

PROJECT FINAL REPORT

Grant Agreement number: 243381

Project acronym: CERCo

Project title: Certified Complexity

Funding Scheme: STREP

Period covered: from 01 February 2010 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>

4.1 Final publishable summary report

This section must be of suitable quality to enable direct publication by the Commission and should preferably not exceed 40 pages. This report should address a wide audience, including the general public.

The publishable summary has to include **5 distinct parts** described below:

- An executive summary (not exceeding 1 page).
- A summary description of project context and objectives (not exceeding 4 pages).
- A description of the main S&T results/foregrounds (not exceeding 25 pages),
- The potential impact (including the socio-economic impact and the wider societal implications of the project so far) and the main dissemination activities and exploitation of results (not exceeding 10 pages).
- The address of the project public website, if applicable as well as relevant contact details.

Furthermore, project logo, diagrams or photographs illustrating and promoting the work of the project (including videos, etc...), as well as the list of all beneficiaries with the corresponding contact names can be submitted without any restriction.

4.1.1 Execute summary

Problem statement: Programs can be specified using both functional constraints (what a program must do) and non functional constraints (how many resources – time, space, energy, etc – is the program allowed to use). Following the current state of the art, functional properties are verified on the high level source code combining user interaction (e.g. to state preconditions and invariants) with a multitude of interoperating techniques (invariant generators, type systems, abstract interpretation, theorem proving, etc.). Non functional properties are independently verified on the object code, but the verification requires knowledge on the functional behaviour, that needs to be reconstructed from the object code via abstract interpretation and user interaction. This situation presents several problems: 1) it may be hard for the user to provide knowledge on the execution of the object code in presence of complex optimizations, which are avoided; 2) techniques that work on the object code are not useful on the early development stages; problems detected by late cost analysis are more expensive to tackle; 3) parametric cost analysis is very hard: how can we reflect a cost that is parametric on the object code state (e.g. the value of a register or a carry bit) to a cost that the user can understand looking at the source code?; 4) the functional analysis performed on the object code using a fixed set of techniques and limited interaction yields results that are less precise than the ones obtained on the source code, affecting the precision of the estimated cost bounds.

CerCo vision and approach: We propose a reconciliation of functional and non functional analysis: both should be performed at the same time on the source code, sharing knowledge. What prevents this approach is the lack of a uniform and precise cost model for high level statements: 1) each statement occurrence is compiled in a different way and compiler optimizations may change the control flow; 2) the cost of an object code instructions depends on the state of execution dependent hardware components like pipelines and caches, which is not visible in the source code. To solve the issue, we envision a new generation of compilers able to keep track of the structure of programs during compilation and able to exploit that information to induce on the source code a non uniform, precise cost model for control blocks which is parametric on the execution history. Once the source level cost model is known, we can simply reduce non functional verification to the functional case and we can exploit any combination of techniques to automate the verification. In particular, the techniques previously used by WCET analyzers on the object code are still available on the source code, but can now be coupled and interact with more complex analysis. When the approach produces precise cost models too complex to reason about, safe approximations can be used to trade complexity with precision. Finally, the analysis on the source code can be performed even during early development stages when the components have been already specified, but not implemented yet: it is sufficient to axiomatize the non functional behaviour of all non implemented components.

Contributions: We have developed a technique, called labelling approach, to implement compilers that induce cost models on the source programs by keeping track of the code evolution in a very lightweight way. We have studied how to formally prove the correctness of compilers that implement that technique. We have implemented and partially certified in an interactive theorem prover such a compiler from C to the 8051 object code. We have implemented a Frama-C plug-in that invokes the compiler on a source program, generates cost invariants from the cost model induced by compilation and finally certifies that the program respects the generated cost invariants by calling automated theorem provers, an innovative technique in the field of cost analysis. As a case study, we show how the plug-in can automatically compute and certify the reaction time of reactive C code obtained compiling Lustre programs to C.

4.1.2 Project context and objectives

Formal methods for verification of functional properties of programs have reached a level of maturity and automation that is allowing a slow but increasing adoption rate in production environments. For safety critical code, it is getting usual to combine rigorous software engineering methodologies and testing with static analysis in order to benefit from the strong points of every approach and mitigate the weaknesses. Particularly interesting are open frameworks for the combination of different formal methods, where the programs can be progressively specified and are continuously enriched with new safety guarantees: every method contributes knowledge (e.g. new invariants) that becomes an assumption for later analysis.

The scenario for the verification of non functional properties (time spent, memory used, energy consumed) is more bleak and the future seems to be getting even worse. Most industries verify that real time systems meet their deadlines simply measuring the time spent in many runs of the systems, computing the maximum time and adding an empirical safety margin, claiming the result to be a bound for the Worst Case Execution Time of the program. Formal methods and softwares to statically analyse the WCET of programs exist, but they often produce bounds that are too pessimistic to be useful. Recent advancements in hardware architectures is all focused on the improvement of the average case performance, not the predictability of the worst case. Execution time is getting more and more dependent from the execution history, that determines the internal state of hardware components like pipelines and caches. Multicores and non uniform memory models are drastically reducing the possibility of performing static analysis in isolation, because programs are less and less time composable. Clock precise hardware models are necessary to static analysis, and getting them is becoming harder as a consequence of the increased hardware complexity.

Despite the latter scenario, the need for reliable real time systems and programs is increasing, and there is an increasing pressure from the research community towards the differentiation of hardware. The aim is the introduction of alternative hardware whose behaviour would be more predictable and more suitable to be statically analyzed, for example decoupling execution time from the execution history by introducing randomization.

In the CerCo project we do not try to address this problem, optimistically assuming that static analysis of non functional properties of programs will return to be feasible in the long term. The main objective of our work is instead to bring together static analysis of functional and non functional properties, which, according to the current state of the art, are completely independent activities with limited exchange of information: while the functional properties are verified on the source code of programs written in high level languages, the analysis of non functional properties is entirely performed on the object code to exploit clock precise hardware models.

There are two main reasons to currently perform the analysis on the object code. The first one is the obvious lack of a uniform, precise cost model for source code instructions (or even basic blocks). During compilation, high level instructions are teared apart and reassembled in context specific ways so that there is no way to identify a fragment of object code with a single high level instruction. Even the control flow of the object and source code can be very different as a result of optimizations. For instance, loop optimizations reduce the number or the order of the iterations of loops, and may assign different object code, and thus different costs, to different iterations. Despite the lack of a uniform, compilation and program independent cost model on the source language, the litterature on the analysis of non asymptotic execution time on high level languages that assumes such a model is growing and getting momentum. Its practical usefulness is doomed to be minimal, unless we can provide a replacement for such cost models. Some hope has been provided by the EMBounded European Project that compositionally compiles high level code to a byte code that is executed by an emulator with guarantees on the maximal execution time spent for each byte code instruction. The approach indeed provides a uniform model, at the price of loosing precision of the model (each cost is a pessimistic upper bound) and performance of the executed code (because the byte code is emulated compositionally instead of performing a fully non compositional compilation).

The second reason to perform the analysis on the object code is that bounding the worst case execution time of small code fragments in isolation (e.g. loop bodies) and then adding up the bounds yields very poor estimations because no knowledge on the hardware state can be assumed when executing the fragment. By analysing longer runs (e.g. by full unrolling loops) the bound obtained becomes more precise because the lack of knowledge on the initial state has less effects on longer computations.

In CerCo we propose a radically new approach to the problem: we reject the idea of a uniform cost model and we propose that the compiler, which knows how the code is translated, must return the cost model for basic blocks of high level instructions. It must do so by keeping track of the control flow modifications to reverse them and by interfacing with static analyzers. By embracing compilation, instead of avoiding it like EMBounded did, a CerCo compiler can at the same time produce efficient code and return costs that are as precise as the static analysis can be. Moreover, we allow our costs to be parametric: the cost of a block can depend on actual program data or on a summary of the execution history or on an approximated representation of the hardware state. For example, loop optimizations assign to a loop body a cost that is a function of the number of iterations performed. For another example, the cost of a loop body may be a function of the vector of stalled pipeline states, which can be exposed in the source code and updated at each basic block exit. It is parametricity that allows to analyze small code fragments without loosing precision: in the analysis of the code fragment we do not have to be ignorant on the initial hardware state. On the contrary, we can assume to know exactly which state (or mode, as WCET literature calls it) we are in.

