INRIA

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team aoste

Models and Methods for the Analysis and Optimization of Systems with Real-time and Embedded Constraints

Sophia Antipolis - Méditerranée, Paris - Rocquencourt

Theme : Embedded and Real Time Systems

Table of contents

Aoste is a joint team with UNS (University of Nice/Sophia-Antipolis) and CNRS UMR I3S. It is also co-located between Sophia-Antipolis and Rocquencourt. Project members originate from the former INRIA Tick and Ostre teams, together with the I3S Sports team.

1. Team

Research Scientist

Robert de Simone [Team Leader, DR, HdR]
Yves Sorel [Team Vice-leader, DR]
Dumitru Potop-Butucaru [CR]
Charles André [UNS Pr, on leave at INRIA, HdR]
Frédéric Mallet [UNS MdC, on leave at INRIA until October]
Laurent Georges [ECE MdC, on leave at INRIA from September, HdR]

Faculty Member

Marie-Agnès Peraldi-Frati [UNS MdC]
Julien Deantoni [UNS MdC]

Technical Staff

Julien Boucaron [IE]
Benoît Ferrero [IE]
Régis Gascon [IE, from June]
Patrick Meumeu Yomsi [IE, from April]
Cécile Stentzel [ODL, until October]
Daniel de Rauglaudre [IR]
Maxence Guesdon [IR, from Development and Experimentation Service]

PhD Student

Anthony Coadou [MESR scholarship]
Omar Kermia [INRIA scholarship, until March]
Jean-François Le Tallec [BDE scholarship with ScaleoChip]
Mohamed Marouf [INRIA scholarship, joint with EPI IMARA]
Aamir Mehmood Khan [Pakistani government scholarship]
Patrick Meumeu Yomsi [INRIA scholarship, until March]

Administrative Assistant

Patricia Lachaume [TCN, CNRS until October, then INRIA]
Christine Anocq [TRS]
Sandra Devauchelle [AI, part-time (10%)]

2. Overall Objectives

2.1. Embedded System Design

Modern embedded systems combine complexity and heterogeneity both at the level of applications (with a mix of control-flow modes and multimedia data-flow streaming), and at the level of execution platforms (with increasing parallelism and multicore architectures). Compilation of the application onto the platform then takes the form of an allocation mapping involving spatial distribution as well as temporal scheduling. Formal models and methods may help to establish the correctness and the efficiency of such transformations. Static and dynamic (model-checking) analyses are also used to provide insights regarding prescribed formal semantics.

The main objective of the Aoste team is thus to promote the formal design of embedded systems, with their intrinsic concurrent, distributed and real-time aspects. For this, we develop a model-based approach, where models here have sound and precise operational semantics. For this we build upon previous experience by team members on synchronous reactive formalisms such as ESTEREL and its graphical SYNCCHARTS version, various GALS or polychronous extensions owing to Concurrency Theory (like Process Networks), and the *Algorithm-Architecture Adequation* methodology (AAA) embodied in the SYNDEX environment.

2.2. Highlights

In 2009 version 1.0 of the UML MARTE profile was adopted at OMG (officially as First Revision).

3. Scientific Foundations

3.1. Models of Computation and Communication (MoCCs)

Participants: Charles André, Julien Boucaron, Anthony Coadou, Liliana Cucu [EPI TRIO], Robert de Simone, Jean-Vivien Millo, Dumitru Potop-Butucaru, Yves Sorel.

Because of their formal semantics, the various Models of Computation and Communication (MoCCs) considered in our team can be used in a true effective design flow methodology based on model transformation to represent compilation, synthesis, analysis and optimization of concurrent embedded applications onto parallel and multicore embedded architectures. Allocation seen in that sense comprises a physical distribution/placement as well as a temporal scheduling aspect. Timing constraints and requirements may be expressed and have to be checked and preserved in the process.

This type of incremental design flow may be applied to represent a number of existing, theoretical or practical approaches to the design of embedded systems.

3.1.1. Synchronous reactive formalisms

In synchronous reactive models the various concurrent processes all run at the speed of a common global logical clock, which sets up the instantaneous reaction step. Synchronous formalisms provide an accurate representation of both hardware and scheduled embedded concurrent software; in both cases, simultaneous behaviors in a single global instant are allowed, and even often required.

Examples include ESTEREL/SYNCCHARTS, LUSTRE/SCADE, and SIGNAL/POLYCHRONY. ESTEREL and SYNCCHARTS are control-oriented state-based formalisms [46], [43], while *Lustre* and *Signal* are declarative data-flow based formalisms. Synchronous formalisms were discussed in many articles and book chapters, amongst which [50], [52], [45], [10], [2], [7].

The INRIA spin-off Esterel-EDA now develops and markets the industrial versions ESTEREL STUDIO and SCADE together with their programming environments.

3.1.2. GALS and multiclock extensions

The purely single-clock synchronous formalisms often prove to be excessively demanding for users to write large systems descriptions, with different clock domains. Independent logical clocks may be used to represent (total or partial) asynchrony amongst concurrent processes. Globally-Asynchronous/Locally-synchronous (GALS) and polychronous/multiclock models are handy extensions to provide flexibility and modularity in system design. The recently proposed theory of latency-insensitive design (LID) with elastic time is a good example of such an approach: specific protocol elements may be inserted between existing “black-box” IP block components, at a subsequent design time, to make them comply with imperative latencies on the global communications.

In any case the basic synchronous model remains the basic semantic level for behaviors, where the reaction step is defined. But natural properties (such as endochrony/asynchrony) allows to view the GALS and multiclock descriptions as higher-level versions with a natural synchronous interpretation provided by simple scheduling. The monoclock version is then obtained by dedicated scheduling techniques known as *clock calculus*.

3.1.3. Process Networks

The previous model extensions were often calling for general results from branches of Theoretical Computer Science and Concurrency Theory such as Process Networks and Process Calculi. Process Networks comprise Petri Nets and Kahn Networks, as well as various specializations and generalizations, such as Event/Marked Graphs, Data-Flow Domains (Synchronous, Boolean, CycloStatic, CycloDynamic,...), while Process Algebras such as CCS and CSP gave rise to extensions with simultaneous events in SCCS and Meije. We held former background in this field, and more generally in the use of formal operational semantics in the design of verification and analysis techniques for such systems. We bridged the gap in using such models to study techniques for optimized placement and (static) scheduling of models. The specific features of hardware targets let us phrase these questions in a specific context, where ad-hoc ultimately periodic regimes can be established.

3.1.4. Static k-periodic scheduling and routing

Following our *time refinement* approach, early untimed causal models may be transformed into multiclock or GALS ones, then precisely scheduled to a uniform single time. This type of approach is used for instance in the static, k-periodic scheduling of dataflow process networks such as Event Graphs [47], [44], or Synchronous DataFlow graphs [54] and various extensions in UC Berkeley's Ptolemy. We extended the approach by providing means for the designers to provide his/her own time constraints on a given modeling framework, and to express the actual refinement from a given (abstract) time frame to another, more concrete one.

This theory of modulo and k-periodic static scheduling for process networks (mostly Marked/Event Graphs) recently got a renewal of interest due to its application in the context of Latency-Insensitive Design [48] of SoCs. The nature of communication channels, used there for interconnect fabric, demands optimal buffer/place sizing, with corresponding flow control. We contributed several results in this direction, with fine characterization of optimal algorithmic techniques to provide such ultimately k-periodic schedules. They are progressively implemented in the K-PASSA prototype software, described in 5.2.

3.1.5. AAA models

The AAA (*Algorithm-Architecture Adequation*) methodology which is intended for optimizing distributed real-time embedded systems relies on three models.

The Algorithm model is an extension of the well known data-flow model from Dennis [49]. It is a directed acyclic hyper-graph (DAG) that we call "conditioned factorized data dependence graph", whose vertices are "operations" and hyper-edges are directed "data or control dependences" between operations. The data dependences defines a partial order on the operations execution. The basic data-flow model was extended in three directions: first infinite (resp. finite) repetition of a sub-graph pattern in order to specify the reactive aspect of real-time systems (resp. in order to specify the finite repetition of a sub-graph consuming different data similar to a loop in imperative languages), second "state" when data dependences are necessary between different infinite repetitions of the sub-graph pattern introducing cycles which must be avoided by introducing specific vertices called "delays" (similar to z^{-n} in automatic control), third "conditioning" of an operation by a control dependence similar to conditional control structure in imperative languages, allowing the execution of alternative subgraphs. Delays combined with conditioning allow the programmer to specify automata necessary for describing "mode changes".

The Architecture model is a directed graph, whose vertices are of two types: "processor" (one sequencer of operations and possibly several sequencers of communications) and "medium" (support of communications), and whose edges are directed connections.

The implementation model [4] is also a directed graph, obtained through an external compositional law, where an architecture graph operates on an algorithm graph in order to give, as a result, a new algorithm graph, which corresponds to the initial algorithm graph, distributed and scheduled according to the architecture graph.

