PROJECT PERIODIC REPORT

PROJECT MANAGEMENT AND USE OF RESOURCES

Grant Agreement number: 243381
Project acronym: CerCo
Project title: Certified Complexity
Funding Scheme: STREP
Date of latest version of Annex I against which the assessment will be made:

Periodic report: 1st □ 2nd □ 3rd X 4th □

Period covered: from 01 February 2012 to 31 March 2013

Project coordinator: Prof. Claudio Sacerdoti Coen
Alma Mater Studiorum – Università di Bologna
Tel: +39 051 2094973
Fax: +39 051 20 9 4510
E-mail: claudio.sacerdoticoen@unibo.it
Project website address: http://cerco.cs.unibo.it

Table of contents

4 Project management

Management activities

CerCo is a relatively small project with only three partners and a small number of work packages. Management activities are included in WP1, while WP2–WP5 correspond to research activities.

The overall objective of WP1, as indicated in Annex I of GA, is the smooth implementation of the CerCo project and in particular:

- to coordinate and supervise activities to be carried out;
- to carry out the overall administrative and financial management of the project;
- to manage contacts with the European Commission;
- to monitor quality and timing of project deliverables;
- to establish effective internal and external communication procedures.

Consortium management tasks and achievements

Consortium structure

The project management and administrative structure is described in Annex I of GA Section B2.1. No changes occurred with respect to the management structure. For convenience, a graph illustrating the project management structure is reproduced below.

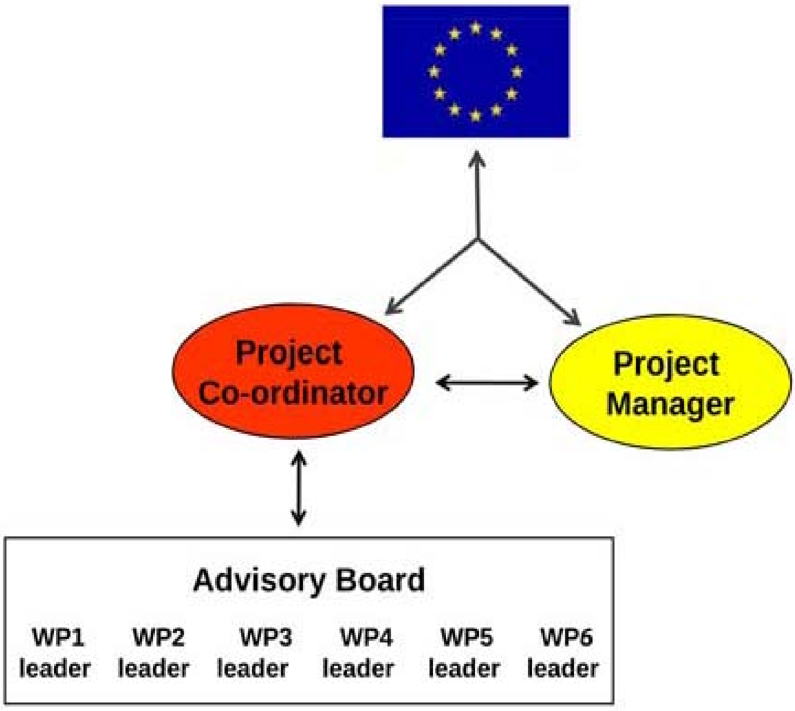

Financial Management and Coordination

The coordinator UNIBO was in charge of the overall administrative and financial tasks for CerCo. The activities during this first period included:

- the general financial management of the EC funding, i.e. transmission pre-financing to all beneficiaries;
- providing partners with detailed information on consortium organisation and FP7 financial management rules during the Kick-off meeting;
- collecting and checking partner's financial information, submission of financial statements through NEF system;
- gathering all necessary information and materials for the preparation and submission of management report, activity report and deliverables.

WP leaders have had an important role for the preparation of project deliverables, acting as intermediary between beneficiaries and coordinator. They have also been responsible for the monitoring of WP activities and the quality check of deliverables and reports

Consortium Agreement

A Consortium Agreement regulating in particular the Intellectual Property Rights and Access Rights, the roles, functioning and responsibilities of consortium governing bodies, the liability of the partners and others, was signed by all CerCo partners in January 2010, before the start date of the project.