The cost of an execution is always the sum of the cost of basic blocks multiplied by the number of times they are executed, which is a functional property of the program. Therefore, in order to perform (parametric) time analysis of programs, it is necessary to combine a cost model with control and data flow analysis. Current state of the art WCET technology performs the analysis on the object code, where the logic of the program is harder to reconstruct and most information available on the source code (e.g. types) has been lost. Imprecision in the analysis leads to useless bounds. To augment precision, the tools ask the user to provide constraints on the object code control flow, usually in the form of bounds on the number of iterations of loops or linear inequations on them. This requires the user to manually link the source and object code, translating his often wrong assumptions on the source code to object code constraints. The task is error prone and, in presence of complex optimizations, may be very hard if not impossible.

The CerCo approach has the potentiality to dramatically improve the state of the art. By performing control and data flow analysis on the source code, the error prone translation of invariants is completely avoided. It is in fact performed by the compiler itself when it induces the cost model on the source language. Moreover, any available technique for the verification of functional properties can be immediately reused and multiple techniques can collaborate together to infer and certify cost invariants for the program. Parametric cost analysis becomes the default one, with non parametric bounds used as last resorts when trading the complexity of the analysis with its precision. A priori, no technique previously used in traditional WCET is lost (e.g. full unrolling for non parametric costs): they can just be applied on the source code.

Traditional techniques for WCET that work on object code are also affected by another problem: they cannot be applied before the generation of the object code. Therefore analysis of functional properties of programs already starts in early development stages, while when analysis of non functional properties becomes possible the cost of changes to the program architecture can already be very high. Our approach already works in early development stages by axiomatically attaching costs to components that are not implemented yet.

All software used to verify properties of programs must be as bug free as possible. The trusted code base for verification is made by the code that needs to be trusted to believe that the property holds. The trusted code base of state-of-the-art WCET tools is very large: one needs to trust the control flow analyzer and the linear programming libraries it uses and also the formal models of the hardware. In CerCo we are moving the control flow analysis to the source code and we are introducing a non standard compiler too. To reduce the trusted code base, we implemented a prototype and a static analyzer in an interactive theorem prover, which was used to certify that the cost computed on the source code is indeed the one actually spent by the hardware. Formal models of the hardware and of the high level source languages were also implemented in the interactive theorem prover. Control flow analysis on the source code has been obtained using invariant generators, tools to produce proof obligations from generated invariants and automatic theorem provers to verify the obligations. If the automatic

provers are able to generate proof traces that can be independently checked, the only remaining component that enters the trusted code base is an off-the-shelf invariant generator which, in turn, can be proved correct using an interactive theorem prover. Therefore we achieve the double objective of allowing to use more off-the-shelf components (e.g. provers and invariant generators) while reducing the trusted code base at the same time.

Summary of the CerCo objectives. To summarize, the goal of CerCo is to reconcile functional with non functional analysis by performing them together on the source code, sharing common knowledge about execution invariants. We want to achieve the goal implementing a new generation of compilers that induce a parametric, precise cost model for basic blocks on the source code. The compiler should be certified using an interactive theorem prover to minimize the trusted code base of the analysis. Once the cost model is induced, off-the-shelf tools and techniques can be combined together to infer and prove parametric cost bounds.

The long term benefits of the CerCo vision are expected to be:

1. the possibility to perform static analysis during early development stages
2. parametric bounds made easier
3. the application of off-the-shelf techniques currently unused for the analysis of non functional properties, like automated proving and type systems
4. simpler and safer interaction with the user, that is still asked for knowledge, but on the source code, with the additional possibility of actually verifying the provided knowledge
5. a reduced trusted code base
6. the increased accuracy of the bounds themselves.

The long term success of the project is hindered by the increased complexity of the static prediction of the non functional behaviour of modern hardware. In the time frame of the European contribution we have focused on the general methodology and on the difficulties related to the development and certification of a cost model inducing compiler.

4.1.3 Main S&T results

We will now review the main S&T results achieved in the CerCo project. We will address them in the following order:

- 1. The (basic) labelling approach.** It is the main technique that underlies all the developments in CerCo. It allows to track the evolution of basic blocks during compilation in order to propagate the cost model from the object code to the source code without losing precision in the process.
- 2. Dependent labelling.** The basic labelling approach assumes a bijective mapping between object code and source code $O(1)$ blocks (called basic blocks). This assumption is violated by many program optimizations (e.g. loop peeling and loop unrolling). It also assumes the cost model computed on the object code to be non parametric: every block must be assigned a cost that does not depend on the state. This assumption is violated by stateful hardware like pipelines, caches, branch prediction units. The dependent labelling approach is an extension of the basic labelling approach that allows to handle parametric cost models. We showed how the method allows to deal with loop optimizations and pipelines, and we speculated about its applications to caches.
- 3. Techniques to exploit the induced cost model.** Every technique used for the analysis of functional properties of programs can be adapted to analyse the non functional properties of the source code instrumented by compilers that implement the labelling approach. In order to gain confidence in this claim, we showed how to implement a cost invariant generator combining abstract interpretation with separation logic ideas. We integrated everything in the Frama-C modular architecture, in order to automatically compute proof obligations from the functional and the cost invariants and to use automatic theorem provers to proof them. This is an example of a new technique that is not currently exploited in WCET analysis. It also shows how precise functional invariants benefits the non functional analysis too. Finally, we show how to fully automatically analyze the reaction time of Lustre nodes that are first compiled to C using a standard Lustre compiler and then processed by a C compiler that implements the labelling approach.
- 4. The CerCo compiler.** This is a compiler from a large subset of the C program to 8051/8052 object code. The compiler implements the labelling approach and integrates a static analyzer for 8051 executables. The latter can be implemented easily and does not require dependent costs because the 8051 microprocessor is a very simple processor whose instructions all have a constant cost. It was picked to separate the issue of exact propagation of the cost model from the target to the source language from the orthogonal problem of the static analysis of object code that requires approximations or dependent costs. The compiler comes in several versions: some are prototypes implemented directly in OCaml, and they implement both the basic and dependent labelling

approaches; the final version is extracted from a Matita partial certification and at the moment implements only the basic approach.

5. A formal partial certification of the CerCo compiler. We implemented the CerCo compiler in the interactive theorem prover Matita and we certified that the cost model induced on the source code correctly and precisely predicts the object code behaviour. We actually induce two cost models, one for time and one for stack space used. We show the correctness of the prediction only for those programs that do not exhaust the available machine stack space, a property that we can statically analyze on the source code using our Frama-C tool. The preservation of functional properties is an assumption of the formal proof that was not proved in CerCo, but that is expected to be unsurprising: other projects have already certified similar compilers considering only the functional semantics preservation. In order to complete the proof for the non functional properties, we have introduced a new semantics for programming languages based on a new kind of structured observables with the relative notions of forward similarity and the study of the intentional consequences of forward similarity. We have also introduced a unified representation for back-end intermediate languages that was exploited to provide a uniform proof of forward similarity.

6. Additional theoretical investigations on computational models and the prediction of program complexity.

- a) The certification of the CerCo compiler formally shows what is common knowledge about computer scientists: it is possible to compilers/interpreters that preserve the asymptotic costs of programs. The property is supposed to hold also for abstract computational models like Turing machines, but there is currently no evidence about it: all known implementations of universal machines do not respect the property. More generally, most of computability theory is focused on extensional (or functional) properties only and little is known on the intensional one. In CerCo we provided the first formal certification of a universal Turing machine where we try to preserve the complexity, reducing the problem to the preservation of the complexity in the encoding/decoding of the machine to be executed. The problem is not settled yet.
- b) We studied extension Implicit Computational Complexity (ICC) techniques. Classical ICC results provide restrictions of purely functional languages that allow to write only programs that belong to a certain complexity class (e.g. elementary, exponential, log-space). They do not provide precise bounds and thus do not need precise cost models. In CerCo we studied how to extend one technique to cover imperative, concurrent languages, achieving elementary or exponential complexity. Then we started the study of a new semantical interpretation that should eventually allow to provide concrete, non exponential, bounds, using precise cost models as the ones that we can provide in CerCo.

