Review of INRIA Embedded and Real-time Systems Theme

ARVIND (Chairman), Béatrice BERARD, Stephen EDWARDS, Alessandro FANTECHI, Emmanuel LEDINOT, Véronique NORMAND, Vijay SARASWAT

Conducted March 2016

It was a pleasure to meet the teams and review the projects under the theme of Embedded and Real-time Systems. Most of the research that was presented is of the highest international standards, and quite a few team members are known internationally for their research. We report on each research team individually in separate appendices attached to this document. Here we will focus on our general observations and recommendations.

The panel thinks that teams should not only do research of the highest caliber but the research should also have impact on industrial practice. We think this view is consistent with INRIA's stated goals. The panelists were of unanimous opinion that the project MuTant represented the biggest success, if "industry" is interpreted broadly. MuTant, under the leadership of Arshia Cont, has produced a world-class interactive computer music system that is able to assist both in music composition and in performance (for accompaniment). It has released Antescofo software system, which is actively used by hundreds if not thousands of users. What is impressive is that Antescofo embodies the group's scientific results that span a wide gamut of disciplines including machine learning, signal processing, and real-time languages. It may help launch a new generation of Karaoke machines one day!

Another project which has had significant impact is Polyhedral which is part of PARKAS. Polyhedral, under the leadership of Albert Cohen, has produced a unified theoretical framework for the compiler analysis of computations involving dense matrix manipulation. Such analysis has been required for decades for generating efficient code for supercomputers but is now needed for embedded systems in a variety of domains. The role of Polyhedral is well recognized by the compiler community world-wide and now the Polyhedral group is engaged in packaging the analytic tool so that it can be incorporated in a variety of compilers in way that does not require any understanding of how the tool works internally.

Several other groups have produced tools that industry has used at least experimentally. For example, CONVECS has produced CADP, an extensive toolbox for the design and formal verification of asynchronous concurrent systems, and AOSTE has produced SynDEx, a system level CAD software to support their *algorithm-architecture adequation* (AAA) methodology. SynDEx is intended to optimize the implementation of embedded control applications onto *multicomponent* architectures under real-time constraints. AOSTE has successfully provided inputs to OMG MARTE standard regarding the insertion of time modeling concepts in industrial modeling language UML. PARKAS has continuous and influential collaborations with two major software publishers - ANSYS/Esterel Technologies and Dassault Systèmes. All such efforts and successes are worth applauding because it often takes time to understand the full impact of any research.

Every group has produced nuggets in the form of research papers and software. However, with the notable exception of MuTant and Polyhedral, the impact of the research has not exceeded the sum of the parts. We saw instances where several papers had been published by a team and yet there was no identifiable theme; there was little sense that these papers had collectively solved some big problem or part of a big problem. It was as if the publication of the papers was enough of an achievement that there was no reason to look for a unifying theme.

We are very aware of the fact that research often takes unexpected turns and generates good results serendipitously. Curiosity driven research should never be discouraged, but that does not mean that a team

cannot have a big theme. Every team should be able to state their goals in few sentences (an 'elevator level pitch' in the American jargon) and should be required to do so before any presentation internal or external. A team's plans must be designed to support such clearly stated goals.

Technology transfer is another area where there is lot of room for improvement. First it has to be recognized that the transfer of technology or new ideas to industry never happens without transfer of people. INRIA researchers can go outside and start a company but most research ideas are not easy to monetize and thus, do not benefit from the start-up culture. Some other ideas spread through consortiums – world-wide web being the most notable example. What is generally not recognized is that real impact on industry often results from successful collaborations where both sides understand the needs of the other. Industry cannot treat researchers at INRIA simply as hired hands; researchers are people with their own innovative ideas. Researchers on their part have to ensure that the problems and the solutions they offer are informed by industrial concerns and worked out at prototypical industrial settings. INRIA administration can do a lot more to bring about such collaborations.

It is always difficult to convince industry to adopt new techniques and some researchers find the task daunting. Industry is always looking for a better 'screwdriver', while a researcher might be trying to sell them the idea that they don't need the 'screwdriver' at all. For example, in the software world, the industry is happy to pay for new tools for finding bugs, but very slow in adopting a new methodology that might actually reduce the number of bugs in the first place. Industry is usually happy to engage with researchers and academicians to solve the specific problems they are facing, provided no significant adoption pain is involved on their part. Of course the problem is not one sided. As researchers, we think our tool or idea should work, but we may not have understood the real problem or tried solving it at the required scale.