3.1.6. Distributed Real-Time Scheduling and Optimization

We adress two main issues: monoprocessor real-time scheduling and multiprocessor real-time scheduling where constraints must mandatorily be met otherwise dramatic consequences may occur (hard real-time) and where resources must be minimized because of embedded features.

In our monoprocessor real-time scheduling work, beside the classical deadline constraint, often equal to a period, we take into consideration dependences beetween tasks and several, possibly related, latencies. A latency is a generalization [3] of the typical "end-to-end" constraint. Dealing with multiple real-time constraints raises the complexity of that issue. Moreover, because the preemption leads to a waste of resources due to its approximation in the WCET (Worst Execution Time) of every task as proposed by Liu and Leyland [55], we first studied non-preemtive real-time scheduling with dependences, periodicities, and latencies constraints. Although a bad approximation may have dramatic consequences on real-time scheduling, there are only few researches on this topic. We have been investigating preemptive real-time scheduling since few years, but seeking the exact cost of the preemption such that it can be integrated in schedulability conditions, and in the corresponding scheduling algorithms. More generally, we are interested in integrating in the schedulability analyses the cost of the RTOS (Real-Time Operating System), for which the exact cost of preemption is the most difficult part because it varies according to the instance of each task [6]. Finally, we investigate also the problem of mixing hard real-time and soft real-time constraints that arises in the most complex applications.

The second research area is devoted to distributed real-time scheduling with embedding constraints. We use the results obtained in the monoprocessor case in order to derive solutions for the problem of multiprocessor (distributed) real-time scheduling. In addition to satisfy the multiple real-time constraints mentioned in the monoprocessor case, we have to minimize the total execution time (makespan) since we deal with automatic control applications involving feedback. Furthermore, the domain of embedded systems leads to solve minimization resources problems. Since these optimization problems are of NP-hard complexity we develop exact algorithms (B & B, B & C) which are optimal for simple problems, and heuristics which are sub-obtimal for realistic problems corresponding to industrial needs. Long time ago we proposed a very fast "greedy" heuristics [51] whose results were regularly improved, and extended with local neighborhood heuristics [5], or used as initial solutions for metaheuristics such as variants of "simulated annealing".

Finally, since real-time distributed architectures are prone to failures we study the possibility to tolerate faults in such systems. We focus on software redondance rather than hardware redondance to guarantee the same real-time behaviour of the system, in the presence of a certain number of faulty processors and of communication media beeing specified by the designer. We investigate fail silent, transient, intermittent, and Byzantine faults.

3.2. Model-Driven Engineering for Real-time and Embedded systems

Participants: Charles André, Julien Deantoni, Frédéric Mallet, Marie-Agnès Peraldi-Frati, Robert de Simone, Yves Sorel.

In the embedded and real-time domain, the behavior of the system is a first-class concern that should enable validations and verifications. However, there are few tentatives to formally express the behavioral semantics of such models. Most of the time, Models bring the syntax while the related (application-specific) analysis tools bring their own behavioral semantics. This forbids a good interoperability between analysis tools and complexifies the understanding of a model behavioral semantics.

Lots of improvements are expected on the research field. To cite a few, the description of the semantics directly at the metamodel level is expected as well as a better definition of the link between the clock instants and the model events.

We promote a model-driven engineering approach for embedded system design based on formal semantics, models and methods. The range of models consists mainly of hierarchical state diagrams and dataflow/activity diagrams for behavior, and component diagrams with connection ports for compositional structure. This brought to light the idea of using the OMG UML formalisms for graphical representation of models, as it contains all these modeling views. We consider these kinds of models as purely syntactic.

In order to provide a formal and explicit behavioral semantics for such models, the idea developped this year was to use the notion of functional time as a first-class concern. Considering functional time makes it possible to specify events and relation between events of the model. While still based on the MARTE Time Model, the relation between the events of the model are now encapsulated in a semantic model.

The semantic model specifies constraints to be respected for the behavioral correctness of a model. It can be considered as the behavioral specification of the model in the same way than an OCL model specifies constraints to be respected for the static correctness of a model. The semantic model is expressed in CCSL and can be input to the TIMESQUARE prototype tool (Section 5.1), to simulate and animate the model according to its formal behavioral semantics.

Then common Models of Computation and Communication (MoCCs) can be built on top of these constructs, to be used directly by the final user. The advanced profile features are aimed primarily at advanced designers and semanticians willing to devise accurate time patterns.

Following the AAA (*Algorithm-Architecture Adequation*) methodology [56], MARTE promotes independent modeling of *applications* called *algorithm* and *embedded platforms* called *architecture* in a first step. The mapping (spatial and temporal) of applications onto embedded platforms is realized only in a subsequent step, through distributed and real-time scheduling analysis and optimizations, relative to the timing constraints and resource costs involved.

MARTE was started as a joint action of Thales, CEA-List and INRIA in their CARROLL collaborative program. The profile RFP (Request For Proposals) was voted early 2005, the initial submission in June 2007, and the (first complete) revised version in middle 2008. After the Finalization Task Force phase (2008–2009), the UML profile MARTE was adopted in November 2009 [37].

3.3. Modeling standards in embedded system design

Participants: Charles André, Frédéric Mallet, Marie-Agnès Peraldi-Frati, Aamir Mehmood Khan, Jean-François Le Tallec, Julien Deantoni, Robert de Simone.

The field of model-driven engineering of hardware/software embedded systems is hosting a number of ad-hoc standards dedicated to specific domains. These standards are instrumental in shaping up the technical and economic activities of model exchange between various industrial and academic partners. The main ones considered in our team are:

AADL for avionic systems;

AutoSar for automotive systems;

IP-Xact for System-on-Chip design;

SystemC, Esterel as Electronic System-Level modeling and programing languages.

These standards may be helpful in performing a number of analyses, such as early component integration, performance/schedulability analysis, and so forth. We conducted a number of comparative studies establishing how generic and specific concepts embodied in these standards could be reflected in MARTE, thereby allowing model transformations and exchanges, in a domain-agnostic fashion.

4. Application Domains

4.1. Multicore System-on-Chip design

Synchronous formalisms and GALS or multiclock extensions are natural model representations of hardware circuits at various abstraction levels. They may compete with HDLs (Hardware Description Languages) at RTL and even TLM levels. The main originality of languages built upon these models is to be based on formal *synthesis* semantics, rather than mere simulation forms.

The flexibility in formal Models of Computation and Communication allows to specify modular Latency-Insensitive Designs, where the interconnect structure is built up and optimized around existing IP components, respecting some mandatory computation and communication latencies prescribed by the system architect. This allows a real platform view development, with component reuse and timing-closure analysis. The design and optimization of interconnect fabric around IP blocks transform at modeling level an (untimed) asynchronous versions into a (scheduled) multiclock timed one.

Also, Network on Chip design may call for computable switching patterns, just like computable scheduling patterns were used in (predictable) Latency-Insensitive Design. Here again formal models, such as Cyclo-static dataflow graphs and extended Kahn networks with explicit routing schemes, are modeling elements of choice for a real synthesis/optimization approach to the desig of systems.

Multicore embedded architecture platform may be represented as Marte UML component diagrams. The semantics of concurrent applications may also be represented as Marte behavior diagrams embodying precise MoCCs. Optimized compilation/synthesis rely on specific algorithms, and are represented as model transformations and allocation (of application onto architecture).

Our current work aims thus primarily at providing Theoretical Computer Science foundations to this domain of multicore embedded SoCs, with possibly efficient application in modeling, analysis and compilation wherever possible due to some natural assumptions. We also deal with a comparative view of Esterel and SystemC TLM for more practical modeling, and the relation between the Spirit IP-Xact interface standard in SoC domain with its Marte counterpart.

4.2. Automotive and avionic embedded systems

Model-Driven Engineering is progressively pertaining to these fields. The formalisms AADL (for avionics) and AutoSar [53] are providing support for this, unfortunately not always with a clean and formal semantics. Yet, some interesting issues are involved there in the mix of event-triggered and time-triggered processing means, the various related protocols, and the coexistence of periodic and aperiodic tasks, with distinct periodicity if ever. The process of scheduling and allocation of multiple heterogeneous and communicating applications onto complex embedded architectural paltforms require adequate model and synthesis/analysis/verification techniques to help the designers converge to acceptable solutions.

5. Software

5.1. TimeSquare

Participants: Charles André [correspondant], Julien Deantoni, Benoît Ferrero, Frédéric Mallet.

TimeSquare is a software environment for modeling and analyzing timed systems. It supports an implementation of the Time Model introduced in the MARTE UML profile (see section 3.2), and its companion Clock Constraint Specification Language (CCSL).

TimeSquare has four main functionalities:

1. interactive clock-based specifications, through dialog boxes,
2. definition/modeling of user-defined clock constraint libraries,
3. simulation and generation of a consistent trace model, using a Boolean solver,
4. attaching call-backs to the trace model to produce domain-specific feedbacks: animation of models, displaying and exploring waveforms, generation of sequence diagrams...