Communication Management

Due to the simple structure and composition of the consortium the communication strategy during the first 12 months required rather simple tools for communication in order to guarantee the exchange of information and data between partners, such as simple e-mail communication, website and consortium meetings.

Organisation of Periodic Meetings

The organisation of meetings was one key aspect for both coordination of future research activities and exchange of information between partners. All partners agreed that it is necessary and helpful to have direct contact between the involved research units especially for the purpose of a better integration among units. The project meetings allowed WP leaders and coordinator to monitor the project activities, prepare deliverables and plan future steps for the implementation of tasks.

During the first period we organised 4 project meetings:

- **Kick-off meeting**, **Bologna, Italy, February 2010**. This first meeting concentrated both on coordination of research activities, planning of future work and financial and administrative organisation of the consortium presented by the Project Manager within UNIBO.

- **Project meeting**, **Edinburgh, Scotland, July 2010**. This meeting was scheduled at the end of the first six months of the project, just after the production of the important deliverable D2.1 (Compiler design and intermediate languages). Having taken the major design decisions, the project was aimed at fixing a more precise scheduling for the activities to be carried out in the remaining months. The meeting has also been scheduled to be co-located with FLOC, the Federated Logic conference, which is the largest venue of logicians. This provided us with an opportunity for informal dissemination of project results at the conference. All beneficiaries were represented.
- **Project meeting**, **Paris, France, September 2010**. The aim of this meeting was mainly to discuss the design of the costing compiler (D2.2) and the memory model and semantics of C (D4.1). In particular, the adoption of the MCS-51 microprocessor suggested to consider in D3.1 some ad-hoc extensions to the C language and required to take in account a non uniform memory model. At the meeting we discussed about the possibility of affecting D2.2 to support these features. Moreover, the choice of memory model will have a direct impact on WP4 too since the front-end and part of the back-end of the compiler will share the same memory model. We also discussed the possibility of having an implementer's meeting at the end of the period to address all the integration problems raised between the two meetings. All beneficiaries were represented.
- **Implementer's meeting, Bologna, Italy, January 2011**. This meeting was scheduled at the end of the tasks active during the first year. The objective of the meeting was to finalize D2.2 by solving the remaining integration problems with D3.1 and D4.1. The one-week-long meeting was used to put in close contact implementers from each project site in order to discuss the integration problems that required a tighter and promptly collaboration. Some of the problems were totally solved during the week. For the remaining one, we decided an integration strategy to be carried out in Task 2.4 during the next period of CerCo. All beneficiaries were represented.

During the second period of CerCo we organised 4 project meetings:

- **Project meeting, Edinburgh, Scotland, February 2011**. This meeting was scheduled at the beginning of the second period to discuss and schedule the activities for the second year and to review the decisions taken during the implementer's meeting in Bologna. In particular, we discussed the global proof strategy and the particular techniques to be adopted in the proofs. The objective was to guarantee the smooth integration of the proofs for the front-end (WP3) with those for the back-end (WP4). We also started a discussion on the tasks to be done in WP5, that had previously not been initiated. All beneficiaries were represented.
- **Project meeting**, **Paris, France, September 2011**. This meeting was a federated meeting of the CerCo and Eternal projects. There were common public dissemination sessions where we presented the results of CerCo to the members of Eternal. These public sessions were followed by private sessions for the discussion of the results described in the deliverables due for the 18th month of the project (D3.2, D3.3, D4.2, D4.3). In addition we discussed progress on WP5 and the problems faced in the development of the Frama-C plugin. An invited speaker from the Frama-C community helped in understanding some of the difficulties. He has also been the target of dissemination. Eternal is an INRIA ARC project that nicely complements CerCo. In CerCo we will provide actual certified computational costs for loop bodies, to be combined by

the user with bounds on loop invariants to certify overall execution costs. However, we do not provide automatic computation of the bounds. The Eternal project studies ways to bound the number of times loops are executed by means of techniques that come from the implicit computational complexity community, married with interactive theorem proving. However, they do not provide any realistic cost model for the bodies of loops. Hence each project provides tools and techniques to give the missing piece of information to the other one. A major topic of discussion in the private CerCo session has been the integration of WCET technologies. In particular, a WCET expert has been invited to give a thorough overview of the AiT tool, confirming the feasibility of a study on the integration of WCET techniques in CerCo. The study had already started in previous months and partial results have also been presented. All beneficiaries were represented.

