



# Activity Report 2022

## Team TEA

Time, Events and Architectures

*Joint team with Centre Inria de l'Université de Rennes*

D4 – Language and Software Engineering





# Contents

<b>Project-Team TEA</b>	<b>1</b>
<b>1 Team members, visitors, external collaborators</b>	<b>2</b>
<b>2 Overall objectives</b>	<b>2</b>
2.1 Introduction . . . . .	2
2.2 Context . . . . .	2
2.3 Motivations . . . . .	3
2.4 Challenges . . . . .	3
<b>3 Research program</b>	<b>4</b>
3.1 Previous Works . . . . .	4
3.2 Timed Modeling . . . . .	4
3.3 Modeling Architectures . . . . .	5
3.4 Scheduling Theory . . . . .	5
3.5 Verified programming for system design . . . . .	6
<b>4 Application domains</b>	<b>6</b>
<b>5 Highlights of the year</b>	<b>7</b>
<b>6 New software and platforms</b>	<b>7</b>
6.1 New software . . . . .	7
6.1.1 CertFC . . . . .	7
6.1.2 ADFG . . . . .	7
6.1.3 POLYCHRONY . . . . .	8
<b>7 New results</b>	<b>9</b>
7.1 Verified programming and secure integration of operating system libraries in RIOT-fp . . . . .	9
7.2 Semantic foundations for cyber-physical systems using higher-order UTP . . . . .	9
7.3 A logical framework to verify requirements of hybrid system models . . . . .	10
7.4 ADFG: Affine data-flow graphs scheduler synthesis . . . . .	10
<b>8 Bilateral contracts and grants with industry</b>	<b>10</b>
8.1 Bilateral contracts with industry . . . . .	10
8.2 Bilateral grants with industry . . . . .	11
<b>9 Partnerships and cooperations</b>	<b>11</b>
9.1 International initiatives . . . . .	11
9.1.1 Visits to international teams . . . . .	12
9.2 National initiatives . . . . .	12
<b>10 Dissemination</b>	<b>12</b>
10.1 Promoting scientific activities . . . . .	12
10.2 Teaching - Supervision - Juries . . . . .	13
<b>11 Scientific production</b>	<b>13</b>
11.1 Major publications . . . . .	13
11.2 Publications of the year . . . . .	13

## Project-Team TEA

*Creation of the Project-Team: 2015 January 01*

### Keywords

#### Computer sciences and digital sciences

- A1.2.5. – Internet of things
- A1.2.7. – Cyber-physical systems
- A1.5.2. – Communicating systems
- A2.1.1. – Semantics of programming languages
- A2.1.4. – Functional programming
- A2.1.6. – Concurrent programming
- A2.1.9. – Synchronous languages
- A2.1.10. – Domain-specific languages
- A2.2.1. – Static analysis
- A2.2.4. – Parallel architectures
- 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. – Formal method for verification, reliability, certification
- A2.4.1. – Analysis
- A2.4.2. – Model-checking
- A2.4.3. – Proofs
- A2.5. – Software engineering
- A2.5.1. – Software Architecture & Design
- A2.5.2. – Component-based Design
- A4.4. – Security of equipment and software
- A4.5. – Formal methods for security
- A7.2. – Logic in Computer Science
- A7.2.3. – Interactive Theorem Proving
- A7.3. – Calculability and computability
- A8.1. – Discrete mathematics, combinatorics
- A8.3. – Geometry, Topology

#### Other research topics and application domains

- B5.1. – Factory of the future
- B6.1.1. – Software engineering
- B6.4. – Internet of things
- B6.6. – Embedded systems

# 1 Team members, visitors, external collaborators

## Research Scientists

- Jean-Pierre Talpin [Team leader, INRIA, Senior Researcher, HDR]
- Thierry Gautier [Inria, Researcher, until Sep 2022]

## Post-Doctoral Fellow

- Benjamin Lion [Inria, from Dec 2022]

## PhD Students

- Stéphane Kastenbaum [MITSUBISHI ELECTRIC]
- Shenghao Yuan [INRIA]

## Administrative Assistant

- Armelle Mozziconacci [CNRS]

# 2 Overall objectives

## 2.1 Introduction

An embedded architecture is an artifact of heterogeneous constituents and at the crossing of several design viewpoints: software, embedded in hardware, interfaced with the physical world. Time takes different forms when observed from each of these viewpoints: continuous or discrete, event-based or time-triggered. Modeling and programming formalisms that represent software, hardware and physics significantly alter this perception of time. Therefore, time reasoning in system design is usually isolated to a specific design problem: simulation, profiling, performance, scheduling, parallelization, simulation. The aim of project-team TEA is to define conceptually unified frameworks for reasoning on composition and integration in cyber-physical system design, and to put this reasoning to practice by revisiting analysis and synthesis issues in real-time system design with soundness and compositionality gained from formalization.

## 2.2 Context

In the construction of complex systems, information technology (IT) has become a central force of revolutionary changes, driven by the exponential increase of computational power. In the field of telecommunication, IT provides the necessary basis for systems of networked distributed applications. In the field of control engineering, IT provides the necessary basis for embedded control applications. The combination of telecommunication and embedded systems into networked embedded systems opens up a new range of systems, capable of providing more intelligent functionalities, thanks to information and communication (ICT). Networked embedded systems have revolutionized several application domains: energy networks, industrial automation and transport systems.