TimeSquare is a plug-in developed with the Eclipse Modeling Tools. It is integrated in the OpenEmbeDD platform and can be downloaded from the team site (http://www-sop.inria.fr/aoste/dev/time_square).. This software is registered by the *Agence pour la Protection des Programmes*, under the number IDDN.FR.001.170007.000.S.P.2009.001.10600, since February 11, 2009.

5.2. K-Passa

Participants: Julien Boucaron [correspondant], Anthony Coadou, Robert de Simone.

This software is dedicated to the simulation, analysis, and static scheduling scheduling of Event/Marked Graphs, SDF and KRG extensions. A graphical interface allows to edit the Process Networks and their time annotations (*latency, ...*). Symbolic simulation allows to statically schedule such graph. Analytic methods allow to compute additional buffers needed to reach maximum achievable throughput of the graph. They can compute also part of the graph that can be slow-down through addition of both integer and fractional latencies. Such extra latencies are used to alter the static-schedule to minimize for instance dynamic power peak. In the case of KRG the (ultimately k-periodic) routing patterns can also be input and transformed. KPASSA can import/export specific UML and IPXACT models compliant with TimeSquare.

K-PASSA currently relies in part on the BOOST GRAPH library for graph algorithms and Qt4 for the GUI. It also uses LP_SOLVE as its underlying integer linear solver, GNU MP for multiprecision arithmetic and Xerces for XML parsing.

This software was developed as a result of researches on Latency-Insensitive Design conducted in the context of the CIM PACA initiative, with the support of industrial partners providing motivations.

KPASSA can be downloaded on AOSTE website. This software is registered by the Agence pour la Protection des Programmes, under the number IDDN.FR.001.310003.000.S.P.2009.000.20700.

5.3. SynDEx

Participants: Maxence Guesdon, Omar Kermia, Yves Sorel [correspondant], Cécile Stentzel.

SynDEx is a system level CAD software implementing the AAA methodology for rapid prototyping and for optimizing distributed real-time embedded applications. Developed in OCAML it can be downloaded free of charge, under the INRIA copyright, at the url: http://www.syndex.org.

The AAA methodology requires the specification of three main ingredients: an application algorithm, an architectural platform, and real-time features or requirements regarding their combination. Given these, SYNDEX will explore the space of possible allocations (distribution and scheduling) from application elements to architecture resources and services to match the real-time requirements, using schedulability analyses and heuristic techniques. It will generate automatically distributed real-time code running on the embedded platform. The last major release of SYNDEX (V7) allows the specification of multi-periodic applications.

Application algorithms can be edited graphically as directed acyclic task graphs (DAG) where each edge represent a data dependence between tasks, or they may be obtained by translation from various sources, such as (formal) synchronous reactive programs for example http://www.scicos.org, http://www.irisa.fr/espresso/Polychrony, and UML2/MARTE models http://www.omg.org/technology/documents/profile_catalog.htm.

Architectures are represented as graphical block diagrams composed of programmable (processors) and non-programmable (ASIC, FPGA) computing components, interconnected by communication media (shared memories, links and busses for message passing). In order to deal with heterogeneous architectures it may feature several components of the same kind but with different characteristics.

Two types of non-functional properties can be specified for each task of the algorithm graph. First, a period that does not depend on the hardware architecture. Second, real-time features that depend on the different types of hardware components, ranging amongst *execution and data transfer time, memory, etc.*. Requirements are generally constraints on deadline equal to period, latency between any pair of tasks in the algorithm graph, dependence between tasks, etc.

Exploration of alternative allocations of the algorithm onto the architecture may be performed manually and/or automatically. The latter possibility is achieved by performing real-time multiprocessor shedulability analyses and optimization heuristics based on the minimization of temporal or resource criteria. For example while satisfying deadlines and latencies constraints they can minimize the total execution time (makespan) of the application onto the given architecture, as well as the amount of memory.

The results of each exploration is visualized as timing diagrams simulating the distributed real-time implementation.

Finally, implementation deployments producing the embedded code use dedicated distributed real-time executives, or general purpose real-time operating systems such as Linux/RTAI or Osek for instance. These executives are deadlock-free, based on off-line scheduling policies. Dedicated executives induce minimal overhead, and are built from processor-dependent executive kernels. Presently, executives kernels are provided for: TMS320C40, PIC18F2680, i80386, MC68332, MPC555, i80C196 and Unix/Linux workstations. Executive kernels for other processors can be ported at reasonable cost following these patterns.

5.4. SAS

Participants: Patrick Meumeu Yomsi, Daniel de Rauglaudre [correspondant], Yves Sorel.

The SAS (Simulation and Analysis of Scheduling) software implementation has started last year. It allows the user to perform the schedulability analysis of periodic otask systems [12] in the monoprocessor case. As we have shown that any periodic task is a particular periodic otask [26], SAS thus also allows the user to perform the schedulability analysis of task systems in the monoprocessor case [25], [14]. The main contribution, compared to other commercial and academic softwares of the same kind, is that SAS takes into account the exact preemption cost during the schedulability analysis. Now, beside the usual real-time constraints (precedence, strict periodicity, latency, etc.) and fixed-priority scheduling policies (Rate Monotonic, Deadline Monotonic, Audsley^{++}, User priorities) that could already be taken into account, SAS has been extended to make it possible to select dynamic scheduling policy algorithms such as Earliest Deadline First (EDF).

The resulting schedule is displayed as a typical Gantt chart with a transient and a permanent phase, or as a disk named "dameid" that clearly shows the idle slots of the processor in the permanent phase. For a schedulable task system under EDF when the exact preemtion cost is considered, the period of the permanent phase may be much longer than the classical one least commun multiple (LCM) of the periods of all tasks in the traditional scheduling theory. A specific effort has been done in SAS in order to improve the display in this case.

The classical utilization factor, the permanent exact utilization factor, the preemption cost in the permanent phase, and the worst response time for each task are displayed when the system is schedulable. In addition, another graphic, showing the response times of each task relative time, can be displayed.

The software is written in OCAML, using CAMLP5 (syntactic preprocessor) and OLIBRT (a graphic toolkit under X). Both are written by Daniel de Rauglaudre.

6. New Results

6.1. K-periodically Routed extended Event Graphs

Participants: Julien Boucaron, Anthony Coadou, Robert de Simone.

The Process Network models based on Event Graphs and varying synchronous or asynchronous interpretations allow powerful results in static scheduling and distribution allocations. Nevertheless they always postulate a uniform data flow. We tried to relax this strong assumption, while preserving the determinism/confluence of computations, as in Kahn Process Networks. Several models were inspirational (such as Boolean and CycloStatic DataFlow Graphs). But the main originality of our KRGs is to rely on two *Select/Merge* operator nodes only (for *mux/demux* effects on data streams). And, *most importantly* the switching patterns for these conditional nodes have to be infinitary periodic binary words, thus using exactly the same description formalisms as our previous *schedule words*.

We have proven a number of powerful algebraic results on such models. They may best be understood by analogy with Boolean Algebra, and the existence of normal forms (such as sum of products, or product of sums). Here, the expansion and factorization of expressions and variables amounts to sharing and unsharing of links and channels in the interconnect fabric representation of the communications across the networks.

Our main results during this year are submitted for publication, and will be published in book chpater form. We believe this model of KRGs to be an important step for the Theoretical Computer Science modeling of modern Networks-on-Chip *NoCs*.

Meanwhile we have also studied further issues of throughput and buffer dimensioning in the context of Latency-Independent Design (LID) [35], [34].

6.2. CCSL: syntax and semantics

Participants: Charles André, Julien Deantoni, Frédéric Mallet, Robert de Simone.

CCSL has been introduced in a non-normative annex of the OMG MARTE specification [37]. It is a language to specify *clock constraints*. The semantics of CCSL given in the specification is informal. A first formal semantics, based on mathematical expressions was published last year.

This year, the expressiveness of CCSL has been compared to two other concurrent models (Signal and Time Petri nets) [21]. Time Petri nets are well-adapted to specify asynchronous clock constraints but cannot deal with reactive aspects like preemption, and reaction to the absence of events. Signal (developed by the Espresso EPI) can easily express synchronous clock constraints and it addresses, like CCSL, polychronous systems. However, while Signal provides a minimal set of operators and focuses on sufficient conditions for deterministic execution (endochronous systems), CCSL proposes high-level time patterns to express time specifications (not necessarily deterministic) closer to UML designer expectations.

To provide a tool support for the analysis of CCSL specifications, we have defined a kernel for CCSL, and given a structural operational semantics to this kernel. The syntax and the semantics of this kernel are described in a research report [32] and have been partially published in the *Journal Européen des Systèmes Automatisés* [13]. This semantics is the reference for the CCSL constraint solver implemented in TimeSquare (Section 5.1).

TimeSquare allows simulation of CCSL specifications but not formal verifications. Relying on the formal semantics of CCSL, we have proposed a methodology [18] to check with formal verification tools that a synchronous implementation conforms to a CCSL specification. An experimentation, using the formal checkers available in Esterel Studio, has been conducted on an Esterel implementation by defining a library of Esterel observers for CCSL constraints.

6.3. UML profiles

Participants: Charles André, Frédéric Mallet, Robert de Simone.