4.1.3.1: The (basic) labelling approach.

Problem statement: given a source program P , we want to obtain an *instrumented* source program P' , written in the same programming language, and the object code O such that: 1) P' is obtained by inserting into P some additional instructions to update global cost informations like the amount of time spent during execution or the maximal stack space required; 2) P and P' must have the same functional behaviour, i.e., they must produce that same output and intermediate observables; 3) P and O must have the same functional behaviour; 4) after execution and in interesting points during execution, the cost information computed by P' must be an upper bound of the one spent by O to perform the corresponding operations (*soundness property*); 5) the difference between the costs computed by P' and the execution costs of O must be bounded by a program dependent constant (*precision property*).

The labeling software components: we solve the problem in four stages, implemented by four software components that are used in sequence.

1. The first component *labels* the source program P by injecting *label emission statements* in the program in appropriate positions. The set of labels with their positions is called *labelling*. The syntax and semantics of the source programming language is augmented with label emission statements. The statement "EMIT l " behaves like a NOP instruction that does not affect the program state or control flow, but it changes the semantics by making the label l observable. Therefore the observables of a run of a program becomes a stream of label emissions: $l_1 \dots l_n$, called the program trace. We clarify the conditions that the labelling must respect later.
2. The second component is a *labelling preserving compiler*. It can be obtained from an existing compiler by adding label emission statements to every intermediate language and by propagating label emission statements during compilation. The compiler is correct if it preserves both the functional behaviour of the program and the generated traces. We may also ask that the function that erases the cost emission statements commute with compilation. This optional property grants that the labelling does not interfere with the original compiler behaviour. A further set of requirements will be added later.
3. The third component is a *static object code analyzer*. It takes in input the object code augmented with label emission statements and it computes, for every such statement, its *scope*. The scope of a label emission statement is the tree of instructions that may be executed after the statement and before a new label emission statement is found. It is a tree and not a sequence because the scope may contain a branching statement. In order to grant that such a finite tree exists, the object code must not contain any loop that is not broken by a label emission statement. This is the first requirement of a *sound labelling*. The analyzer fails if the labelling is unsound. For each scope, the analyzer computes an upper bound of the execution time required by the scope, using the maximum of the costs of the two branches in case of a conditional statement. Finally, the

analyzer computes the cost of a label by taking the maximum of the costs of the scopes of every statement that emits that label.

4. The fourth and last component takes in input the cost model computed at step 3 and the labelled code computed at step 1. It outputs a source program obtained by replacing each label emission statement with a statement that increments the global cost variable with the cost associated to the label by the cost model. The obtained source code is the *instrumented source code*.

Correctness: Requirements 1, 2 and 3 of the program statement obviously hold, with 2 and 3 being a consequence of the definition of a correct labelling preserving compiler. It is also obvious that the value of the global cost variable of an instrumented source code is at any time equal to the sum of the costs of the labels emitted by the corresponding labelled code. Moreover, because the compiler preserves all traces, the sum of the costs of the labels emitted in the source and target labelled code are the same. Therefore, to satisfy the 4th requirement, we need to grant that the time taken to execute the object code is equal to the sum of the costs of the labels emitted by the object code. We collect all the necessary conditions for this to happen in the definition of a *sound labelling*: a) all loops must be broken by a cost emission statement; b) all program instructions must be in the scope of some cost emission statement. To satisfy also the 5th requirement, additional requirements must be imposed on the object code labelling to avoid all uses of the maximum in the cost computation: the labelling is *precise* if every label is emitted at most once and both branches of a conditional jump start with a label emission statement.

The correctness and precision of the labelling approach only rely on the correctness and precision of the object code labelling. The simplest, but not necessary, way to achieve them is to impose correctness and precision requirements also on the initial labelling produced by the first software component, and to ask the labelling preserving compiler to preserve these properties too. The latter requirement imposes serious limitations on the compilation strategy and optimizations: the compiler may not duplicate any code that contains label emission statements, like loop bodies. Therefore several loop optimizations like peeling or unrolling are prevented. Moreover, precision of the object code labelling is not sufficient per se to obtain global precision: we also implicitly assumed the static analysis to be able to associate a precise constant cost to every instruction. This is not possible in presence of stateful hardware whose state influences the cost of operations, like pipelines and caches. In the next section we will see an extension of the basic labelling approach to cover this situation.

The labelling approach described in this section can be applied equally well and with minor modifications to imperative and functional languages. In the CerCo project, we have tested it on a simple imperative language without functions (a While language), to a subset of C and to two compilation chains for a purely functional higher order language. The two main changes required to deal with functional languages are: 1) because global variables and updates are not available, the instrumentation phase produces monadic code to “update” the global costs; 2) the requirements for a sound and precise labelling of the source code must be changed when the compilation is

based on CPS translations. In particular, we need to introduce both labels emitted before a statement is executed and labels emitted after a statement is executed. The latter capture code that is inserted by the CPS translation and that would escape all label scopes.

Phases 1, 2 and 3 can be applied as well to logic languages (e.g. Prolog). However, the instrumentation phase cannot: in standard Prolog there is no notion of (global) variable whose state is not retracted during backtracking. Therefore, the cost of executing computations that are later backtracked would not be correctly counted in. Any extension of logic languages with non backtrackable state should support the labelling approach.

4.1.3.2: Dependent labelling.

The core idea of the basic labelling approach is to establish a tight connection between basic blocks executed in the source and target languages. Once the connection is established, any cost model computed on the object code can be transferred to the source code, *without affecting the code of the compiler or its proof*. In particular, it is immediate that we can also transport cost models that associate to each label functions from hardware state to natural numbers. However, a problem arises during the instrumentation phase that replaces cost emission statements with increments of global cost variables. The global cost variable must be incremented with the result of applying the function associated to the label to the hardware state at the time of execution of the block.

The hardware state comprises both the “functional” state that affects the computation (the value of the registers and memory) and the “non functional” state that does not (the pipeline and caches content for example). The former is in correspondence with the source code state, but reconstructing the correspondence may be hard and lifting the cost model to work on the source code state is likely to produce cost expressions that are too hard to reason on. Luckily enough, in all modern architectures the cost of executing single instructions is either independent of the functional state or the jitter --- the difference between the worst and best case execution times --- is small enough to be bounded without losing too much precision. Therefore we can concentrate on dependencies over the “non functional” parts of the state only.

The non functional state has no correspondence in the high level state and does not influence the functional properties. What can be done is to expose the non functional state in the source code. We just present here the basic intuition in a simplified form: the technical details that allow to handle the general case are more complex and can be found in the CerCo deliverables. We add to the source code an additional global variable that represents the non functional state and another one that remembers the last labels emitted. The state variable must be updated at every label emission statement, using an update function which is computed during static analysis too. The update function associates to each label a function from the recently emitted labels and

old state to the new state. It is computed composing the semantics of every instruction in a basic block and restricting it to the non functional part of the state.

Not all the details of the non functional state needs to be exposed, and the technique works better when the part of state that is required can be summarized in a simple data structure. For example, to handle simple but realistic pipelines it is sufficient to remember a short integer that encodes the position of bubbles (stuck instructions) in the pipeline. In any case, the user does not need to understand the meaning of the state to reason over the properties of the program. Moreover, at any moment the user or the invariant generator tools that analyze the instrumented source code produced by the compiler can decide to trade precision of the analysis with simplicity by approximating the parametric cost with safe non parametric bounds. Interestingly, the functional analysis of the code can determine which blocks are executed more frequently in order to approximate more aggressively the ones that are executed less.

The idea of dependent labelling can also be applied to allow the compiler to duplicate blocks that contain labels (e.g. to allow loop optimizations). The effect of duplication is to assign a different cost to the different occurrences of a duplicated label. For example, loop peeling turns a loop into the concatenation of a copy the loop body (that executes the first iteration) with the conditional execution of the loop (for the successive iterations). Because of further optimizations, the two copies of the loop will be compiled differently, with the first body usually taking more time. By introducing a variable that keep tracks of the iteration number, we can associate to the label a cost that is a function of the iteration number. The same technique works for loop unrolling without modifications: the function will assign a cost to the even iterations and another cost to the odd ones. The actual work to be done consists in introducing in the source code and for each loop a variable that counts the number of iterations. The loop optimization code that duplicate the loop bodies must also modify the code to propagate correctly the update of the iteration numbers. The technical details are more complex and can be found in the CerCo reports and publications. The implementation, however, is quite simple and the changes to a loop optimizing compiler are minimal.

