by the expertise of David Nowak in 2XS team at CNRS Lille. The integration of abstract and powerful mathematical ideas to understand and analyse programs is at the core of work packages in Objective 1. Weekly meeting are currently set up to discuss ideas with David.

Objective 2 is practical and the aim is the extension of the semantic preservation theorem of CompCert, and will be performed in collaboration with Frédéric Besson in Epicure Team at the Inria Centre at Rennes University. Bridging the gap between theoretical model and practical tools is key to improve current software design, which is what Objective 2 fulfils. The consideration of a system language such as the C language makes the solution useful for programming engineers as well.

Objective 3 is an application of the theoretical and practical results on a UART, a widely used protocol, with the expertise of Jean-Pierre Talpin in the TEA team at the Inria Centre at Rennes University. The TEA team is specialised in programming engineering and application to embedded systems.

My research background focuses on the interaction between digital and physical systems, and fits therefore the interdisciplinary spectrum of the project.

Gender dimension

No gender dimension is foreseen, as the project deals purely with physical data and its mathematical treatment.

Open Science practices and research data management and management of research outputs

The project promotes open science practices by choosing for publication open access conferences and journals, and by using appropriate tools for reusable and accessible implementation.

Results of the projects will be Findable, as I will use Inria's gitlab for continuous development of the Coq framework, and zenodo to periodically and permanently store milestones achievements. The gitlab containing the source code repository will also be uploaded on the software heritage platform for archiving. Theoretical results will be hosted in the open access Hyper Articla en Ligne (HAL) library, and conferences and journals will be chosen for their open access policy. The results will be accessible through the team's webpage and my own webpage. More particularly, a strong importance will be given to make the library interoperable, by choosing the appropriate licenses (e.g., GNU LGPL the same as for the Coq proof assistant), and integrate the library within the standard library for Coq. Finally, the code will be published on the gitlab of Inria with a detailed explanation on how to install the tool in order to make it reusable.

1.3 Quality of the supervision and of the two-way transfer of knowledge between the researcher and the host

Qualifications and experience of the supervisors

The supervisor at Inria will be Jean-Pierre Talpin who has been team leader of the TEA team for 25 years. Jean-Pierre has an internationally established network with teams in Germany, China, USA. Jean-Pierre has been in the program committee of more than 75 international conferences (e.g., EMSOFT, FACS, MEMOCODE). Jean-Pierre TALPIN actively participates, as a key member, in the RIOT-FP (Future-proof IoT, ~200k€) project, in collaboration with other Inria teams. The TEA team collaborates with the Institute for Software of the Chinese Academy of Science (ISCAS) at Beijing on Compositional Verification. Moreover, the TEA team has a strong collaboration with Mitsubishi Electric Europe (MERCE) on contract for cyber-physical systems. The proposed objective in this project are also answering challenges in Mitsubishi, such as the certification of cyber-physical properties (e.g., time guarantees) on critical systems.

The main collaboration on the first objective will be with David Nowak, a permanent researcher at 2XS team in CNRS Lille. David has a strong expertise in theoretical computer science, and applied category theory. Moreover, he is proficient with the Coq proof system, as shown by his latest works to prove properties of an EDF (Earliest Deadline First) scheduler in Coq[22]. David contributed to the development of the dx tool[23] and used CompCert to generate a provably correct implementation of the EDF scheduler.

The main collaboration on the second objective will be with Frédéric Besson, a permanent research at the Epicure team at Inria. Frédéric is expert in CompCert as shown by his latest work on a proof for memory isolation of an rBPF virtual machine[24].

Further collaboration with CWI will be developed during the project, and more precisely with a young researcher Hans-Dieter Hiep. Hans-Dieter will defend his PhD end of 2023 and is already on board for long term collaborations.

Planned training activities for the researcher

I am very proponent of the learning-by-doing approach. I will use the basic knowledge I have in Coq[25] and its underlying type system to learn advanced techniques and gain expertise. I will provide an introductory course for using the Coq proof assistant at Inria, and I will follow the advanced topics in the online software foundation courses.

With the project, I will become expert in applied category for computer science, and therefore contribute to solve current challenges in the field of functional programming. On an educational topic, Inria has a strong link with universities, and I intend to teach an introductory lecture on the basis of category theory. My teaching ambition is to share exciting topics with young and promising students. Also, teaching advanced mathematics is an opportunity to share advanced problems, collaborate with students on research questions, and eventually motivate students to do research in theoretical computer science.