Administration can be very helpful in connecting INRIA teams to the relevant groups in industry. Sometimes such contacts are needed so that the researchers can find out a realistic scenario where their ideas can be tried. If a group is developing a new language for controlling cyber- physical systems, like a power grid, or real time control in an automobile, they need to demonstrate that they can solve a realistic problem. For such a demonstration, meaningful data from industry is often essential, and it can take significant effort to convince an industrial development group that they should engage with the INRIA group and provide the relevant information. In some other situations, a research team may actually have a working solution and may need to convince the industry to try their solution inside the company. This latter situation faces many more challenges than the first one. At the same time everyone has to recognize that researchers cannot be expected to pursue the adoption beyond some reasonable effort. This problem is often compounded by the fact that not all brilliant researchers are good in interfacing with industry. In such cases help from administration would be essential to make any inroads.

There is recognition in the INRIA administration that projects should be encouraged and rewarded for industrial impact. We are told that there is an office to promote INRIA research in Industry. But as far as we could tell it has had minimal impact. We don't think that researchers look towards that office for help in selling their ideas to industry.

We also want to reiterate a concern that was expressed in the 2012 report on this research theme:

> "Our only areas of concern are the ways in which new projects are initiated where, it seems to us, the opportunities for change and renewal (of groupings and leadership, as well as research topics) are not fully exploited, ... The policy that projects have a maximum duration is an excellent one, but it seems that many of the new projects are formed by rearranging the components of expiring ones and selecting a new goal that is compatible with the skills and interests of the participants. While some degree of such bottom-up assembly is desirable and inevitable, ...
>
> We think it would be useful for INRIA researchers and management periodically to survey the entire field of embedded systems as a "clean slate" to identify promising and important research

areas, without regard to existing activities within the theme, and to use this as one of the inputs in formulating new projects, and possibly in recruiting new researchers. ”

We don’t think much has been done along the lines suggested above. If such an exercise was undertaken by administration, the Embedded theme as a whole, as well as each research team, would have a clearer idea of what they are trying to accomplish. We strongly encourage INRIA to undertake such an approach before forming new teams. It may cause one to think if it was worthwhile for so many groups to pursue competing ideas all within the confines of synchronous languages, or if new projects should be launched to exploit the technology developed by the MuTant project. To the panel it seemed that even the “machine listening algorithms” could apply to audio compression, audio surveillance, videogame interactions, audio content classification and searching, and perhaps even speech processing.

Finally, we think the primary reward structure in INRIA is centered on publishing papers in top conferences, and researchers, young and old, have internalized this. This reward structure has to be made to serve a new reward structure, which recognizes the impact of ideas on industry and society. Even if the idea has impact years after the project ceases to exist, it should always be publicized in media and INRIA literature. The impact of such recognition on researchers, especially the young and upcoming ones, cannot be over emphasized.

In conclusion, we saw lot of high quality research under the theme of Embedded and Real-Time Systems and think it can have even greater impact if more effort is put into appropriate industrial collaborations. We would also like to thank INRIA for the opportunity it provided us to contribute our observations for the future directions of the theme.

Project-team AOSTE

Scientific Leader: Robert De Simone

Primary Topics and Objectives:
The team AOSTE focuses on "Models and methods of analysis and optimization for systems with real-time and embedded constraints". More precisely, following the approach Adequation of Algorithm / Architecture (AAA), its goal is the optimization of the mapping from application models to architecture models. In this aspect, the team stands at the heart of the INRIA theme of Embedded and Real-Time Systems.

International Standing and Reputation in the Field:
The AOSTE team maintains a wide range of interactions both with other academic groups in the French and European communities and with several industrial partners. The publications of the team in highly ranked journals and conferences, as well as the participation of members in conference organizations and scientific boards, also attest to its reputation. AOSTE is also a major partner in LIAMA, a French/Chinese laboratory common with ECNU Shanghai and INRIA.

Major Achievements and Impact (Theory, Research Software, etc.):
The wide scope of the group, from models to architectures, has produced several important theoretical results and software. For instance, the Clock Constraint Specification Language (CCSL) has evolved in the period to attain a significant maturity level, with expressiveness results and applications in various domains. Other achievements concern real-time scheduling, with performance improvements for the uniprocessor case and probabilistic extensions.