- **Implementer's meeting, Edinburgh, Scotland, December 2011**. The one week long meeting was used to draw together implementers from each project site in order to discuss integration problems that required a tight and prompt collaboration. It was also used for a long session of brain-storming on one critical problem faced in the proof of correctness for the static prediction of costs. The brain-storming resulted in the notion of structured traces to be used to provide the operational semantics of well behaved back-end programs. The new semantics will have an impact all over the proof. Finally, the untrusted compiler was patched to integrate the feedback gained during the formalization of the compiler. All beneficiaries were represented.

During the third period we organised 4 project meetings:

- **Project meeting, Paris, France, March 2012.** This short meeting was held in Paris the day before the second review of the project. We discussed and scheduled activities for the third year and we reviewed decisions taken during the implementer's meeting in Edinburgh. All beneficiaries were represented.
- **Project meeting, Tallin, Estonia, March 2012.** This short meeting was held during the ETAPS 2012 multiconference. There was a particular focus on refining plans for dissemination events, and arrangements for an ETAPS 2013 workshop were agreed with the conference organisers. There was detailed discussion on progress with the compiler proof, as highlighted by the reviewers. In response, partners agreed to review their remaining site budget in preparation for requesting a project extension. All beneficiaries were represented.
- **Project meeting, Bologna, Italy, May 2012**. This meeting lasted three days, with extended discussion of two major scientific topics. The first was the actual shape of the final correctness statement that needed to be changed to incorporate stack usage. The notion of *measurable subtrace* was identified and the top level layer of the correctness proof was re-discussed. The second topic was the decision on how to accommodate in the global schedule the work towards pipeline and cache prediction, which was not planned in the original project plan. We decided to reuse part of the budget of a Post-Doc that was leaving for a lectured position to hire a new Post-Doc to work part-time on the topic. We also decided to have him focus on the analysis of the conditions for dependent labels, without attempting an implementation that would have not been possible with the resources available. A whole session was devoted to brainstorming on the organization of the two events planned for the third period. Finally, we performed a review of the remaining budget and we decided to ask for a project extension of two months. All beneficiaries were represented physically or via teleconference.

- **Implementer's meeting, Edinburgh, Scotland, August 2012.** The meeting was used to draw together implementers from UNIBO and UEDIN in order to discuss integration problems between the formal proofs of the front-end and the back-end. Some additional invariants have been negotiated at the front/back-end interface. UNIBO and UEDIN were represented.

Changes in the consortium

No major changes have occurred with regard to the organisation and management structure of the consortium.

The legal statuses of the beneficiaries have remained unchanged.

Project planning and status

During the third period we asked for a project extension of two months. The deadlines for most deliverables have been shifted according to the new end of period.

During the third period, the project has followed the project plan (both technical and management) described in the updated version of the Grant Agreement. Additional work unplanned in the GA has been performed in WP5 and WP4 to extend the labelling approach technique to cover loop optimizations and cost models that are sensitive to the hardware state.

We failed to complete the formal proof of correctness to be delivered in D3.4 and D4.4. At the end of the second period we presented a detailed estimation of the effort required to complete the proofs, as required by the reviewers. According to the estimation, we should have been able to complete the proof. The estimation was based on the amount of code to be proved and the ratios of proof lines/code lines taken from the CompCert project. It turned out that the ratios were greatly under-estimated because the additional effort required for the parts of the simulation that deal with non functional invariants was more significant than expected. In particular, we had to perform stack space analysis in addition to execution time analysis. While we were able to reuse the methodology we already developed for time, this was not on its own sufficient because running out of stack space has also effects on functional properties, which is not the case for time. Another cause of delay has been the back-porting of optimizations implemented for the first time in the untrusted prototype branch that deals with loop optimizations. The ported optimizations were necessary to obtain object code with reasonable performance when compiling 32 bit code on 8 bit hardware.