The TEA team has developed a strong collaboration with MERCE, a Mitsubishi research centre at Rennes. Monthly meetings are taking place to discuss possible collaborations with the team. I will train on sharing my research by presenting the work and progress at least twice a year in those meetings.

The career of a researcher is not only about his or her strong technical skills, but also the ability to access funds and broadcast its results. To that end, I will apply to the workshops given by Inria about the different licensing possibilities for open access, free and open source software. On similar topics, I intend to gain knowledge in IP property by following the workshop provided by Inria. Moreover, I will follow the workshops given by Inria to get familiar with the submission of an ERC *starting grant* and national funding for young researchers (such as ANR JCJC). Then, I will be ready, at the end of the MSCA scholarship, to submit a proposal and become a recognized expert in the domain of compiler correctness for time certification. Finally, the technical and soft skills (e.g., for communication) trained during the fellowship will be of great value for applying to a permanent research position at Inria or CNRS in France.

Two ways transfer of knowledge between the researcher and host organisation.

Inria hosts research teams actively involved in the development of state of the art tools in formal methods, and, more specifically, on certified compilation. The transfer of knowledge through the collaboration with the Epicure team, and Frédéric especially, is fundamental for the success of Objective 2. Moreover, the many collaborations that the TEA team has with, for instance, the 2XS team at CNRS Lille allows me, in collaboration with David Nowak, to quickly get the knowledge about advanced mathematical techniques to fulfill Objective 1. On that respect, the collaboration has already started, and I share and discuss progress with David every week.

During my PhD, I studied formal models and practical applications to design *correct by construction* communication protocols. The advanced algebraic concepts (co-inductive data types, formal models for cyber-physical coordination) developed during my PhD are valuable in the environment of the TEA team, in order to abstract real programming

[22] F. Vanhems, V. Rusu, D. Nowak, G. Grimaud, *A Formal Correctness Proof for an EDF Scheduler Implementation*

[23] https://gitlab.univ-lille.fr/samuel.hym/dx

[24] S Yuan, F Besson, JP Talpin, et al, *End-to-end Mechanized Proof of an eBPF Virtual Machine for Micro-controllers*

[25] I read part of *Type Theory and Formal Proof*, and followed the Volume 1 of https://softwarefoundations.cis.upenn.edu/

challenges into mathematical description. Developing a semantics that captures real time as an effect in a programming language requires an advanced algebraic model, in the same line as I worked on during my PhD. The project will allow me to apply powerful theoretical results to practical settings within the Coq proof assistant. I also plan to contribute to the regional ecosystem by teaching at university in Rennes.
Finally, I join the TEA team with an active network of Dutch computer scientists, as I did my PhD thesis at CWI in Amsterdam. I intend to consolidate a collaboration with the Computer Security group at CWI in Amsterdam, with Farhad Arbab and Hans-Dieter Hiep, and with Marcello Bonsangue at Leiden University. The collaboration consists of two visits per year, for discussing results, and of monthly online meetings for presentation of research.

1.4 Quality and appropriateness of the researcher's professional experience, competences and skills

I defended my PhD on an algebra for interaction of cyber-physical components. I showed in my PhD how interaction between cyber and physical systems can be modelled explicitly in an algebraic operator. The approach taken in the thesis has both a theoretical algebraic dimension and a practical framework to simulate and verify properties of those systems in Maude, a functional language developed at Stanford Research Institute. The research has lead to 3 journal papers in the Journal for Logic and Algebraic Methods for Programming (JLAMP) and papers in selective conferences. Moreover, I applied and received a three months grant at the East China Normal University (ECNU) at Shanghai, in the group of Min Zhang from February to April 2023, in order to collaborate on an extension of the results of my PhD thesis.