4.1.3.3: Techniques to exploit the induced cost model.

We review the cost synthesis techniques developed in the project.

The starting hypothesis is that we have a certified methodology to annotate blocks in the source code with constants which provide a sound and possibly precise upper bound on the cost of executing the blocks after compilation to object code.

The principle that we have followed in designing the cost synthesis tools is that the synthetic bounds should be expressed and proved within a general purpose tool built to reason on the source code. In particular, we rely on the Frama – C tool to reason on C code and on the Coq proof-assistant to reason on higher-order functional programs.

This principle entails that:

- The inferred synthetic bounds are indeed correct as long as the general purpose tool is.
- There is no limitation on the class of programs that can be handled as long as the user is willing to carry on an interactive proof.

Of course, automation is desirable whenever possible. Within this framework, automation means writing programs that give hints to the general purpose tool. These hints may take the form, say, of loop invariants/variants, of predicates describing the structure of the heap, or of types in a light logic. If these hints are correct and sufficiently precise the general purpose tool will produce a proof automatically, otherwise, user interaction is required.

The Cost plug-in and its application to the Lustre compiler

Frama – C is a set of analysers for C programs with a specification language called ACSL. New analyses can be dynamically added through a plug-in system. For instance, the Jessie plug-in allows deductive verification of C programs with respect to their specification in ACSL, with various provers as back-end tools.

We developed the CerCo Cost plug-in for the Frama – C platform as a proof of concept of an automatic environment exploiting the cost annotations produced by the CerCo compiler. It consists of an ocaml program which in first approximation takes the following actions: (1) it receives as input a C program, (2) it applies the CerCo compiler to produce a related C program with cost annotations, (3) it applies some heuristics to produce a tentative bound on the cost of executing the C functions of the program as a function of the value of their parameters, (4) the user can then call the Jessie tool to discharge the related proof obligations.

In the following we elaborate on the soundness of the framework and the experiments we performed with the Cost tool on the C programs produced by a Lustre compiler.

Soundness The soundness of the whole framework depends on the cost annotations added by the CerCo compiler, the synthetic costs produced by the Cost plug-in, the verification conditions (VCs) generated by Jessie, and the external provers discharging the VCs. The synthetic costs being in ACSL format, Jessie can be used to verify them. Thus, even if the added synthetic costs are incorrect (relatively to the cost annotations), the process in its globality is still correct: indeed, Jessie will not validate incorrect costs and no conclusion can be made about the WCET of the program in this case. In other terms, the soundness does not really depend on the action of the Cost plug-in, which can in principle produce any synthetic cost. However, in order to be able to actually prove a WCET of a C function, we need to add correct annotations in a way that Jessie and subsequent automatic provers have enough information to deduce their validity. In practice this is not straightforward even for very simple programs composed of branching and assignments (no loops and no recursion) because a fine analysis of the VCs associated with branching may lead to a complexity blow up.

Experience with Lustre Lustre is a data-flow language to program synchronous systems and the language comes with a compiler to C. We designed a wrapper for supporting Lustre files.

The C function produced by the compiler implements the step function of the synchronous system and computing the WCET of the function amounts to obtain a bound on the reaction time of the system. We tested the Cost plug-in and the Lustre wrapper on the C programs generated by the Lustre compiler. For programs consisting of a few hundreds loc, the Cost plug-in computes a WCET and Alt – Ergo is able to discharge all VCs automatically.

Handling C programs with simple loops

The cost annotations added by the CerCo compiler take the form of C instructions that update by a constant a fresh global variable called the cost variable. Synthesizing a WCET of a C function thus consists in statically resolving an upper bound of the difference between the value of the cost variable before and after the execution of the function, i.e. find in the function the instructions that update the cost variable and establish the number of times they are passed through during the flow of execution. In order to do the analysis the plugin makes the following assumptions on the programs:

- No recursive functions.
- Every loop must be annotated with a variant. The variants of ‘for’ loops are automatically inferred.

The plugin proceeds as follows.

- First the callgraph of the program is computed. If the function f calls the function g then the function g is processed before the function f .
- The computation of the cost of the function is performed by traversing its control flow graph. The cost at a node is the maximum of the costs of the successors.
- In the case of a loop with a body that has a constant cost for every step of the loop, the cost is the product of the cost of the body and of the variant taken at the start of the loop.
- In the case of a loop with a body whose cost depends on the values of some free variables, a fresh logic function f is introduced to represent the cost of the loop in the logic assertions. This logic function takes the variant as a first parameter. The other parameters of f are the free variables of the body of the loop. An axiom is added to account the fact that the cost is accumulated at each step of the loop.

- The cost of the function is directly added as post-condition of the function

The user can influence the annotation by different means:

- By using more precise variants.

- By annotating function with cost specification. The plugin will use this cost for the function instead of computing it.

C programs with pointers

When it comes to verifying programs involving pointer-based data structures, such as linked lists, trees, or graphs, the use of traditional first-order logic to specify, and of SMT solvers to verify, shows some limitations. Separation logic is then an elegant alternative. Designed at the turn of the century, it is a program logic with a new notion of conjunction to express spatial heap separation. Separation logic has been implemented in dedicated theorem provers such as Smallfoot or VeriFast. One drawback of such provers, however, is to either limit the expressiveness of formulas (e.g. to the so-called symbolic heaps), or to require some user-guidance (e.g. open/close commands in Verifast).

In an attempt to conciliate both approaches, we introduced the notion of separation predicates. The approach consists in reformulating some ideas from separation logic into a traditional verification framework where the specification language, the verification condition generator, and the theorem provers were not designed with separation logic in mind. Separation predicates are automatically derived from user-defined inductive predicates, on demand. Then they can be used in program annotations, exactly as other predicates, i.e., without any constraint. Simply speaking, where one would write $P * Q$ in separation logic, one will here ask for the generation of a separation predicate sep and then use it as $P \wedge Q \wedge sep(P, Q)$. We have implemented separation predicates within the Jessie plug-in and tested it on non-trivial case studies (e.g. the composite pattern from the VACID-0 benchmark). In these cases, we achieve a fully automatic proof using three existing SMT solver.

We have also used the separation predicates to reason on the cost of executing simple heap manipulating programs such as an in-place list reversal.

The cost of memory management for functional programs

Modern high level programming languages implicitly handle memory. In particular, all functional programming languages do that. We already said how we were able to simply modify the labelling approach to deal with an higher order functional language without any implicit form of memory management. We address here how to cope with the latter.

In a realistic implementation of a functional programming language, the runtime environment usually includes a garbage collector. In spite of considerable progress in real-time garbage collectors it seems to us that such collectors do not offer yet a viable path to a certified and usable WCET prediction of the running time of functional programs. As far as we know, the cost predictions concern the amortized case rather than the worst case and are supported more by experimental evaluations than by formal proofs.

The approach we have adopted instead, following the seminal work of Tofte et al., is to enrich the last calculus of the compilation chain : (1) with a notion of memory region, (2) with operations to allocate and dispose memory regions, and (3) with a type and effect system that guarantees the safety of the dispose operation. This allows to further extend the compilation chain mentioned above and then to include the cost of safe memory management in our analysis. Actually, because effects are intertwined with types, what we have actually done, following the work of Morrisett et al., is to extend a typed version of the compilation chain. An experimental validation of the approach is left for future work and it would require the integration of region-inference algorithms such as those developed by Aiken et al. in the compilation chain.

4.1.3.4 The CerCo Compiler

In CerCo we have developed a certain number of cost preserving compilers based on the labelling approach. Excluding an initial certified compiler for a While language, all remaining compilers target realistic source languages --- a pure higher order functional language and a large subset of C with pointers, gotos and all data structures --- and real world target processors --- MIPS and the Intel 8051/8052 processor family. Moreover, they achieve a level of optimization that ranges from moderate (comparable to gcc level 1) to intermediate (including loop peeling and unrolling, hoisting and late constant propagation). The so called Trusted CerCo Compiler is the only one that was implemented in the interactive theorem prover Matita and partially certified. The code distributed is obtained extracting OCaml code from the Matita implementation. In the rest of this section we will only focus on the Trusted CerCo Compiler, that only targets the C language and the 8051/8052 family, and that does not implement the advanced optimizations yet. Its user interface, however, is the same as the one of the other versions, in order to trade safety with additional performances. In particular, the Framac CerCo plug-in can work without recompilation with all compilers.