The Finalization Task Force (FTF) of the UML profile for MARTE has issued its report in May 2009. Aoste actively contributes to MARTE and SysML standardizations and Frédéric Mallet was a voting member of the MARTE FTF and SysML RTF. Frédéric Mallet and Charles André have proposed resolutions of issues related to the Time and Allocation chapters of MARTE. The UML profile MARTE has been adopted by the OMG in November 2009 [37]. MARTE and its Time model of MARTE were presented at *Neptune 2009* (Paris, May 2009) and published in the journal of *Génie Logiciel* [16], [17].

Besides this standardization effort, Aoste has proposed other profiles following MDE approaches. A first proposal concerns the *multilevel modeling* (i.e., not restricted to the Class-Instance relationship). A profile, called *DomainSpecification*, has been defined and allows automated construction of domain specific UML profiles. This study extends preliminary results proposed in the François Lagarde's thesis (November 2008) and was presented at the Conference on Software Language Engineering [22]. The second contribution is a UML profile for IP-Xact, a standard of the Electronic Design Automation (EDA) domain (see section 6.5).

6.4. A semantic model for domain-specific languages

Participants: Charles André, Julien Deantoni, Frédéric Mallet, Robert de Simone.

This year, we started an ambitious work focused on the explicitation of a model behavioral semantics. We proposed a general model-based framework to describe a formal and explicit behavioral semantics as a separate model [38]. This model is based on an abstract syntax of CCSL (Section 6.2). While CCSL specifies explicitly the logical relations between clock instants, the model links the clock to events in a specific model. The association between clocks and model events is for now in its very crude form but the prototype tool [39] integrated in TimeSquare opens promising perspectives.

Specifically addressing the avionics domain 4.2, we have shown [20] how to use CCSL to build a semantic model for AADL, the Architecture and Analysis Description Language adopted by the Society of Automotive Engineers (SAE). UML structured classifiers and activities are used to model AADL application software elements and execution platforms. MARTE allocation replaces the AADL binding mechanism. CCSL clock constraints give an explicit timed execution semantics. TimeSquare (Section 5.1) then interprets the operational semantics (see Section 6.2) of the CCSL constraints to animate the UML/MARTE model.

6.5. Interoperability and system-level validation of IP component blocks

Participants: Charles André, Benoît Ferrero, Aamir Mehmood Khan, Frédéric Mallet, Robert de Simone.

In the context of the project Sys2RTL of the CIMPACA platform *Conception* and of the project ID-TLM in collaboration with ST, we are investigating ways to use UML as a modeling framework for the design and integration of Hardware and Software IPs. For interoperability of IPs, the Spirit Consortium promotes the IP-Xact format. IP-Xact only captures the interface of IPs and relies on Hardware Description Languages for the description of the behavior: SystemC for TLM descriptions and VHDL/Verilog for RTL descriptions. To build a virtual platform for System-On-Chips that allows the early analysis and verification of systems, an abstract specification of the IP behavior is required. We propose to use UML as a graphical front-end to capture both structural and behavioral aspects. The UML Profile for MARTE is used to specify the time aspects and capture non-functional properties. Following a Model-Driven Approach, we have developed transformation models in ATL to transform UML models to IP-Xact specifications back and forth. The transformation requires the annotation of UML/Marte models with IP-Xact specific stereotypes. These stereotypes have been gathered in a UML Profile for IP-Xact. Beyond the transformation of structural elements, we are also considering behavioral and non-functional aspects. The CCSL specification provides a golden model, at the timed communicating process level, and against which implementations at different levels must be checked. As a first step, we have built a library of VHDL observers to validate RTL implementations [24]. We are also building a library of SystemC observers that should be available soon to verify the conformance of implementations at the transaction level.

6.6. Virtual Platform modeling

Participants: Charles André, Julien Deantoni, Jean-François Le Tallec, Robert de Simone.

In the context of the PhD thesis of Jean-François Le Tallec, funded under the CIM PACA programme jointly with Scaleo Chip, a local SME company, we studied the system-level representation of full embedded platform at virtual levels. Case studies were provided by former such platforms, but not virtual and lower-level, previously implemented by the industrial partner. Models of a dedicated flash memory controller were realized in several fashions (Verilog VHDL, SystemC, Esterel), and then inserted in prototype platform models. The relevance of specific bus protocols (AMBA mainly), and the difficulties encountered with current proposals for such modeling (IP-XACT standards, existing environments such as the Synopsys tools available in CIM PACA), were instrumental in the definition of new approaches that use the formalisms of MARTE Time model and CCSL, to provide abstract yet precise semantics to the virtual components.

Preliminary results were presented in [41], in which systematic connection rules are proposed.

6.7. Executable Timed Requirements for automotive applications

Participants: Charles André, Frédéric Mallet, Marie-Agnès Peraldi-Frati.

The UML profile for East-ADL2 was defined by the ATESST project and provides a support to define requirements for automotive applications. It is currently integrated in the first phases of the AutoSAR methodology promoted by the AutoSAR consortium; AutoSAR covers the whole design flow whereas East-ADL2 only considers the requirement aspects.

In the context of the project MeMVaTEx (Section 8.2.2), we have defined a CCSL relation definition library for East-ADL2 Timing Requirements [23]. It contains four relation definitions, one for each East-ADL2 Timing Requirement: repetition rate, delay requirement, input/output synchronizations.

Another contribution [15] addresses a larger scope and focuses on methological aspects including traceability and validation of requirements. The knock control of 4-stroke engines is studied as an example. The methodology combines UML, MARTE and SysML. This was a first step of our effort to ensure the convergence of MARTE and SysML, which is one of the major problem addressed in the Lambda project (Section 8.2.5), and that we conduct within the MARTE revision task force at the OMG (Section 6.3).

6.8. Monoprocessor real-time scheduling

Participants: Liliana Cucu-Grosjean, Laurent George, Patrick Meumeu Yomsi, Daniel de Rauglaudre, Yves Sorel.

Last year, we introduced a new model, called the *otask model*, in order to solve the general scheduling problem of hard real-time systems with various kinds of constraints. This new model allowed us to perform the schedulability analysis of a task set while taking the exact RTOS cost into account for any scenario of first release for all tasks and when priorities are given according to a fixed-priority scheduling policy. In this case, the schedulability analysis is based on an algebraic approach through the definition of a binary scheduling operation denoted by $\oplus$ which is associative and not commutative. For a given set of otasks, operation $\oplus$ is used as many times as there are tasks in the system.

This year, we have closely studied the impact of the scenario of first release times of all tasks [25], [27], [14], then the impact of the priority assignment and finally the impact of the scheduler cost [26], all on both the schedule and the schedulability analysis for a given system [12]. First, when tasks are scheduled according to any fixed-priority scheduling policy, we have proposed an algorithm to improve the sensitivity of the deadlines and then the quality of control of the system by using our approach for the computation of the worst response times. Here, the worst response time of a task is the maximum time elapsed between the activation times and the completion times of the task amongst all its instances. Second, because a dynamic-priority scheduling policy may produce a valid schedule for a system that can not be schedulable by any fixed-priority scheduling policy, we have extended our approach to tackle the scheduling problem where priorities are assigned to tasks (resp. otasks) according to a dynamic-priority scheduling policy such as Earliest Deadline First (EDF). This has been done while taking real-time constraints such precedence, strict periodicity and latency into account. As our main goal is to build a predictable RTOS, we have performed an extensive state of the art work on the mechanisms which are used to implement a scheduler. While doing so, we have come up with two main types of schedulers: the *timer-driven schedulers* and the *event-driven schedulers*. We have given the advantages and disadvantages of each type of schedulers.

Meanwhile, we continued our work on multiple latency constraints. We provided a schedulability test for dependant periodic task systems scheduled using non-preemptive policies, that must satisfy latency constraints [19].

6.9. Multiprocessor real-time scheduling

Participants: Omar Kermia, Mohamed Marouf, Yves Sorel.

We performed a lot of algorithmic and experimental work on the different theoretical results proposed last years. In particular we tested the proposed multiprocessor real-time schedulability conditions for sets of non preemptive tasks with deadline equal to period, precedence, and multiple latency constraints. Using these schedulability conditions, we proposed a back-tracking optimization heuristic which, in addition to verify the schedulability of tasks, minimizes on the one hand the total execution time (makespan) of the set of tasks running onto a multiprocessor architecture, and on the other hand the amount of memory. This heuristic which obviously gives sub-optimal results was compared to an exact algorithm giving an optimal result. All the theoretical results as well as the algorithmic and experimental results are described in [11]. Furthermore, they have been exploited to develop in Ocaml a new major version of SynDEx (V7) which allows the designer to specify now multiperiod applications.

Since distributed architectures are prone to failures we started a new work on fault tolerance for multiperiodic set of tasks running on a multiprocessor. This work is a follow-up of H. Kalla's former PhD thesis, which handled only monoperiodic set of tasks and simple model of fault (fail silent, transient), without sensor and actuator faults. A new PhD thesis started at the beginning of the year in collaboration with the IMARA team which wants to develop automatic control applications running on the various CyberCars platforms they develop, while providing fault tolerance. In order to tackle these complex embedded applications, we plan to study for multiperiodic set of tasks intermittent and Byzantine faults not only for the processors and the communication media, but also for the sensors and the actuators.