Industry Transfer and Partnership:
Many partnerships exist, with large companies like Thales, Airbus, or car constructors, ST MicroElectronics, as well as with small ones like Adacore or ClearSy. The various softwares developed by AOSTE (SynDEx, TimeSquare, K-Passa, LoPhT, EVT-Kopernic) have all been used experimentally in industrial context. Some industrial transfer was achieved, for instance for the AAA methodology and the tool SynDEx, but the report is not really focused on this aspect of the work. On the other hand, several PhD students were offered positions of software engineer in companies like Kontron and ANSYS Esterel technologies, which had collaborated with the AOSTE team. This is a natural way to transfer technology to industry.

Training of Personnel:
The team has been training and graduating PhD students (6 in the period), who are currently on post-doctorate positions, or have become software engineers or assistant professors.

Principal Strengths and Weaknesses of the Project:
The large size of the project, developed in two geographical areas, and its wide scope, from applications to architectures, imply both strengths and weaknesses. On one hand, the project benefits from a global view of embedded systems and timing constraints, and the various members can cooperate within the group: people involved in execution models and scheduling problems can provide execution platforms to those wanting to experiment higher level constructs. Conversely, a view on the targeted applications is necessary for researchers trying to improve scheduling efficiency. This was achieved to some degree during the life span of the project but maintaining cohesion and consistency is not easy, especially when the team is large and distributed.

Plan for the next period (4 years):

The proposal for the next period is to split the project into two new groups: one in the Paris area and the other in Nice/Sophia-Antipolis. The first one, ASTRE (for Analysis and Synthesis of multiprocessor real-Time possibly Probabilistic Embedded systems) would retain a global view of the engineering process, with a focus on scheduling problems. The second one, Kairos, would extend the activities on model based design and investigate the application of reactive programming to the context where short-distance contactless communication can occur. Although the objectives are not yet completely fixed, the proposition seems like a reasonable one.

Opportunities and risks/difficulties faced by the project:
For the project as a whole, which is now terminating, this point is not really relevant. It will be more interesting to discuss it when the two new projects have defined clearer objectives. Still, real-time scheduling and model-based design are two important topics with a great potential impact.

Recommended actions and suggested measures of success:
The results produced by the AOSTE team must certainly be pushed forward to be applied in actual systems. This perspective would be worth pursuing by both new teams, even in the difficult context of French industry. Since it was already the case up to now, the communication between the two new groups should go on for mutual benefit.

Project CONVECS

Scientific Leader: Radu Mateescu

Primary Topics and Objectives:

The CONVECS project-team addresses the rigorous design of concurrent asynchronous systems using formal methods and automated analysis. At this regard, the team is focused on verification methods based on state space exploration (reachability analysis, model checking, equivalence checking, etc.). However, state space exploration typically does not scale well, while the complexity of designs is ever increasing, and requires considerable computing power (both storage capacity and execution speed). These are the challenges that CONVECS seeks to address. In particular, the objectives that guide the team are:

1. New formal languages and their concurrent implementations
2. Parallel and distributed verification
3. Timed, probabilistic, and stochastic extensions
4. Component-based architectures for on-the-fly verification
5. Real-life applications and case studies

International Standing and Reputation in the Field:

The CONVECS team has an excellent reputation within its scientific community, as witnessed by many publications in top-notch journals and conferences, and exhibits an impressive record of world-wide diffusion of the toolset that the team has produced over the years. The CONVECS team maintains a wide range of interactions both with other academic groups in the French and European communities, and participates to several European projects and interest groups.

Major Achievements and Impact (Theory, Research Software, etc*):

The team's work is centered around its extensive CADP software toolset, which the team has released to the public and is actively used by many academic and industrial sites. CADP is the outcome of a long-running effort by the predecessor teams of CONVECS, and it has been continuously maintained and upgraded, especially by providing new front-ends to cope with new challenges and for deeper industrial penetration. Indeed, the ability to incorporate fully developed theoretical concepts into high-level computer languages and associated software tools which are usable by industry is the most distinctive trait of the CONVECS team.

Industry Transfer and Partnership:

Some partnerships exist, with companies like Crouzet Automatismes and STMicroelectronics, also with entities inside an umbrella that pulls together French academic centers and industries, including SMEs. One PhD student has been funded by STMicroelectronics. These connections have allowed the toolkit developed by CONVECS to be experimented in different industrial contexts.

Training of Personnel:

The team has trained and graduated five PhD students in the period.

Principal Strengths and Weaknesses of the Project:

CONVECS exhibits good synergy between theory and practice. The team is a remarkably compact and has been striving to apply the developed tool base in advanced research and industrial projects. They have done so by providing new ways to exploit the potential of the toolset in new domains. The team has a lot of possibilities to exploit their competence given the need of powerful formal verification tools in a variety of different areas.