The (trusted) CerCo compiler implements the following optimizations: cast simplification, constant propagation in expressions, liveness analysis driven spilling of registers, dead code elimination, branch displacement, tunnelling. The two latter optimizations are performed by the optimizing assembler which is part of the compiler. The back-end of the compiler works on three addresses instructions, preferred to static single assignment code for the simplicity of the formal certification.

The CerCo compiler is loosely based on the CompCert compiler, a recently developed certified compiler from C to the PowerPC, ARM and x86 microprocessors. Contrarily to CompCert, both the CerCo code and its certification are open source. Some data structures and language definitions for the front-end are directly taken from CompCert, while the back-end is a redesign and reimplementaion of a didactic compiler from Pascal to MIPS used by Francois Pottier for a course at the Ecole Polytechnique.

Following CompCert tradition, the compiler is organized in an unusually large number of intermediate passes, all responsible for just one change in the semantics of the source and target languages. Introducing multiple passes has minor performance implications on modern hardware and it allows to simplify the simulation proofs. The

first three intermediate languages for the front-end. They are syntactically and semantically quite different between each other. For example, in the first language we find the traditional looping structures of C, in the second all loops are infinite loops (built with GOTOs) interrupted using BREAKs and in the third one the code is organized as a graph of statements where loops become loops in the graph. The four back-end languages, instead, have a more similar syntax.

Departing from CompCert, we do not provide a stand alone syntax and semantics for every back-end language. Instead, we developed a generic representation of back-end languages as a parametric data type that can be instantiated to the wanted language. The generic representation allows to multiply the number of passes without increasing too much the code size. For example, we also provide a single generic semantics for the generic representation, parameterized over pass specific details.

Other departures from CompCert are:

1. all of our intermediate languages include label emitting instructions to implement the labelling approach, and the compiler preserves execution traces.
2. the target language of CompCert is an assembly language with additional macro-instructions to be expanded before assembly; for CerCo we need to go all the way down to object code in order to perform the static analysis. Therefore we developed also an optimizing assembler and a static analyzer, all integrated in the compiler.
3. to avoid implementing a linker and a loader, we do not support separate compilation and external calls. Adding a linker and a loader is a transparent process to the labelling approach and should create no unknown problem
4. we target an 8-bit processor. Targeting an 8 bit processor requires several changes and an increased code size, but it is not fundamentally more complex. The proof of correctness, however, becomes much harder.
5. we target a microprocessor that has a non uniform memory model, which is still often the case for microprocessors used in embedded systems and that is becoming common again in multi-cores processors. Therefore the compiler has to keep track of the position of data and it must move data between memory regions in the proper way. Also the size of pointers to different regions is not uniform. In our case, an additional difficulty was that the space available for the stack in internal memory in the 8051 is tiny, allowing only a minor number of nested calls. To support full recursion in order to test the CerCo tools also on recursive programs, the compiler manually manages a stack in external memory.
6. while there is a rough correspondence between CompCert and CerCo back-end passes, the order of the passes is permuted. In the future we want to explore how to exploit our generic back-end language representation to try to freely compose and permute passes.

4.1.3.5 A formal partial certification of the CerCo compiler

The Trusted CerCo Compiler has been implemented and certified using the interactive theorem prover Matita. In this section we briefly hint at the exact correctness statement and at the main ingredients of the proof. Details on the proof techniques employed and further informations can be collected from the CerCo deliverables and papers.

The statement

The most natural statement of correctness for our complexity preserving compiler is that the time spent for execution by a terminating object code program should be the time predicted on the source code by adding up the cost of every label emission statement. This statement, however, is too naïve to be useful for real world real time programs like those used in embedded systems.

Real time programs are often reactive programs that loop forever responding to events (inputs) by performing some computation followed by some action (output) and the return to the initial state. For looping programs the overall execution time does not make sense. The same happens for reactive programs that spend an unpredictable amount of time in I/O. What is interesting is the reaction time that measure the time spent between I/O events. Moreover, we are interested in predicting and ruling out programs that crash running out of space on a certain input.

Therefore we need to look for a more complex statement that talks about sub-runs of a program. The most natural statement is that the time predicted on the source code and spent on the object code by two corresponding sub-runs are the same. The problem to solve to make this statement formal is how to identify the corresponding sub-runs and how to single out those that are meaningful.

The solution we found is based on the notion of measurability. We say that a run has a measurable sub-run when *both the prefix of the sub-run and the sub-run do not exhaust the stack space, the number of function calls and returns in the sub-run is the same, the sub-run does not perform any I/O and the sub-run starts with a label emission statement and ends with a return or another label emission statements*. The stack usage is estimated using the stack usage model that is computed by the compiler.

The statement that we want to formally prove is that *for each C run with a measurable sub-run there exists an object code run with a sub-run such that the observables of the pairs of prefixes and sub-runs are the same and the time spent by the object code in the sub-run is the same as the one predicted on the source code using the time cost model generated by the compiler*.

We briefly discuss the constraints for measurability. Not exhausting the stack space is a clear requirement of meaningfulness of a run, because source programs do not crash for lack of space while object code ones do. The balancing of function calls/returns is a requirement for precision: the labelling approach allows the scope of

label emission statements to extend after function calls to minimize the number of labels. Therefore a label pays for all the instructions in a block, excluding those executed in nested function calls. If the number of calls/returns is unbalanced, it means that there is a call we have not returned to that could be followed by additional instructions whose cost has already been taken in account. To make the statement true (but less precise) in this situation, we could only say that the cost predicted on the source code is a safe bound, not that it is exact. The last condition on the entry/exit points of a run is used to identify sub-runs whose code correspond to a sequence of blocks that we can measure precisely. Any other choice would start/end the run in the middle of a block and we would be forced again to weaken the statement taking as a bound the cost obtained counting in all the instructions that precede the starting one in the block/follow the final one in the block.

I/O operations can be performed in the prefix of the run, but not in the measurable sub-run. Therefore we prove that we can predict reaction times, but not I/O times, as it should be.

Dealing with function calls

For obvious complexity reasons, the proof of correctness must be modular in the compilation passes. We are looking for independent proofs of correctness for the single passes that, composed together, should yield the global proof. Achieving modularity for a language that does not have function calls and function pointers is quite simple. The most difficult part of the proof is a forward simulation lemma that says that for each pass and each (sub)-run in the source program, there is a corresponding (sub)-run in the target program that emits the same observables.

The previous lemma holds also when function calls and function pointers are introduced, but make the lemma too weak. As we already explained, the labelling approach assumes that the cost of a function call will be paid by the labels of the function body, while the instructions after the call are paid by the label before the call. In order to predict correctly the costs, we therefore need to assume that: 1) every function body starts with a label; 2) every terminating function returns just after the call. The problem with function pointers is that property 1 is no longer a syntactic property that can be checked at compile time and that can be propagated by compilation: a function pointer allows to jump everywhere in the code. The problem with function calls in general is that property 2 is no longer valid and cannot be checked at compile time for every back-end language where the RET statement passes control to the address on top of the stack, which is not necessarily the one after the last function call. It is even simple to design compilers that intentionally change the order of instructions and introduce code to manipulate the stack so that function do not return to the call point, but the compiler preserves all functional properties anyway.

In order to strengthen the statement of the lemma, we need identify those runs that satisfy the assumption above and prove the statement for those runs only (both in the source and target language). Note that both properties are automatically guaranteed by the semantics of the high level source language, which is structured and more

compositional. Therefore what we are doing is to try to identify the low-level runs that keep part of the structure of the original program (a clearly intensional or non functional property).