6.10. Clock-driven real-time implementation of synchronous specifications

Participants: Dumitru Potop-Butucaru, Robert de Simone, Yves Sorel, Jean-Pierre Talpin [EPI ESPRESSO].

One important line of work in our project concerns the model-based mapping of the computations and communications of the functional specification onto corresponding resources of the implementation architecture. This mapping comprises both temporal scheduling and spatial allocation aspects. Therefore, we promote an approach which starts from loosely-timed/asynchronous models and proceeds by refining them to fully synchronized ones, using so-called clock calculus techniques under the architecture constraints.

This year, we provided a modeling framework [28] based on an intermediate representation format, called clocked graphs, for polychronous endochronous specifications, which are the ones that can be safely considered for deterministic distributed real-time implementation using static scheduling techniques. Our formalism allows the specification of both "intrinsic" correctness properties of the specification, such as causality and clock consistency, and "external" correctness properties, such as endochrony, which ensure compatibility with the desired implementation architecture, including both hardware and software aspects. Using this formalism, we define a new method for distributed real-time implementation of synchronous specification, where the move from (endochronous) synchronous specification to realtime scheduled implementation is a seamless sequence of model decorations.

When compared with current state-of-the-art, represented by the AAA/SynDEx methodology, this new approach has the advantage of providing a seamless transformation all the way from specification to implementation models. Our approach also has the advantage of promoting activation conditions (known as clocks) as first class citizens, which is not presently the case for SynDEx, where they can be defined only through structured dataflow constructs (which often results in pessimization of both the specification and the implementation).

6.11

6.11. From Concurrent Multiclock Programs to Deterministic Asynchronous Implementations

Participants: Dumitru Potop-Butucaru, Robert de Simone, Yves Sorel, Jean-Pierre Talpin (EPI ESPRESSO).

Current techniques for the compilation of multi-clock synchronous programs often produce implementations that are over-synchronized. In such implementations, all the clocks (activation conditions) are forced to derive from a single base clock to allow a simple, hierarchical code generation. This approach is well-suited when the target is a sequential processor. For distributed implementations, however, it results in unnecessary inter-processor synchronizations with may result in important performance losses.

We proposed this year a general method to characterize and synthesize correctness-preserving, asynchronous wrappers for synchronous processes on a globally asynchronous locally synchronous (GALS) architecture. Our technique is mathematically founded on the theory of weakly endochronous systems, due to Potop, Caillaud, and Benveniste. Weak endochrony gives a compositional sufficient condition establishing that a concurrent synchronous specification exhibits no behavior where information on the absence of an event is needed. Thus, the synchronous specification can safely be executed with identical results in any asynchronous environment (where absence cannot be sensed). Weak endochrony thus gives a latency-insensitivity and scheduling-independence criterion.

We defined the first general method to check weak endochrony on multi-clock synchronous programs. The method is based on the construction of so-called generator sets. Generator sets contain minimal synchronization patterns that characterize all possible reactions of a multi-clocked program. These sets are used to check that a specification is indeed weakly endochronous, in which case they can be used to generate the GALS wrapper. In case the specification is not weakly endochronous, the generators can be used to generate intuitive error messages. Thus, we provide an alternative to classical compilation schemes for multi-clock programs, such as the clock hierarchization techniques used in Signal/Polychrony.

We are currently working on the application of our technique in the compilation of the Signal language, and on the generation of simpler communication protocols in the SynDEx tool.

7. Contracts and Grants with Industry

7.1. ID/TL-M project with ST MIcroelectronics

Participants: Charles André, Julien Boucaron, Robert de Simone, Benoît Ferrero, Aamir Mehmood Khan.

ID/TL-M is a project launched as part of the larger NANO2012 programme conducted by ST Microlectronics in Rhône-Alpes. Its main goal is to study the potential use of model-driven engineering techniques (MDE) for Electronic System-Level Design (ESL) of Systems-on-Chip (SoC). While SystemC is a de-facto standard in this domain, it suffers a number of lacks. One is the absence of clear formal semantics (unlike Esterel), which largely proscribes high-level synthesis; the second is the absence, at the transaction-level modeling (TLM) level, of a clear associated interface, as an Architecture Description Language (ADL), which could support annotations for extra non-functional properties.

The ongoing standard IP-XACT is a candidate for being such an ADL but, being developed by a consortium of industrial partners with too low-level intents, it does not fit with all our aim perfectly. As a result we studied the specialization of MARTE to encompass most existing features of IP-XACT, while being truly extensible to non-functional properties with genericity. As a by-product, formal semantic definition using our CCSL language and MARTE Time model are becoming feasible, with synchronous artefacts borrrowed from Esterel as guidelines.

This year we have provided the specification of translations in both directions between IP-XACT 1.5 and a proper subset of MARTE. Links to formal Models of Computation and Communication (MoCCs) have been drawn, andwe expect to implement them soon to allow interpretation and analysis using our K-PASSA tool. Logical time multiclock semantics has also been studied, which resulted in the specification of interpretation amenable to TIME-SQUARE. The PAPYRUS UML modeler by CEA, and ECLIPSE environments shall be considered to support the forthcoming implementations.

We held two meetings in Grenoble this year. Both partners (INRIA Aoste and STM STG) are members of the recently started ANR project HELP, which adds the low-power and thermal dimensions to the range of non-functional properties.

7.2. Thales TRT

Participants: Patrick Meumeu Yomsi, Yves Sorel.

Last year, we started a collaboration with the Thales Research and Technology group on the design of new multi-core architectures dedicated to critical embedded systems. In order to achieved this, we performed a first study on predictable RTOS that are needed for such systems.

This year, we continued this work along three main axes. First, after an extensive state of the art work on the existing RTOS, we have set up the foundations of the design of an RTOS which would allow us to incorporate its exact cost into the schedulability conditions. Second, we have shown the impact of the RTOS cost on both the schedule and the schedulability analysis of a system. Finally, we have focused our attention on the implementation of the scheduler which is a central component of the RTOS. We came to a first structure for the algorithm implementing such a scheduler in the case of a simple processor.

8. Other Grants and Activities

8.1. Regional collaborations

8.1.1. CIM PACA

Participants: Jean-François Le Tallec, Julien Boucaron, Aamir Mehmood Khan, Robert de Simone.

This ambitious regional initiative is intended to foster collaborations between local PACA industry and academia partners on the topics of microelectronic design, though mutualization of equipments, resources and R&D concerns. We are actively participating in the Design Platform (one of the three platforms launched in this context). Other participants are UNS, CNRS (I3S and LEAT laboratories), L2MP Marseille, CMP-ENSE Gardanne on the academic side, and Texas Instruments, Philips, ST Microelectronics, ATMEL, and Esterel-EDA on the industrial side.

Inside this platform we are coordinating a dedicated project, named Spec2RTL, on methodological flows for high-level SoC synthesis. Participants are Texas Instruments, NXP, ST Microelectronics, Synopsys, Esterel-EDA, and Scaleo Chip as industrial partners, INRIA, I3S (CNRS/UNSA) and ENST on the academic side. A pool of PhD students are funded on a par basis between industrial partners and local PACA PhD grants under the BDI programme.

Jean-François Le Tallec started his PhD thesis in connection with Scaleo Chip, a local SME company developing SoC platform simulators. The PhD topic is to investigate new *virtual* platform environments at ESL TLM level, and their relation to formal modeling in multiclock ESTEREL.

Sadly this year we had to cope with large lay-offs inside our Texas Instruments partner, which indirectly contributed to the shut-down of the Esterel-EDA INRIA spin-off company. All this was due to reasons fully external to the project, but impacted it negatively. The Esterel Studio compiler has recently been bought by the Synfora company, with which we are seeking further collaborations.

In the prospects of the upcoming second phase of the programme, named CIM PACA 2011, we are considering possible future extensions leading to a better and more realistic use of Virtual Hardware Platforms (VHP), In this direction we seriously considered the Synopsys tools, which were acquired as mutualized softwares by the Design Platform: INNOVATOR, SYSTEM STUDIO, COREASSEMBLER, COREBUILDER. Other similar tools are also under evaluation.

8.2. Nation-wide collaborations

8.2.1. RNTL platform OpenEmbeDD

Participants: Charles André, Benoît Ferrero, Fadoi Lakhal, Robert de Simone, Yves Sorel.

OpenEmbeDD was a large platform project aimed at connecting several formalisms with model-driven engineering tools, in the embedded domain. Its focus was on the use of model-driven approaches to combine various specification formalisms, analysis and modeling techniques, into an interoperable framework. Partners were: INRIA, CEA-List, Thales, Airbus, France Telecom, CS, LAAS, and VERIMAG. Four INRIA teams were involved (ATLAS, Triskell, Aoste and DaRT).