20th-century science and technology brought us effective methods and tools for designing both computational and physical systems, such as for instance Simulink and Matlab. But the design of cyber-physical systems (CPS) is much more than the union of those two fields. Traditionally, information scientists only have a hazy notion of requirements imposed by the physical environment of computers. Similarly, mechanical, civil, and chemical engineers view computers strictly as devices executing algorithms. CPS design is, to date, mostly executed in this ad-hoc manner, without sound, mathematically grounded, integrative methodology. A new science of CPS design will allow to create machines with complex dynamics and high control reliability, and apply to new industries and applications, such as IoT or edge devices, in a reliable and economically efficient way. Progress requires nothing less than the

construction of a new science and technology foundation for CPS that is simultaneously physical and computational.

### 2.3 Motivations

Beyond the buzzword, a CPS is a ubiquitous object of our everyday life. CPSs have evolved from individual independent units (e.g. an ABS brake) to more and more integrated networks of units, which may be aggregated into larger components or sub-systems. For example, a transportation monitoring network aggregates monitored stations and trains through a large scale distributed system with relatively high latency. Each individual train is being controlled by a train control network, each car in the train has its own real-time bus to control embedded devices. More and more, CPSs are mixing real-time low latency technology with higher latency distributed computing technology.

A common feature found in CPSs is the ever presence of concurrency and parallelism in models. Large systems are increasingly mixing both types of concurrency. They are structured hierarchically and comprise multiple synchronous devices connected by buses or networks that communicate asynchronously. This led to the advent of so-called GALS (Globally Asynchronous, Locally Synchronous) models, or PALS (Physically Asynchronous, Logically Synchronous) systems, where reactive synchronous objects are communicating asynchronously. Still, these infrastructures, together with their programming models, share some fundamental concerns: parallelism and concurrency synchronization, determinism and functional correctness, scheduling optimality and calculation time predictability.

Additionally, CPSs monitor and control real-world processes, the dynamics of which are usually governed by physical laws. These laws are expressed by physicists as mathematical equations and formulas. Discrete CPS models cannot ignore these dynamics, but whereas the equations express the continuous behavior usually using real (irrational) variables, the models usually have to work with discrete time and approximate floating point variables.

### 2.4 Challenges

A cyber-physical, or reactive, or embedded system is the integration of heterogeneous components originating from several design viewpoints: reactive software, some of which is embedded in hardware, interfaced with the physical environment through mechanical parts. Time takes different forms when observed from each of these viewpoints: it is discrete and event-based in software, discrete and time-triggered in hardware, continuous in mechanics or physics. Design of CPS often benefits from concepts of multiform and logical time(s) for their natural description. High-level formalisms used to model software, hardware and physics additionally alter this perception of time quite significantly.

In model-based system design, time is usually abstracted to serve the purpose of one of many design tasks: verification, simulation, profiling, performance analysis, scheduling analysis, parallelization, distribution, or virtual prototyping. For example in non-real-time commodity software, timing abstraction such as number of instructions and algorithmic complexity is sufficient: software will run the same on different machines, except slower or faster. Alternatively, in cyber-physical systems, multiple recurring instances of meaningful events may create as many dedicated logical clocks, on which to ground modeling and design practices.

Time abstraction increases efficiency in event-driven simulation or execution (i.e SystemC simulation models try to abstract time, from cycle-accurate to approximate-time, and to loosely-time), while attempting to retain functionality, but without any actual guarantee of valid accuracy (responsibility is left to the model designer). Functional determinism (a.k.a. conflict-freeness in Petri Nets, monotonicity in Kahn PNs, confluence in Milner's CCS, latency-insensitivity and elasticity in circuit design) allows for reducing to some amount the problem to that of many schedules of a single self-timed behavior, and time in many system studies is partitioned into models of computation and communication (MoCCs). Multiple, multiform time(s) raises the question of combination, abstraction or refinement between distinct time bases. The question of combining continuous time with discrete logical time calls for proper discretization in simulation and implementation. While timed reasoning takes multiple forms, there is no unified foundation to reason about multi-form time in system design.

The objective of project-team TEA is henceforth to define formal models for timed quantitative reasoning, composition, and integration in embedded system design. Formal time models and calculi should

allow us to revisit common domain problems in real-time system design, such as time predictability and determinism, memory resources predictability, real-time scheduling, mixed-criticality and power management; yet from the perspective gained from inter-domain timed and quantitative abstraction or refinement relations. A regained focus on fundamentals will allow to deliver better tooled methodologies for virtual prototyping and integration of embedded architectures.

### 3 Research program

#### 3.1 Previous Works

The challenges of team TEA support the claim that sound Cyber-Physical System design (including embedded, reactive, and concurrent systems altogether) should consider multi-form time models as a central aspect. In this aim, architectural specifications found in software engineering are a natural focal point to start from. Architecture descriptions organize a system model into manageable components, establish clear interfaces between them, collect domain-specific constraints and properties to help correct integration of components during system design. The definition of a formal design methodology to support heterogeneous or multi-form models of time in architecture descriptions demands the elaboration of sound mathematical foundations and the development of formal calculi and methods to instrument them.