It should be pointed out, however, that the team has not addressed different approaches to the verification that could prove more advantageous in some cases. An example of alternative approach is SAT/SMT-

based techniques, which do not ask for explicit state-space exploration. Such approaches are currently exploited by competing tools. The CONVECS activities appear to be focused much more on the "front-end" ability of the toolset to cope with different areas of formal verification, than on the "back-end" verification engine, which is considered as a consolidated asset.

Plan for the next period (4 years):
The CONVECS team is now 4 years old and hence it naturally will continue its activity into the next 4 years, by developing the already set objectives.

Opportunities and risks/difficulties faced by the project:
The competence of the team allows for a lot of opportunities to exploit the developed tool base for a variety of different research and industrial applications. In particular, the recent studies on distributed verifications open the way to the usage of large distributed networks to support the verification of complex systems.

The only difficulties of the team may be related to generating enough industrial interest for a continuous flow of support and collaboration.

Recommended actions and suggested measures of success:
The committee has essentially no criticism of what the CONVECS team has been doing and how it has been doing it. In order to develop further the impact of the team, we make the following recommendations:

- Strengthen the industrial collaboration (in the panel's opinion, industrial collaboration should at least double in the next period)
- Investigate improving the verification engines by adopting, where possible and convenient, other techniques than on-the-fly verification

Project MuTant

Scientific Leader: Arshia Cont

Primary Topics and Objectives:
MuTant addresses real-time machine listening and timed real-time programming for computer music. Its goal is the production of a world-class interactive computer music system able to assist with both music composition and performance (e.g., for accompaniment), and in this it has succeeded admirably.

International Standing and Reputation in the Field:
The MuTant group is remarkable in that it has both an excellent reputation within its scientific community, as witnessed by many publications in top-notch journals and conferences, as well as with non-technical users of its Antescofo software. To the committee, this is perfect: their work is appreciated by both those who appreciate the nuances of its technical underpinnings as well as those who merely want to use it to accomplish an end goal.

Major Achievements and Impact (Theory, Research Software, etc*):
The group's work is embodied in its extensive Antescofo software system, which they have released to the public and is actively used by hundreds if not thousands of users. While many pieces of software have popularity of this magnitude, Antescofo is distinguished because it is directly driven by new, significant, scientific results by the group that span a wide gamut of disciplines including machine learning, signal processing, and real-time languages.

Industry Transfer and Partnership:
While the MuTant group has advised and/or licensed some of their work to several startup companies, its main path for technology transfer has been through public dissemination of the Antescofo system itself, which has put this group's findings in the hands of musicians: its primary user base.

In March 2016, the team leader and two PhD students have created a startup focused on further developing the Antescofo product.

Training of Personnel:
The team has been training and graduating PhD students, who have gone on to take postdoc positions.

Principal Strengths and Weaknesses of the Project:
This project embodies a near perfect synergy between theory and practice. For example, the new machine learning algorithms the group has developed are subtle, complex, and highly theoretical, yet they have direct application and benefit to the problem of computer listening, e.g., segmenting waveforms into distinct instrument sounds. To the committee, we find such work a superlative example of the power of computer science: high-powered mathematics harnessed to produce algorithms of practical utility.

The only weakness of the project is that its application domain, music, is not as important an economic activity as, say, automobile or aircraft design, and might initially appear to some as being frivolous.
However, the sheer scientific quality of the work coupled with its potential applications outside the somewhat narrow realm of live music performance, easily overcomes this.

Plan for the next period (4 years):
The committee is actually a little sad to learn that the MuTant group, as it stands, is coming to an end because its leader and two PhD students are leaving their academic environment to found a startup focused on further commercializing the technology. Given the project's success to date in its current setting, it seems almost a shame to dismantle such a success.

Opportunities and risks/difficulties faced by the project:
The technology developed by the group seems like it could be used for a much wider range of applications than it has currently addressed. For example, the machine listening algorithms seem like they could apply to audio compression, audio surveillance, videogame interactions, audio content classification and searching, and perhaps even speech processing. The committee encourages the group to consider applying the wealth of technology it has developed to new application areas.

Recommended actions and suggested measures of success:
The committee has essentially no criticism of what the MuTant group has been doing and how it has been doing it. If anything, we encourage other project groups to examine and attempt to emulate the MuTant group's success to the extent possible.