The solution we found to single out those runs consisted in proposing a new style for describing the semantics of programs. We call it the structured traces semantics. A structured traces semantics maps programs to structured traces, that are streams of observables where additional structure is imposed on the stream. For example, we impose that every converging function call should return immediately after the call itself and that every call must be followed by a cost emission statement. We also integrate in the structure additional requirements that capture on traces all the syntactic requirements that soundness and precision of labelling impose on the code. For example, we drop the precision requirement that every branch in the code must be followed by a cost emission statement and we put add it to the structure of the trace: a conditional branch statement in a structured trace must be followed by a cost emission statement. The aim is to drop the proof that the compiler preserves precision and correctness, incorporating (and generalizing) it in the proof of forward simulation of structured traces. The generalization is actually genuine because, for example, with the new approach dead code does not need to be soundly labelled since it produces no runs.

The proof sketch

The formal proof, based on the structured traces just described, is achieved combining several stand alone results. Those starred are standard proofs that show that functional properties are preserved. Not all the starred proofs have been completed during the project, while we completed the ones that are issued by the labelling methodology and deal with non functional properties:

- 1 Correctness of the instrumentation algorithm: the source code augmented with label emission statements in proper places simulates the original program and the labelling obtained is sound and precise
- 2 * Existence of a standard functional forward simulation between the source code and the last code of the front-end.
- 3 Correctness of an algorithm that checks if the program generated by the front-end is soundly and precisely labelled. Compilation is aborted otherwise. The algorithm can be dropped replacing the proof with the longer and harder proof of preservation of soundness and precision of labelling by the front-end.
- 4 Existence of structured traces: we prove that every measurable sub-run of a soundly and precisely labelled program written in the last front-end language admits a corresponding structured trace such that the sub-run and the trace have the same observables and their prefixes have the same property too.
- 5 The definition of a forward simulation relation between structured trace.

- 6 The proof that similar structured traces produce that same set of observables.
- 7 The existence of a forward simulation (in the sense of 5) for every back-end pass. The proof relies on the hypothesis that the stack usage predicted for a structured trace does not exceed the available machine stack. The proofs have been simplified by factorizing out some common results:
 - 7.1 A complex proof that reduces the existence of a forward simulation (which is a global condition) to a set of local conditions to be checked for every source step. For example, in a structured trace all RETs must return control just after the corresponding CALL statement (the global condition). This can only happen if the CALL pushes the right address on the stack and if all instructions in the function body manipulate the stack in the expected way (the local conditions). The local conditions combine together standard functional forward simulation conditions with additional ones that still mention structured traces.
 - 7.2 Thanks to our uniform and generic representation of back-end languages, we were able to provide a generic representation of back-end passes that covers most of them. A second complex but generic proof allows to reduce the local proof obligations generated by 7.1 to a new set of local proof obligations that are exactly the functional ones with additional very simple ones that do not mention structured traces any longer.
 - 7.3 * Finally, one needs to inhabit all the proof obligations generated by 7.2, i.e. it must show the preservation of the functional properties.
- 8 Correctness of static analysis: the actual execution cost of a structured object code trace is the sum of the statically computed cost of each label in the trace.

The correctness statement is then proved as follow: by 1,3 and 4 there exists an intermediate structure trace (and prefix) that have the same observables of the source program and the same functional behaviour; by 6 and 7 there exists a structured object code (and prefix) that have the same observables of the source program and the same functional behaviour. Therefore any cost model that is correct for the object code can be used to predict the cost on the source code. Finally, by 8 the cost model computed by the compiler is correct for the object code and thus any cost analysis performed on the source code is also correct.

Additional work

The executable semantics of C and of the 8051 object code are part of the trusted code base of CerCo. Errors in the semantics may hide errors in the compiler. Therefore it is of maximal importance that the semantics are as accurate and error free as possible.

Executability is an important key to the minimization of errors because it allows to run in parallel tests on the semantics and on the machine (or a C emulator) and

compare the executions. Debugging the two semantics took a considerable time in the project.

For the C semantics we had another way to raise confidence: we have formally proved in Matita the equivalence of our executable semantics for C with the non-executable one developed in CompCert and ported to Matita. Non executable semantics are often picked in formal verifications because they allow to simplify the simulation proof. This was also the case for CompCert.

A very surprising discovery had been bugs in the CompCert semantics, one of which actually masking a bug in the CompCert compiler itself.

The discover represents further evidence of the benefits of an executable semantics that can be actually tested. Contrary to CompCert, all the semantics in use in CerCo are executable.

4.1.3.6 Additional theoretical investigations on computational models and the prediction of program complexity

In the introduction of Section 4.1.3 we have already hinted at additional theoretical investigations that have been performed in the CerCo project and that are loosely connected to the labelling approach. Detailed informations are provided in the CerCo papers. We only want here to recall the motivations for the investigation.

Complexity preserving universal Turing machines

Most computational models that are studied theoretically are considered to be realistic up to the “minor” detail of allowing an infinite amount of resources. For example, Turing machines and lambda calculi allow infinite memory, process algebras allow infinite channels, RAMs allow for unbounded integers to be stored in registers, etc. These unrealistic assumptions make the problems interesting. For example, termination is trivially decidable when program states are bounded. Moreover, often in practice a large number of resources only do not make practical any analysis that is rooted in the finiteness of the resources. Continuing the example, the number of different states of a real world computer is no huge that it is unfeasible to build the transition matrix between all pair of states to determine termination this way.

Functional properties are, almost by definition, insensitive to resources. Therefore most computability theory, that deals only with functional properties, continues to hold when applied to realistic models that have limited resources. Curiously, almost nothing is known theoretically for non functional properties. The object of study of CerCo is the] preservation of the exact computational cost during compilation. As expected, we have formally certified that such preservation is possible when working with real world languages, hardware and compilers. The natural question is if the same holds for the usual theoretical computational models with unbounded resources. An even simpler problem, which is the one we started to address in the timeframe of CerCo, is the existence of universal interpreters that preserve the computational costs.

Feasible bounds by light typing

In our experience, the cost analysis of higher-order programs requires human intervention both at the level of the specification and of the proofs. One path to automation consists in devising programming disciplines that entail feasible bounds (polynomial time). The most interesting approaches to this problem, even if maybe not the most practical ones at the moment, build on light versions of linear logic. Our main contribution is to devise a type system that guarantees feasible bounds for a higher-order call-by-value functional language with references and threads.

The first proof of this result relies on a kind of standardisation theorem and it is of a combinatorial nature. More recently, we have shown that a proof of a similar result can be obtained by semantic means building on the so called quantitative realizability models proposed by Dal Lago and Hofmann. We believe this semantic setting is particularly appropriate because it allows to reason both on typed and untyped programs. Thus one can imagine a framework where some programs are feasible 'by typing' while others are feasible as a result of an 'interactive proof' of the obligations generated by quantitative realizability. Beyond building such a framework, an interesting issue concerns the certification of concrete bounds at the level of the compiled code. This has to be contrasted with the current state of the art in implicit computational complexity where most bounds are asymptotic and are stated at the level of the source code.

4.1.4 Potential impact, main dissemination activities and exploitation of results

4.1.5. Further informations

Further informations can be obtained from the project Web site:

<http://cerco.cs.unibo.it>

or contacting the project coordinator:

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

4.2 Use and dissemination of foreground

A plan for use and dissemination of foreground (including socio-economic impact and target groups for the results of the research) shall be established at the end of the project. It should, where appropriate, be an update of the initial plan in Annex I for use and dissemination of foreground and be consistent with the report on societal implications on the use and dissemination of foreground (section 4.3 – H).

The plan should consist of:

- Section A

This section should describe the dissemination measures, including any scientific publications relating to foreground. **Its content will be made available in the public domain** thus demonstrating the added-value and positive impact of the project on the European Union.

- Section B

This section should specify the exploitable foreground and provide the plans for exploitation. All these data can be public or confidential; the report must clearly mark non-publishable (confidential) parts that will be treated as such by the Commission. Information under Section B that is not marked as confidential **will be made available in the public domain** thus demonstrating the added-value and positive impact of the project on the European Union.

Section A (public)

This section includes two templates

- Template A1: List of all scientific (peer reviewed) publications relating to the foreground of the project.
- Template A2: List of all dissemination activities (publications, conferences, workshops, web sites/applications, press releases, flyers, articles published in the popular press, videos, media briefings, presentations, exhibitions, thesis, interviews, films, TV clips, posters).