2. Impact #@IMP-ACT-IA@#

2.1 Credibility of the career enhancement measures and their contribution to his skills development

Skill development.
The Coq proof assistant is a state of the art tool with solid mathematical foundations based on dependent type theory. The formalization of mathematical theories in Coq is a crucial step in the verification and validation of a proof. I want to provide a theoretical library for categorical semantics, formalized in the Coq proof assistant, to reason about reactive programs. This first experience will be valuable for future formalization of theories in other fields of mathematics. The gap to master such tool is high, and the skill is attractive for collaboration in other research projects, as I will be able, using the Coq proof assistant, to provide certifications that the theory developed with my collaborators is correct. Inria is pioneer in this domain, and currently leading the development of the Coq tool. The TEA team of the Inria Centre at Rennes University is at the forefront of the research to use the Coq proof assistant for end-to-end verification of exokernel services and reactive systems.

The CompCert compiler has received one of the most prestigious ACM awards in computer science in the past two years. The development and improvement of CompCert is a very active field of research. Through the project, I will acquire the knowledge to interface CompCert in order to increase the class of programs that can be certified. CompCert is used in the industry, for critical systems, where the need to preserve properties through compilation is crucial. The knowledge that I will acquire in the project, and the extension that I will build on CompCert will be of great academic and industrial value, for possibly building other extensions.

Category theory is a prolific research area applied to abstract and formalize programming challenges. For instance, the gap between functional and imperative programming is often bridged with categorical structures, such as monads. The representation of time as an effect, and the categorical semantics is a new perspective that requires new knowledge. The project will give me a strong background in applied category theory, which is a skill required in both fundamental research and programming engineering.

Expected impact on the research career.
The project I propose focuses on a formal and compositional semantics that captures execution time as an effect. This project is a unique opportunity to build an expertise in the intersection of theorem proving and cyber-physical systems. The effects of physical impacts of programs are increasingly important in today's digital infrastructure (energy, time consumption) and mechanisms to formally quantify the physical effects are still lacking. Results of the research project can lead to breakthrough that span new avenues to long term research projects, such as the development of formal languages with timing supports, and verification of compilers for timing or energy usage guarantees. The project will give me the opportunity to master and contribute to the development of the Coq proof assistant (possibly, as a next step, to extend the underlying type system with resource management) and the CompCert certified compiler. I will also take the opportunity given by the grant to start applying to a fixed position as a permanent researcher in Inria Centre at Rennes University, as shown in the Gantt diagram. Moreover, I will follow a workshop for writing a starting grant ERC proposal, and organize some time to already start a draft during year two of the project. The results and discoveries of the project will strengthen the skills required for an ambitious five years

research plan to include more physical effects into the semantics of system languages, and eventually extend the type system of Coq to include cyber-physical effects.

2.2 Suitability and quality of the measures to maximise expected outcomes and impacts #@COM-DIS-VIS-CDV@#

Dissemination, communication, and exploitation activities.

Dissemination. The theoretical results will lead to four high quality papers, that will be peer reviewed and presented at top conferences on programming and verification, e.g., POPL, PLDI, OOPSLA, EMSOFT, FM, CAV. Simultaneously, the project leads to artefacts that will be made open source, and provided as a Coq library. I also plan to present my results at the Coq workshop (https://coq-workshop.gitlab.io/). As the TEA team at Inria is also in contact with industrial partners (MERCE Mitsubishi), I will present the results over the years and motivate my approach with the certification of a real world use cases, illustrated by the UART protocol.

Communication. The project also focuses on leveraging formal methods to a wider audience. The library will be accompanied by a tutorial and a detailed procedure to start using the tool. The project will lead to a presentation to the general audience of the Inria Centre at Rennes University. Finally, I will stay in contact with Dutch research institutes, such as Leiden University and CWI. Common interests have been created with the researchers at CWI. In that sense, close collaboration with the CWI for dissemination of research results will be ensured by annual visits and joint research works. More generally, I have submitted a proposal for a monthly broadcast emission on local radio in Rennes to discuss the research subject of young researchers. This will be an opportunity to bridge the gap with general audience and specialized researchers. I have myself a personal website (https://benjaminlion.fr/) that I will update regularly with practical and theoretical news of the project, together with the team website. I also plan to present the result to the seminar Sci-Rennes Inria seminar in Rennes. This seminar is meant for general scientific audience to present ongoing research projects.

Exploitation. The theoretical results of the project will be published and advertised in the theoretical computer science community. Since the theory will be formalized in Coq, the research community can directly use the library and extend the framework to more expressive language constructs. As a result, this project can lead to a new research avenue about real time guarantees programs through certified compilation. The practical results, such as the extension of the CompCert semantic preservation theorem, will also be reusable by the research community to improve on the class of real time properties that can be certified. Moreover, this is a step towards leveraging formal method tools to time critical software in industry and critical societal infrastructure. The open access of the library makes the result reusable, and favours the construction of an open source community.