System design based on the “synchronous paradigm” has focused the attention of many academic and industrial actors on abstracting non-functional implementation details from system design. This elegant design abstraction focuses on the logic of interaction in reactive programs rather than their timed behavior, allowing to secure functional correctness while remaining an intuitive programming model for embedded systems. Yet, it corresponds to embedded technologies of single cores and synchronous buses from the 90s, and may hardly cover the semantic diversity of distribution, parallelism, heterogeneity, of cyber-physical systems found in 21st century Internet-connected, true-time-synchronized clouds, of tomorrow’s grids.

By contrast with a synchronous hypothesis, yet from the same era, the polychronous MoCC is inherently capable of describing multi-clock abstractions of GALS systems. Polychrony is implemented in the data-flow specification language Signal, available in the Eclipse project POP (**Polychrony on Polarsys**) and in the CCSL standard **Clock Constraints in CCSL** available from the TimeSquare project. Both provide tooled infrastructures to refine high-level specifications into real-time streaming applications or locally synchronous and globally asynchronous systems, through a series of model analysis, verification, and synthesis services. These tool-supported refinement and transformation techniques can assist the system engineer from the earliest design stages of requirement specification to the latest stages of synthesis, scheduling and deployment. These characteristics make polychrony much closer to the required semantic for compositional, refinement-based, architecture-driven, system design.

While polychrony was a step ahead of the traditional synchronous hypothesis, CCSL is a leap forward from synchrony and polychrony. The essence of CCSL is “multi-form time” toward addressing all of the domain-specific physical, electronic and logical aspects of cyber-physical system design.

#### 3.2 Timed Modeling

To formalize timed semantics for system design, we shall rely on algebraic representations of time as clocks found in previous works and introduce a paradigm of “time system”: refinement types that represent timed behaviors. Just as a type system abstracts data carried along operations in a program, a “time system” abstracts the causal interaction of that program module or hardware element with its environment, its pre- and post-conditions, its assumptions and guarantees, either logical or numerical, discrete or continuous. Some fundamental concepts we envision are present in the clock calculi found in data-flow synchronous languages like Signal or Lustre, yet bound to a particular model of timed concurrency.

In particular, the principle of refinement type systems<sup>1</sup>, is to associate information (data-types) inferred from programs and models with properties pertaining, for instance, to the algebraic domain

<sup>1</sup>*Abstract Refinement Types*. N. Vazou, P. Rondon, and R. Jhala. European Symposium on Programming. Springer, 2013.

on their value, or any algebraic property related to its computation: effect, memory usage, pre-post condition, value-range, cost, speed, time, temporal logic<sup>2</sup>. Being grounded on type and domain theories, such type systems system should naturally be equipped with program analysis techniques based on type inference (for data-type inference) or abstract interpretation (for program properties inference) to help establish formal relations between heterogeneous component “types”.

Gaining scalability requires the capacity to modularly decompose systems which can be obtained using Abadi and Lamport’s “*Composing Specifications*” and implemented by the notion of assume-guarantee contracts or Dijkstra monads. Verification problems encompassing heterogeneously timed specifications are common and of great variety: checking correctness between abstract (e.g. the synchronous hypothesis) and concrete time models (e.g. real-time architectures) relates to desynchronisation (from synchrony to asynchrony) and scheduling analysis (from synchronous data-flow to hardware). More generally, they can be perceived from heterogeneous timing viewpoints (e.g. mapping a synchronous-time software on a real-time middleware or hardware).

This perspective demands capabilities to use abstraction and refinement mechanisms for time models (using simulation, refinement, bi-simulation, equivalence relations) but also to prove more specific properties (synchronization, determinism, endochrony). All this formalization effort will allow to effectively perform the tooled validation of common cross-domain properties (e.g. cost v.s. power v.s. performance v.s. software mapping) and tackle problems such as these integrating constraints of battery capacity, on-board CPU performance, available memory resources, software schedulability, to logical software correctness and plant controllability.

### 3.3 Modeling Architectures

To address the formalization of such cross-domain case studies, modeling the architecture formally plays an essential role. An architectural model represents components in a distributed system as boxes with well-defined interfaces, connections between ports on component interfaces, and specifies component properties that can be used in analytical reasoning about the model.

In system design, an architectural specification serves several important purposes. First, it breaks down a system model into components of manageable size and complexity, to establish clear interfaces between components. In this way, complexity becomes manageable by hiding details that are not relevant at a given level of abstraction. Clear, formally defined, and semantically rich component interfaces facilitate integration by allowing most validation efforts to be conducted modularly. Connections between components, which specify how components interact with each other, help propagate the guaranteed effects of a component to the assumptions of linked components.

Most importantly, an architectural model is a repository to share knowledge about the system being designed. This knowledge can be represented as requirements, design artifacts, component implementations, held together by a structural backbone. Such a repository enables automatic generation of analytical models for different aspects of the system, such as timing, reliability, security, performance, energy, etc. Since all the models are generated from the same source, the consistency of assumptions w.r.t. guarantees, of abstractions w.r.t. refinements, used for different analyses becomes easier, and can be properly ensured in a design methodology based on formal verification and synthesis methods.

### 3.4 Scheduling Theory

Based on sound formalization of time and CPS architectures, real-time scheduling theory provides tools for predicting the timing behavior of a CPS which consists of many interacting software and hardware components. Expressing parallelism among software components is a crucial aspect of the design process of a CPS. It allows for efficient partition and exploitation of available resources.

The literature about real-time scheduling<sup>3</sup> provides very mature schedulability tests regarding many scheduling strategies, preemptive or non-preemptive scheduling, uniprocessor or multiprocessor scheduling, etc. Scheduling of data-flow graphs has also been extensively studied in the past decades.