The project was concluded at mid-year, and the final version of the integrated platform by delivered by then on INRIA's forge. It contains amongst other things our contributions on MARTE Time model and CCSL, with its TIMESQUARE implementation. It also contains model-based translations from OPENEMBEDDD metamodel to and from SYNDEX, and various connections to POLYCHRONY/SME (from the Espresso EPI) that are also related to our work.

The outcome of this platform are intended to be one of the primary inputs to the large ARTEMIS CESAR project.

8.2.2. RNTL project MeMVaTEx

Participants: Marie-Agnès Peraldi-Frati, Dumitru Potop-Butucaru, Yves Sorel.

The partners of the MeMVaTEx project are: Continental, INRIA, CEA-List, CNRS-UTC, and Sherpa Engineering. The project focuses on developing a design methodology, centered on the requirements, their traceability and their validation. The application domain is the design of complex real-time automotive systems. The methodology is based on the standards EAST-ADL2, AUTOSAR, SYSML and MARTE. The last phase of the project is currently centered on the integration of heterogeneous model and tools for validation and verification of requirements. This phase integrates the Simulink and SynDEx tools in order to provide the validation of requirements and models.

During this year, we have been working on the integration of the SynDEx tool in a timing flow based on EAST_ADL2/Autosar and MARTE. A gateway has been developed that exports temporal informations from the hardware and software models to a SynDEx model. SynDEx is used as a test cases verificator linked to the requirements. We worked also on the MeMVaTEx demonstrator for illustrating the previous results.

We held the project final meeting in September 2009, and presented its results in a journal article [15], and to the ANR yearly symposium.

8.2.3. ANR RT-Simex

Participants: Frédéric Mallet, Julien Deantoni.

This new project is dedicated to the reverse engineering of analysis traces of simulation and execution back up to the source code, or in our case most likely into the original models in a MARTE profile representation. The prime contractor is the OBEO company. We held a mid-year meeting in Sophia-Antipolis.

8.2.4. ANR HeLP

Participants: Julien Boucaron, Dumitru Potop-Butucaru, Robert de Simone.

HELP is a recently accepted ANR project, with ST MicroElectronics, Verimag, LEAT, and Docea Power (a SME from Grenoble) as other partners. The project name stands for "High-Level/Low-Power" design and analysis.

There are two main focuses inside the project. The first one deals with the comparison of efficient System-Level modeling, between multiple logical clocks originated from synchronous languages on the one hand, loosely timed and approximate-time simulation schemes of SystemC on the other hand. The second focus goes with the modeling of energy consumption and thermal levels for SoCs, their representation in a model-driven engineering approach, and their interaction with time modeling.

We held the kick-off meeting for this project in December in Grenoble. There are links between this project and the ID/TL-M collaboration with ST Microelectronics.

8.2.5. FUI Lambda

Participants: Charles André, Julien Deantoni, Robert de Simone, Frédéric Mallet.

The Lambda project is headed by Thales, with ST Microelectronics, Airbus, Esterel-EDA, CoFluent, CEA-LIst, and several other partners.

Our contribution in this project remains rather light. We bring expertise to help with the definition of a model transformation between SYNCCHARTS and UML State Diagrams. We then contribute to the combination of SYSML and MARTE artefacts, in the context of SoC design as well as the SPIRIT IP-XACT standard.

8.2.6. FUI PARSEC

Participants: Dumitru Potop-Butucaru, Yves Sorel.

The Parsec project was accepted within the 2009 FUI for projects, and is scheduled to last from 2010 to 2012. This a large project with The partners of the project are Thales, CEA, Elidiss, INRIA, Systerel, OpenWide, Alstom, and TelecomParisTech. The project aims at defining a framework for the development of distributed real-time embedded systems that are subject to strict certification standards such as DO-178B (for avionics), IEC 61508 (for transportation systems), or ISO/IEC 15408 (the Common Criteria for information technology security evaluation).

The AOSTE team will use its expertise in the modeling and distributed real-time implementation of embedded applications using synchronous formalisms and associated tools. The two main scientific challenges of the project are (1) a better modeling of the distributed implementation architectures, allowing code generation for novel architectures and better code generation for architectures we currently handle, and (2) the modeling and efficient implementation of mode changes, as they are specified in an industrial context.

8.2.7. ARC Triade – Combining models of computation for the design of real-time and embedded applications

Participants: Dumitru Potop-Butucaru, Yves Sorel, Robert de Simone.

The Triade Cooperative Research Action (ARC) is a partnership between the AOSTE, DaRT, and ESPRESSO teams of INRIA. Triade aims at using formal models with structuring programmatic constructs as means to translate programs and descriptions written in formalisms widely used in Embedded System and SoC design, and provide a seamless flow of increasingly time-defined and time-accurate models, so as to progressively obtain the final mapped implementation through provably correct steps from the early description elements.

Triade funded regular meetings with our colleagues from Rennes (ESPRESSO) and Lille (DaRT). Two publications resulted [29], [28]. We intend a common implementation initiative of these topics.

8.3. European collaborations

8.3.1. ARTEMIS CESAR

Participants: Charles André, Julien Deantoni, Régis Gascon, Frédéric Mallet, Yves Sorel, Robert de Simone.

CESAR is a large project with over 60 participants, most of which important industrial partners from avionics, railways, and automotive domains. It is currently the flagship project of the European ARTEMIS programme. It aims at defining a common integrated environment to organize collaborative use of many tools and methods relevant to embedded system design, positioned according to effective design flows validated by industrial needs and experience.

The project is split between 6 main subprojects:

SP1 is devoted to the construction of a common reference technological platform (RTP), on which all tools and models are to be plugged. It borrows partly on experience gained in the former OPENEMBEDD RNTL platform;

SP2 deals with requirements engineering and the issues of consistency when conducting large development projects inside multiple organizations (evolving product lines, contracts, non-functional property requirements, and so on);

SP3 deals with individual tools and methods, both from classical approaches and from innovative academic sources that are found useful to support prospective design flows and thus be connected onto the RTP. EPI Aoste is most specifically leading this part for INRIA;

SP5/6/7 are devoted to industrial domains, where industrial partners provide use cases and potential scenarios, showing needs (and evaluating solutions in a latter phase).

This year we attended the kick-off meeting in Vienna in March, then two technical SP3 meetings in Stockholm and Toulouse, and finally a plenary meeting in Torino. This project is funding the postdoctoral position of Régis Gascon.

8.3.2. ITEA OPENPROD

Participants: Dumitru Potop-Butucaru, Yves Sorel.

The OPENPROD project was accepted within the 2008 ITEA call for projects, and is scheduled to last from 2009 to 2011. The partners of the project are Bosh, Siemens, SKF, Nokia, IFP, EDF, PSA, EADS, LMS Imagine, VTT, CEA, Fraunhofer, etc. The project aims at providing an open whole-product model-driven rapid systems development, modeling, and simulation environment integrating the leading open industrial software development platform (Eclipse) with open-source tools (OpenModelica, etc.), and industrial tools and applications.

The AOSTE team participates to the workpackage devoted to code generation and run-times issues. More precisely it adresses the automatic translation of the discrete part of high-level specifications written in Modelica into a format able to be implemented onto a distributed architecture while satisfying real-time constraints. In addition, models translated into this format should be implemented onto multi-core architectures. IFP will provide a test case based on a engine control models intended to be implemented onto their real-time co-simulation platform running on a multi-core.

8.3.3. IST Network of Excellence ARTIST2 & ARTIST-Design

Participants: Julien Boucaron, Robert de Simone, Frédéric Mallet.

We attended this year several events sponsored by the ARTISTDESIGN Network of Excellence, including he perioidcal review meeting and its preparation workshop in Brussels. Other such events were the Synchron seminar in Dagstuhl, in late november, which brings together the whole community of Synchronous languages, and the EMSOFT conference, itself surrounded by a numbers of satellite events gathering the main european leaders on Embedded System designs for brainstorming informal meetings.

9. Dissemination

9.1. Leadership within scientific community

- Robert de Simone was programme committee member for MemoCode'09, FDL'09 and EmSoft'09. He was on the Selection Board of experts for the ANR programme ARPEGE 2009. He is one of the two INRIA representative to the University of Nice/Sophia-Antipolis (UNS) Doctoral School Counsil. He represents INRIA in the CA board of ARCSIS, the ruling association for the CIM PACA, as well being a member of its Strategic Council. He holds similar positions in CIM PACA Design Platform non-profit organization. He was reviewer for the PhD theses of Yann Glouche (IRISA) and Florence Plateau (LRI), and jury member for the thesis of Farooq Muhammad (LEAT).

- Charles André was jury member for the PhD thesis of Charlotte Seidner (IrCCyN).
- Dumitru Potop-Butucaru was program committee member for the FMGALS'09 workshop on globally asynchronous, locally synchronous systems. He is program committee member of ACSD 2010. Within INRIA, he is a member of the scientific recruitement committee *Détachements, délégations, post-doctorants*.
- Yves Sorel is programme committee member of the following conferences: RTSS, RNTS, DATE, DASIP, EUSIPCO, GRESTSI. He is member of the OCDS/SYSTEM@TIC Paris-Region Cluster Steering Committee. He was reviewer for the PhD theses of Isabelle Perseil (ENST) and Jean-Paul Perez-Seva (UNS), and jury member for the theses of Julien Forget (ONERA), Omar Kermia (Paris 11), Patrick Meumeu Yomsi (Paris 11).
- Marie-Agnès Peraldi-Frati is member of the CNRS/I3S conseil de laboratoire and member of the CERTEC (conseil d'études et de la recherche technologique) of the IUT of Nice-Sophia Antipolis.