Concretely, we recommend other groups consider how the MuTant project has been centered around a single grand, but not impossible, goal: computer-assisted music performance. Such a goal is less grand than, say, Hoare's challenge of a verifying compiler or of simulating the workings of a complete cell, but this makes it more attainable and has worked extremely well. Moreover, it is clear that the MuTant group's interaction with users has helped guide their work and kept it focused on delivering results that remain relevant. The committee suspects other groups could successfully mimic such a single-minded focus on a particular goal.

Project PARKAS

Scientific Leader: Marc Pouzet

Primary Topics and Objectives:
The primary stated objective of the project is the design of deterministic parallel programming languages and their implementation on multi-core architectures. A particular focus is the design of a synchronous programming language supporting hybrid computation. The project is organized along three loosely coupled themes (a) programming language definition (b) compilation and run-time for compute-intensive programs running on multicores, (c) validation and proof techniques for compilers, particularly focused on understanding weak memory models.

International Standing and Reputation in the Field:
All the PIs on this project are well known in the research field and have significant international reputations

Major Achievements (Theory, Research Software, etc.).
We wish to highlight the following achievements:
(a) Work on hybrid systems' semantics that lead to the development and implementation of Zelus. Zelus extends a Lucid Synchrone like language with continuous time variables, whose dynamics is expressed using differential equations. The compiler is an incremental modification to the Lucid Synchrone compiler and involved extending the underlying type system. The ideas were further developed in an extension of the SCADE compiler, SCADE Hybrid. Particularly notable is that the extension to the SCADE compiler required only 5% extra lines of code (LOC). This highlights the clean design of the language as an extension of an existing synchronous language. It is also a good example of theoretical work (non-standard analysis for causality analysis) with true impact on practice. Extending these results to Differential Algebraic Equations (DAE) would be a great achievement of PARKAS and Hycomes for the next period.

(b) Studies of the Chase-Lev deque algorithm, a critical algorithm for the implementation of dynamic task-parallel libraries. The paper provides a correctness proof of the algorithm for relaxed memory models, and discusses the performance of multiple implementations, in the context of work-stealing schedulers. An interesting aspect of the paper is a demonstration that the most efficient implementation cannot be written using C11 low-level atomics.

Also notable is the work presented at POPL'15, in collaboration with Vafeiadis and colleagues demonstrating problems with common compiler optimizations and the C11 memory models.

Industry transfer and Partnership:
The project has an effective way of transferring some of its results to industry, even if the evidence is not clear when viewed through normal channels like software licenses, start-up creations or people transfers.

Marc Pouzet was key in the design of major evolutions of the SCADE compiler at ANSYS/Esterel Technologies when synchronous hierarchical and parallel state machines were added to data-flow programming. He helped crafting an appropriate trade-off between expressivity and causality analysis which was pivotal for the success of SCADE 6. He also gave scientific consultancy to Dassault Systèmes on the LCM compiler of CATIA Systems.

These collaborations have been continued over a period of time. They led to SCADE Hybrid at ANSYS and to Modelica 3.3 at standard level (Modelica Association) and product level (Dymola of Dassault Systèmes). In both cases transfer included compilation of synchronous data-flow programs into parallel code for multi-cores.

Albert Cohen has closely collaborated with Kalray, a start-up on energy-efficient manycores. He contributed OpenMP and OpenCL support (LLVM development) for the MPPA. He is also committed to transferring polyhedral compilation in the production-quality free platform supported by ARM (Polly Labs).

In the middle of its profusion of research activities and prototypes, PARKAS has also managed to have seminal contributions to industry in a significant number of places.

Training of Personnel:
The project has a very good track record of training researchers, with 12 PhDs being graduated from 2012-2015. The record of positions in industry and academia of the former PhD students is also very good.

Principal Strengths and Weaknesses of the Project:
The principal strength of the project is the individual strength of the principal investigators, and the choice of hybrid programming as the focus. Particularly strong results have been obtained in compilation and validation. The work on programming language design for hybrid computing is still in progress, with fruitful ongoing work with the Hycomes team.

We would highlight one area in particular where the team can improve. Developing hybrid synchronous programming -- the concept, languages, type systems, compilation schemes, static analysis frameworks, correct program design frameworks, reasoning frameworks, and application domains -- is a very important scientific task. This team is capable of "moving the needle" in this space -- it has all the right skills, background and brain power.