These tables are cumulative, which means that they should always show all publications and activities from the beginning until after the end of the project. Updates are possible at any time.

TEMPLATE A1: LIST OF SCIENTIFIC (PEER REVIEWED) PUBLICATIONS, STARTING WITH THE MOST IMPORTANT ONES										
NO.	Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers ¹ (if available)	Is/Will open access ² provided to this publication?
1	<i>Economic transformation in Hungary and Poland'</i>		<i>European Economy</i>	<i>No 43, March 1990</i>	<i>Office for Official Publications of the European Communities</i>	<i>Luxembourg</i>	<i>1990</i>	<i>pp. 151 - 167</i>		yes/no
2										
3										

¹ A permanent identifier should be a persistent link to the published version full text if open access or abstract if article is pay per view) or to the final manuscript accepted for publication (link to article in repository).

² Open Access is defined as free of charge access for anyone via Internet. Please answer "yes" if the open access to the publication is already established and also if the embargo period for open access is not yet over but you intend to establish open access afterwards.

TEMPLATE A2: LIST OF DISSEMINATION ACTIVITIES

NO.	Type of activities ³	Main leader	Title	Date/Period	Place	Type of audience ⁴	Size of audience	Countries addressed
1	Conference		European Conference on Nanotechnologies	26 February 2010				
2								
3								

³ A drop down list allows choosing the dissemination activity: publications, conferences, workshops, web, press releases, flyers, articles published in the popular press, videos, media briefings, presentations, exhibitions, thesis, interviews, films, TV clips, posters, Other.

⁴ A drop down list allows choosing the type of public: Scientific Community (higher education, Research), Industry, Civil Society, Policy makers, Medias, Other ('multiple choices' is possible).

Section B (Confidential⁵ or public: confidential information to be marked clearly)

Part B1

The applications for patents, trademarks, registered designs, etc. shall be listed according to the template B1 provided hereafter.

The list should, specify at least one unique identifier e.g. European Patent application reference. For patent applications, only if applicable, contributions to standards should be specified. This table is cumulative, which means that it should always show all applications from the beginning until after the end of the project.

TEMPLATE B1: LIST OF APPLICATIONS FOR PATENTS, TRADEMARKS, REGISTERED DESIGNS, ETC.					
Type of IP Rights ⁶ :	Confidential Click on YES/NO	Foreseen embargo date dd/mm/yyyy	Application reference(s) (e.g. EP123456)	Subject or title of application	Applicant (s) (as on the application)

⁵ Note to be confused with the "EU CONFIDENTIAL" classification for some security research projects.

⁶ A drop down list allows choosing the type of IP rights: Patents, Trademarks, Registered designs, Utility models, Others.

Part B2

Please complete the table hereafter:

Type of Exploitable Foreground ⁷	Description of exploitable foreground	Confidential Click on YES/NO	Foreseen embargo date dd/mm/yyyy	Exploitable product(s) or measure(s)	Sector(s) of application ⁸	Timetable, commercial or any other use	Patents or other IPR exploitation (licences)	Owner & Other Beneficiary(s) involved
	<i>Ex: New superconductive Nb-Ti alloy</i>			<i>MRI equipment</i>	<i>1. Medical 2. Industrial inspection</i>	<i>2008 2010</i>	<i>A materials patent is planned for 2006</i>	<i>Beneficiary X (owner) Beneficiary Y, Beneficiary Z, Poss. licensing to equipment manuf. ABC</i>

In addition to the table, please provide a text to explain the exploitable foreground, in particular:

- Its purpose
- How the foreground might be exploited, when and by whom
- IPR exploitable measures taken or intended
- Further research necessary, if any
- Potential/expected impact (quantify where possible)

⁷ A drop down list allows choosing the type of foreground: General advancement of knowledge, Commercial exploitation of R&D results, Exploitation of R&D results via standards, exploitation of results through EU policies, exploitation of results through (social) innovation.

⁸ A drop down list allows choosing the type sector (NACE nomenclature) : http://ec.europa.eu/competition/mergers/cases/index/nace_all.html

4.3 Report on societal implications

Replies to the following questions will assist the Commission to obtain statistics and indicators on societal and socio-economic issues addressed by projects. The questions are arranged in a number of key themes. As well as producing certain statistics, the replies will also help identify those projects that have shown a real engagement with wider societal issues, and thereby identify interesting approaches to these issues and best practices. The replies for individual projects will not be made public.

A General Information <i>(completed automatically when Grant Agreement number is entered.</i>	
Grant Agreement Number:	243381
Title of Project:	Certified Complexity
Name and Title of Coordinator:	Prof. Claudio Sacerdoti Coen
B Ethics	
1. Did your project undergo an Ethics Review (and/or Screening)? <ul style="list-style-type: none"> If Yes: have you described the progress of compliance with the relevant Ethics Review/Screening Requirements in the frame of the periodic/final project reports? <p>Special Reminder: the progress of compliance with the Ethics Review/Screening Requirements should be described in the Period/Final Project Reports under the Section 3.2.2 'Work Progress and Achievements'</p>	No
2. Please indicate whether your project involved any of the following issues (tick box) :	YES
RESEARCH ON HUMANS	
• Did the project involve children?	
• Did the project involve patients?	
• Did the project involve persons not able to give consent?	
• Did the project involve adult healthy volunteers?	
• Did the project involve Human genetic material?	
• Did the project involve Human biological samples?	
• Did the project involve Human data collection?	
RESEARCH ON HUMAN EMBRYO/FOETUS	
• Did the project involve Human Embryos?	
• Did the project involve Human Foetal Tissue / Cells?	
• Did the project involve Human Embryonic Stem Cells (hESCs)?	

• Did the project on human Embryonic Stem Cells involve cells in culture?	
• Did the project on human Embryonic Stem Cells involve the derivation of cells from Embryos?	
PRIVACY	
• Did the project involve processing of genetic information or personal data (eg. health, sexual lifestyle, ethnicity, political opinion, religious or philosophical conviction)?	
• Did the project involve tracking the location or observation of people?	
RESEARCH ON ANIMALS	
• Did the project involve research on animals?	
• Were those animals transgenic small laboratory animals?	
• Were those animals transgenic farm animals?	
• Were those animals cloned farm animals?	
• Were those animals non-human primates?	
RESEARCH INVOLVING DEVELOPING COUNTRIES	
• Did the project involve the use of local resources (genetic, animal, plant etc)?	
• Was the project of benefit to local community (capacity building, access to healthcare, education etc)?	
DUAL USE	
• Research having direct military use	No
• Research having the potential for terrorist abuse	

C Workforce Statistics

3. Workforce statistics for the project: Please indicate in the table below the number of people who worked on the project (on a headcount basis).

Type of Position	Number of Women	Number of Men
Scientific Coordinator	0	1
Work package leaders	0	5
Experienced researchers (i.e. PhD holders)	0	
PhD Students	0	
Other	2	0

4. How many additional researchers (in companies and universities) were recruited specifically for this project?

Of which, indicate the number of men:

D Gender Aspects

5. Did you carry out specific Gender Equality Actions under the project?	+	Yes No
---	---	-----------

6. Which of the following actions did you carry out and how effective were they?		
	Not at all effective	Very effective
Design and implement an equal opportunity policy		<input type="radio"/>
Set targets to achieve a gender balance in the workforce		<input type="radio"/>
Organise conferences and workshops on gender		<input type="radio"/>
Actions to improve work-life balance		<input type="radio"/>
Other:		

7. Was there a gender dimension associated with the research content – i.e. wherever people were the focus of the research as, for example, consumers, users, patients or in trials, was the issue of gender considered and addressed?	
Yes- please specify	
No	

E Synergies with Science Education

8. Did your project involve working with students and/or school pupils (e.g. open days, participation in science festivals and events, prizes/competitions or joint projects)?	
Yes- please specify	
No	

9. Did the project generate any science education material (e.g. kits, websites, explanatory booklets, DVDs)?	
Yes- please specify	
No	

F Interdisciplinarity

10. Which disciplines (see list below) are involved in your project?		
Main discipline ⁹ :		
Associated disciplineError: Reference source not found:	<input type="radio"/>	Associated disciplineError: Reference source not found:

G Engaging with Civil society and policy makers

11a Did your project engage with societal actors beyond the research community? (if 'No', go to Question 14)	<input type="radio"/>	Yes
	<input type="radio"/>	No

11b If yes, did you engage with citizens (citizens' panels / juries) or organised civil society (NGOs, patients' groups etc.)?	
No	
Yes- in determining what research should be performed	
Yes - in implementing the research	
Yes, in communicating /disseminating / using the results of the project	

⁹ Insert number from list below (Frascati Manual).