---

<sup>2</sup>*LTL types FRP*. A. Jeffrey. Programming Languages meets Program Verification.

<sup>3</sup>*A survey of hard real-time scheduling for multiprocessor systems*. R. I. Davis and A. Burns. *ACM Computing Survey* 43(4), 2011.



A milestone in this prospect is the development of abstract affine scheduling techniques<sup>4</sup>. It consists, first, of approximating task communication patterns (e.g. between Safety-Critical Java threads) using cyclo-static data-flow graphs and affine functions. Then, it uses state of the art ILP techniques to find optimal schedules and to concretize them as real-time schedules in the program implementations<sup>5 6</sup>.

Abstract scheduling, or the use of abstraction and refinement techniques in scheduling borrowed to the theory of abstract interpretation<sup>7</sup> is a promising development toward tooled methodologies to orchestrate thousands of heterogeneous hardware/software blocks on modern CPS architectures (just consider modern cars or aircrafts). It is an issue that simply defies the state of the art and known bounds of complexity theory in the field, and consequently requires a particular focus.

### 3.5 Verified programming for system design

The IoT is a network of devices that sense, actuate and change our immediate environment. Against this fundamental role of sensing and actuation, design of edge devices often considers actions and event timings to be primarily software implementation issues: programming models for IoT abstract even the most rudimentary information regarding timing, sensing and the effects of actuation. As a result, applications programming interfaces (API) for IoT allow wiring systems fast without any meaningful assertions about correctness, reliability or resilience.

We make the case that the "API glue" must give way to a logical interface expressed using contracts or refinement types. Interfaces can be governed by a calculus – a refinement type calculus – to enable reasoning on time, sensing and actuation, in a way that provides both deep specification refinement, for mechanized verification of requirements, and multi-layered abstraction, to support compositionality and scalability, from one end of the system to the other.

Our project seeks to elevate the "function as type" paradigm to that of "system as type": to define a refinement type calculus based on concepts of contracts for reasoning on networked devices and integrate them as cyber-physical systems<sup>8</sup>. An invited paper<sup>9</sup> outlines our progress with respect to this aim and plans towards building a verified programming environment for networked IoT devices: we propose a type-driven approach to verifying and building safe and secure IoT applications.

Accounting for such constraints in a more principled fashion demands reasoning about the composition of all the software and hardware components of the application. Our proposed framework takes a step in this direction by (1) using refinement types to make physical constraints explicit and (2) imposing an event-driven programming discipline to simplify the reasoning of system-wide properties to that of an event queue. In taking this approach, a developer could build a verified IoT application by ensuring that a well-typed program cannot violate the physical constraints of its architecture and environment.

## 4 Application domains

In collaboration with Mitsubishi R&D, we explore another application domain where time and domain heterogeneity are prime concerns: factory automation. In factory automation alone, a system is conventionally built from generic computing modules: PLCs (Programmable Logic Controllers), connected to the environment with actuators and detectors, and linked to a distributed network. Each individual, physically distributed, PLC module must be timely programmed to perform individually coherent actions and fulfill the global physical, chemical, safety, power efficiency, performance and latency requirements of the whole production chain. Factory chains are subject to global and heterogeneous (physical, electronic, functional) requirements whose enforcement must be orchestrated for all individual components.

Model-based analysis in factory automation emerges from different scientific domains and focuses on different CPS abstractions that interact in subtle ways: logic of PLC programs, real-time electro-mechanical processing, physical and chemical environments. This yields domain communication

<sup>4</sup> *Buffer minimization in EDF scheduling of data-flow graphs*. A. Bouakaz and J.-P. Talpin. LCTES, ACM, 2013.

<sup>5</sup> *ADFG for the synthesis of hard real-time applications*. A. Bouakaz, J.-P. Talpin, J. Vitek. ACS D, IEEE, June 2012.

<sup>6</sup> *Design of SCJ Level 1 Applications Using Affine Abstract Clocks*. A. Bouakaz and J.-P. Talpin. SCOPES, ACM, 2013.

<sup>7</sup> *La vérification de programmes par interprétation abstraite*. P. Cousot. Séminaire au Collège de France, 2008.

<sup>8</sup> Refinement types for system design. Jean-Pierre Talpin. FDL'18 keynote.

<sup>9</sup> Steps toward verified programming of embedded computing systems. Jean-Pierre Talpin, Jean-Joseph Marty, Deian Stefan, Shrawan Nagarayan, Rajesh Gupta, DATE'18.

problems that render individual domain analysis useless. For instance, if one domain analysis (e.g. software) modifies a system model in a way that violates assumptions made by another domain (e.g. chemistry) then the detection of its violation may well be impossible to explain to either the software or chemistry experts. As a consequence, cross-domain analysis issues are discovered very late during system integration and lead to costly fixes. This is particularly prevalent in multi-tier industries, such as avionic, automotive, factories, where systems are prominently integrated from independently-developed parts.

## 5 Highlights of the year

We released the first fully verified implementation of "femto-containers": a virtual machine executing eBPF scripts (extended Berkeley Packet Filters) to provide kernel extensibility in the RIOT operating system [10]. The workflow of our implementation consisted of 1) modelling an executable interpreter of the rBPF virtual machine in Coq; 2) formally verifying that the rBPF interpreter satisfies software faults isolation; 3) refining this proof model into an equivalent monadic representation; 4) transpiling this monadic representation as an imperative program in CompCert Clight using  $\delta x$ ; 4) proving a forward simulation relation between the monadic Coq model and the Clight implementation. The resulting artifact: CertFC [6.1.1] was benchmarked, validated, and integrated to RIOT's-featured femto-containers as an open-source library. A second publication reports this evaluation and the integration of CertFC in RIOT OS [11].