Yet the team does not seem to have rallied together around this singular mission. While the work on compilation and weak memory models (e.g. C++ Memory models) is scientifically important, significant and stands on its own, it is not clear that this work would be on top of the list *if* the goal was to properly develop the hybrid synchronous programming story.

Plan for the next period (4 years):
We would prefer to see a tighter focus on designing and implementing hybrid synchronous languages, with tools to generate code for multicores and distributed architectures, to reason about program behavior (e.g semantic equivalence between sequential and parallel executions), and tools for "sketching" programs, cf [Solar-Lezama, Bodik]. If tools for debugging causality loops and distributed executions were added to the research agenda, there might be opportunities for collaboration with SPADES.

There are elements that indicate a bigger compiler agenda in future, but this does not seem to be in line with the general goal of PARKAS, which is to develop a framework for hybrid synchronous programming.

Opportunities and risks/difficulties faced by the project:
Since the research leads for the three areas – M.Pouzet's, A.Cohen's, F.Zappa Nardelli's -- are very strong and have large, vibrant research agendas, there is a tendency towards breadth and getting into other related areas, vs focusing on really nailing one area. Each subgroup are perceived as doing their own thing.

Clearly the decision here should be made by INRIA management and team leads on the opportunity to be more focused for achieving highly visible "game changing" results. Such a discussion might also include Hycomes' and SPADES' leads.

Recommended actions and suggested measures of success:
(a) Bring all three parts of the project together in realizing productive hybrid synchronous programming languages, with parallel and distributed implementations on (collections) of multi-cores.

(b) Identify areas of application that highlight the range of possibilities for hybrid synchronous programming, even if they step outside the current bounding box of embedded programming research in France. An example is animation and movie-making -- the field naturally supports the need for a high-level, compositional design framework within which complex story lines need to be pulled together, with hybrid visualization / animation techniques being at the core.

For this, suggested measures of success are application frameworks built on top of the hybrid programming framework that capture the imagination of application developers in that field, "move the needle" in the application field, enable new kinds of applications that the current set of tools simply cannot support, and become the new de facto standard. The work of MUTANT on Antescofo is an example.

Project SPADES

Scientific Leader: Alain Girault

Primary Topics and Objectives:
SPADES covers a broad spectrum of topics: distributed computing, schedulability analysis, machine-checked correctness proofs of algorithms, fault-tolerance at middleware and hardware levels, causality analysis in distributed systems, etc. This rich portfolio of competencies originates from the merge in 2013 of two significantly different threads of research: that of POPART (application level programming of hard real-time situated systems), and that of SARDES (middleware level theoretical modeling and programming for large scale component-based distributed systems). The overall objective of SPADES team is to provide formal methods to program systems that either feature dynamic structures, or execute on multicores, or need mixed criticality and fault-tolerance.

International Standing and Reputation in the Field:
The team members have good to very good international visibility, as evidenced by the record of participations to program committees of tier 1 conferences, including some participations as chair or co-chair, and by some forefront editorial responsibilities or working group leaderships.

Major Achievements (Theory, Research Software, etc.).
The objectives for the period 2013-2016, and for the next one alike, are:

- Components and contracts,
- Real-time multicore programming,
- Language-based fault-tolerance.

The record of the team publications is very good, qualitatively and quantitatively speaking. During the evaluation seminar the team put forward three particular achievements that were selected from about ten possible ones: Coq-verified time redundant fault tolerance mechanisms for hardware, typical worst case schedulability analysis for the automotive industry, and logical causality analysis for fault ascription to components in distributed executions.

All the research activities, presented or not, address some practical concerns and require high technical capability. Regarding multicore programming however, and more specifically certification-oriented control of contentions on mainstream marketed multicores, the global positioning of the team and its perspectives of true impact on industrial practice remained unclear to the panelists.

There was also some perception of lack of cohesiveness. The team appears more as an aggregate of bright researchers, all sharing inclination towards formality and possibly benefiting from one another, but pursuing mainly personal and isolated threads of research.

As far as software is concerned, three contributions are listed in the synthesis report: LDDL (a Coq library for hardware description), COSYMA (a tool for controller synthesis) and pyCPA_TWCA (typical worst case response time analysis for weakly hard guarantees).

Industry transfer and Partnership:
The team was created only 2 years ago. A transfer of libraries dedicated to abstract interpretation of models or progams was signed with the start-up ArgoSim in 2013 when Bertrand Jeannet left POPART-SPADES. There are partnerships with STMicroelectronics and Thales (CIFRE PhD thesis), and collaborations with Daimler and Bosch on so-called “weak-hard” real-time schedulability analysis.