9.2. Teaching

- Julien Deantoni gives courses at different cursus levels of the polytechnic school of the UNS: A course and labs on micro-controller and Real-Time operating system programming in the last year of the engineering cursus. He also teaches object oriented programmation through C++ in the 4^{th} year of the mathematical and modeling cursus as well as linux shell programming in the second year of the engineering cursus.
- Robert de Simone gave a course on Formal Methods and Models for Embedded Systems in the STIC Master program of the university of Nice-Sophia Antipolis (UNS), for approximately 15h.
- Frédéric mallet gives courses at the "Université de Nice Sophia Antipolis" and the "Ecole Polytechnique Universitaire" from the first year of the Bachelor Degree (Licence) to the master. He teaches software engineering, computer architecture, object-oriented programming and model-driven engineering. He is also in charge of the Master 1 MIAGE. In July, he was invited by the Software Engineering Institute in Shanghaï to give a talk on "Model-Driven Engineering for Real-Time and Embedded systems" in its annual summer school.
- Marie-Agnès Peraldi-Frati gives courses at different cursus levels of UNS: A course and labs on UML for real-time in the TSM master (Telecommunication, System and Microelectronics) at the University of Nice. She gives different courses (Systems and networks, Programming, Web development, Computer architecture) at the L1 level of the IUT Informatique.
- Yves Sorel gives courses in the Research Master SETI (Systèmes Embarqués et Traitement de l'Information) at the University of Orsay Paris 11, in the last year cursus of the ESIEE (Engineering School located in Noisy-le-Grand), and in the last year cursus of the ENSTA (Engineering School located in Paris), on topics comprising the AAA methodology and the optimization of distributed real-time embedded systems.

10. Bibliography

Major publications by the team in recent years

[1] C. ANDRÉ, F. MALLET, R. DE SIMONE. *Modeling Time(s)*, in "MoDELS'2007 10th Intern. Conf. on Model Driven Engineering Languages and Systems", 2007.

[2] A. BENVENISTE, P. CASPI, S. EDWARDS, N. HALBWACHS, P. LE GUERNIC, R. DE SIMONE. *Synchronous Languages Twelve Years Later*, in "Proceedings of the IEEE", January 2003.

[3] L. CUCU, N. PERNET, Y. SOREL. *Periodic real-time scheduling: from deadline-based model to latency-based model*, in "Annals of Operations Research", 2007, http://www-rocq.inria.fr/syndex/publications/pubs/aor07/aor07.pdf.

[4] T. GRANDPIERRE, Y. SOREL. *From Algorithm and Architecture Specification to Automatic Generation of Distributed Real-Time Executives: a Seamless Flow of Graphs Transformations*, in "Proceedings of First ACM and IEEE International Conference on Formal Methods and Models for Codesign, MEMOCODE'03, Mont Saint-Michel, France", June 2003, http://www-rocq.inria.fr/syndex/publications/pubs/memocode03/memocode03.pdf.

[5] O. KERMIA, Y. SOREL. *A Rapid Heuristic for Scheduling Non-Preemptive Dependent Periodic Tasks onto Multiprocessor*, in "Proceedings of ISCA 20th International Conference on Parallel and Distributed Computing Systems, PDCS'07, Las Vegas, Nevada, USA", September 2007, http://www-rocq.inria.fr/syndex/publications/pubs/pdcs07/pdcs07.pdf.

[6] P. MEUMEU YOMSI, Y. SOREL. *Extending Rate Monotonic Analysis with Exact Cost of Preemptions for Hard Real-Time Systems*, in "Proceedings of 19th Euromicro Conference on Real-Time Systems, ECRTS'07, Pisa, Italy", July 2007, http://www-rocq.inria.fr/syndex/publications/pubs/ecrts07/ecrts07.pdf.

[7] D. POTOP-BUTUCARU, S. EDWARDS, G. BERRY. *Compiling Esterel*, Springer, 2007.

[8] D. POTOP-BUTUCARU, R. DE SIMONE. *Optimizations for Faster Execution of Esterel Programs*, in "MEMOCODE'03", 2003.

[9] E. VECCHIÉ, R. DE SIMONE. *Syntax-driven optimizations for Reachable State Space construction of Esterel programs*, in "International Journal of Embedded Systems", April 2005.

[10] R. DE SIMONE, D. POTOP-BUTUCARU, JEAN-PIERRE. TALPIN. *8*, in "The Synchronous Hypothesis and Synchronous Languages", CRC Press, 2005.

Year Publications

Doctoral Dissertations and Habilitation Theses

[11] O. KERMIA. *Ordonnancement temps réel multiprocesseur de tâches non préemptives avec contraintes de précédence, de périodicité stricte et de latence*, Université de Paris Sud, 2009, http://www-rocq.inria.fr/syndex/publications/pubs/theses/THOK.pdf, Spécialité Physique, Ph. D. Thesis.

[12] P. MEUMEU YOMSI. *Prise en compte du coût exact de la préemption dans l'ordonnancement temps réel monoprocesseur avec contraintes multiples*, Université de Paris Sud, 2009, http://www-rocq.inria.fr/syndex/publications/pubs/theses/THPM.pdf, Spécialité Physique, Ph. D. Thesis.

Articles in International Peer-Reviewed Journal

[13] C. ANDRÉ, F. MALLET. *Modèles de contraintes temporelles pour systèmes polychrones*, in "JESA", vol. 43, n^{o} 7–8–9, 2009, p. 725–739.

[14] P. MEUMEU YOMSI, L. GEORGE, Y. SOREL, D. DE RAUGLAUDRE. *Improving the Quality of Control of Periodic Tasks Scheduled by FP with an Asynchronous Approach*, in "International Journal on Advances in

Systems and Measurements", vol. 2, n^{o} 2, 2009, http://www-rocq.inria.fr/syndex/publications/pubs/ijasm09/ijasm09.pdf, to appear.

Articles in National Peer-Reviewed Journal

[15] A. ALBINET, H. DUBOIS, MARIE-AGNÈS. PERALDI-FRATI. *MeMVaTEx : des exigences aux modèles dans le domaine automobile*, in "Revue de l'électricité et de l'électronique", n^{o} 2, February 2009, p. 52–59.

[16] C. ANDRÉ, M. BELAUNDE, B. BERTHOMIEU, C. BRUNETTE, A. CANALS, H. GARAVEL, S. GRAF, F. LANG, V. MAHÉ, M. NAKHLÉ, R. SCHNEKENBURGER, R. DE SIMONE, J.-P. TALPIN, F. VERNADAT. *Présentation des résultats du projet OpenEmbeDD*, in "Génie Logiciel", 2009, p. 25–30, http://hal.inria.fr/inria-00381639/en/.

[17] C. ANDRÉ, F. MALLET. *Les modèles de temps de MARTE et CCSL*, in "Génie logiciel", n^{o} 89, 2009, p. 44–49, http://hal.inria.fr/inria-00416597/en/.

International Peer-Reviewed Conference/Proceedings

[18] C. ANDRÉ, F. MALLET. *Specification and Verification of Time Requirements with CCSL and Esterel*, in "Languages, Compilers, and Tools for Embedded Systems ACM SIGPLAN Notices, Irlande Dublin", C. KIRSCH, M. KANDEMIR (editors), vol. 44, ACM SIGPLAN/SIGBED, 2009, p. 167-176, http://hal.inria.fr/inria-00416654/en/.

[19] L. CUCU-GROSJEAN, Y. SOREL. *A schedulability test for real-time dependant periodic task systems with latency constraints*, in "Proceedings of conference Models and Algorithms for Planning and Scheduling Problems, MAPSP'09, Abbey Rolduc, The Netherlands", July 2009, http://www-rocq.inria.fr/syndex/publications/pubs/mapsp09/mapsp09.pdf.

[20] F. MALLET, C. ANDRÉ, J. DEANTONI. *Executing AADL models with UML/Marte*, in "Int. Conf. Engineering of Complex Computer Systems - ICECCS'09, Allemagne Potsdam", IEEE Computer Society, 2009, p. pp. 371-376, http://hal.inria.fr/inria-00416592/en/.

[21] F. MALLET, C. ANDRÉ. *On the semantics of UML/Marte Clock Constraints*, in "Int. Symp. on Object/component/service-oriented Real-time distributed Computing (ISORC'09), Japon Tokyo", IEEE, 2009, p. 305–312, http://hal.inria.fr/inria-00383279/en/.

[22] F. MALLET, F. LAGARDE, C. ANDRÉ, S. GÉRARD, F. TERRIER. *An automated process for implementing multilevel domain models*, in "2nd Int. Conf. on Software Language Engineering, SLE 2009", October 2009.