IP management. The code written during the project will be owned by Inria and will be shared through open source licences. If any result is created jointly with David Nowak from CNRS, in the first work package, and it is not possible to distinguish between the contribution of each contributor, such work will be jointly owned by Inria and CNRS. Details concerning jointly owned results, joint inventions, and joint patent applications will be addressed in a separate agreement.

2.3. The magnitude and importance of the project's contribution to the expected scientific, societal and economic impacts

Scientific impact. In the past years, model of computations have switched from being purely functional (such as lambda calculus) to introducing effects (e.g., with monads) and modeling more real world programs. The theoretical question of "what can a computer *compute*?" can be reexamined in light of new physical interactions that a computer accesses through its sensors and actuators. Mostly, this project focuses on the time effect of a program, and extends current model of computation with a notion of time. As a consequence, a future step would consider an extension of the underlying calculus of inductive construction of Coq to include data structures that are time sensitive. In that sense, dependent types that are now first class concepts in Coq could also be extended with physically aware data types, that change over time. As an implication, several theories for cyber-physical systems that lack formal support in proof systems (because of being too complex) could now be expressed using a time sensitive type system.

Industrial impact. The use of the UART protocol is mainstream in industrial application, to flash a micro-controller. The reference implementation that I deliver in objective 3 will be the first implementation that provably comply to the UART protocol. The same methodology or the use of the reference implementation can therefore provide guarantees for critical industrial use cases. More generally, the extension of CompCert with time guarantees in the semantic preservation theorem can also benefit industries that use CompCert as a compiler. New certifications can be delivered for systems with hard real time constraints. The project therefore leverages mathematical tools and formal verification to enable programming engineers to increase safety, reliability, and predictability of sensitive reactive application. As a sign of the impact of such tool, the market volume in IoT is expected to double by 2028 and reach US$602.00bn[26], which includes new critical industrial and societal applications.

[26] https://www.statista.com/outlook/tmo/internet-of-things/europe

Societal impact. Finally, the project contributes **to larger societal impacts**. A program needs energy during its execution, and minimizing the energy consumption is a current societal challenge (computers are responsible of 2% of global emissions[27]). By providing a formal description of timing effects of programs, the energy consumption becomes apparent as related to the time a program takes to run. The certification framework is therefore an important step towards quantifying the societal contribution of computation to, for instance, energy consumption. More generally, the societal deployment of digital systems with physical actuation requires certifications that those systems are safe. The consequence of this project is to provide the formal first steps to understand and reduce the risks of digital systems with cyber-physical effects. #§COM-DIS-VIS-CDV§# #§IMP-ACT-IA§#

3. Quality and Efficiency of the Implementation #@QUA-LIT-QL@# #@WRK-PLA-WP@# #@CON-SOR-CS@# #@PRJ-MGT-PM@#

3.1 Quality, effectiveness of the work plan, risks, and appropriateness of the effort assigned to WPs

The work plan is divided into five work packages (WP) with three technical (WP2, WP3, and WP4), and two non-technical (WP1 and WP5). The two work packages WP2 and WP3 are respectively the theoretical and the practical ground works required to extract a reference implementation for the UART protocol in WP4.