Training of Personnel:

It has significantly decreased over the 3-year period: from 4 PhD students and 2 post-docs in 2013, to 2 PhD students and no post-doc as of evaluation. Most of former team members (6 over 9) found positions in major tool-vendor companies or co-founded a start-up, which is a good transfer record

Principal Strengths and Weaknesses of the Project:
The principal strength of the project is its wide spectrum of competencies and its scientific excellence. The counterpart of this strength is a perceived lack of cohesiveness at team level. Some joint work encompassing a significant subgroup of the team was not put forward during the seminar.

Some works are original and likely to be rich of potential applications (*e.g.*, reversible distributed computation, fault ascription). Others seem closer to research carried out since the 1990s and close to that of AOSTE.

Plan for the next period (4 years):
The research plan for the next four years is basically the continuation of the three main areas defined at team creation in 2013. The foundational part (pi-calculus, location graphs, reversibility, logical causality, fault ascription, etc.) of objectives 1 and 3 is rich and promising for future tier 1 scientific publications or for deep insights in middleware design, if any.

How objective 2 "Real-time multicore programming" is addressed remained partly unclear to the committee. We share that the stringent economic constraints on the automotive sector motivate continued research on schedulability analysis (Typical Worst Case Analysis). But the part named "synchronous programming for multicores" looks more like new variations on the old theme of parallel or distributed execution of synchronous programs, than a true intent to provide comprehensive prototype solutions to program full-fledged software on marketed multicores. The positioning of ForeC for instance, and that of the related activities (*e.g.*, WCET estimation), do not seem to address contention control at hypervisor and OS level, nor in the bare metal case. We probably missed something in SPADES' research rational, but it seems to the panel that deterministic programming at application level is only part of what is needed for true industrial impact on the safety-critical multicore problem.

Opportunities and risks/difficulties faced by the project:
The multi-faceted nature of the research objectives of the team makes sense: it provides opportunities of crosspollination among the team members, and today systems of systems feature all these facets, their development requires multi-disciplinary teams like that of SPADES.

However, the synthesis report on the evaluation period and the presentations in private session made the team appear as a weakly cohesive aggregate of bright person-scale research activities. What will be at team-scale over the next 4-year period? What will contribute to institution-level challenges with some potential to contribute to INRIA's world class visibility in the field? We found no answer to these questions, possibly we missed something, but if so, it should have been put forward more clearly in such a *team*-evaluation and a *theme*-evaluation seminar. There is a need for higher vision of SPADES' research agenda as far as higher impact scale is concerned.

A suggestion of opportunity for SPADES within the "embedded and real-time systems" theme would be to define a world visible "grand challenge" (scientific impact and transfer potential) based on the unique portfolio of competencies INRIA has on formal semantics of Modelica (DAE integration, typed objects, switched and dynamic structures), on theoretical models of distributed processes, and on logical causality analysis for debugging of distributed software.

PARKAS and HYCOMES have spontaneously focused this way and engaged in-depth collaborations with the Swedish academic and industrial core partners of the technology. SPADES could join this thread of research for part of its agenda (e.g., fast multicore and distributed *guaranteed* simulations, with breakpoints, rollbacks, and causality-based debugging capabilities).

Recommended actions and suggested measures of success:
The suggested opportunity allies scientific excellence (*e.g.,* precision timed semantics based on non-standard analysis, Coq-proved semantical results on models of distributed processes, etc.) and great industrial impact potential: such a middleware, or compilation/deployment/execution platform, is required to meet the digital *functional* mock-up objective, especially for systems of systems.

The digital functional mock-up objective was formulated a decade ago in the model-based CPS development roadmaps of nearly all industrialists (energy, transportation, health-care, IoT, etc.). Current state of the art of commercial Modelica offering suggests that it will stay in these roadmaps for the next decade, at the very least. It is likely to be hard, even for majors like Siemens, Wolfram, ANSYS and Dassault-Systèmes, to develop the kind of certifiable, efficient, and user-friendly CPS and CPSoS simulation middleware mentioned previously. INRIA has the portfolio of competencies these tool vendors miss to tackle such a challenge, to develop this kind of ideal component featuring at the same time all these characteristics users need.

Project TEA

Scientific Leader: Jean-Pierre Talpin

Primary Topics and Objectives:
Project TEA is a newly created project, formally started in January 2015.

Within the INRIA theme on real-time embedded systems, this project focuses on systems architecture and integration. The project aims at developing formal foundations for time-related reasoning in software and systems architecture design, early verification, and integration.