## 6 New software and platforms

### 6.1 New software

#### 6.1.1 CertFC

**Name:** End-to-end Mechanized Proof of an eBPF Virtual Machine for Micro-controllers

**Keywords:** Virtualization, Network Function Virtualization, Proof, Isolation, Code generation

**Functional Description:** CertrBPF includes a verified C verifier and interpreter. The verifier performs static checks to verify that the rBPF instructions are syntactically correct. If the verification succeeds, the program is run by the interpreter. The static verification and the dynamic checks ensure software fault isolation of the running rBPF program. Namely, we have the guarantee that:

The interpreter never crashes. More precisely, the proved C program is free of undefined behaviours such as division by zero or invalid memory accesses. All the memory accesses are performed in designated memory regions which are an argument of the interpreter.

The development of CertrBPF follows a refinement methodology with three main layers:

The proof model: an executable Coq specification of the rBPF virtual machine  
The synthesis model: a refined and optimised executable Coq program that is close in style to a C program. Eventually, we have the synthesis model (named dx model) which is compliant with the dx C extraction tool.

The implementation model: the extracted C implementation in the form of CompCert Clight AST.

**URL:** <https://github.com/future-proof-iot/CertFC>

**Contact:** Shenghao Yuan

#### 6.1.2 ADFG

**Name:** Affine data-flow graphs schedule synthesizer

**Keywords:** Code generation, Scheduling, Static program analysis

**Functional Description:** ADFG is a synthesis tool of real-time system scheduling parameters: ADFG computes task periods and buffer sizes of systems resulting in a trade-off between throughput

maximization and buffer size minimization. ADFG synthesizes systems modeled by ultimately cyclo-static dataflow (UCSDF) graphs, an extension of the standard CSDF model.

Knowing the WCET (Worst Case Execute Time) of the actors and their exchanges on the channels, ADFG tries to synthesize the scheduler of the application. ADFG offers several scheduling policies and can detect unschedulable systems. It ensures that the real scheduling does not cause overflows or underflows and tries to maximize the throughput (the processors utilization) while minimizing the storage space needed between the actors (i.e. the buffer sizes).

Abstract affine scheduling is first applied on the dataflow graph, that consists only of periodic actors, to compute timeless scheduling constraints (e.g., relations between the speeds of two actors) and buffering parameters. Then, symbolic schedulability policies analysis (i.e., synthesis of timing and scheduling parameters of actors) is applied to produce the scheduler for the actors.

ADFG, initially defined to synthesize real-time schedulers for SCJ/L1 applications, may be used for scheduling analysis of AADL programs.

**URL:** <http://polychrony.inria.fr/ADFG/>

**Authors:** Thierry Gautier, Jean-Pierre Talpin, Adnan Bouakaz, Alexandre Honorat, Loïc Besnard

**Contact:** Loïc Besnard

### 6.1.3 POLYCHRONY

**Keywords:** Code generation, AADL, Proof, Optimization, Multi-clock, GALS, Architecture, Cosimulation, Real time, Synchronous Language

**Functional Description:** Polychrony is an Open Source development environment for critical/embedded systems. It is based on Signal, a real-time polychronous data-flow language. It provides a unified model-driven environment to perform design exploration by using top-down and bottom-up design methodologies formally supported by design model transformations from specification to implementation and from synchrony to asynchrony. It can be included in heterogeneous design systems with various input formalisms and output languages. The Polychrony tool-set provides a formal framework to: validate a design at different levels, by the way of formal verification and/or simulation, refine descriptions in a top-down approach, abstract properties needed for black-box composition, compose heterogeneous components (bottom-up with COTS), generate executable code for various architectures. The Polychrony tool-set contains three main components and an experimental interface to GNU Compiler Collection (GCC):

\* The Signal toolbox, a batch compiler for the Signal language, and a structured API that provides a set of program transformations. It can be installed without other components and is distributed under GPL V2 license.

\* The Signal GUI, a Graphical User Interface to the Signal toolbox (editor + interactive access to compiling functionalities). It can be used either as a specific tool or as a graphical view under Eclipse. It has been transformed and restructured, in order to get a more up-to-date interface allowing multi-window manipulation of programs. It is distributed under GPL V2 license.

\* The POP Eclipse platform, a front-end to the Signal toolbox in the Eclipse environment. It is distributed under EPL license.

**URL:** <https://www.polarsys.org/projects/polarsys.pop>

**Contact:** Loïc Besnard

**Participants:** Loïc Besnard, Paul Le Guernic, Thierry Gautier

**Partners:** CNRS, Inria

## 7 New results

### 7.1 Verified programming and secure integration of operating system libraries in RIOT-fp

**Participants:** Shenghao Yuan, Frédéric Besson, Jean-Pierre Talpin, Samuel Hym, Koen Zandberg, Emmanuel Bachelli.

Our project aims at the formal verification of safety and security properties for embedded operating system libraries in the RIOT operating system, focusing on a virtual machine to implement kernel-extensibility by hosting runtime applications and services using the virtual instruction set architecture (ISA) of eBPF (extended Berkeley packet filter).