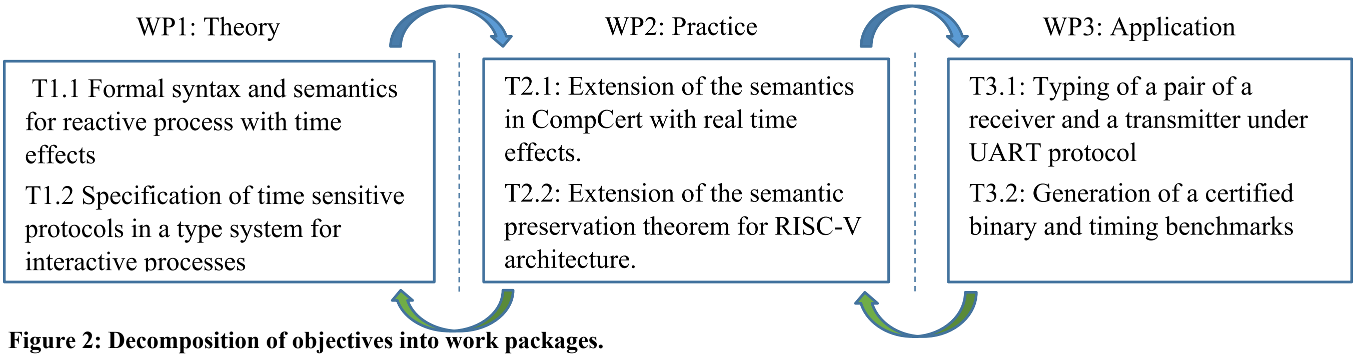

WP1: Theory

T1.1 Formal syntax and semantics for reactive process with time effects

T1.2 Specification of time sensitive protocols in a type system for interactive processes

WP2: Practice

T2.1: Extension of the semantics in CompCert with real time effects.

T2.2: Extension of the semantic preservation theorem for RISC-V architecture.

WP3: Application

T3.1: Typing of a pair of a receiver and a transmitter under UART protocol

T3.2: Generation of a certified binary and timing benchmarks

Figure 2: Decomposition of objectives into work packages.

<u>**WP 1** (Management and career planning - *1pers. month*):</u> The two deliverable D1.1 and D1.2 respectively correspond to the management and career plan and constitute the first milestone M1. The work package takes the duration of the whole project, with weekly meeting with my main supervisor, and weekly meeting with the collaborator of each work package. Assuming the project starts in May, the diagonal hashed period (▨) is a two months period in which I will revise an application for a fixed position as a researcher at Inria or CNRS in France. In the second year, I will follow an Inria workshop for ERC writing (🟩). The horizontal hashed period (▤) is a three months period in which I will write an ERC *starting grant* proposal that generalizes the results of this project to include and certify more of the physical effects of computation.

<u>**WP 2** (Theory):</u> Develop a formal and compositional semantics that capture time of execution as an effects.

- *Task 2.1 (Process) - (5 pers. months):* give a compositional semantics of a subset of a system language (e.g., subset of the C language) in order to capture time as an effect. I will use categorical constructs, such as monads and co-monads, to define functions with effects and their compositions. I will deliver a formalization in Coq (D2.1) of the theory as an open access library.

- *Task 2.2 (Types) - (4 pers. months)*: formalize time properties as a type systems for the subset of the programming language defined in Task 1.1. The type system is implemented in Coq and delivered as a library (D2.2).

<u>**WP 3** (Practice) :</u> Extension of the verified C compiler CompCert with timing guarantees.

- *Task 3.1 (Extension) - (4 pers. months)*: Extension of the semantics of intermediate language in CompCert to include hard real time information (WP 2.1), using the formal semantics defined in Task 1.1. Next, The certification ensures that the compiled binaries, running on the targeted architecture, satisfies the specification of an interaction protocol. The existence of tool-chains using CompCert demonstrates the feasibility of such approach. This task will lead to an extension of the semantic preservation theorem of CompCert as a library (D3.1).

- *Task 3.2 (Certification) - (4 pers. months)*: I will define refinement techniques to provably certify the safety of runtime effects of reactive programs within their environment and architecture, and extract a runnable binary with safety properties (WP 2). This task will lead to a documentation that explains each refinement techniques (D3.2).

<u>**WP 4** (Application):</u> Certification of the UART protocol.

- *Task 4.1 (Specification) - (2 pers. months):* Specify, within the formal framework of T1.1, a receiver and a

[27] https://circularcomputing.com/news/carbon-footprint-laptop/

transmitter for the UART protocol; and specify the protocol as a logical proposition on the observable behaviors.

- *Task 4.2 (Extraction) - (2 pers. months):* Certify that the composition of the two processes defined in T3.1 satisfies the UART protocol, and extract a reference implementation for the receiver and transmitter, using dx and CompCert. This task will lead to a verified binary for the receiver and transmitter of the UART protocol (D3.1).