Because of the delay, we have therefore decided to concentrate on the parts of the formal proof that present novel research material. In particular, we have focused on:

1. the top-down decomposition of the proof into three layers that deal with non functional properties plus one remaining layer whose proof obligations are essentially the same of a standard functional simulation. The formal proof of the novel layers has been completed, while the proof of the standard layer has not been completed for several compiler passes. Completing those parts is necessary to claim that the compiler is 100% correct, but the proofs are not likely to add new insights to the research community.
2. the proof of passes and generic results that deal with our unified representation of back-end languages. The unified representation is a departure from the approach taken in other

projects. Completing the parts of the formal proof that is mostly influenced by this design choice allows to assess the benefits and drawbacks of it.

3. the formal proof of the functional properties of our assembler, which implements the branch displacement optimization. Branch displacement is an hard problem and implementations are known to be error prone. Our is the first formal verification of a branch displacement implementation and it should be easy to port the formal proof to other microprocessors.

At month 28 we submitted two addenda for D2.1 and D6.2 and revised versions for D4.2, D4.3 and D5.1, in accordance with requests made by the reviewers during the second Project Review.

Deliverable D6.6, expected at month 33, was ready only at month 36 and delivered at the end of the project. The deliverable provides Debian packages and a Live CD for the software developed in CerCo. A strong requirement by the reviewers at the end of the second period has been to produce an extracted version of the compiler code. At the time of the review, however, Matita did not implement code extraction, which was made available only at month 35. Therefore we were able to extract the code to make the Debian packages only at month 36. No other deliverable or milestone depends on D6.6, so the delay had no noticeable effects on the schedule.

Deliverable D6.4 (Organization of an Event Targeted to Potential Industrial Stakeholders) was scheduled for month 35 and organized respectively at months 36 to co-locate it with HiPEAC 2013.

All remaining scientific deliverables (D3.4, D4.4, D5.2, D5.3, D6.3, D6.5) were expected at the end of the project (month 38) and have been delivered at month 39.

Deliverables D1.3 (Periodic Activity Report and Financial Statements) and D1.4 (Final Report) will be submitted within 60 days after the end of the period.

5 Dissemination and use of knowledge

Use of foreground and dissemination activities

List of presentations given to international conferences:

- **Certified Complexity**. talk given by Claudio Sacerdoti Coen at Types 2010, Warsaw (15 October 2010).
- **Realizability Models for Cost-Preserving Compiler Correctness**, talk to be given by Marco Gaboardi at DICE 2010.
- **A canonical locally nameless formalisation of an algebraic logical framework,** talk given by Wilmer Ricciotti at Types 2010, Warsaw (13 October 2010).
- **An Elementary Affine λ-Calculus with Multithreading and Side Effects**, talk given by A. Madet at TLCA 2011 (2 June 2011)
- **A type system for positivity,** talk given by Claudio Sacerdoti Coen at Types 2011 (10 September 2011)
- **The Matita Interactive Theorem Prover**, talk given by Wilmer Ricciotti at CADE 2011 (3 August 2011)
- **A Web Interface for Matita**, talk given by W. Ricciotti at Conference on Intelligent Computer Mathematics 2012, (10 July 2012)

- **Certifying and Reasoning on Cost Annotations in C Programs**, talk given by Y. Régis-Gianas at Formal Methods for Industrial Critical Systems, Paris 2012 (28 August 2012)
- **A polynomial time lambda-calculus with multithreading and side effects**, talk given by A. Madet at Principles and Practice of Declarative Programming, Leuven 2012 (20 September 2012)
- **Separation Predicates: A Taste of Separation Logic in First-Order Logic**, talk given by F. Bobot at International Conference on Formal Engineering Methods, Kyoto 2012 (15 November 2012)
- **On the correctness of an optimising assembler for the MCS-51 microcontroller**, talk given by Claudio Sacerdoti Coen at International Conference on Certified Programs and Proofs, (13 December 2012)
- **Rating Disambiguation Errors**, talk given by W. Ricciotti at International Conference on Certified Programs and Proofs, (15 December 2012)
- **An executable semantics for CompCert C,** talk given by B. Campbell at International Conference on Certified Programs and Proofs, (14 December 2012)
- **Indexed Labels for Loop Iteration Dependent Costs,** talk given by Paolo Tranquilli at Quantitative Aspects of Programming Languages and Systems (23 March 2013)
- **CerCo: Certified Complexity,** talk given by I. Stark at Laboratory for Foundations of Computer Science, University of Edinburgh,(24 July 2012)

List of posters presented at international conferences:

- **Certified Complexity**. Poster presented by Claudio Sacerdoti Coen and Dominic Mulligan at FET 2011, Budapest (4-6 May 2011). Long abstract published in the proceedings.