11c In doing so, did your project involve actors whose role is mainly to organise the dialogue with citizens and organised civil society (e.g. professional mediator; communication company, science museums)?	<input type="radio"/> <input type="radio"/>	Yes No
12. Did you engage with government / public bodies or policy makers (including international organisations)		
No Yes- in framing the research agenda Yes - in implementing the research agenda <hr/> Yes, in communicating /disseminating / using the results of the project		
13a Will the project generate outputs (expertise or scientific advice) which could be used by policy makers? Yes – as a primary objective (please indicate areas below- multiple answers possible) Yes – as a secondary objective (please indicate areas below - multiple answer possible) No		
13b If Yes, in which fields?		
Agriculture Audiovisual and Media Budget Competition Consumers Culture Customs Development Economic and Monetary Affairs Education, Training, Youth Employment and Social Affairs	Energy Enlargement Enterprise Environment External Relations External Trade Fisheries and Maritime Affairs Food Safety Foreign and Security Policy Fraud Humanitarian aid	Human rights Information Society Institutional affairs Internal Market Justice, freedom and security Public Health Regional Policy Research and Innovation Space Taxation Transport

13c If Yes, at which level?		
Local / regional levels		
National level		
European level		
International level		
H Use and dissemination		
14. How many Articles were published/accepted for publication in peer-reviewed journals?		
To how many of these is open access¹⁰ provided?		
How many of these are published in open access journals?		
How many of these are published in open repositories?		
To how many of these is open access not provided?		
Please check all applicable reasons for not providing open access:		
<input type="checkbox"/> publisher's licensing agreement would not permit publishing in a repository <input type="checkbox"/> no suitable repository available <input type="checkbox"/> no suitable open access journal available <input type="checkbox"/> no funds available to publish in an open access journal <input type="checkbox"/> lack of time and resources <input type="checkbox"/> lack of information on open access <input type="checkbox"/> other ¹¹ :		
15. How many new patent applications ('priority filings') have been made? <i>("Technologically unique": multiple applications for the same invention in different jurisdictions should be counted as just one application of grant).</i>		
16. Indicate how many of the following Intellectual Property Rights were applied for (give number in each box).	Trademark	
	Registered design	
	Other	
17. How many spin-off companies were created / are planned as a direct result of the project?		
<i>Indicate the approximate number of additional jobs in these companies:</i>		
18. Please indicate whether your project has a potential impact on employment, in comparison with the situation before your project:		
Increase in employment, or Safeguard employment, or Decrease in employment, Difficult to estimate / not possible to quantify	<input type="checkbox"/> In small & medium-sized enterprises <input type="checkbox"/> In large companies <input type="checkbox"/> None of the above / not relevant to the project	
19. For your project partnership please estimate the employment effect resulting directly from your participation in Full Time Equivalent (FTE = one person working fulltime for a year) jobs:	<i>Indicate figure:</i>	

¹⁰ Open Access is defined as free of charge access for anyone via Internet.

¹¹ For instance: classification for security project.

Difficult to estimate / not possible to quantify		<input type="checkbox"/>
I Media and Communication to the general public		
20. As part of the project, were any of the beneficiaries professionals in communication or media relations?		
Yes	No	
21. As part of the project, have any beneficiaries received professional media / communication training / advice to improve communication with the general public?		
Yes	No	
22. Which of the following have been used to communicate information about your project to the general public, or have resulted from your project?		
Press Release	<input type="checkbox"/>	Coverage in specialist press
Media briefing	<input type="checkbox"/>	Coverage in general (non-specialist) press
TV coverage / report	<input type="checkbox"/>	Coverage in national press
Radio coverage / report	<input type="checkbox"/>	Coverage in international press
Brochures /posters / flyers	<input type="checkbox"/>	Website for the general public / internet
DVD /Film /Multimedia	<input type="checkbox"/>	Event targeting general public (festival, conference, exhibition, science café)
23. In which languages are the information products for the general public produced?		
Language of the coordinator	<input type="checkbox"/>	English
Other language(s)		

Question F-10: Classification of Scientific Disciplines according to the Frascati Manual 2002 (Proposed Standard Practice for Surveys on Research and Experimental Development, OECD 2002):

FIELDS OF SCIENCE AND TECHNOLOGY

1. NATURAL SCIENCES

- 1.1 Mathematics and computer sciences [mathematics and other allied fields: computer sciences and other allied subjects (software development only; hardware development should be classified in the engineering fields)]
- 1.2 Physical sciences (astronomy and space sciences, physics and other allied subjects)
- 1.3 Chemical sciences (chemistry, other allied subjects)
- 1.4 Earth and related environmental sciences (geology, geophysics, mineralogy, physical geography and other geosciences, meteorology and other atmospheric sciences including climatic research, oceanography, vulcanology, palaeoecology, other allied sciences)
- 1.5 Biological sciences (biology, botany, bacteriology, microbiology, zoology, entomology, genetics, biochemistry, biophysics, other allied sciences, excluding clinical and veterinary sciences)

2. ENGINEERING AND TECHNOLOGY

- 2.1 Civil engineering (architecture engineering, building science and engineering, construction engineering, municipal and structural engineering and other allied subjects)
- 2.2 Electrical engineering, electronics [electrical engineering, electronics, communication engineering and systems, computer engineering (hardware only) and other allied subjects]

- 2.3. Other engineering sciences (such as chemical, aeronautical and space, mechanical, metallurgical and materials engineering, and their specialised subdivisions; forest products; applied sciences such as geodesy, industrial chemistry, etc.; the science and technology of food production; specialised technologies of interdisciplinary fields, e.g. systems analysis, metallurgy, mining, textile technology and other applied subjects)

3. MEDICAL SCIENCES

- 3.1 Basic medicine (anatomy, cytology, physiology, genetics, pharmacy, pharmacology, toxicology, immunology and immunohaematology, clinical chemistry, clinical microbiology, pathology)
3.2 Clinical medicine (anaesthesiology, paediatrics, obstetrics and gynaecology, internal medicine, surgery, dentistry, neurology, psychiatry, radiology, therapeutics, otorhinolaryngology, ophthalmology)
3.3 Health sciences (public health services, social medicine, hygiene, nursing, epidemiology)

4. AGRICULTURAL SCIENCES

- 4.1 Agriculture, forestry, fisheries and allied sciences (agronomy, animal husbandry, fisheries, forestry, horticulture, other allied subjects)
4.2 Veterinary medicine

5. SOCIAL SCIENCES

- 5.1 Psychology
5.2 Economics
5.3 Educational sciences (education and training and other allied subjects)
5.4 Other social sciences [anthropology (social and cultural) and ethnology, demography, geography (human, economic and social), town and country planning, management, law, linguistics, political sciences, sociology, organisation and methods, miscellaneous social sciences and interdisciplinary, methodological and historical S1T activities relating to subjects in this group. Physical anthropology, physical geography and psychophysiology should normally be classified with the natural sciences].

6. HUMANITIES

- 6.1 History (history, prehistory and history, together with auxiliary historical disciplines such as archaeology, numismatics, palaeography, genealogy, etc.)
6.2 Languages and literature (ancient and modern)
6.3 Other humanities [philosophy (including the history of science and technology) arts, history of art, art criticism, painting, sculpture, musicology, dramatic art excluding artistic "research" of any kind, religion, theology, other fields and subjects pertaining to the humanities, methodological, historical and other S1T activities relating to the subjects in this group]

2. FINAL REPORT ON THE DISTRIBUTION OF THE EUROPEAN UNION FINANCIAL CONTRIBUTION

This report shall be submitted to the Commission within 30 days after receipt of the final payment of the European Union financial contribution.

Report on the distribution of the European Union financial contribution between beneficiaries

Name of beneficiary	Final amount of EU contribution per beneficiary in Euros
1.	
2.	
n	
Total	