WP 5 (Dissemination, Exploitation, Communication - *2pers. month*): targeted conferences (●) for submission of theoretical results (around D2.1 and D2.2), and practical and applied results (around D3.1 and D4.1). The two milestones M2 and M3 correspond to the writing, as a conference paper, of the theoretical and practical results. Such publication demonstrates the theoretical soundness of the approach. In parallel, I intend to communicate my work at the Interactive Theorem Prover workshop and seminars at CWI and Inria (▲). Finally, the deliverable from WP2, WP3, and WP4 will be followed with blog posts on my personal website and on the TEA team website, that explain the results obtained and show how to reproduce them for exploitation. I plan a three months visit in the 2XS team of David Nowak, in month 13, 14, and 15, for dissemination of the results, and direct collaboration on the next steps.

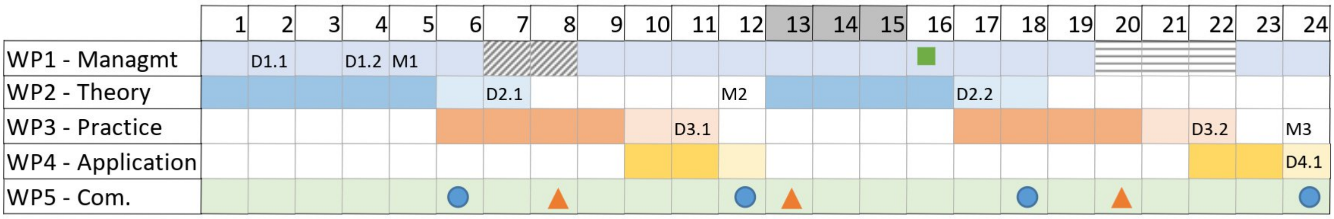

Risk description, (P)robability of occurrence, (I)mpact	Mitigation strategy
Theory is slower to develop than foreseen (WP2, Task 2.1) P: moderate, I: low	Time buffers (in light colors) are allocated at the end of each task to implement the feedback loops as displayed in green in Figure 2.
Formalization in Coq is harder than foreseen due to constraints in the tool (WP2, Tasks 2.1 and 2.2) P: moderate, I: low	The second task of WP2 can be simplified to tailor the use case example of the UART protocol only, instead of the general type system I envision. This way, some of the months planned for Task 2.2 can be allocated to Task 2.1.
The CompCert semantics is difficult to extend with time effect (WP3, Tasks 3.1 and 3.2) P: moderate, I: high	The two months buffer time allocated at the end of each task makes it possible to mitigate this risk by continuous development of the semantics and the integration in CompCert. If this is still not enough, Task 2.2 of WP2 will be simplified to tailor to the application of the UART protocol only. It would cost a bit on the generality of the theory, but still achieve all objectives.
Collaborations with David or Frédéric cannot be continued, P: moderate, I: medium	The collaboration with David and Frédéric speeds up the learning phase. If the collaboration with David ends, I already have some background due to the current collaboration to continue alone. If the collaboration with Frédéric stops, I will get in touch with the Epicure team and find new collaborations to work on the CompCert compiler.
Difficulty in applying the framework to the UART protocol. P: low, I: medium	If the underlying physical layers cannot be formalized, we will simplify the physical model to still be able to provably certify the timing properties.

3.2 Quality and capacity of the host institutions and participating organisations, hosting arrangements

Besides scientific excellence, Inria provides significant support for the researchers . In particular, the researchers can count on the Operational Committee for the Evaluation of Legal and Ethical Risks (OCELER) to deal with ethics and the legal evaluation of their projects. The Innovation Transfer and Partnerships Department supports the researchers in exploitation-related duties. The experimentation and development department provides assistance in software development. The scientific information and publishing department helps the researchers to disseminate their results as well as open science related tasks. The communication department supports communication and dissemination activities, and the HR department supports the researchers in their administrative tasks. In May 2019, Inria obtained the Human Resources Excellence in Research from the European Commission (HRS4R). Gender equality is also strictly respected through the “Committee on Gender Equality and Equal Opportunities. *The* Inria *Centre at Rennes University has strong collaborations with academic and research institutions in Rennes, for instance,* Rennes University, University of Rennes 2, CentraleSupelec, INSA Rennes and ENS Rennes. It has also a branch in Nantes, which is developing alongside the Nantes University. These agreements offer the Inria researchers the opportunity to assist in courses within these universities as well as collaborating with external researchers and scientists. #§CON-SOR-CS§# #§PRJ-MGT-PM§#