List of seminars given to inviting institutions:

- **Certifying cost annotations in compilers**, talk given by Nicolas Ayache at the LIP6 laboratory of the University Paris 6 (18 November 2010);
- **A canonically nameless representation of binders**, talk given by Randy Pollack at U. Penn at Philadelphia (11 June 2010);
- **A canonically nameless representations of binders**, talk given by Randy Pollack at Lab PPS (Paris) to the INRIA-Rocquencourt team π⍰r2 (26 November 2010);
- **Certifying cost annotations in compilers**, talk given by Roberto Amadio at ENS/Lyon.
- **Certified Complexity,** talk given by Claudio Sacerdoti Coen at the University of Birmingham (20 December 2011).
- **Certifying cost annotations in compilers**, talk given by Nicolas Ayache at ENSTA (École Nationale Supérieure de Techniques Avancées) (13 December 2011)
- **Certifying and reasoning on cost annotations of functional programs**, talk given by Roberto Amadio at the University Paris XIII, Paris, 2012.
- **Certified Complexity**, talk given by Dominic Mulligan at the Semantics Group seminar, University of Cambridge, (31 May 2012)

Other seminars:

- **Operational Semantics in Specifications**, talk given by Brian Campbell at the Laboratory for Foundations of Computer Science, University of Edinburgh (16 March 2010)
- **An executable semantics for CompCert C**, talk given by Brian Campbell at the Scottish Programming Language seminar at Heriot-Watt University (11 November 2011).

List of papers published in international journals or proceedings of international conferences:

- **A Canonical Locally Named Representation of Binding,** R. Pollack, M. Sato and W. Ricciotti, *Journal of Automated Reasoning,* DOI:10.1007/s10817-011-9229-y.
- **Certifying and reasoning on cost annotations of functional programs,** R.M. Amadio, Y. Regis-Gianas, In proceedings of Foundations and Practical Aspects of Resource Analysis (FOPARA 2011), LNCS, 71277:72-89, DOI:10.1007/978-3-642-32495-6_5.
- **An Elementary Affine λ-Calculus with Multithreading and Side Effects**, A. Madet, R.M. Amadio., In Proceedings Typed Lambda Calculi and Applications (TLCA 2011), Springer LNCS 6690:138-152, 2011, DOI:10.1007/978-3-642-21691-6_13.
- **Certified Complexity**, R. Amadio, A. Asperti, N. Ayache, B. Campbell, D. Mulligan, R. Pollack, Y. Regis-Gianas, C. Sacerdoti Coen, I. Stark, in Procedia Computer Science, Volume 7, 2011, Proceedings of the 2nd European Future Technologies Conference and Exhibition 2011 (FET 11), 175-177, DOI:10.1016/j.procs.2011.09.054.
- **The Matita Interactive Theorem Prover**, A. Asperti, W. Ricciotti, C. Sacerdoti Coen, E. Tassi, In Automated Deduction – CADE-23, LECTURE NOTES IN COMPUTER SCIENCE, ISBN:978-3-642-22437-9, Pages 64-69, Volume 6803, Year 2011, DOI:10.1007/978-3-642-22438-6_7.
- **Certifying and Reasoning on Cost Annotations in C Programs**, N. Ayache, R.M. Amadio, Y. Régis-Gianas, in Proc. FMICS, Springer LNCS 7437: 32-46, 2012, DOI:10.1007/978-3-642-32469-7_3.
- **Rating Disambiguation Errors,** A. Asperti, W. Ricciotti, In *Proceedings of the 2nd International Conference on Certified Programs and Proofs (CPP 2012)*, LNCS, Volume 7679, Pages 240--255, Springer, 2012, DOI:10.1007/978-3-642-35308-6_19.
- **Certifying and reasoning on cost annotations of functional programs,** R.M. Amadio, Y. Régis-Gianas, in *Higher-order and Symbolic Computation,* to appear, 2012.
- **Separation Predicates: A Taste of Separation Logic in First-Order Logic**, F. Bobot, J.-C. Filliatre, in Proc. IFCEM, Springer LNCS 7635:167-181, 2012, DOI:10.1007/978-3-642-34281-3_14.
- **A Web Interface for Matita**, A. Asperti, W. Ricciotti, In Proceedings of the Conference on Intelligent Computer Mathematics (CICM 2012), LNAI, Volume 7362/2012, Pages 417-421, DOI:10.1007/978-3-642-31374-5_28, ISBN 978-3-642-31373-8.
- **Indexed Realizability for Bounded-Time Programming with References and Type Fixpoints**, A. Brunel, A. Madet, In Proc. APLAS, Springer LNCS 7705:264-279 , 2012, DOI:10.1007/978-3-642-35182-2_19.
- **A polynomial time λ-calculus with multithreading and side effects,** A. Madet in Proc. ACM-PPDP, pages 55-66, 2012, DOI:10.1145/2370776.2370785.
- **An executable semantics for CompCert C,** B. Campbell, in Proceedings of Certified Proofs and Programs 2012, LNCS 7679:60-75, 2012, DOI:10.1007/978-3-642-35308-6_8.