We implemented a fully verified rBPF virtual machine for RIOT's "femto-containers" in the proof assistant Coq [10]. The workflow of our implementation consists of 1) modelling an executable interpreter of the rBPF virtual machine in Coq; 2) formally verifying that the rBPF interpreter satisfies software faults isolation; 3) refining this proof model into a monadic form; 4) transpiling this monadic representation as an imperative program in CompCert Clight; 4) proving a forward simulation relation between the monadic Coq model and the Clight implementation.

The resulting artifact: CertFC [6.1.1] was then benchmarked, validated, and integrated to RIOT's-featured femto-containers as an open-source library. A second publication reports this evaluation and its integration to RIOT OS [11].

### 7.2 Semantic foundations for cyber-physical systems using higher-order UTP

**Participants:** Xiong Xu, Jean-Pierre Talpin, Naijun Zhan, Shuling Wang, Bohua Zhan.

The associate-team CONVEX yielded a number of important results in the development of theorem-proving models for cyber-physical systems verification, under the joint supervision of Xiong Xu with team TEA and Naijun Zhan's group at ISCAS, Beijing. In [8], we define an extension of Hoare&He's Unifying Theories of Programming (UTP) with higher-order quantification and provide a formal semantics for modeling and verifying hybrid systems. Higher-order UTP (HUTP) provides a semantics foundation for cyber-physical systems (CPSs). It separates the concerns in CPS design into time, state and trace, which support the specification of discrete, real-time and continuous dynamics, concurrency and communication, and higher-order quantification. Within HUTP, we defined a calculus of normal hybrid designs to model, analyze, compose, refine and verify heterogeneous hybrid system models. In addition, we defined respective formal semantics for Hybrid Communicating Sequential Processes (HCSP) and Simulink using HUTP and justified the correctness of the translation from Simulink to HCSP by translation validation as an application of HUTP. We have started applying this framework to co-modeling real-time hardware platforms in AADL and hybrid discrete-continuous models in Simulink/Stateflow (AADL+S/S) to uniformly analyze and verify this combination in HUTP. In this aim, [9] provides, to our knowledge, the simplest formalization of a semantic for Simulink.

We are furthering our study and application of the HUTP in the Isabelle/HOL theorem prover by developing a hybrid extension of Milner's  $\pi$ -calculus. Our aim is to model, analyze and verify mobile hybrid systems, or more generally mobile and extensible hybrid systems as found in the IoT in the form of intermittent, mobile, distributed cyber-physical systems such as transportation grids in general (data, resources, vehicles).

First, based on our previous work on higher-order unifying theories of programming (HUTP), we will investigate the mechanism of mobile hybrid systems, i.e., an HUTP theory extended with mobility. Then, we will focus on the operational semantics of a hybrid extension to a generic  $\pi$ -calculus. Finally, we will refine this theory to a protocolar session calculus and build a session type system for hybrid mobile processes, to characterize important safety and security properties, such as determinism and isolation.

### 7.3 A logical framework to verify requirements of hybrid system models

**Participants:** Stéphane Kastenbaum, Benoit Boyer, Jean-Pierre Talpin.

The goal of this PhD project is to build on the previous work done in Simon Lunel's PhD thesis. The goal is to ensure the correctness-by-design of cyber-physical system models. It deals with cyber-physical systems (CPS), which are assemblies of networked, heterogeneous, hardware, and software components sensing, evaluating, and actuating a physical environment. This heterogeneity induces complexity that makes CPSs challenging to model correctly. Since CPSs often have critical functions, it is however of utmost importance to formally verify them in order to provide the highest guarantees of safety. Faced with CPS complexity, model abstraction becomes paramount to make verification attainable. To this end, assume/guarantee contracts enable component model abstraction to support a sound, structured, and modular verification process.

While abstractions of models by contracts are usually proved sound, none of the related contract frameworks themselves have, to the best of our knowledge, been formally proved correct so far. In this aim, we present the formalization of a generic assume/guarantee contract theory in the Coq proof assistant. We identify and prove theorems that ensure its correctness. Our theory is generic, or parametric, in that it can be instantiated and used with any given logic, in particular hybrid logics, in which highly complex cyber-physical systems can uniformly be described.

We are pursuing the development of this model by instantiating this parametric theory with differential dynamic logic, a language to specify cyber-physical systems. This instantiation provides a use case for the meta-theory of contracts. To this end, we are using CoqDL, a formalization of differential dynamic logic that was developed by the Logical Systems Lab of Carnegie Mellon University.

### 7.4 ADFG: Affine data-flow graphs scheduler synthesis

**Participants:** Thierry Gautier, Jean-Pierre Talpin, Shuvra Bhattacharyya, Alexandre Honorat, Hai Nam Tran.

We continued our research involving the development of advanced signal processing dataflow methods for the ADFG tool. This research is collaboratively developed with the TEA Team, Hai Nam Tran (Lab-STICC/UBO), Alexandre Honorat (INSA) and Shuvra Bhattacharyya (UMD/INSA/INRIA). Our emphasis during this reporting period is the public release of the ADFG infrastructure on Gitlab [6.1.2](#).