[23] F. MALLET, M.-A. PERALDI-FRATI, C. ANDRÉ. *Marte CCSL to execute East-ADL Timing Requirements*, in "Int. Symp. on Object/component/service-oriented Real-time distributed Computing (ISORC'09), Japon Tokyo", IEEE, 2009, p. 249–253, http://hal.inria.fr/inria-00383262/en/.

[24] A. MEHMOOD KHAN, F. MALLET, C. ANDRÉ, R. DE SIMONE. *IP-XACT components with abstract time characterization*, in "Forum on specification, verification and design languages, FDL 2009", September 2009.

[25] P. MEUMEU YOMSI, L. GEORGE, Y. SOREL, D. DE RAUGLAUDRE. *Improving the Sensitivity of Deadlines with a Specific Asynchronous Scenario for Harmonic Periodic Tasks scheduled by FP*, in "Proceedings of the 4th International Conference on Systems, ICONS'09, Cancun, Mexico", March 2009, http://www-rocq.inria.fr/syndex/publications/pubs/icons09/icons09.pdf.

[26] P. MEUMEU YOMSI, Y. SOREL. *A New Model for Hard Real-Time Systems*, in "Proceedings of the 3rd Junior Researcher Workshop on Real-Time Computing, JRWRTC'09, in conjunction with the 17th International conference on Real-Time and Network Systems, RTNS'09, Paris, France", Oct 2009, http://www-rocq.inria.fr/syndex/publications/pubs/rtns09/rtns09.pdf.

[27] P. MEUMEU YOMSI, Y. SOREL. *Schedulability Analysis for non Necessarily Harmonic Real-Time Systems with Precedence and Strict Periodicity Constraints using the Exact Number of Preemptions and no Idle Time*, in "Proceedings of the 4th Multidisciplinary International Scheduling Conference, MISTA'09, Dublin, Ireland", August 2009, http://www-rocq.inria.fr/syndex/publications/pubs/mista09/mista09.pdf.

[28] D. POTOP-BUTUCARU, R. DE SIMONE, Y. SOREL, J.-P. TALPIN. *Clock-driven distributed real-time implementation of endochronous synchronous programs*, in "Proceedings of the 7th International Conference on Embedded Software (EMSOFT), Grenoble", October 2009, http://www-rocq.inria.fr/syndex/publications/pubs/emsoft09/emsoft09.pdf.

[29] D. POTOP-BUTUCARU, R. DE SIMONE, Y. SOREL, J.-P. TALPIN. *From Concurrent Multiclock Programs to Deterministic Asynchronous Implementations*, in "Proceedings of the 9th Application of Concurrency to System Design Conference, ACSD'09, Augsburg, Germany", July 2009, http://www-rocq.inria.fr/syndex/publications/pubs/acsd09/acsd09.pdf.

Scientific Books (or Scientific Book chapters)

[30] F. MALLET, R. DE SIMONE. *MARTE vs. AADL for Discrete-Event and Discrete-Time Domains*, in "Languages for Embedded Systems and their Applications", M. RADETZKI (editor), vol. 36, Springer, 2009, p. 27-41, http://hal.inria.fr/inria-00416656/en/.

[31] R. DE SIMONE, D. POTOP-BUTUCARU, JEAN-PIERRE. TALPIN. *6*, in "The Synchronous Hypothesis and Polychronous Languages", CRC Press, 2009.

Research Reports

[32] C. ANDRÉ. *Syntax and Semantics of the Clock Constraint Specification Language (CCSL)*, INRIA, 05 2009, http://hal.inria.fr/inria-00384077/en/, RR-6925, Rapport de recherche.

[33] C. ANDRÉ, F. MALLET. *Combining CCSL and Esterel to specify and verify time requirements*, INRIA, 2009, http://hal.inria.fr/inria-00360528/en/, RR-6839, Rapport de recherche.

[34] J. BOUCARON, A. COADOU. *Dynamic Variable Stage Pipeline: an Implementation of its Control*, INRIA, 2009, http://hal.inria.fr/inria-00381563/en/, RR-6918, Rapport de recherche.

[35] J. BOUCARON, A. COADOU, R. DE SIMONE. *Throughput and FIFO Sizing: an Application to Latency-Insensitive Design*, INRIA, 2009, http://hal.inria.fr/inria-00381644/en/, RR-6919, Rapport de recherche.

Scientific Popularization

[36] J. DEANTONI, JEAN-PHILIPPE. BABAU. *SAIA : une utilisation conjointe du génie logiciel et des méthodes formelles*, 2010, Journal du département IF de l'INSA de Lyon.

Patents and standards

[37] OMG. *UML Profile for MARTE, v1.0*, OMG (Object Management Group), November 2009, http://www.omg.org/spec/MARTE/1.0/PDF/, Document number: formal/2009-11-02.

Other Publications

[38] J. DEANTONI, F. MALLET, C. ANDRÉ. *On the Formal Execution of UML and DSL models*, April 2009, http://www.mdd4dres.info/_media/mdd4dreswip09_submission_10.pdf?id=wip&cache=cache, 4th Int. School on Model-Driven Development for Distributed, Realtime, Embedded Systems.

[39] J. DEANTONI, F. MALLET, C. ANDRÉ, B. FERRERO. *TimeSquare, on the Formal Execution of UML and DSL models*, April 2009, http://www.mdd4dres.info/_tools/1f3d0539532e6396ad1ecadc4d363a9a, 4th Int. School on Model-Driven Development for Distributed, Realtime, Embedded Systems.

[40] B. FERRERO, C. ANDRÉ, F. MALLET, R. DE SIMONE. *TimeSquare: a software environment for timed systems*, April 2009, http://www-sop.inria.fr/oasis/SAFA/Abstracts_SAFA2009.html, Design, Automation & Test in Europe (DATE'09), U-Booth session.

[41] JEAN-FRANÇOIS. LE TALLEC, J. DEANTONI. *Toward a TLM to RTL refinement: a formal approach*, October 2009, http://rtns09.ece.fr, 3rd Junior Researcher Workshop on Real-Time Computing.

[42] F. MALLET. *MARTE: The OMG UML2 Profile for Modeling and Analysis of Real-Time and Embedded systems*, September 2009, Keynote speaker.

References in notes

[43] C. ANDRÉ. *Representation and Analysis of Reactive Behavior: a Synchronous Approach*, in "Computational Engineering in Systems Applications (CESA)", IEEE-SMC, 1996, p. 19–29.

[44] F. BACCELLI, G. COHEN, GEERT JAN. OLSDER, JEAN-PIERRE. QUADRAT. *Synchronization and Linearity: an algebra for discrete event systems*, John Wiley & Sons, 1992, http://cermics.enpc.fr/~cohen-g/SED/book-online.html.

[45] A. BENVENISTE, G. BERRY. *The Synchronous Approach to Reactive and Real-Time Systems*, in "Proceedings of the IEEE", vol. 79, n^{o} 9, September 1991, p. 1270-1282.

[46] F. BOUSSINOT, R. DE SIMONE. *The Esterel Language*, in "Proceedings of the IEEE", September 1991.

[47] J. CARLIER, P. CHRÉTIENNE. *Problèmes d'ordonnancement*, Masson, 1988.

[48] L. CARLONI, K. MCMILLAN, A. SANGIOVANNI-VINCENTELLI. *Theory of Latency-Insensitive Design*, in "IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems", 2001.

[49] J.B. DENNIS. *First Version of a Dataflow Procedure Language*, in "Lecture Notes in Computer Sci.", vol. 19, Springer-Verlag, 1975, p. 362-376.

[50] S. EDWARDS. *Languages for Digital Embedded Systems*, Kluwer, 2000.

[51] T. GRANDPIERRE, C. LAVARENNE, Y. SOREL. *Optimized Rapid Prototyping For Real-Time Embedded Heterogeneous multiprocessors*, in "Proceedings of 7th International Workshop on Hardware/Software Co-Design, CODES'99", 1999.

[52] N. HALBWACHS. *Synchronous Programming of Reactive Systems*, in "Computer Aided Verification", 1998, p. 1-16, http://citeseer.ist.psu.edu/10686.html.

[53] H. HEINECKE. *AUTOSAR, an industrywide initiative to manage the complexity of emerging Automotive E/E-Architecture*, in "Electronic Systems for Vehicles 2003, VDI Congress, Baden-Baden", 2003.

[54] EDWARD A. LEE, D. G. MESSERSCHMITT. *Static Scheduling of Synchronous Data Flow Programs for Digital Signal Processing*, in "IEEE Trans. Computers", 1987.

[55] C.L. LIU, J.W. LAYLAND. *Scheduling Algorithms for Multiprogramming in a Hard-Real-Time Environment*, in "Journal of the ACM", 1973.

[56] Y. SOREL. *Massively Parallel Systems with Real Time Constraints, the Algorithm Architecture Adequation Methodology*, in "Proceedings of Conference on Massively Parallel Computing Systems, MPCS'94, Ischia, Italy", May 1994.