- **On the Correctness of an Optimising Assembler for the Intel MCS-51 Microprocessor,** D. Mulligan, C. Sacerdoti Coen, in Proceedings of Certified Proofs and Programs 2012, LNCS 7679:43-59, 2012, DOI:10.1007/978-3-642-35308-6_7.
- **Indexed Labels for Loop Iteration Dependent Costs,** P. Tranquilli, in Proceedings of the 11th International Workshop on Quantitative Aspects of Programming Languages and Systems (QAPL 2013), Rome, 23rd-24th March 2013, Electronic Proceedings in Theoretical Computer Science, 2013, to appear.

List of papers currently submitted or to be submitted to international conferences or journals:

- **An executable formalisation of the MCS-51 microprocessor in Matita**, Dominic Mulligan and Claudio Sacerdoti Coen. The paper has been previously unsuccessfully submitted to Interactive Theorem Proving 2011 and VSTTE 2012. The main motivation for rejection was that the formalization was not accompanied by further work built on top, like a certified assembler. Therefore it lacked validation and evidence of usefulness. We plan to resubmit as part of a journal paper certifying the optimizing assembler for MCS-51 developed in CerCo.
- **On the correctness of a branch displacement algorithm**, Jacob-Pieter Boender and Claudio Sacerdoti Coen. The paper has previously been submitted to Certified Proofs and Programs (CPP 2012) together with a companion paper on the certification of an assembly language that uses the branch displacement algorithm. The reviewers suggested some improvements to the presentation and pushed us to try to merge the two papers, that share some material. After trying, we came to the conclusion that there was not enough space to present all the relevant material. We plan to resubmit the paper to CPP 2013 after improving the presentation. After receiving feedback from CPP, we also plan to combine the two papers and the executable formalization of the MCS-51 microprocessor in a single journal paper.
- **Polymorphic variants in dependent type theory,** Dominic Mulligan and Claudio Sacerdoti Coen. Polymorphic variants are structural algebraic types that allow type inference in ML-like languages. They provide an efficient and elegant solution to the expression problem and they can be used to impose strong typing guarantees on data structures that present a high frequency of irregularity in the typing constraints. The latter situation was faced in CerCo in the description of the syntax of the assembly language. Polymorphic variants were exploited to solve the problem in the untrusted CerCo prototype. In order to develop the trusted CerCo prototype, we had to look for an encoding of polymorphic variants in dependent type theory. In this scenario we are not interested in type inference, but we look for structural algebraic types (opposed to inductive types which are nominal) that admit efficient and elegant notions of subset types. In CerCo we found an encoding of polymorphic variants (and of their duals, polymorphic records) that satisfy the required properties. The encoding is also a nice application of dependent types, as it is strongly based on recursively defined type formers. The encoding will be presented on 25 April 2013 at the next “Types for Proofs and Programs” meetings in Toulouse, and we plan to submit a paper to the post-proceedings. We also have in preparation a long version of the paper to be submitted in 2013 to the Journal of Formalized Reasoning. On completion the formal proof will be submitted to the Journal of Formalized Reasoning.