Four distinct objectives are proposed:

1. Time modeling and formal reasoning framework for system design;
2. System architecture-based approaches for system property analyses;
3. Real-time dynamic scheduling leveraging abstract interpretation or probabilistic approaches;
4. Virtual prototyping (hardware modeling and simulation) theory and software.

International Standing and Reputation in the Field:
The project has a small team lead by two senior researchers of international standing, with well-developed connections especially in the US and China.

Major Achievements (Theory, Research Software, etc.):
At this preliminary stage (18-month existence), the project has produced a number of intermediary results around SAT/SMT solver usage for timing analysis (objective 1) and abstract affine scheduling techniques (objective 3).
Beyond those intermediary results, two main achievements are reported, that leverage previous work:

- The definition of formal semantics for the AADL, with a reference implementation in the Polychrony tool; and the publication of this work into the AADL standardization body.
- Additions to the SimSoc project including the formal proof of the ARM instruction set simulator.

Impact analysis, industry transfer:
Architecture languages achievements: The AADL is a standard in the automotive and aerospace industry domain (SAE International standardization body). The user community for the AADL has slowly developed over the years and remains essentially academic, with unfortunately few reported usages in industrial settings as of today. The TEA work on formal semantics is thus expected to feed the applied research projects in the community (e.g. Airbus-led research in IRT Saint-Exupery), further developing the scientific reach and technical capability of the AADL to address software architecture modeling, simulation and analysis.

In relation to this work, the team reports a specific collaboration with Toyota US for exploring the insertion of contract-based software architecture modeling and analysis techniques into the Toyota engineering processes, leading to two patents.

Hardware simulation achievements: Regarding the SimSoc area of achievements, it is difficult for us to assess the impact of this work, as this topic could not be discussed during the meeting (the appropriate team member was on holidays). Collaboration is reported with ST Microelectronics; our assumption is that this work is liable to be transferred to the ST tool set.

Plan for the next period (4 years):
The proposed plan includes the following:

- Applied research for Mitsubishi R&D addressing the modeling and analysis of factory automation systems. Our understanding is that this applied research would orient and leverage more

fundamental research listed below.

- Software and system architecture modeling for multi-domain property analysis – temporal, mechanical, power, etc.; supporting contract concepts and refinement theory.
- Time synchronization protocols, refinement-types theory for time.
- A collaboration with AOSTE and DIVERSE to integrate the TEA time reasoning framework into the GEMOC environment.
- Probabilistic scheduling theory, syntax guided scheduling
- Multi-domain, multi-scale simulation for large physical systems.

Principal Strengths and Weaknesses of the Project:

The intended scope of the TEA project and its specific positioning in the INRIA theme appear to us as strengths, for two reasons:

1) Balance within the INRIA theme organization: the overlap with other projects is reduced, with good complementarities with several other projects, especially AOSTE and SPADE.
2) Industrial relevance: architecture modeling and analysis as well as system integration support are key industrial concerns. Breathing scientific foundations and formal analysis capabilities into current component-based embedded systems design practices does seem to us as a useful and challenging objective for INRIA researcher.

Other strengths include: the software development capability of the team; the connections to industry (Mitsubishi in particular); the scientific reputation of the leaders; etc.

Weaknesses of the project involve the following:

- Incomplete overarching vision. What is the grand challenge tackled by the project, and how do the four objectives inter-relate and contribute to solving the challenge? How is the team interacting on a common overarching vision? The presentation made by the team during the private session did not help answering these questions. The feeling of the panel is that the variety of scientific and technical activities does not convey a strong sense of focus and team integration.
- Size of the team. The team is quite small, while the project plan features many different areas of work, which may look unrealistic.

Recommended actions and suggested measures of success

- Consolidate the overarching vision; strengthen the focus of the roadmap accordingly with the vision.
- In particular, make sure to fully address the end to end research issues, from system specification and architectural design down to actual implementation and integration. In that context, a particular industrial concern is the one of reuse – how to ensure sound reuse of components, leveraging qualification credits attached to a component. This may be a dimension to explore.
- Further develop connections with industry. The Mitsubishi collaboration looks fine. The Toyota collaboration looks a little tenuous. Strengthen or add collaborations with embedded industry.
- Develop collaborations with other INRIA projects to compensate for the small size of the team and ensure a good coverage of the architecture and integration area in the INRIA theme. For example, with SPADES on contract-based design.