## 8 Bilateral contracts and grants with industry

### 8.1 Bilateral contracts with industry

#### Inria – Mitsubishi Electric framework program (2018+)

Title: Inria – Mitsubishi Electric framework program

INRIA principal investigator: Jean-Pierre Talpin

International Partner: Mitsubishi Electric R&D Europe (MERCE)

Duration: 2018+

Abstract: Following up the fruitful collaboration of TEA with the formal methods group at MERCE, Inria and Mitsubishi Electric signed an Inria-wide collaboration agreement, which currently hosts projects with project-teams Sumo and Tea, as well as Toccata.

## 8.2 Bilateral grants with industry

**Participants:** Jean-Pierre Talpin, Stéphane Kastenaum, Benoit Boyer.

### Mitsubishi Electric R&D Europe (2019-2022)

**Title:** A logical framework to verify requirements of hybrid system models

**INRIA principal investigator:** Jean-Pierre Talpin, Stéphane Kastenaum

**International Partner:** Mitsubishi Electric R&D Europe

**Duration:** 2019 - 2022

**Abstract:** The goal of this doctoral project is to verify and build cyber-physical systems (CPSs) with a correct-by-construction approach in order to validate system requirements against the two facets of the cyber and physical aspects of such designs. Our approach is based on components augmented with formal contracts that can be composed, abstracted or refined. It fosters on the proof of system-level requirements by composing individual properties proved at component level. While semantically grounded, the tooling of this methodology should be usable by regular engineers (i.e. not proof theory specialists).

## 9 Partnerships and cooperations

**Participants:** Jean-Pierre Talpin, Shenghao Yuan, Thierry Gautier.

### 9.1 International initiatives

#### Associate Teams in the framework of the Inria International Program with the Chinese Academy of Science

**Title:** **Convex** Compositional Verification of Cyber-Physical Systems (**Convex**)

**Partner Institution:** Institute of Software, Chinese Academy of Science (ISCAS)

**Date/Duration:** **3 years** Formal modeling and verification methods have successfully improved software safety and security in vast application domains in transportation, production and energy. However, formal methods are labor-intensive and require highly trained software developers. Challenges facing formal methods stem from rapid evolution of hardware platforms, the extensibility and increasing cost of software infrastructures, and from the interaction between software, hardware and physics in mobile and networked cyber-physical systems. Automation and expressivity of formal verification tools must be improved not only to scale functional verification to very large software stacks, but also verify non-functional properties from models of hardware (time, energy) and physics (domain). Abstraction, compositionality and refinement are essential properties to provide the necessary scalability to tackle the complexity of system design with methods able to scale heterogeneous, concurrent, networked, timed, extensible, discrete and continuous models of cyber-physical systems. The aim of project Convex is to define a verified CPS architecture design methodology that takes advantage of existing time and concurrency modeling standards, yet focuses on interfacing heterogeneous and exogenous models using simple, mathematically-defined structures, to achieve the single goal of correctly integrating CPS components.

### 9.1.1 Visits to international teams

**Jean-Pierre Talpin**

**Visited institution:** ISCAS

**Country:** China

**Dates:** 15/12/22 - 28/01/23

**Context of the visit:** Convex

**Mobility program/type of mobility:** definition and development of a hybrid  $\pi$ -calculus in the HOL4 theorem prover.

## 9.2 National initiatives

**Title:** RIOT-fp: Future-proof IoT

**Duration:** 4 years

**Coordinator:** Emmanuel Bachelli

**Partners:** Tribe, Eva, Grace, Prosecco, Tea, Freie Universität Berlin and Fujitsu.

**Summary:** RIOT-fp is a research project on cyber-security targeting low-end, microcontroller-based IoT devices, on which operating systems such as RIOT run, and the development of a low-power network stack. Taking a global and practical approach, RIOT-fp gathers partners planning to enhance RIOT with an array of security mechanisms. The main challenges tackled by RIOT-fp are: 1/ developing high-speed, high-security, low-memory IoT crypto primitives, 2/ providing guarantees for software execution on low-end IoT devices, and 3/ enabling secure IoT software updates and supply-chain, over the network. Beyond academic outcomes, the output of RIOT-fp is open source code published, maintained and integrated in the open source ecosystem around RIOT. As such, RIOT-fp strives to contribute usable building blocks for an open source IoT solution improving the typical functionality vs. risk tradeoff for end-users. The goal of project-team TEA in RIOT-fp is to build verified operating system libraries for IoT devices using proof-oriented programming techniques in Coq and F\*, such as the actual bootloader of RIOT and its femto-containers: virtual machines isolating kernel-level execution of user-supplied applications and services using the eBPF virtual ISA.

## 10 Dissemination

Jean-Pierre Talpin

### 10.1 Promoting scientific activities

Jean-Pierre Talpin chairs the steering committee of the ACM-IEEE International Conference on Formal Methods and Models for System Design (MEMOCODE) which, for its 20th anniversary, joined the ESWEEK event.

He also cochaired the program committee of the Symposium on Dependable Software Engineering: Theories, Tools and Applications (SETTA'22).

He participated in the program committee of the International Conference on Embedded Computer Systems: Architectures, Modeling and Simulation (SAMOS XXII) and of the 1st International Workshop on Formal Engineering of Cyber-Physical Systems (FE-CPS@TASE'22).

## 10.2 Teaching - Supervision - Juries

Jean-Pierre Talpin served as rapporteur and president for the Ph.D. defense committee of Frédéric Fort, entitled "Programming adaptive real-time systems", at the University of Lille.