6 Explanation of the use of resources

Justification of major cost items and resources

Beneficiary 1 — UNIBO

In the third reporting period UNIBO was involved in WP1 and WP6 on management and dissemination activities. With regards scientific work, UNIBO was mostly active on WP4, also contributing to WP3 and WP5.

WP3: UNIBO and UEDIN collaborated at the development of the semantics based on structured traces. The requirements on structured traces were imposed in the back-end in order to allow the static prediction of costs of object code programs and also to prove that preservation of structure implies preservation of the labelled traces. Later on, more requirements were imposed to grant simulation in presence of a finite amount of stack memory. The proof that structured traces exist for all RTLabs programs is part of the front-end, and had to be updated at every change in the requirements. The last requirements and changes to the proof came almost at the end of the project.

WP4: UNIBO lead WP4, whose only task during the third period was aimed at the formalization of the proof of correctness of the back-end of the compiler. An OCaml version of the compiler is obtained by extracting the code from the formalization. The extracted code is not self contained: it needs to be linked with untrusted code that performs graph colouring and dataflow analysis, whose results may be later verified. Moreover, additional untrusted code is required to pretty-print the final and intermediate compilation results, and to provide the extracted code with the same interface of the untrusted prototype. The development or adaptation of the untrusted code has also been performed in WP4. The formal proof of the back-end proceeds in layers and a Post-Doc was assigned to each one of the layers.

WP5: UNIBO has collaborated with UPD in the development of the Frama-C plugin. Moreover, during the third version, UNIBO has carried out tests to verify the compatibility of the trusted prototype with the Frama-C plug-in, initially developed for the untrusted prototype. UNIBO has also extended the dependent labelling approach developed in WP5 during the second period to capture cost models that depend on the status of hardware components. The latter work has been only studied theoretically, and deserves an implementation.

Beneficiary 2 — UPD

In the third reporting period UPD was mainly involved in WP5 with a minor role in WP1 on management activities.

WP5: UPD worked on the following contributions:

- A method to generate and automatically certify synthetic cost bounds for C programs with simple loops.
- A Frama-C plug-in to handle C programs with pointers, inspired by ideas coming from separation logic.
- An account of the cost of memory management in a higher-order functional language which relies on a region based memory management system.
- A type system that guarantees feasible bounds for a higher-order functional language with references and threads.

WP6: UPD participated in the organisation of the dissemination events in Berlin and Rome.

Beneficiary 3 — UEDIN

In the third reporting period UEDIN was mainly involved in WP3 research activities, and some contribution to WP4 with UNIBO. UEDIN also undertook dissemination activities WP6 and standard contributions to project management WP1.

WP3: UEDIN led work on completing the formalization and proof of correctness of the compiler front-end. This was done in close collaboration with UNIBO work on the back-end, which crucially depends on structures and invariant properties provided at the interface between the compiler stages.

Two particular innovations developed in this period were essential in enabling this component of the large-scale collaborative correctness proof: formal identification of *measurable* fragments of execution traces; and certified intermediate checks of invariants. These intermediate checks are in the spirit of *translation validation*, and mean that correctness proofs of different compiler components — such as the front-end and back-end — can proceed independently in the knowledge that they will be compatible when completed.

WP4: UEDIN collaborated with UNIBO on the interface between front-end and back-end correctness proofs, and in making precise the proof requirements which the front-end must establish for the back-end to succeed. The key vehicle for this is the innovative concept of *structured traces*, proposed in the previous project period and refined to useful maturity in this period.

WP6: UEDIN contributed to the two dissemination events in this period: the CerCo workshop at HiPEAC, targeted to potential industrial stakeholders, and the joint CerCo/PROARTIS workshop at ETAPS, targeted to the scientific community.