He also served as rapporteur for the Ph.D. defense committee of Nicolas Dejon, entitled "Conception d'un noyau sécurisé pour objets contraints", at the University of Lille.

Jean-Pierre Talpin supervises the Ph.D. Thesis of Shenghao Yuan and Stéphane Kastenbaum.

## 11 Scientific production

### 11.1 Major publications

- [1] L. Besnard, T. Gautier, P. Le Guernic, C. Guy, J.-P. Talpin, B. Larson and E. Borde. 'Formal Semantics of Behavior Specifications in the Architecture Analysis and Design Language Standard'. In: *Cyber-Physical System Design from an Architecture Analysis Viewpoint*. Cyber-Physical System Design from an Architecture Analysis Viewpoint. Springer, Jan. 2017. DOI: [10.1007/978-981-10-4436-6\\_3](https://doi.org/10.1007/978-981-10-4436-6_3). URL: <https://hal.inria.fr/hal-01615143>.
- [2] L. Besnard, T. Gautier, P. Le Guernic and J.-P. Talpin. 'Compilation of Polychronous Data Flow Equations'. In: *Synthesis of Embedded Software*. Ed. by S. K. Shukla and J.-P. Talpin. Springer, 2010, pp. 1–40. DOI: [10.1007/978-1-4419-6400-7\\_1](https://doi.org/10.1007/978-1-4419-6400-7_1). URL: <https://hal.inria.fr/inria-00540493>.
- [3] A. Bouakaz and J.-P. Talpin. 'Design of Safety-Critical Java Level 1 Applications Using Affine Abstract Clocks'. In: *International Workshop on Software and Compilers for Embedded Systems*. St. Goar, Germany, June 2013, pp. 58–67. DOI: [10.1145/2463596.2463600](https://doi.org/10.1145/2463596.2463600). URL: <https://hal.inria.fr/hal-00916487>.
- [4] A. Honorat, H. N. Tran, L. Besnard, T. Gautier, J.-P. Talpin and A. Bouakaz. 'ADFG: a scheduling synthesis tool for dataflow graphs in real-time systems'. In: *International Conference on Real-Time Networks and Systems*. Grenoble, France, Oct. 2017, pp. 1–10. DOI: [10.1145/3139258.3139267](https://doi.org/10.1145/3139258.3139267). URL: <https://hal.inria.fr/hal-01615142>.
- [5] S. Lunel, S. Mitsch, B. Boyer and J.-P. Talpin. 'Parallel Composition and Modular Verification of Computer Controlled Systems in Differential Dynamic Logic'. In: *FM 2019 - 23rd International Symposium on Formal Methods*. Long version of an article accepted to the conference FM'19. Porto, Portugal, Oct. 2019, pp. 1–22. URL: <https://hal.inria.fr/hal-02193642>.
- [6] S. Nakajima, J.-P. Talpin, M. Toyoshima and H. Yu. *Cyber-Physical System Design from an Architecture Analysis Viewpoint*. Communications of NII Shonan Meetings. Springer, Jan. 2017. DOI: [10.1007/978-981-10-4436-6](https://doi.org/10.1007/978-981-10-4436-6). URL: <https://hal.inria.fr/hal-01615144>.
- [7] H. Yu, J. Prashi, J.-P. Talpin, S. K. Shukla and S. Shiraishi. 'Model-Based Integration for Automotive Control Software'. In: *Digital Automation Conference*. ACM. San Francisco, United States, June 2015. URL: <https://hal.inria.fr/hal-01148905>.

### 11.2 Publications of the year

#### International journals

- [8] X. Xu, J.-P. Talpin, S. Wang, B. Zhan and N. Zhan. 'Semantics Foundation for Cyber-Physical Systems Using Higher-Order UTP'. In: *ACM Transactions on Software Engineering and Methodology* (23rd Apr. 2022), pp. 1–47. DOI: [10.1145/3517192](https://doi.org/10.1145/3517192). URL: <https://hal.inria.fr/hal-03888055>.
- [9] X. Xu, B. Zhan, S. Wang, J.-P. Talpin and N. Zhan. 'A denotational semantics of Simulink with higher-order UTP'. In: *Journal of Logical and Algebraic Methods in Programming* 130 (Jan. 2023), p. 100809. DOI: [10.1016/j.jlamp.2022.100809](https://doi.org/10.1016/j.jlamp.2022.100809). URL: <https://hal.inria.fr/hal-03888092>.



**International peer-reviewed conferences**

- [10] S. Yuan, F. Besson, J.-P. Talpin, S. Hym, K. Zandberg and E. Baccelli. ‘End-to-end Mechanized Proof of an eBPF Virtual Machine for Micro-controllers’. In: CAV 2022 - 34th International Conference on Computer Aided Verification. Haifa, Israel, 7th Aug. 2022, pp. 1–23. URL: <https://hal.inria.fr/hal-03888082>.
- [11] K. Zandberg, E. Baccelli, S. Yuan, F. Besson and J.-P. Talpin. ‘Femto-Containers: Lightweight Virtualization and Fault Isolation For Small Software Functions on Low-Power IoT Microcontrollers’. In: Middleware 2022 - 23rd ACM/IFIP International Conference Middleware. quebec, Canada, 7th Nov. 2022, pp. 1–12. DOI: [10.1145/3528535.3565242](https://doi.org/10.1145/3528535.3565242). URL: <https://hal.inria.fr/hal-03888109>.